

IBM TS7700 Release 4.0 Guide

Larry Coyne
Katja Denefleh
Derek Erdmann
Joe Hew
Sosuke Matsui
Aderson Pacini
Michael Scott
Chen Zhu



Storage



International Technical Support Organization

IBM TS7700 Release 4.0 Guide

January 2017

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

First Edition (January 2017)

This edition applies to Version 4, Release 0, Modification 0 of IBM TS7700 (product number 3957-AGK0).

© Copyright International Business Machines Corporation 2017. 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	xv
Trademarks	xvi
Preface	xvii
Summary of contents	xviii
Authors	xix
Now you can become a published author, too	xxi
Comments welcome	xxii
Stay connected to IBM Redbooks	xxii
Part 1. Architecture and planning	1
Chapter 1. Introducing the IBM TS7700	3
1.1 Overview	4
1.2 New capabilities	5
1.3 Concepts of storage virtualization	5
1.4 Benefits of tape virtualization	11
1.5 Managing the IBM TS7700	12
1.6 Data storage values	12
Chapter 2. Architecture, components, and functional characteristics	15
2.1 TS7700 architecture	16
2.1.1 Monolithic design of a Virtual Tape Server	16
2.1.2 Modular design of the TS7700	16
2.1.3 Previous Peer-to-Peer Virtual Tape Server design	20
2.1.4 Principles of grid design	20
2.1.5 TS7700 Models	21
2.1.6 Introduction of the TS7700T	21
2.1.7 Management of the TS7700	25
2.2 Stand-alone cluster: Components, functions, and features	29
2.2.1 Views from the Host: Library IDs	29
2.2.2 Tape Volume Cache	31
2.2.3 Virtual volumes and logical volumes	32
2.2.4 Mounting a scratch virtual volume	33
2.2.5 Mounting a specific virtual volume	34
2.2.6 Logical WORM support and characteristics	35
2.2.7 Virtual drives	35
2.2.8 Selective Device Access Control	36
2.2.9 Physical drives	37
2.2.10 Stacked volume	37
2.2.11 Selective Dual Copy function	38
2.2.12 General TVC management in a stand-alone cluster	39
2.2.13 TVC Cache management in a TS7740 stand-alone cluster	41
2.2.14 About TVC cache management in a TS7700D and TS7700T CP0 stand-alone cluster	42
2.2.15 TVC Cache management in a TS7700T CPx stand-alone cluster	42
2.2.16 Expired virtual volumes and the Delete Expired function	44
2.2.17 TVC management processes for TS7740 or TS7700T CPx	45
2.2.18 TVC handling in outage situations	46

2.2.19	Copy Consistency Point: Copy policy modes in a stand-alone cluster	46
2.2.20	TVC selection in a stand-alone cluster	46
2.2.21	TVC encryption	46
2.2.22	Physical volume pools	47
2.2.23	Logical and stacked volume management	50
2.2.24	Secure Data Erase function	52
2.2.25	Copy Export function.	53
2.2.26	Encryption of physical tapes	53
2.2.27	User Management: Roles and profiles	56
2.2.28	Security identification by using Lightweight Directory Access Protocol	57
2.2.29	Service preparation mode.	57
2.2.30	Service mode	57
2.3	Multi-cluster grid configurations: Components, functions, and features	58
2.3.1	Rules in a multi-cluster grid.	58
2.3.2	Required grid hardware	59
2.3.3	Data integrity by volume ownership	61
2.3.4	I/O TVC selection	63
2.3.5	Copy Consistency Points	63
2.3.6	Cluster family concept.	65
2.3.7	Override settings concept.	66
2.3.8	Host view of a multi-cluster grid and Library IDs.	68
2.3.9	Tape Volume Cache	69
2.3.10	Virtual volumes and logical volumes.	69
2.3.11	Mounting a scratch virtual volume.	69
2.3.12	Mounting a specific virtual volume	69
2.3.13	Logical WORM support and characteristics	70
2.3.14	Virtual drives	70
2.3.15	Allocation assistance	71
2.3.16	Selective Device Access Control	74
2.3.17	Physical drives	74
2.3.18	Stacked volume	74
2.3.19	Selective Dual Copy function	74
2.3.20	General TVC management in multi-cluster grids	74
2.3.21	Expired virtual volumes and the Delete Expired function	75
2.3.22	TVC management for TS7740 and TS7700T CPx in a multi-cluster grid	75
2.3.23	TVC management for TS7760 or TS7720 in a multi-cluster grid	76
2.3.24	TVC management processes in a multi-cluster grid	80
2.3.25	Copy Consistency Point: Copy policy modes in a multi-cluster grid	80
2.3.26	TVC (I/O) selection in a multi-cluster grid	88
2.3.27	TVC handling in an unavailability condition.	89
2.3.28	Remote (cross) cluster mounts	89
2.3.29	TVC encryption	89
2.3.30	Logical and stacked volume management	90
2.3.31	Secure Data Erase	90
2.3.32	Copy Export	90
2.3.33	Encryption of physical tapes	90
2.3.34	Autonomic Ownership Takeover Manager	90
2.3.35	Selective Write Protect for disaster recovery testing.	91
2.3.36	FlashCopy for disaster recovery testing R3.1	92
2.3.37	Service preparation mode.	94
2.3.38	Service mode	95
2.4	Grid configuration examples	95
2.4.1	Homogeneous versus hybrid grid configuration	95

2.4.2	Planning for high availability or disaster recovery in limited distances	96
2.4.3	Disaster recovery capabilities in a remote data center	97
2.4.4	Configuration examples	98
Chapter 3.	IBM TS7700 usage considerations	103
3.1	Introduction	104
3.1.1	A short look at history	104
3.1.2	Challenges of today's businesses	104
3.1.3	Challenges of technology progress	105
3.2	Gather your business requirements	106
3.2.1	Requirement types	106
3.2.2	Environment: Source of data	107
3.2.3	Backup data, active data, and archive data	108
3.2.4	IBM DB2 archive log handling	110
3.2.5	DFSMSHsm Migration Level 2	110
3.2.6	Object access method: Object processing	111
3.2.7	Batch processing: Active data	111
3.2.8	Data type and cache control	112
3.3	Features and functions for all TS7700 models	112
3.3.1	Stand alone versus grid environments	113
3.3.2	Sharing a TS7700	113
3.3.3	Tape Volume Cache selection	114
3.3.4	Copy Consistency policy	115
3.3.5	Synchronous mode copy	116
3.3.6	Override policies	116
3.3.7	Cluster family	116
3.3.8	Logical Volume Delete Expire Processing versus previous implementations	117
3.3.9	Encryption	117
3.3.10	z/OS Allocation assistance	118
3.3.11	25 GB logical volumes	118
3.4	Features and functions available only for the TS7700T	119
3.5	Operation aspects: Monitoring and alerting	119
3.5.1	Message handling	120
3.5.2	Regularly scheduled performance monitoring	120
3.5.3	Optional checks	120
3.6	Choosing a migration method	121
3.6.1	Host-based migration	121
3.6.2	TS7700 internal data migration	121
3.6.3	Tape drive technology behind a TS7700	122
Chapter 4.	Preinstallation planning and sizing	125
4.1	Hardware installation and infrastructure planning	126
4.1.1	System requirements	127
4.1.2	TS7700 specific limitations	134
4.1.3	TCP/IP configuration considerations	136
4.1.4	Factors that affect performance at a distance	143
4.1.5	Host attachments	144
4.1.6	Planning for LDAP for user authentication in your TS7700 subsystem	149
4.1.7	Cluster time coordination	149
4.2	Planning for a grid operation	150
4.2.1	Autonomic Ownership Takeover Manager (AOTM) Considerations	150
4.2.2	Defining grid copy mode control	151
4.2.3	Defining scratch mount candidates	153

4.2.4	Retain Copy mode	154
4.2.5	Defining cluster families	154
4.2.6	TS7720 and TS7760 cache thresholds and removal policies	154
4.2.7	Data management settings (TS7740/TS7700T CPx in a multi-cluster grid)	158
4.3	Planning for software implementation	160
4.3.1	Host configuration definition	161
4.3.2	Software requirements	163
4.3.3	System-managed storage tape environments	163
4.3.4	Sharing and partitioning considerations	164
4.3.5	Sharing the TS7700 by multiple hosts	165
4.3.6	Partitioning the TS7700 between multiple hosts	166
4.3.7	Logical path considerations	166
4.4	Planning for logical and physical volumes	167
4.4.1	Volume serial numbering	167
4.4.2	Logical volumes	168
4.4.3	Logical WORM	171
4.4.4	Physical volumes for TS7740, TS7720T, and TS7760T	171
4.4.5	Data compression	173
4.4.6	Secure Data Erase function	173
4.4.7	Planning for tape encryption in a TS7740, TS7720T, and TS7760T	174
4.4.8	Planning for cache disk encryption in the TS7700	176
4.5	Tape analysis and sizing the TS7700	178
4.5.1	IBM tape tools	178
4.5.2	BatchMagic	181
4.5.3	Workload considerations	181
4.5.4	Education and training	185
4.5.5	Implementation services	186
	Chapter 5. Disaster recovery	189
5.1	TS7700 disaster recovery principles	190
5.1.1	Data availability	190
5.1.2	Deferred Copy Queue	191
5.1.3	Volume ownership	191
5.2	Failover scenarios	193
5.3	Planning for disaster recovery	193
5.3.1	Disaster recovery site connectivity IODF considerations	194
5.3.2	Grid configuration	194
5.3.3	Planning guidelines	195
5.4	High availability and disaster recovery configurations	196
5.4.1	Example grid configurations	196
5.4.2	Restoring the host and library environments	204
5.5	Disaster recovery testing basics	205
5.5.1	Selective write protect for disaster recovery testing	205
5.6	A real disaster	206
5.7	Geographically Dispersed Parallel Sysplex for z/OS	208
5.7.1	Geographically Dispersed Parallel Sysplex considerations in a TS7700 grid configuration	208
5.7.2	Geographically Dispersed Parallel Sysplex functions for the TS7700	209
5.7.3	Geographically Dispersed Parallel Sysplex implementation	210
	Part 2. Implementation and migration	211
	Chapter 6. IBM TS7700 implementation	213
6.1	TS7700 implementation	214

6.1.1	Implementation tasks	214
6.2	TS4500/TS3500 tape library definitions	215
6.3	Setting up the TS7700	216
6.3.1	Definitions for TS7760T TS7740, or TS7720T	216
6.3.2	TS7700 definitions	216
6.4	Hardware configuration definition	216
6.4.1	Defining devices through HCD	218
6.4.2	Activating the I/O configuration	221
6.5	Setting values for the Missing Interrupt Handler	222
6.6	TS7700 software definitions	223
Chapter 7. Hardware configurations and upgrade considerations		227
7.1	TS7700 hardware components	228
7.1.1	Common components for the TS7700 models	230
7.1.2	TS7760 components	234
7.1.3	TS7720 components	240
7.1.4	TS7740 components	245
7.1.5	TS7700 tape library attachments, drives, and media	248
7.1.6	TS3000 Total System Storage Console	249
7.1.7	Cables	249
7.2	TS7700 component upgrades	251
7.2.1	TS7700 concurrent system component upgrades	251
7.2.2	TS7700 non-concurrent system component upgrades	253
7.2.3	TS7760 Cache upgrade options	256
7.2.4	TS7720 Cache upgrade options	259
7.2.5	TS7740 Tape Volume Cache upgrade options	259
7.2.6	Upgrading drive models in an existing TS7740 or TS7700T	260
7.2.7	Frame replacement of old hardware with new hardware	268
7.3	TS7700 upgrade to Release 4.0	268
7.3.1	Planning for the upgrade	269
7.4	Adding clusters to a grid	269
7.4.1	TS7700 grid upgrade concept	269
7.4.2	Considerations when adding a cluster to the existing configuration	271
7.4.3	Considerations for merging an existing cluster or grid into a grid	275
7.5	Removing clusters from a grid	280
7.5.1	Reasons to remove a cluster	280
7.5.2	High-level description of the process	281
Chapter 8. Migration		283
8.1	Migration to a TS7700	284
8.1.1	Grid to Grid Migration	285
8.2	Moving data in and out of the TS7700	288
8.2.1	Phased method of moving data	288
8.2.2	Quick method of moving data	289
8.2.3	Products to simplify the task	292
8.2.4	Combining methods to move data into the TS7700	293
8.2.5	Moving data out of the TS7700	293
8.3	Migration of DFSMSHsm-managed data	296
8.3.1	Volume and data set sizes	297
8.3.2	TS7700 implementation considerations	300
8.3.3	DFSMSHsm task-related considerations	302
8.4	DFSMSRmm and other tape management systems	304
8.5	IBM Spectrum Protect	306

8.6 DFSMSdss	309
8.6.1 Full volume dumps	309
8.6.2 Stand-Alone Services	310
8.7 Object access method	311
8.8 Database backups	312
8.8.1 DB2 data	312
8.8.2 CICS and IMS	314
8.8.3 Batch data	315
Part 3. Operation	317
Chapter 9. Operation	319
9.1 User interfaces	320
9.1.1 The tape library management GUI	321
9.1.2 Call Home and Electronic Customer Care	323
9.2 TS7700 Management Interface	326
9.2.1 Connecting to the Management Interface	326
9.2.2 Using the TS7700 Management Interface	329
9.2.3 The Systems icon	334
9.2.4 The Monitor icon	358
9.2.5 Performance	362
9.2.6 The Virtual icon	373
9.2.7 The Physical icon	410
9.2.8 The Constructs icon	437
9.2.9 The Access icon	446
9.2.10 The Settings icon	469
9.2.11 The Service icon	496
9.3 Common procedures	509
9.3.1 The tape library with the TS7700T cluster	509
9.3.2 TS7700T definitions	529
9.3.3 TS7700 definitions	546
9.3.4 TS7700 multi-cluster definitions	571
9.4 Basic operations	584
9.4.1 Clock and time setting	584
9.4.2 Library in Pause mode	585
9.4.3 Preparing a TS7700 for service	586
9.4.4 The Tape Library inventory	587
9.4.5 Inventory upload	589
9.5 Tape cartridge management	589
9.5.1 3592 tape cartridges and labels	589
9.5.2 Manual insertion of stacked cartridges	590
9.6 Cluster intervention scenarios	592
9.6.1 Hardware conditions	592
9.6.2 TS7700 LIC processing failure	597
9.7 TS7700 Management Interface considerations	597
Chapter 10. Host Console operations	601
10.1 System-managed tape	602
10.1.1 DFSMS operator commands	602
10.1.2 MVS system commands	605
10.1.3 Host Console Request function	608
10.1.4 Library LMPOLICY command	621
10.1.5 Useful DEVSERV QUERY commands	622
10.1.6 Scratch volume recovery for logical volumes	624

10.1.7	Ejecting logical volumes	626
10.2	Messages from the library	627
10.2.1	CBR3750I Console Message	627
10.2.2	Alert setting messages	627
10.2.3	TS7700 Host Console messages	628
10.3	EXPIRE HOLD and scratch processing considerations	630
10.4	Scratch count mismatch	631
10.5	Effects of changing categories	631
10.6	Library messages and automation	632
10.7	Return-to-scratch enhancement	632
10.8	Deleting virtual volumes	632
Chapter 11.	Performance and monitoring	635
11.1	Overview	636
11.2	TS7700 performance characteristics	637
11.3	Basic performance overview	640
11.3.1	TS7700 components and task distribution	640
11.3.2	Grid considerations and replication modes	642
11.3.3	Workload profile from your hosts	644
11.3.4	Lifecycle Management of your data	644
11.3.5	Parameters and customization of the TS7700	644
11.3.6	Terminology of throughput	645
11.3.7	Throttling in the TS7700	646
11.4	Monitoring TS7700 performance	648
11.4.1	Base information: Types of statistical records	649
11.4.2	Using the TS4500 Management GUI	651
11.4.3	Using the TS3500 Tape Library Specialist for monitoring	652
11.4.4	Using the TS7700 Management Interface to monitor IBM storage	655
11.5	Cache capacity	665
11.5.1	Interpreting Cache Usage: MI	666
11.5.2	Interpreting Cache Usage: VEHSTATS	666
11.5.3	Interpreting Cache Usage: LI REQ,distlib,CACHE	666
11.5.4	Tuning cache usage - Making your cache deeper	667
11.5.5	Tuning cache usage - Management of unwanted copies	668
11.6	Cache throughput / Cache bandwidth	669
11.6.1	Interpreting Cache throughput: Performance graph	669
11.6.2	Interpreting cache throughput: VEHSTATS HOURFLOW	669
11.6.3	Tuning Cache bandwidth: Premigration	670
11.6.4	Premigration and premigration throttling values	670
11.7	TS7700 throughput: Host I/O increments	672
11.7.1	HOST I/O in the performance graphs	673
11.7.2	HOST I/O in the VEHSTATS	674
11.7.3	Host Throughput Feature Codes	674
11.7.4	Tuning for HOST I/O	676
11.8	Grid link and replication performance	676
11.8.1	Installed grid link hardware: Mixing of different Grid link adapters	676
11.8.2	Bandwidth and quality of the provided network	676
11.8.3	Selected replication mode	677
11.8.4	Tuning possibilities for copies: COPYCOUNT Control	682
11.8.5	Tuning possibilities for copies: Deferred Copy Throttling	682
11.8.6	Grid link performance monitoring	685
11.9	Considerations for the backend TS7740 / TS7700T	686
11.9.1	Amount of Back-end drives	686

11.9.2	Monitor Backend drives in the MI	687
11.9.3	Monitor Backend drives in the VEHSTATS	687
11.9.4	Monitor Backend drives with a LI REQ command	689
11.9.5	Tune the usage of Back-end drives	689
11.9.6	Amount of Back-end cartridges	692
11.9.7	Monitor the usage of Back-end cartridges on the MI	692
11.9.8	Monitor the usage of Back-end cartridges with VEHSTATS	693
11.9.9	Tuning of the usage of Back-end cartridges with VEHSTATS	694
11.10	Throttling the TS7700	694
11.10.1	Monitoring throttling with the MI	694
11.10.2	Monitoring throttling with VEHSTATS	695
11.10.3	Tuning to avoid the throttling	695
11.11	Adjusting parameters in the TS7700	696
11.12	Monitoring after service or outage	697
11.13	Performance evaluation tool: Plotting cache throughput from VEHSTATS	697
11.14	Bulk Volume Information Retrieval	699
11.14.1	Overview of the BVIR function	700
11.14.2	Prerequisites	702
11.14.3	Request data format	702
11.14.4	Response data format	705
11.14.5	Interpreting the BVIR response data	706
11.15	Alerts and exception and message handling	713
11.15.1	Alerting of specific events	713
11.15.2	Handling Replication Exceptions	715
11.16	IBM Tape Tools	719
11.16.1	Introduction to IBM Tape Tools	719
11.16.2	Tools download and installation	721
11.16.3	IBM Tape Tools for TS7700 monitoring	723
11.17	Using Volume Mount Analyzer	724
11.18	Using VEHSTATS and VEHGRXCL for monitoring and reporting	725
11.18.1	VEHSTATS tool overview	725
11.18.2	Running the VEHSTATS jobs	726
11.18.3	VEHSTATS reports	728
11.18.4	VEHGRXCL tool overview	732
11.18.5	VEHAUDIT overview	735
11.19	IBM z/OS commands for monitoring	736
11.19.1	DISPLAY SMS commands	736
11.19.2	LIBRARY command	738
11.20	What to look for and where	739
11.21	Virtual Device Allocation in z/OS with JES2	741
11.21.1	EQUAL allocation	743
11.21.2	BYDEVICES allocation	745
11.21.3	Allocation and Copy Consistency Point setting	747
11.21.4	Allocation and device allocation assistance	749
11.21.5	Allocation and scratch allocation assistance	752
Chapter 12	Copy Export	757
12.1	Copy Export overview and considerations	758
12.1.1	General considerations for Copy Export	759
12.1.2	Copy Export grid considerations	763
12.1.3	Reclaim process for Copy Export physical volumes	765
12.1.4	Copy Export process messages	767
12.2	Implementing and running Copy Export	770

12.2.1	Setting up data management definitions.	770
12.2.2	Validating before activating the Copy Export function.	771
12.2.3	Running the Copy Export operation	773
12.2.4	Canceling a Copy Export operation	778
12.2.5	Host completion message.	778
12.3	Using Copy Export Recovery	780
12.3.1	Planning and considerations for testing Copy Export Recovery	780
12.3.2	Performing Copy Export Recovery	781
12.3.3	Restoring the host and library environments.	786
Chapter 13.	Disaster Recovery Testing	787
13.1	DR Testing Overview	788
13.2	DR Testing Methods	788
13.2.1	Method 1: DR Testing using FlashCopy	788
13.2.2	Method 2: DR Testing using Write Protect Mode on DR clusters	789
13.2.3	Method 3: DR testing without using Write Protect Mode on DR clusters	790
13.2.4	Method 4: Breaking the interconnects between the TS7700 grid	791
13.3	DR General Considerations	792
13.3.1	The z/OS test environment represents a point in time	792
13.3.2	The data that is available in the DR cluster.	792
13.3.3	Write Protect Mode	792
13.3.4	Protection of your production data	793
13.3.5	Separating production and disaster recovery hosts: Logical volumes	793
13.3.6	Creating data during the disaster recovery test from the DR host: Selective Write Protect	794
13.3.7	Creating data during the disaster recovery test from the disaster recovery host: Copy policies	795
13.3.8	Restoring the DR host from a production host	796
13.3.9	Scratch runs during the disaster recovery test from the production host	796
13.3.10	Scratch runs during the disaster recovery test from the DR host	796
13.3.11	Cleanup phase of a disaster recovery test	796
13.3.12	Considerations for DR tests without Selective Write Protect mode	797
13.3.13	Returning to scratch without using Selective Write Protect.	799
13.4	DR for FlashCopy Concepts and Command Examples	801
13.4.1	Livecopy enablement in a DR Family	803
13.4.2	Stopping FlashCopy and Write Protect Mode for a DR Family	805
13.5	DR Testing Methods Examples.	810
13.5.1	Method 1: DR Testing using FlashCopy	811
13.5.2	Method 2: Using Write Protect Mode on DR clusters	814
13.5.3	Method 3: DR Testing without Write Protect Mode	816
13.5.4	Method 4: Breaking the grid link connections	818
13.6	Expected failures during a DR test	820
Part 4.	Appendixes	821
Appendix A.	Feature codes and RPQ	823
RPQ		824
3952 F06 RPQ		824
Feature code lists.		824
3952 F05 features.		824
3952 F06 features.		825
Server features for 3957-V07, 3957-VEB and 3957-VEC.		826
Cache Controller features for 3956-CC9, 3956-CS9 and 3956-CSA		828
TS7700 Cache Drawer features 3956-CX9, 3956-XS9 and 3956-XSA		828

Appendix B. IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments	831
Software requirements	832
Software implementation in z/VM and z/VSE.	832
General support information	832
z/VM native support that uses DFSMS/VM.	833
Native z/VSE.	835
VM/ESA and z/VM guest support	836
z/VSE as a z/VM guest using a VSE Guest Server	837
Software implementation in z/OS Transaction Processing Facility	839
Usage considerations for TS7700 with z/TPF.	840
Performance considerations for TS7700 multi-cluster grids with z/TPF	841
Implementing Outboard Policy Management for non-z/OS hosts	843
Appendix C. JES3 examples and information	847
JES3 support for system-managed tape	848
Library device groups	848
Updating the JES3 INISH deck.	850
Example with two separate tape libraries.	852
LDG definitions necessary for the first example	852
Device statements that are needed for this configuration	853
SETNAME statements that are needed for this configuration.	853
HWSNAME statement that is needed for this configuration	854
Example with three Tape Libraries.	854
LDG definitions that are needed for the second configuration example	855
Device statement needed	856
SETNAME statements needed	857
High-watermark setup name statements	857
Additional examples	857
Processing changes.	858
JES3/DFSMS processing	859
Selecting UNITNAMEs	860
New or modified data sets	860
Old data sets.	860
DFSMS catalog processing.	860
DFSMS VOLREF processing	861
Fetch messages	861
JES3 allocation and mounting	861
Multi-cluster grid considerations	862
Scratch allocation assistance and device allocation assistance	863
Appendix D. DEVSERV QLIB command	867
Appendix E. Sample job control language	871
BVIR jobs to obtain historical data	872
BVIRHSTS	872
BVIRHSTU	874
BVIRHSTV	875
Extra BVIR reporting	877
Volume Map report	877
Cache Contents report	879
Copy Audit report	879
Volume Status report	879
Physical volume status	881

Physical Volume Status report	882
Physical Volume Pool Status report	883
Physical Volume and Pool Status Report Writer.	884
VEHSTATS reports	886
Export list volume sample JCL.	891
JCL for TS7700 migration scenarios	892
Using EDGUTIL to validate tape configuration database inconsistencies	892
IDCAMS example to delete a library definition in the TCDB.	892
IDCAMS example to list all entries in the TCDB.	892
IDCAMS example to change the TCDB	893
JCL to change volumes in RMM.	893
REXX EXEC to update the library name.	893
Appendix F. Library Manager volume categories.	895
Appendix G. IBM TS7700 parameter examples	905
General example setup	906
Example 1: Two-cluster grid for HA and DR.	907
Example 2: Two-cluster grid for HA and DR.	910
Example 3: Three-cluster grid for HA and DR.	913
Example 4: Four-cluster grid for HA and DR.	916
Example 5: Four-cluster grid for HA and DR by using cluster families	919
General example setup for Tape partitions	922
Basic considerations how to find the best configuration and setup.	922
Example 1: All data in cache.	923
Example 2: All data on physical tape (premigrated ASAP), and with only one tape partition 923	
Example 3: HSM ML2 will be kept in cache only, all other data will be premigrated, and tape partitions will be used	923
Example 4: Delay premigration will be used to expire data in cache	924
Appendix H. Extra IODF examples	925
General IODF principles.	926
Using switches to connect to the control unit	926
Directly connecting	927
Upgrading to 8-Gb channels.	927
Adding more devices	927
Sharing ports.	935
LIBPORT-IDs in the MVSCP.	936
Appendix I. Case study for logical partitioning of a two-cluster grid.	937
Overview of partitioning	938
Definitions and settings in z/OS	939
Definitions in HCD.	940
PARMLIB definitions	941
DFSMSrmm definitions	943
JES2 definitions	943
SMS constructs and definitions.	944
RACF definitions.	945
Automation activities	945
Definitions on the TS7700 Management Interface.	946
Physical volume pools	946
Scratch (Fast Ready) categories.	947
Defining constructs	947

Library Port Access Groups	950
Logical volume ranges or insert volumes connected to defined ranges	955
User Management on the Management Interface	955
Verification of changes	956
Related publications	959
IBM Redbooks publications	959
Other publications	959
Technical documents on the IBM Techdocs website	961
Help from IBM	961

Notices

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

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®	IBM Spectrum Protect™	Redbooks (logo)  ®
CICS®	IBM z™	S/390®
DB2®	IBM z Systems®	System i®
DS8000®	IBM z13®	System Storage®
EnergyScale™	IMS™	System z®
FICON®	MVS™	Tivoli®
FlashCopy®	OS/400®	WebSphere®
GDPS®	Parallel Sysplex®	z Systems®
Geographically Dispersed Parallel Sysplex™	POWER®	z/OS®
Global Technology Services®	POWER7®	z/VM®
IBM®	POWER8®	z/VSE®
IBM Spectrum™	RACF®	z13™
	Redbooks®	

The following terms are trademarks of other companies:

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

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

Microsoft, 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

This IBM® Redbooks® publication highlights IBM TS7700 Release 4.0. The IBM TS7700 is part of a family of IBM Enterprise tape products. This book is intended for system architects and storage administrators who want to integrate their storage systems for optimal operation.

The IBM TS7700 offers a modular, scalable, and high-performance architecture for mainframe tape virtualization for the IBM z Systems® environment. It is a fully integrated, tiered storage hierarchy of disk and tape. This storage hierarchy is managed by robust storage management microcode with extensive self-management capability. It includes the following advanced functions:

- ▶ Policy management to control physical volume pooling
- ▶ Cache management
- ▶ Redundant copies, including across a grid network
- ▶ Copy mode control

The IBM TS7700 offers enhanced statistical reporting. It also includes a standards-based Management Interface (MI) for IBM TS7700 management. IBM TS7700 R4.0 continues the next generation of IBM TS7700 for z Systems tape:

- ▶ The IBM TS7760 is an all new hardware refresh and features Encryption Capable, high-capacity cache that uses 4 TB serial-attached Small Computer System Interface (SAS) HDDs in arrays that use dynamic disk pool configuration. This setup can scale to large capacities with the highest level of data protection.
- ▶ Release 4.0 introduces the option to attach to a TS4500 tape library, and to the previous TS3500 tape library, which contains back-end physical tape drives and policies to manage up to eight of the disk repositories in a tape-attached TS7760T. This TS7760T (Tape Attached) configuration mimics the behavior of a TS7740, with additional features that go beyond what a TS7740 can provide.

The TS7760T writes data by policy to physical tape through attachment to high-capacity, high-performance IBM TS1150 and IBM TS1140 tape drives installed in an IBM TS4500 or TS3500 tape library.

The TS7760 models are based on high-performance and redundant IBM POWER8® technology. They provide improved performance for most z Systems tape workloads when compared to the previous generations of IBM TS7700.

Summary of contents

This book contains valuable information about the IBM TS7700 for anyone who is interested in this product. The following summary helps you understand the structure of this book, and to decide which of the chapters are of the most interest.

In addition to the material in this book, other IBM publications are available to help you better understand the IBM TS7700.

If you have limited knowledge of the IBM TS7700, see the documentation for TS7700:

<http://www.ibm.com/support/knowledgecenter/STFS69/welcome>

A series of technical documents and white papers that describe many aspects of the IBM TS7700 are available. Although the basics of the product are described in this book, more detailed descriptions are provided in these documents. For that reason, most of these detailed record descriptions are not in this book, although you are directed to the appropriate technical document. For these additional technical documents, go to the IBM Techdocs Technical Sales Library website and search for TS7700:

<http://www.ibm.com/support/techdocs/atmastr.nsf/Web/TechDocs>

Familiarize yourself with the contents of Chapter 1, “Introducing the IBM TS7700” on page 3, Chapter 2, “Architecture, components, and functional characteristics” on page 15, and Chapter 3, “IBM TS7700 usage considerations” on page 103. These chapters provide a functional description of all of the major features of the product, and they are a prerequisite for understanding the other chapters.

If you are planning for the IBM TS7700, see Chapter 4, “Preinstallation planning and sizing” on page 125 for hardware information. Information on planning for Software begins in Chapter 4.3, “Planning for software implementation” on page 160. Chapter 6, “IBM TS7700 implementation” on page 213 describes the implementation and installation tasks to set up an IBM TS7700.

If you already have an IBM TS7700 or even an IBM 3494 Virtual Tape Server (VTS) installed, see Chapter 7, “Hardware configurations and upgrade considerations” on page 227. Chapter 8, “Migration” on page 283 describes migrating to a TS7700 environment.

Chapter 9, “Operation” on page 319 provides information about the operational aspects of the IBM TS7700. This information includes the layout of the MI windows to help with daily operational tasks. Chapter 9 “Host console operations” provides information on commands and procedures that are initiated from the host operating system.

If you have a special interest in the performance and monitoring tasks as part of your operational responsibilities, see Chapter 11, “Performance and monitoring” on page 635. Although this chapter gives a good overview, more information is available in the technical documents on the Techdocs website.

For availability and disaster recovery specialists, and those individuals who are involved in the planning and operation that is related to availability and disaster recovery, see Chapter 12, “Copy Export” on page 757.

Information that is related to disaster recovery can be found in Chapter 5, “Disaster recovery” on page 189 and Chapter 13., “Disaster Recovery Testing” on page 787.

In addition, the following appendixes conclude this book:

- ▶ For information about feature codes and requests for price quotation (RPQ), see Appendix A, “Feature codes and RPQ” on page 823, which lists all of the features available for the IBM TS7700.
- ▶ For information about implementation with various IBM systems, such as IBM z/VM®, IBM z/VSE®, the IBM TPF Operations Server, and IBM z/Transaction Processing Facility (IBM z/TPF), see Appendix B, “IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments” on page 831. This appendix gives a short overview and scheme for the IBM TS7700 implementation.
- ▶ For information about job entry subsystem 3 (JES3), an operating system component, see Appendix C, “JES3 examples and information” on page 847. This appendix provides additional information to assist you if you are running an IBM z/OS® system with JES3.
- ▶ For information about the layout of a new command that can be helpful with the IBM TS7700 configuration in z/OS, see Appendix D, “DEVSERV QLIB command” on page 867.
- ▶ For information about job control language, see Appendix E, “Sample job control language” on page 871, which gives you examples of jobs that are needed for installation and operational tasks.
- ▶ For information about categories, see Appendix F, “Library Manager volume categories” on page 895, which gives you a full list of all category codes that are used in both the IBM TS7700 and the IBM 3494 VTS.
- ▶ For information about parameters, see Appendix G, “IBM TS7700 parameter examples” on page 905, which provides parameter examples in different grid configurations.
- ▶ For information about the input/output definition file (IODF) and the input/output configuration program (IOCP), see Appendix H, “Extra IODF examples” on page 925.
- ▶ For information about a partitioning case study, see Appendix I, “Case study for logical partitioning of a two-cluster grid” on page 937, which provides a scenario about a partitioned IBM TS7700 hardware configuration.

Authors

This book was produced by a team working at IBM Tucson, Arizona.



Larry Coyne is a Project Leader for the IBM International Technical Support Organization (ITSO) at Tucson, Arizona, in the US. He has 34 years of IBM experience, with 23 years in IBM storage software management. He holds degrees in software engineering from the University of Texas at El Paso and in project management from George Washington University. His areas of expertise include client relationship management, quality assurance, development management, and support management for IBM Tivoli® Storage Software.




Katja Deneleh works in the Advanced Technical Skill group in Germany. She is responsible for providing second-level support for high-end tape products for Europe, the Middle East, and Africa (EMEA). Katja has worked more than 15 years as an IBM System z® systems programmer, and more than 10 years as a Mainframe Architect for outsourcing clients. Her areas of expertise cover all System z hardware, IBM Parallel Sysplex®, and operations aspects of large mainframe installations. Before joining IBM in 2003, she worked for companies using IBM systems and storage in Germany.



Derek Erdmann is a DFSMS Software Technical Support Engineer specializing in the OAM product area, where he has been the Team Lead for 4 years. He graduated from Northern Illinois University in 2009 with a Master's degree in Computer Science with an emphasis in Enterprise Computing. He has spent the last 7 years with IBM working with customers and developers to enhanced the quality of the DFSMS product set.




Joe Hew works in the Tucson, Arizona product field engineering group, supporting the IBM TS7700. With many years in the information technology (IT) field, Joe has worked in system-level test on various products, such as storage controllers, tape libraries, adapters, Serial Storage Architecture (SSA), and storage area networks (SANs). Joe is a Microsoft Certified Professional and a Certified Level 2 Fibre Channel Practitioner (awarded by the Storage Networking Industry Association).




Sosuke Matsui is a software development engineer in Tokyo, Japan. He joined IBM Japan in 2009, and worked on the development and testing of asynchronous replication of IBM Scale Out Network Attached Storage for 5 years. Currently, he is responsible for the development and testing of TS7700 hierarchical storage management (HSM) component. He is a member of the Association for Computing and Machinery (ACM) and Information Processing Society of Japan (IPSJ).



Aderson 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 extensive experience servicing a broad range of IBM products. He has installed, implemented, and supported all of the IBM Tape Virtualization Servers, from the IBM VTS B16 to the IBM TS7700 Virtualization Engine. Aderson joined IBM in 1976 as a Service Representative, and his entire career has been in IBM Services.



Michael Scott is an IBM Senior Accredited IT Specialist at the ESCC Tape Solution Center in Mainz, Germany. He has 17 years of experience in tape storage systems and tape solutions as a product field engineer. In his current job role as OEM Product Application Engineer (PAE), he is responsible for all EMEA OEM storage clients in terms of technical solutions and sales opportunities.



Chen Zhu is a Senior Data Facility Storage Management Subsystem (DFSMS) Technical Support Engineer in the IBM Systems, Client Enablement, and Systems Assurance team. He has 16 years of experience in DFSMS technical support. He holds a Masters in Business Administration and a Bachelor of Sciences in Mathematics. Michael has six patents that are issued in the computer sciences field, and is a DFSMS Technical Advocate. Currently, he is the team lead for the DFSMS technical support education program.

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

Norbert Schlumberger

IBM SO Delivery, Server Systems Operations

Felipe Barajas, Erika Dawson, Lawrence M. (Larry) Fuss, Charles House, Katsuyoshi Katori, Khanh Ly, Takeshi Nohta, Sam Smith, Joe Swingler

IBM Systems

Tom Koudstaal

E-Storage B.V.

Thanks to the authors of the previous edition, which was published in May 2016:

Larry Coyne, Katja Deneleh, Joe Hew, Sosuke Matsui, Aderson Pacini, Michael Scott, Chen Zhu

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 2 - 4 weeks in length, and you can participate either in person or as a remote resident working from your home base.

Learn more about the residency program, browse the residency index, and apply online:

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:

ibm.com/redbooks

- ▶ Send your comments in an email:

redbooks@us.ibm.com

- ▶ Mail your comments:

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:

<https://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:

<http://w3.itso.ibm.com/itsoapps/redbooks.nsf/subscribe?OpenForm>

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

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



Part 1

Architecture and planning

This part introduces the IBM TS7700 R4.0 (IBM TS7700) family. The family consists of the IBM TS7760 disk-only virtualization solution, the IBM TS7760T (Tape Attached with TS4500 or TS3500 support). This part provides a detailed technical description of the architecture and components of the TS7700. In addition, the information that is needed to plan for the implementation is addressed. The information covers the TS7760, TS7760T, TS7720, TS7720T, and TS7740.

This part includes the following chapters:

- ▶ Introducing the IBM TS7700
- ▶ Architecture, components, and functional characteristics
- ▶ IBM TS7700 usage considerations
- ▶ Preinstallation planning and sizing
- ▶ Disaster recovery



Introducing the IBM TS7700

The IBM TS7700, which was introduced in 2006, is now in its fifth generation of IBM Tape Virtualization products for mainframes. It replaces the highly successful IBM TotalStorage Virtual Tape Server (VTS).

This chapter includes the following sections:

- ▶ Overview
- ▶ New capabilities
- ▶ Concepts of storage virtualization
- ▶ Benefits of tape virtualization
- ▶ Managing the IBM TS7700
- ▶ Data storage values

1.1 Overview

With cloud infrastructures on the rise and data volumes that are expanding exponentially, organizations need a cost effective way to manage both primary and backup data that is active, inactive, or even archived. Long-term retention of data is a business priority, as is continuous availability from anywhere at any time, and the storage solution must also fit within today's budget constraints. Storing infrequently accessed data on costly disk storage simply does not make sense. At the same time, physical tape libraries can require long access times, making the use cost-prohibitive in transactional storage infrastructures. That's where virtualized tape storage comes in.

This publication explains the all-new hardware that is introduced with IBM TS7700 release 4.0, (R4.0) and the concepts associated with it. TS7700 R4.0 can be installed only on the IBM TS7720, TS7740, and the all-new, hardware-refreshed TS7760 Models. A request for quote (RFQ) must be submitted for R4.0 to be installed on previous generations of TS7700 models: TS7740 Model V06 and TS7720 Model VEA. The IBM TS7720T and TS7760T (tape attach) partition mimics the behavior of the previous TS7740, but with higher performance and capacity.

The fifth-generation TS7700 consists of the following models:

- ▶ Current models:
 - TS7760T (tape-attached)
 - TS7760 VEC (disk only, upgradeable to TS7760T tape-attached)
- ▶ Previous models:
 - TS7740 V07 (tape-attached)
 - TS7740 V06 (tape-attached)
 - TS7720 VEA (disk only)
 - TS7720T (tape-attached)
 - TS7720 VEB (disk only)

The TS7700 is a modular, scalable, and high-performance architecture for mainframe tape virtualization. This is a fully integrated, tiered storage hierarchy of disk and tape. It incorporates extensive self-management capabilities consistent with IBM Information Infrastructure initiatives.

These capabilities can improve performance and capacity. Better performance and capacity help lower the total cost of ownership for tape processing, and help avoid human error. A TS7700 can improve the efficiency of mainframe tape operations by efficiently using disk storage, tape capacity, and tape speed. It can also improve efficiency by providing many tape addresses.

TS7700 provides tape virtualization for the IBM z environment. Tape virtualization can help satisfy the following requirements in a data processing environment:

- ▶ Improved reliability
- ▶ Reduction in the time that is needed for the backup and restore process
- ▶ Reduction of services downtime that is caused by physical tape drive and library outages
- ▶ Reduction in cost, time, and complexity by moving primary workloads to virtual tape
- ▶ More efficient procedures for managing daily backup and restore processing
- ▶ Infrastructure simplification through reduction of the number of physical tape libraries, drives, and media

1.2 New capabilities

TS7700 R4.0 delivers the following capabilities:

- ▶ IBM POWER8 server technology in the TS7760
- ▶ Increased cache capacity and performance in the TS7760
- ▶ Higher cache availability and rebuild times with dynamic disk partitioning in the TS7760
- ▶ Supports attachment to the IBM TS4500 and TS3500 tape libraries
- ▶ Improvements to reliability, availability, and serviceability

The new IBM TS7760T combines the best features from both IBM TS7760 models. It supports not only large cache sizes to improve hit ratios, or even a 100% cache hit ratio, but also has physical tape that is attached to store additional copies or add capacity. In addition, seldom-used or never-used data can be placed on physical tape.

The ability to define multiple cache partitions (CP0-CP7) enables separate workloads in an TS7760T.

The delay premigration is a perfect addition to control the cache content and reduce unnecessary back-end activities. Using this feature, you can ensure that data is kept long enough in cache to support fast access, while ensuring that data with a high reuse factor does not pre-migrate unnecessarily, and back-end activities can be reduced.

1.3 Concepts of storage virtualization

A virtual tape subsystem presents emulated tape drives to the host, and stores tape data on emulated tape volumes. These volumes are in a disk-based cache rather than physical tape media. The TS7700 emulates the function and operation of IBM 3490 Enhanced Capacity (3490E) tape drives. It uses a Redundant Array of Independent Disks (RAID) technology disk subsystem to store volumes that are written by the host. The disk space that is provided is called a *Tape Volume Cache (TVC)*.

The main components of the IBM TS7700 are shown in Figure 1-1.

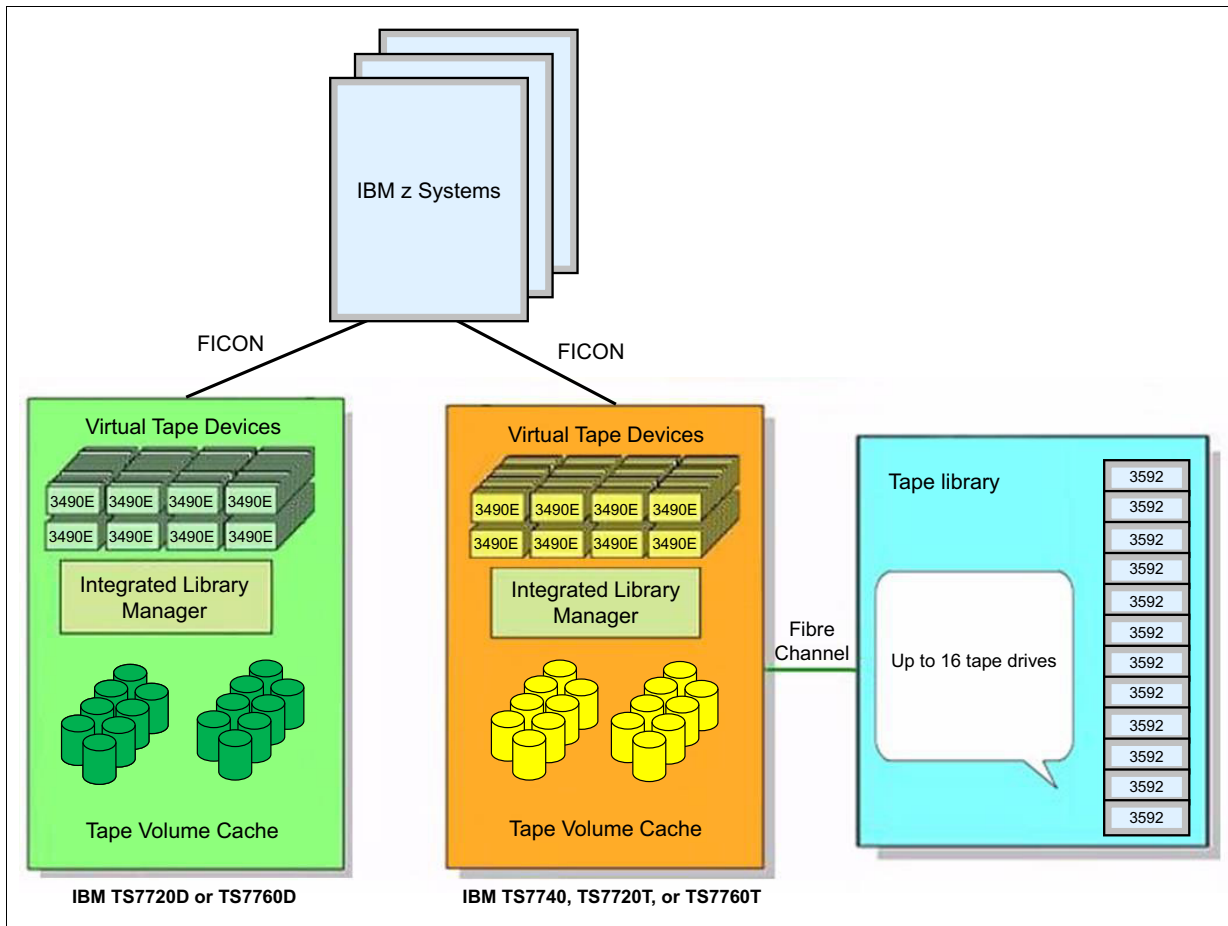


Figure 1-1 Main components of the TS7700

Emulated tape drives are also called *virtual drives*. To the host, virtual IBM 3490E tape drives look the same as physical 3490E tape drives. Emulation is not apparent to the host and applications. The host always writes to and reads from virtual tape drives. It never accesses the physical tape drives (commonly referred to as the back-end tape drives) attached to TS7740, TS7720T, and TS7760T configurations. In fact, it does not need to identify that these tape drives exist.

Even an application that supports only 3490E tape technology can use the TS7700 without any changes. Therefore, the application benefits from the high capacity and high-performance tape drives in the back end. For TS7720 VEB and TS7760 VEC (disk-only) configurations, no physical tape attachment exists. However, the virtual tape drives work the same for the host.

Because the host exclusively accesses the virtual tape drives, all data must be written to or read from emulated volumes in the disk-based TVC. These emulated tape volumes in the TVC are called *virtual volumes*.

When the host requests a volume that is still in disk cache, the volume is virtually mounted. No physical mount is required. After the virtual mount is complete, the host can access the data at disk speed. Mounting 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 only the space in cache that the compressed data on the volume requires. For this reason, tape virtualization makes efficient use of disk capacity. In IBM TS7740, IBM TS7720T and TS7760T (tape-attached) configurations, the virtual volumes are copied from disk to tape. They also need only the amount of tape capacity that is occupied by the data that is stacked end-to-end, making efficient use of both disk and tape capacity.

Another benefit of tape virtualization is the large number of drives available to applications. Each IBM TS7700 can support up to a maximum of 496 virtual tape devices. Often, applications contend for tape drives, and jobs must wait because no physical tape drive is available. Tape virtualization efficiently addresses these issues by providing many virtual tape drives. The TS7740, TS7720T, and TS7760T manage the physical tape drives and physical volumes in the tape library. It also controls the movement of data between physical and logical volumes.

In the TS7740, TS7720T, and TS7760T data that is written from the host into the TVC is scheduled for copying to tape later. The process of copying data to tape that exists only in cache is called *premigration*. When a volume is copied from cache to tape, the volume on the tape is called a *logical volume*.

A physical volume can contain many logical volumes. The process of putting several logical volumes on one physical tape is called *stacking*. A physical tape that contains logical volumes is referred to as a *stacked volume*. This concept does not apply to TS7720 VEB and TS7760 VEC because no physical tape devices are attached to it.

Without a TS7740, TS7720T, and TS7760T, many applications would be unable to fill the high capacity media of modern tape technology, and you might end up with many under-used cartridges. This wastes much space, and requires an excessive number of cartridge slots.

Tape virtualization eliminates any unused volume capacity, and fully uses physical tape capacity when present. Also, you can use tape virtualization to use the full potential of modern tape drive and tape media technology. In addition, it does so without changes to your applications or job control language (JCL).

When space is required in the TVC of a TS7740, TS7720T, and TS7760T for new data, volumes that were copied to tape are removed from the cache. By default, removal is based on a *least recently used* (LRU) algorithm. Using this algorithm ensures that no new data or recently accessed data is removed from cache. The process of deleting volumes in cache that were premigrated to tape is called *migration*. Volumes that were deleted in the cache and exist only on tape are called *migrated volumes*.

In a TS7720 and TS7760 (disk-only) configuration, no migrated volumes exist 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 and TS7760 is larger than the capacity for the TS7740.

When a TS7720 and TS7760 is a member of a multicluster hybrid grid, virtual volumes in the TS7720 and TS7760 cache can be automatically removed. This removal is done by using a Volume Removal Policy if another valid copy exists elsewhere in the grid. A *TS7700 grid* refers to two or more physically separate TS7700 clusters that are connected to one another by using a customer-supplied Internet Protocol network.

On the TS7740, TS7720T, and TS7760T, a previously migrated volume must be copied back from tape into the TVC to be accessed. It must be copied because the host has no direct access to the physical tapes. When the complete volume is copied back into the cache, the host can access the data. The process of copying data back from tape to the TVC is called *recall*.

Figure 1-2 shows IBM TS7740, IBM TS7720T, and TS7760 TVC processing.

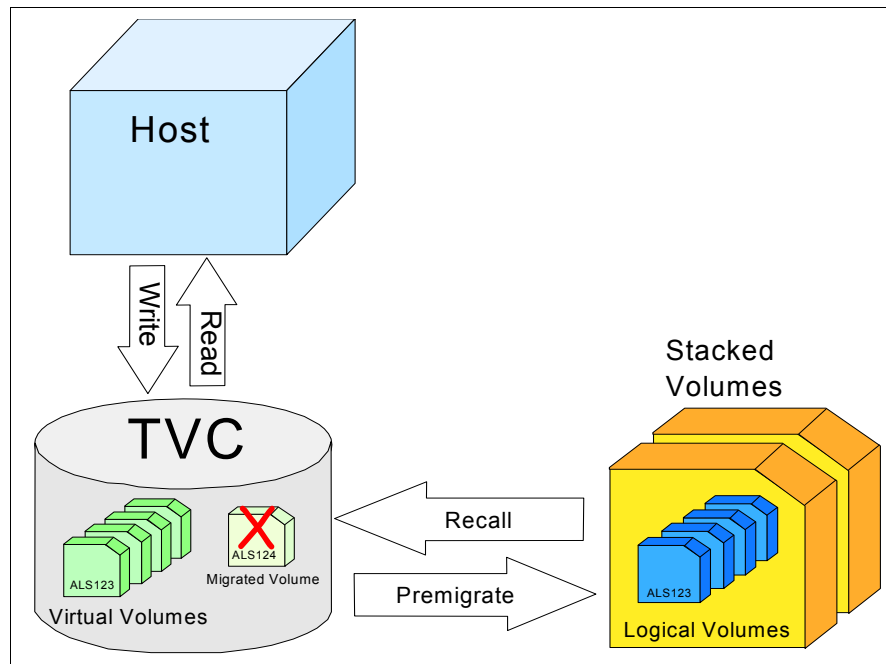


Figure 1-2 TS7740, TS7720T, and TS7760T Tape Volume Cache processing

With a TS7720 VEB and TS7760 VEC (disk-only), the virtual volumes are accessed by the host within the TVC.

Figure 1-3 shows the IBM TS7720 and IBM TS7760 TVC processing.




Figure 1-3 TS7720 and TS7760 Tape Volume Cache processing

Another benefit of tape virtualization is the data replication functions. Two, three, four, five, and six IBM TS7700 tape products can be interconnected. The connections can be through one of the following sets of links:

- ▶ Two or Four 1-gigabit (Gb) Ethernet links (copper or shortwave (SW) fiber)
- ▶ Two or Four 10-gigabit per second (Gbps) Ethernet links (longwave (LW) fiber)

These sets of links form a *multi-cluster grid configuration*. Adapter types cannot be mixed in a cluster. They *can* vary within a grid, depending on your network infrastructure. Logical volume attributes and data are replicated across the clusters in a grid. Any data that is replicated between the clusters is accessible through any other cluster in the grid configuration. Through remote volume access, you can reach any virtual volume through any virtual device. You can reach volumes even if a replication has not been made.

Setting policies on the TS7700 defines where and when you have multiple copies of your data. You can also specify for certain kinds of data, such as test data, that you do not need a secondary or tertiary copy.

You can group clusters within a grid into *families*. Grouping enables the TS7700 to make improved decisions for tasks, such as replication or TVC selection.

Depending on the configuration, multiple TS7700 tape products that form a grid provide the following types of solutions:

- ▶ High availability (HA)
- ▶ Disaster recovery (DR)
- ▶ HA and DR
- ▶ Metro and global business continuance

Before R3.2, a multi-cluster grid configuration presented itself to the attached hosts as one large library with the following maximums:

- ▶ 512 virtual devices for a two-cluster grid
- ▶ 768 virtual tape devices for a three-cluster grid
- ▶ 1024 virtual tape devices for a four-cluster grid
- ▶ 1536 virtual devices for a six-cluster grid

These numbers can now be exceeded in steps by 16 virtual drives, up to 496 virtual devices per cluster.

The copying of the volumes in a grid configuration is handled by the clusters, and it is not apparent to the host. By intermixing TS7720, TS7740, and TS7760 Models you can build a hybrid two, three, four, five, or six-cluster grid.

Figure 1-4 shows multiple IBM TS7700 tape products in an example of possible host and grid connections.

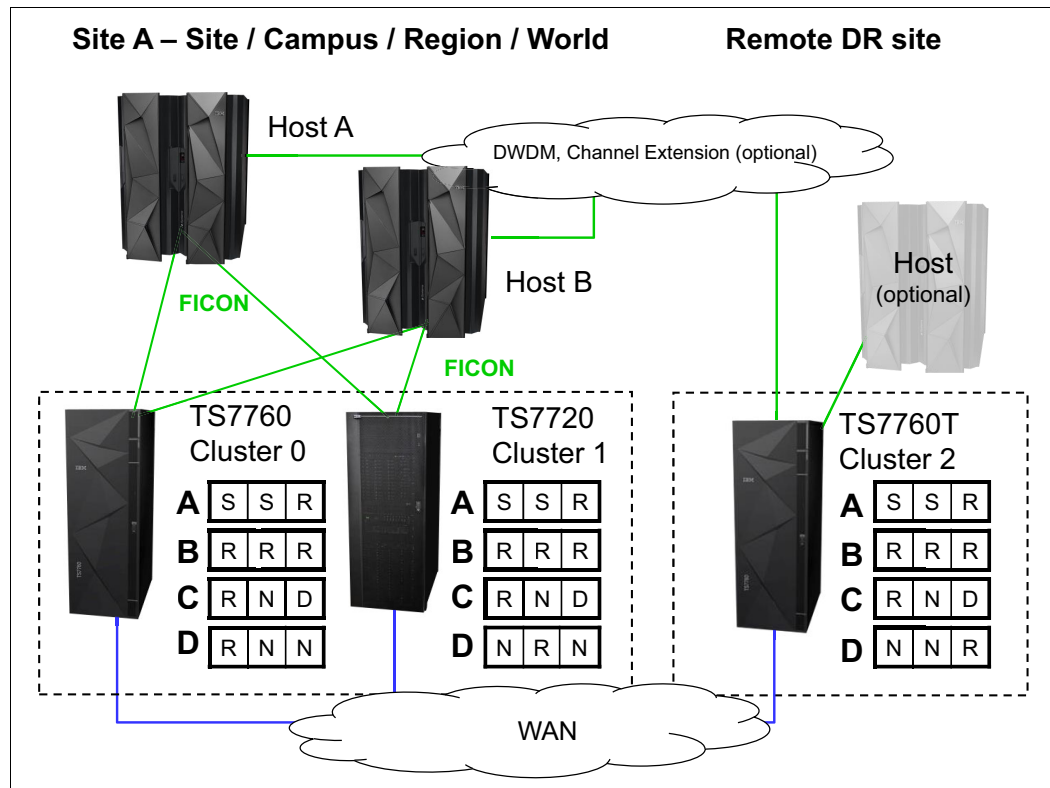


Figure 1-4 Multiple TS7700 tape products that depict possible host and grid connections

For TS7740, TS7720T, and TS7760T grid configuration, each TS7740, TS7720T, and TS7760T manages its own set of physical volumes. Each maintains the relationship between logical volumes and the physical volumes on which they are located.

The clusters in a TS7700 grid can be, but do not need to be, geographically dispersed. In a multiple cluster grid configuration, two TS7700 clusters are often located within 100 kilometers (km) or 62 miles of each other, whereas the remaining clusters can be located more than 1000 km (621.37 miles) away. This configuration provides both a highly available and redundant regional solution. It also provides a remote DR solution outside of the region. A multi-cluster grid supports the concurrent growth and reduction of cluster counts.

1.4 Benefits of tape virtualization

The current global marketplace is increasingly information-oriented, which has far-reaching implications for businesses. The ability to rapidly and securely access information can create a competitive advantage. The following information-related business trends are causing an explosion of information and complexity in data centers:

- ▶ 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 offers an extraordinary range of systems, storage, software, and services that are based on decades of innovation. This range is designed to help you get the best solutions for your business requirements. It also manages challenges, such as exploding data growth, new applications, dynamic workloads, and new regulations. IBM Information Infrastructure intelligently stores, retrieves, protects, and distributes information to help you get a competitive advantage. Converting data centers to service-oriented architectures (SOAs) helps you identify and support multiple service levels, including information services.

Certain information services must be high speed to support websites and databases. In some cases, information services must be multiplexed to multiple locations, or require extra encryption and overwrite protection. IBM Information Infrastructure helps you apply the correct services and service levels so that vital information can be delivered. IBM Information Infrastructure solutions are designed to help you manage this information explosion. They also address challenges of information compliance, availability, retention, and security.

This approach helps your company move toward improved productivity and reduced risk without driving up costs. The IBM TS7700 is part of the IBM Information Infrastructure. This strategy delivers information availability, supporting continuous and reliable access to data. It also delivers information retention, supporting responses to legal, regulatory, or investigatory inquiries for information.

The TS7700 can be the answer to the following challenges:

- ▶ Enterprise storage platform to support business in the cloud era
- ▶ Growing storage requirements
- ▶ Shrinking backup windows
- ▶ The need for continuous access to data

You can expect the following types of benefits from tape virtualization:

- ▶ Brings efficiency to the tape operation environment
- ▶ Reduces the batch window
- ▶ Provides HA and DR configurations
- ▶ Provides fast access to data through caching on disk
- ▶ Provides optional use of current tape drive, tape media, and tape automation technology
- ▶ Provides optional use of filling high capacity media to 100%
- ▶ Provides many tape drives for concurrent use
- ▶ Provides data consolidation, protection, and sharing
- ▶ Requires no additional software
- ▶ Reduces the total cost of ownership

1.5 Managing the IBM TS7700

The TS7700 uses a Management Interface (MI) to manage key management functions:

- ▶ Grid configuration
- ▶ Logical and volume cartridge management
- ▶ Constructs management
- ▶ Monitoring and utilization overview
- ▶ Monitoring and defining partitions and premigration queues
- ▶ Ownership takeover mode
- ▶ User access and roles management
- ▶ Stand-alone volume mount support
- ▶ Pool encryption setting modification
- ▶ Library Request Command panel

The TS7700 also includes a set of commands and enhanced statistical reporting.

1.6 Data storage values

The IBM TS7700 R4.0 documentation displays data storage values that use both decimal (base-10) prefixes and binary (base-2) units of measurement. Decimal units, such as kilobytes (KB), megabytes (MB), GB, and TB, are commonly used to express certain values. However, the base of the units can be misleading.

To prevent confusion, IBM uses a convention to differentiate between binary and decimal units. At the kilobyte level, the difference between decimal and binary units of measurement is relatively small (2.4%). This difference grows as data storage values increase. When values reach terabyte levels, the difference between decimal and binary units approaches 10%.

Both decimal and binary units are available throughout the TS7700 Tape Library documentation. Table 1-1 compares the names, symbols, and values of the binary and decimal units.

Table 1-1 Names, symbols, and values of the binary and decimal units

Decimal			Binary		
Name	Symbol	Value (base-10)	Name	Symbol	Value (base-2)
kilo	K	10 ³	kibi	Ki	2 ¹⁰
mega	M	10 ⁶	mebi	Mi	2 ²⁰
giga	G	10 ⁹	gibi	Gi	2 ³⁰
tera	T	10 ¹²	tebi	Ti	2 ⁴⁰
peta	P	10 ¹⁵	pebi	Pi	2 ⁵⁰
exa	E	10 ¹⁸	exbi	Ei	2 ⁶⁰

Table 1-2 shows the increasing percentage of difference between binary and decimal units.

Table 1-2 Increasing percentage of difference between binary and decimal units

Decimal value	Binary value	Percentage difference
100 kilobytes (KB)	97.65 kibibytes (KiB)	2.35%
100 megabytes (MB)	95.36 mebibytes (MiB)	4.64%
100 gigabytes (GB)	93.13 gibibytes (GiB)	6.87%
100 terabytes (TB)	90.94 tebibytes (TiB)	9.06%
100 petabytes (PB)	88.81 pebibytes (PiB)	11.19%
100 exabytes (EB)	86.73 exbibytes (EiB)	13.27%



Architecture, components, and functional characteristics

This chapter provides a description of the architecture of the IBM TS7700. The description includes general virtualization concepts, new concepts, and functions that were introduced with TS7700 R4.0. In addition, configuration examples are addressed.

First, stand-alone and clustered configuration features are explained, followed by features that apply only to multi-cluster grid configurations.

The following topics are described:

- ▶ Terms and expressions that are used to describe the TS7700
- ▶ Architecture of the TS7700
- ▶ Underlying concepts of tape virtualization within the TS7700
- ▶ Differences of the TS7700 models
- ▶ Hardware components for TS7700 Release 4.0
- ▶ Attachment of the TS7740 and TS7700T to an IBM TS3500 or TS4500 tape library
- ▶ Tape drive support
- ▶ Multi-cluster grid examples
- ▶ Functional characteristics of the TS7700:
 - Replication policies
 - Tape Partitions and delay premigration concepts on a TS7700T
 - Cluster families
 - Logical Write Once Read Many (LWORM) support
 - Enhanced cache removal policies for grids that contain one or more TS7700D clusters
 - IBM FlashCopy® and Selective write protect for disaster recovery (DR) testing
 - Device allocation assistance (DAA)
 - Scratch allocation assistance (SAA)
 - Selective Device Access Control (SDAC)
 - On-demand support of up to 4,000,000 logical volumes in a grid
 - Support up to 496 devices per cluster
- ▶ User security and user access enhancements
- ▶ Grid network support

This chapter includes the following sections:

- ▶ TS7700 architecture
- ▶ Stand-alone cluster: Components, functions, and features
- ▶ Multi-cluster grid configurations: Components, functions, and features
- ▶ Grid configuration examples

2.1 TS7700 architecture

The architectural design of the IBM TS7700 and many of its capabilities are addressed. A short description of the original IBM Virtual Tape Server (VTS) architecture is included to help you understand the differences.

The TS7700 family now includes three different models:

- ▶ IBM TS7760D and IBM TS7760T
- ▶ IBM TS7720D and IBM TS7720T
- ▶ IBM TS7740

Though there are some differences between these models, the underlying architecture is the same. If a function or feature is unique or behaves differently for a given model, it is clearly stated. If not, you can assume that it is common across all models.

When the TS7700 is referenced, it implies all models and types, including the TS7760D, TS7760T, TS7720D, TS7720T, and TS7740. When the function is only applicable to models that are disk-only, then TS7700D is used, if they are only applicable to tape attached models, then TS770T is used. If the function is only applicable to a specific version of the TS7700, TS7760D, TS7760T, TS7720D, TS7720T, TS7740 or a subset, the product name or names are referenced.

2.1.1 Monolithic design of a Virtual Tape Server

The previous IBM 3494 VTS performed all functions within a single IBM System p server. The previous generation VTS also serves as the Redundant Array of Independent Disks (RAID) disk controller. The RAID was tightly integrated into the system. Peer-to-peer (PTP) functions were created with additional hardware components, and limited to two sites. The complex design had reached an architectural limit that made it difficult to further enhance. A fundamental architecture change was required.

IBM decided that it was time to create a next-generation solution with a focus on scalability and business continuance. Many components of the original VTS were retained, although others were redesigned. The result was the TS7700 Virtualization Engine.

2.1.2 Modular design of the TS7700

The modular design of the TS7700 separates the functions of the system into smaller components. These components have well-defined functions that are connected by open interfaces. The platform enables components to be scaled up from a small configuration to a large one. This provides the capability to grow the solution to meet your business objectives.

The TS7700 is built on a multi-node architecture. This architecture consists of nodes, clusters, and grid configurations. The elements communicate with each other through standard-based interfaces. In the current implementation, a virtualization node (vnode) and hierarchical data storage management node (hnode) are combined into a general node (gnode), running on a single System p server.

The Tape Volume Cache (TVC) module in a R4.0 is a high-performance Dynamic Disk pooling (DDP) protected disk controller or a set of controllers. Before R4.0 the TVC was a high-performance RAID 6 disk controller or set of controllers. The TVC has redundant components for high availability (HA) and attaches through Fibre Channel (FC) to the TS7700.

A TS7700 and the previous VTS design are shown in Figure 2-1.

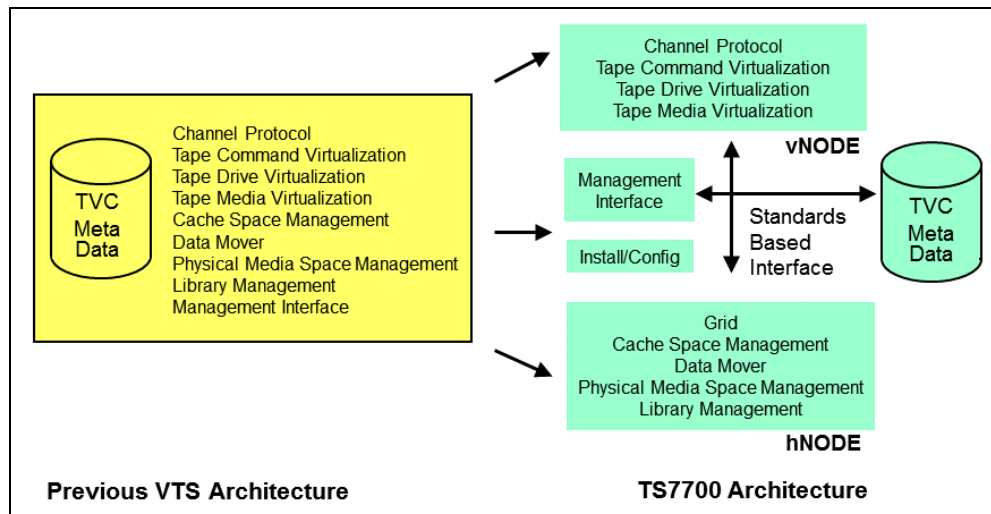


Figure 2-1 TS7700 virtualization design compared to a VTS design

Nodes

Nodes are the most basic components in the TS7700 architecture. A node has a separate name, depending on the role that is associated with it. There are three types of nodes:

- ▶ Virtualization nodes
- ▶ Hierarchical data storage management nodes
- ▶ General nodes

Virtualization node

A *virtualization node (vnode)* is a code stack that presents the virtual image of a library and drives to a host system. When the TS7700 is attached as a virtual tape library, the vnode receives the tape drive and library requests from the host. The vnode then processes them as real devices process them. It then converts the tape requests through a virtual drive and uses a file in the cache subsystem to represent the virtual tape image. After the logical volume is created or altered by the host system through a vnode, it is in disk cache.

Hierarchical data storage management node

An *hierarchical data storage management node (hnode)* is a code stack that performs management of all logical volumes that are in disk cache or physical tape. This management occurs *after* the logical volumes are 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 between clusters. An hnode uses standardized interfaces, such as Transmission Control Protocol/Internet Protocol (TCP/IP), to communicate with external components.

General node

A *general node (gnode)* can be considered a vnode and an hnode sharing a physical controller. The current implementation of the TS7700 runs on a gnode. The engine has both a vnode and hnode that are combined in an IBM POWER8 processor-based server.

Figure 2-2 shows a relationship between nodes.




Figure 2-2 Node architecture

Cluster

The TS7700 cluster combines the TS7700 server with one or more external (from the server's perspective) disk subsystems. This subsystem is the TS7700 cache controller. This architecture enables expansion of disk cache capacity.

Figure 2-3 shows the TS7700 configured as a cluster.




Figure 2-3 TS7700 cluster

A TS7700 cluster provides Fibre Channel connection (IBM FICON®) host attachment, and a default count of 256 virtual tape devices. Features are available that enable the device count to reach up to 496 devices per cluster. The IBM TS7740 and IBM TS7700T cluster also includes the assigned TS3500 or TS4500 tape library partition, fiber switches, and tape drives. The IBM TS7720 and TS7760 can include one or more optional cache expansion frames.

Figure 2-4 shows the components of a TS7740 cluster.

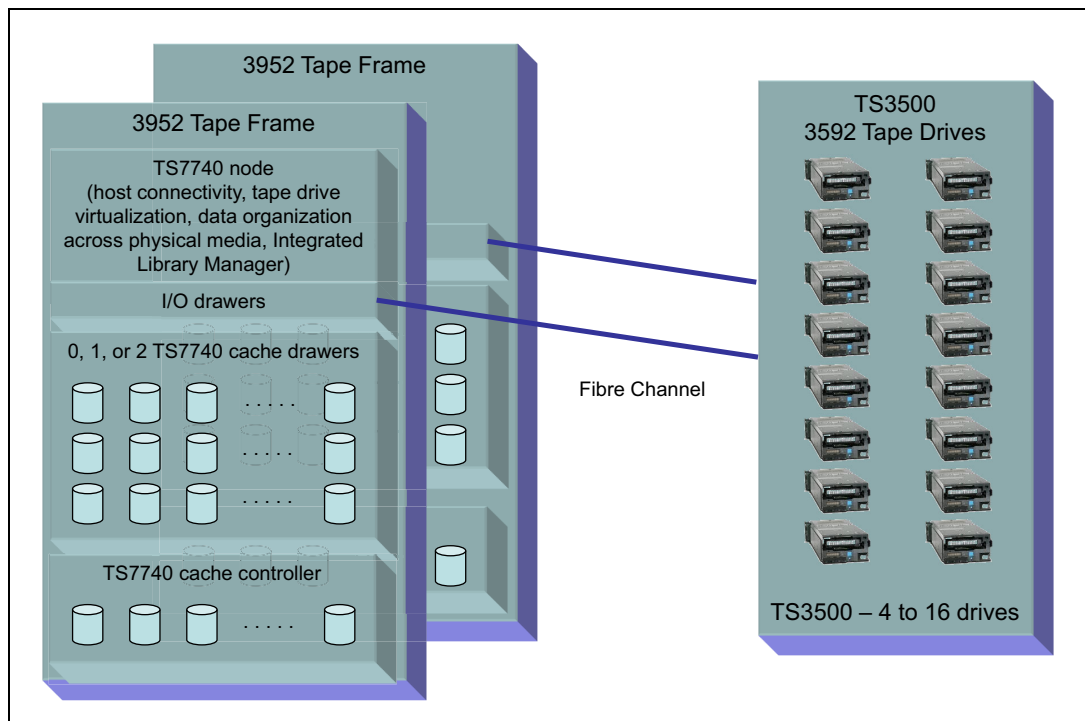


Figure 2-4 TS7740 cluster components

The TS7700 Cache Controller consists of a RAID controller and associated disk storage media. These items act as cache storage for data. The capacity of each disk drive module (DDM) depends on your configuration.

The TS7700 Cache Drawer acts as an expansion unit for the TS7700 Cache Controller. One or more controllers and their expansion drawers are collectively referred to as the TS7700 Cache, or often named the *TVC*. The amount of cache available per TS7700 Cache Drawer depends on your configuration.

The TS7760 Cache (CSA, CXA) provides a new TVC protection, called Dynamic Disk Pooling (DDP).

The TS7740 Cache provided a RAID 6 (since CC9) and RAID 5 (up to CC8) protected TVC to temporarily contain compressed virtual volumes before they are offloaded to physical tape.

The TS7720 and TS7720T CS9/CS9 use RAID 6 protection. If an existing installation is upgraded, the existing cache is protected by RAID6, where the new CSA/CXA cache uses DDP for protection.

2.1.3 Previous Peer-to-Peer Virtual Tape Server design

In the IBM 3494 PtP VTS, you needed external Virtual Tape Controller (VTC) hardware to present two VTSs as a single library to the host. The VTCs were connected to the host through IBM Enterprise Systems connection (ESCON) or FICON. Each VTC was connected to both VTSs. Only two VTSs were supported in P2P configuration.

This limited P2P design was one of the main reasons that the previous VTS needed to be redesigned. The new TS7700 replaced the P2P concepts with an industry-leading new technology referred to as a *grid*.

2.1.4 Principles of grid design

The TS7700 R4.0 *grid configuration* is a series of two, three, four, five, or six clusters. These clusters are connected by grid links to each other by a *grid network* to constitute resilient DR and HA solutions.

Fast path: Five-cluster grid and six-cluster grid configurations are available with a request for price quotation (RPQ).

A grid configuration and all virtual tape drives emulated in all configured clusters appear as one large library to the attached z Systems hosts.

Logical volumes that are created within a grid can be selectively replicated to one or more peer clusters by using a selection of different replication policies. Each replication policy or Copy Consistency Point provides different benefits, and can be intermixed. The grid architecture also enables any volume that is located within any cluster to be accessed remotely, which enables ease of access to content anywhere in the grid.

In general, any data that is initially created or replicated between clusters is accessible through any available cluster in a grid configuration. This concept ensures that data can still be accessed even if a cluster becomes unavailable. In addition, it can reduce the need to have copies in all clusters because the adjacent or remote cluster's content is equally accessible.

A grid can be of all one TS7700 model type, or any mixture of models types, including TS7760D, TS7760T, TS7720D, TS7720T, and TS7740. When a mixture of models is present within the same grid, it is referred to as a *hybrid grid*.

The term *multi-cluster grid* is used for a grid with two or more clusters. For a detailed description, see 2.3, “Multi-cluster grid configurations: Components, functions, and features” on page 58.

2.1.5 TS7700 Models

With R4.0 a new model that is called TS7760 is introduced in the TS7700 family. A disk only model is available (referenced as TS7760D). The TS7760T provides a tape attachment to a physical library, either IBM TS3500 or IBM TS4500. In an IBM TS3500, all types of tape drives of the TS1100 family are supported. In an IBM TS4500, the TS1140 and TS1150 can be used. Both models provide up to 1.3 petabytes (PB) of usable data in cache.

The TS7760 provides a disk only model as well as an option to attach to a physical tape library with TS1100 tape drives, as did its predecessor the TS7720. To support the IBM TS4500, R4.0 needs to be installed on a TS7720T. Both models deliver a maximum of 1 PB usable data in cache.

The TS7740 provides up to 28 terabytes (TB) of usable disk cache space, and supports the attachment to the IBM TS3500 tape library and TS1100 family of physical tape drives.

2.1.6 Introduction of the TS7700T

When the TS7700 was first created, the product family’s first model (TS7740) was designed with disk cache and physical tape concepts similar to the original VTS. The disk cache is primarily used for temporary storage, and most of the solution’s capacity was provided by using back-end physical tape. As the IBM z Systems® tape industry evolved, there became a need for large capacity, disk-only tape solutions, enabling more primary data workloads to move to tape with minimal penalty in performance.

The TS7720 was introduced as a response to this need. Hybrid grid configurations combined the benefits of both the TS7720 (with its large disk cache) and the TS7740 (with its economical and reliable tape store). Through Hybrid grids, large disk cache repositories and physical tape offloading were all possible. The next evolution of the TS7700 was combining the benefits of both TS7720 and TS7740 models into one solution.

Through the combination of the technologies, the TS7760, TS7720, TS7740, and Hybrid grid benefits can now be achieved with a single product. All features and functions of the TS7720, the TS7740, and hybrid grid have been maintained, although additional features and functions have been introduced to further help with the industry’s evolving use of z Systems virtual tape.

In addition to the features and functions that are provided on the TS7700D and TS7740, two key, unique features were introduced as part of the R3.2 TS7720T product release:

- ▶ Disk Cache Partition, which provides better control of how workloads use the disk cache
- ▶ Delay Premigration, or the ability to delay movement to tape

TS7700T disk cache partitioning

With the TS7700T supporting up to 1 PB of disk cache, the traditional TS7740 disk cache management style might not be adequate for many workload types. Different workloads can have different disk cache residency requirements, and treating all of the types with one *recently used* algorithm isn't always sufficient. A method to manage disk cache usage at workload granularity might be needed.

The TS7700T supports the ability to create 1 - 7 tape-managed partitions. Each partition is user-defined in 1 TB increments. Workloads that are directed to a tape-managed partition are managed independently concerning disk cache residency. After you create 1 - 7 tape-managed partitions, the disk cache capacity that remains is viewed as the resident-only partition. Partitions can be created, changed, and deleted concurrently from the Management Interface (MI).

Within this document, the tape-managed partitions are referred to as CP1 - CP7, or generically as *CPx*. The resident-only partition is referred to as *CP0*. The partitions are logical, and have no direct relationship to one or more physical disk cache drawers or types. All CPx partitions can use back-end physical tape, but the CP0 partition has no direct access to back-end physical tape. In addition, CPx partitions have no direct relationship to physical tape pools. Which partition and which pool are used for a given workload is independent.

Storage Class (SC) is used to direct workloads to a given partition. There is no automatic method to have content move between partitions. However, it can be achieved through mount/demount sequences, or through the **LIBRARY REQUEST** command.

TS7700T tape-managed partitions (CP1-CP7, CPx)

At least one tape-managed partition must exist in a TS7700T configuration. The default is CP1. However, you can configure new partitions and delete CP1 if you like, if at least one other CPx partition exists. Each CPx partition can be a unique customized size in 1 TB increments. The minimum size is 1 TB, and the maximum size is the size of the TS7700T disk cache minus 2 TB. CPx partitions support the movement of content to tape (premigration), and the removal of content from disk cache (migration).

Workloads that are directed to a given CPx partition are handled similarly to a TS7740, except that the hierarchal storage management of the CPx content is only relative to workloads that target the same partition. For example, workloads that target a particular CPx partition do not cause content in a different CPx partition to be migrated. This enables each workload to have a well-defined disk cache residency footprint.

Content that is replicated through the grid accepts the SC of the target cluster, and uses the assigned partition. If more than one TS7700T exists in a grid, the partition definitions of the two or more TS7700Ts do not need to be the same.

TS7700T CPx premigration queue

All CPx partitions share a premigration queue. The maximum amount of content that can be queued in the premigration queue is limited by a new TS7700T feature code FC5274. Each feature provides 1 TB of premigration queue. The minimum is one feature for 1 TB, and the maximum is of 10 features for 10 TB of premigration queue.

Content queued for premigration is already compressed, so the premigration queue size is based on post-compressed capacities. For example, if you have a host workload that compresses at 3:1, 6 TB of host workload results in only 2 TB of content queued for premigration.

PMPRIOR and PMTHLVL are **LIBRARY REQUEST**-tunable thresholds that are used to help manage and limit content in the premigration queue. As data is queued for premigration, and premigration activity is minimal until the PMPRIOR threshold is crossed. When crossed, the premigration activity increases based on the defined premigration drive count.

If the amount of content in the premigration queue continues to increase, the PMTHLVL threshold is crossed, and the TS7700T intentionally begins to throttle inbound host and copy activity into all CPx partitions to maintain the premigration queue size. This is when the TS7700T enters the sustained state of operation. The PMPRIOR and PMTHLVL thresholds can be no larger than the FC5274 resulting premigration queue size. For example, if three FC5274 features are installed, PMTHLVL must be set to a value of 3 TB or smaller.

After a logical volume is premigrated to tape, it is no longer counted against the premigration queue. The volume exists in both disk cache and physical tape until the migration policies determine whether and when the volume should be deleted from disk cache.

How many FC5274 features should be installed is based on many factors. The IBM tape technical specialists can help you determine how many are required based on your specific configuration.

TS7700T CPx delay premigration

With more workloads benefiting from a larger disk cache, you might determine that copying data to tape isn't necessary unless the data has aged a certain amount of time. This provides a method to retain data only in disk cache until a delay criteria is met, and only then queuing it for premigration. If the logical volume expires before this delay period, the data is never moved to tape. This reduces physical tape activity to only the workload that is viewed as archive content. It also can greatly reduce the back-end physical tape reclamation processing that can result from data that expires quickly.

Another reason that you might want to delay premigration is to run the TS7700T longer in the peak mode of operation, which can help reduce your job run times. By delaying premigration, the amount of content in the premigration queue can be reduced, which helps eliminate any throttling that can occur if the PMTHLVL threshold is crossed while running your workloads.

The delay normally is enough to get you through your daily job window. However, this is only valid for environments that have a clearly defined window of operation. The delayed premigration content is eventually queued, and any excessive queuing past the PMTHLVL threshold might result in heavy throttling. If workloads continue throughout the day, this might not be a feasible option.

The delay period is in hours, and is an attribute of the SC. Independent of which CPx partition that the data is assigned to, the delay period can be unique per workload.

TS7700T CPx Migrations

TS7700T migration operates similarly to the TS7740, except that each CPx-configured partition is treated independently concerning space management. Migration is the process of removing a logical volume in disk cache after first putting a copy on physical tape and meeting other criteria. *When* a migration of a logical volume takes place depends on a few factors:

- ▶ A copy of the logical volume must already be premigrated to primary physical tape, and if configured, secondary physical tape.
- ▶ Peer clusters in a TS7700 grid configuration have completed copies of the logical volume.
- ▶ This prerequisite can be lifted when an excessive backlog of copies exists.
- ▶ The preference group criteria has been met.

- ▶ PG0: Volumes are removed from disk cache immediately, independent of which CPx partition it is contained in.
- ▶ PG1: Volumes are removed from disk cache based on a *least recently used* algorithm. Only when space is required for a specific CPx partition are these logical volumes migrated.

TS7700T resident-only partition (CP0)

Logical volumes that are stored in CP0 are treated the same as volumes in a TS7700 disk-only cluster. The logical volumes in CP0 cannot directly move to tape. Auto-removal policies are applicable to the content assigned to the CP0 partition, including pinned, prefer keep, prefer remove, and retention policies. Content that is assigned to CPx partitions is never a candidate for auto removal. If CP0 is less than 10 TB in usable size, auto removal is disabled.

The CP0 usable size is determined by the remaining configured capacity after defining one or more CPx partitions. The CP0 partition must be at least 2 TB, and can be as large as the configured cache size minus 3 TB. As CPx partitions are created or increased in size, CP0 loses usable capacity. As CPx partitions are deleted or decreased in size, CP0 gains usable capacity. Other than workloads directly targeting CP0, the CP0 usable capacity is also used for FlashCopy processing, and for overcommit or overspill, as described in the next section.

Overcommit and overspill

When a CPx partition contains more content than its configured size, the partition is moved to the *overcommit* state. The partition remains in this state until the excess can be migrated. There are a few ways to have a partition enter the overcommitted state:

- ▶ An existing partition is decreased if configured by a user to a new size value that is smaller than the total amount of data currently resident in the CPx partition that is being resized.
- ▶ An excess of volume content is moved from one partition to another as part of a policy change.
- ▶ CPx receives more content during a workload than can be premigrated before the partition fills. This is referred to as *overspill*.
- ▶ In each of these cases, the excess space is taken from CP0's usable capacity, and the TS7700T is not viewed as degraded. It is by design that CP0 lends capacity for these expected use cases.

If CP0 has no remaining free space, further overspill is prevented. The CPx partitions are not allowed to overcommit any further. A new LI REQUEST command was introduced in R4.0 to reserve space for the CP0, which cannot be used for overspill purposes.

Moving logical volumes between CP partitions

In the following scenarios, a logical volume can be moved from one partition to another:

- ▶ A virtual volume's policy changes during a mount/demount sequence, and a new SC rule is applied. The movement occurs when the volume is closed or unmounted. While mounted, the volume remains in its originally assigned partition.
- ▶ The LI REQ PARTRFSH command was run, which enables a partition assignment change to occur without a mount/demount sequence. When using PARTRFSH, no other construct changes are refreshed other than the assigned partition. For example, pool properties that are assigned to the volume during its last mount/demount sequence are retained. If more construct changes are required, such as moving data from one pool to another, use a mount/demount sequence instead.

In either case, logical volumes can be moved from CP0 to CPx, from CPx to CP0, and from CPx to a different CPx partition. Movement rules are as described.

The following rules apply for CPx to CPx movements:

- ▶ Virtual volumes that are still in disk cache are reassigned to the new partition, and adjust the active content of both the source and target partition.
- ▶ Any delay in premigration continues to be respected, assuming that the target partition can accommodate the delay.
- ▶ The movement is part of a mount/demount sequence, and any delay relative to the last access is refreshed.
- ▶ Any other changes in constructs, such as preference group, premigration delay rules, and pool attributes, are only kept if the movement is the result of a mount/demount sequence.

The following rules apply for CPx to CP0 movements:

- ▶ Virtual volumes only in CPx cache are reassigned to CP0.
- ▶ Virtual volumes currently in CPx cache and on tape have the cache copy that is reassigned to CP0, and all copies on tape are invalidated.
- ▶ Virtual volumes currently in CPx and only on tape have the partition reassigned, but a copy is not automatically moved to CP0 disk cache. If a recall occurs later, the instance recalled into CP0 disk cache becomes the only copy, and all instances on physical tape become invalid. Until then, the content remains only on physical tape.
- ▶ Any other changes in constructs, such as removal properties, are only kept if the movement is the result of a mount/demount sequence.

The following rules apply for CP0 to CPx movements:

- ▶ The partition assignment of the logical volume in disk cache is immediately reassigned.
- ▶ If no delay premigration is active for the assigned SC, the volume is immediately queued for premigration.
- ▶ If a delay premigration is active for the assigned SC, the delay criteria based on last access or creation time is accepted. The movement itself does not alter the last access or creation time reference point.
- ▶ If the logical volume was previously migrated in a CPx partition, moved to CP0, and then moved back to a CPx partition before it was recalled into CP0 disk cache, it operates the same as though it is a CPx to CPx move.
- ▶ Any other changes in constructs, such as preference group, premigration delay rules, and pool attributes, are only kept if the movement is the result of a mount/demount sequence.

PARTRFSH command to move volumes between partitions

Similar to the **LI REQ COPYRFSH** command, a **LI REQ** command is supported by the TS7700T. It mimics a mount/demount regarding the SC's partition assignment. Other construct attributes are not updated. The command must be issued to a TS7700T distributed library, and the command supports only a single volume at a time within the current release.

2.1.7 Management of the TS7700

The management of the TS7700 is based on the following key components:

- ▶ TS7700 MI
- ▶ TS3500 or TS4500 web interface

- ▶ Advanced (outboard) policy management
- ▶ Data Facility Storage Management Subsystem (DFSMS) integration with the TS7700 to provide the storage management subsystem (SMS) constructs' names for policy management
- ▶ Host commands to control the TS7700
- ▶ Messages for automated alerting and operations
- ▶ Tools
- ▶ Call home support

TS7700 Management Interface

The TS7700 MI is a web-based graphical user interface (GUI). It is used to configure the TS7700, set up outboard policy management behavior, monitor the systems, and perform many other customer-facing management functions.

TS3500/ TS4500 web interface

The TS3500 and TS4500 web interface is used to configure and operate the tape library, particularly for the management of physical drives and media.

Advanced (outboard) policy management

Policy management enables you to better manage your logical and stacked volumes through the usage of the SMS construct names. With IBM z/OS and DFSMS, the SMS construct names that are associated with a volume (Storage Class (SC), Storage Group (SG), Management Class (MC), and Data Class (DC)) are sent to the library.

When a volume is written from load point, the eight-character SMS construct names (as assigned through your automatic class selection (ACS) routines) are passed to the library. At the library's MI, you can then define policy actions for each construct name, enabling you and the TS7700 to better manage your volumes. For the other z Systems platforms, constructs can be associated with the volumes, when the volume ranges are defined through the library's MI.

DFSMS constructs in z Systems platform and their equivalents in TS7700

In z Systems platform, the following DFSMS constructs exist:

- ▶ Storage Class
- ▶ Storage Group
- ▶ Management Class
- ▶ Data Class

Each of these constructs is used to determine specific information about the data that must be stored. All construct names are also presented to the TS7700. They need to have an equivalent definition at the library. You can define these constructs in advance on the TS7700 MI. For more information, see "Defining TS7700 constructs" on page 551. If constructs are sent to the TS7700 without having predefined constructs on the TS7700, the TS7700 creates the construct with default parameters.

Tip: Predefine your SMS constructs on the TS7700. The constructs that are created automatically might not be suitable for your requirements.

Storage class in SMS

SCs perform three functions. They decide whether data is SMS-managed. They decide the level of performance of a data set. They decide whether you can override SMS and place data on specific volumes.

Storage class in TS7700

The SC in TS7700 is used to set the cache preferences for the logical volume. This definition is cluster-based.

Storage group in SMS

SGs are the fundamental concept of DFSMS. DFSMS groups disks together into storage pools, so you allocate by storage pool. Storage pools can also consist of tape volumes. This enables SMS to direct tape allocations to a VTS or automated library. For tape SGs, one or more tape libraries can be associated with them.

Connectivity is defined at both the library level and the SG level. If an SG is connected to certain systems, any libraries that are associated with that SG must be connected to the same systems. You can direct allocations to a local or remote library, or to a specific library by assigning the appropriate SG in the SG ACS routine.

Storage Group in TS7700

The SG in the TS7700 is used to map the logical volume to a physical pool and to the primary pool number. This definition is cluster-based.

Management Class in SMS

MCs are used to determine backup and migration requirements. When assigned to data sets, MCs replace and expand attributes that otherwise are specified on job control language (JCL) data definition (DD) statements, IDCAMS DEFINE commands, and DFSMS Hierarchical Storage Manager (DFSMSHsm) commands. An MC is a list of data set migration, backup, and retention attribute values. An MC also includes object expiration criteria, object backup requirements, and class transition criteria for the management of objects.

Management Class in TS7700

From the TS7700 side, the MC is used for functions, such as Copy Policy, Selective Dual Copy Pool (depending on the physical pool, this function might be used for Copy Export), Retain Copy Mode, and Scratch Mount Candidate for Scratch Allocation assistance. This definition is cluster-based.

DATACLASS in SMS

The DATACLASS construct defines what a file looks like. The DATACLASS ACS routine is always started, even if a file is not SMS-managed. A DATACLASS is only ever assigned when a file is created and cannot be changed. A file is described by its data set organization, its record format, its record length, its space allocation, how many volumes it can span, its data compaction, its media type, and its recording information.

DATACLASS in TS7700

DATACLASS in the TS7700 is used for the definition of the virtual volume size, and whether it must be treated as an LWORM volume. This definition is shared on the grid. If you define it on one cluster, it is propagated in all other clusters in the grid.

Important: DATACLASS assignment is applied to all clusters in a grid when a volume is written from beginning of tape. Given that SG, SC, and MC can be unique per cluster, they are independently recognized at each cluster location for each mount/demount sequence.

Host commands

Several commands to control and monitor your environment are available. They are described in detail in Chapter 6, “IBM TS7700 implementation” on page 213, Chapter 8, “Migration” on page 283, Chapter 9, “Operation” on page 319, and Appendix F, “Library Manager volume categories” on page 895. These major commands are available:

D SMS,LIB	Display library information for composite and distributed libraries.
D SMS,VOLUME	Display volume information for logical volumes.
LI REQ	<p>The LIBRARY REQUEST command, also known as the Host Console Request function, is initiated from a z/OS host system to a TS7700 composite library or a specific distributed TS7700 library within a grid. Use the LIBRARY REQUEST command to request information that is related to the current operational state of the TS7700, its logical and physical volumes, and its physical resources.</p> <p>The command can also be used to run outboard operations at the library, especially setting alerting thresholds. Because all keyword combinations are passed to the TS7700 and all responses are text-based, the LIBRARY REQUEST command is a primary means of adding management features with each TS7700 release without requiring host software changes.</p> <p>The LIBRARY REQUEST command can be issued from the MI for TS7700 clusters that are running R3.2 or later. When settings are changed, the TS7700 behavior can change for all of the hosts that use the TS7700, which you need to consider when changing settings by using the LI REQ command. For more information, see the white paper found on the following website:</p> <p>http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091</p>
DS QLIB	Use the DEVICE SERVICES QUERY LIBRARY command to display library and device-related information for the composite and distributed libraries.

There is a subtle difference, but it is important to understand. The **DS QLIB** command can return different data, depending on which host it is entered. An **LI** command returns the same data without regard to the host if both hosts have full accessibility.

Automation handling and messages

Mainly, consider the content-based retrieval (CBRxxxx) messages. For more information, see the following document:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101689>

Tools

There are many helpful tools that are provided for the TS7700. For more information, see Chapter 11, “Performance and monitoring” on page 635.

Call Home support

The Call Home function automatically generates a service alert when a problem is detected within the subsystem, such as a problem in the following components:

- ▶ Inside the TS7700 components themselves
- ▶ In the associated TS3500 or TS4500 library and tape drives
- ▶ In the cache disk subsystem

Status information is transmitted to the IBM Support Center for problem evaluation. An IBM Service Support Representative (IBM SSR) can be dispatched to the installation site if maintenance is required. Call Home is part of the service strategy that is adopted in the TS7700 family. It is also used in a broad range of tape products, including VTS models and tape controllers, such as the IBM System Storage® 3592-C07.

The Call Home information for the problem is transmitted with the appropriate information to the IBM product support group. This data includes the following information:

- ▶ Overall system information, such as system serial number and Licensed Internal Code level
- ▶ Details of the error
- ▶ Error logs that can help to resolve the problem

After the Call Home is received by the assigned IBM support group, the associated information is examined and interpreted. Following analysis, an appropriate course of action is defined to resolve the problem. For instance, an IBM SSR might be sent to the site location to take the corrective actions. Alternatively, the problem might be repaired or resolved remotely by IBM support personnel through a broadband (if available) or telephone (if necessary) connection.

The TS3000 Total Storage System Console (TSSC) is the subsystem component responsible for placing the service call or Call Home when necessary. Since model 93p and release TSSC V4.7, only broadband connection is supported.

2.2 Stand-alone cluster: Components, functions, and features

In general, any cluster can be used as a stand-alone cluster. The TS7700 has several internal characteristics for High Availability (DDP or RAID 6 protection, dual power supplies, and so forth). However, a grid configuration can be configured for both additional HA and DR functions with different levels of business continuance. See Chapter 3, “IBM TS7700 usage considerations” on page 103.

Next, general information is provided about the components, functions, and features used in a TS7700 environment. The general concepts and information are also in 2.2, “Stand-alone cluster: Components, functions, and features” on page 29. Only deviations and additional information for multi-cluster grid are in 2.3, “Multi-cluster grid configurations: Components, functions, and features” on page 58.

2.2.1 Views from the Host: Library IDs

All host interaction with tape data in a TS7700 is through virtual volumes and virtual tape drives.

You must be able to identify the logical entity that represents the virtual drives and volumes, but also address the single entity of a physical cluster. Therefore, two types of libraries exist, a composite library and a distributed library. Each type is associated with a library name and a Library ID.

Composite library

The *composite library* is the logical image of the stand-alone cluster or grid that is presented to the host. All logical volumes and virtual drives are associated with the composite library. In a stand-alone TS7700, the host sees a logical tape library with up to 31 3490E tape CUs.

These CUs each have 16 IBM 3490E tape drives, and are connected through 1 - 8 FICON channels. The composite library is defined through the Interactive Storage Management Facility (ISMF). A composite library is made up of one or more distributed libraries.

Figure 2-5 illustrates the host view of a stand-alone cluster configuration.

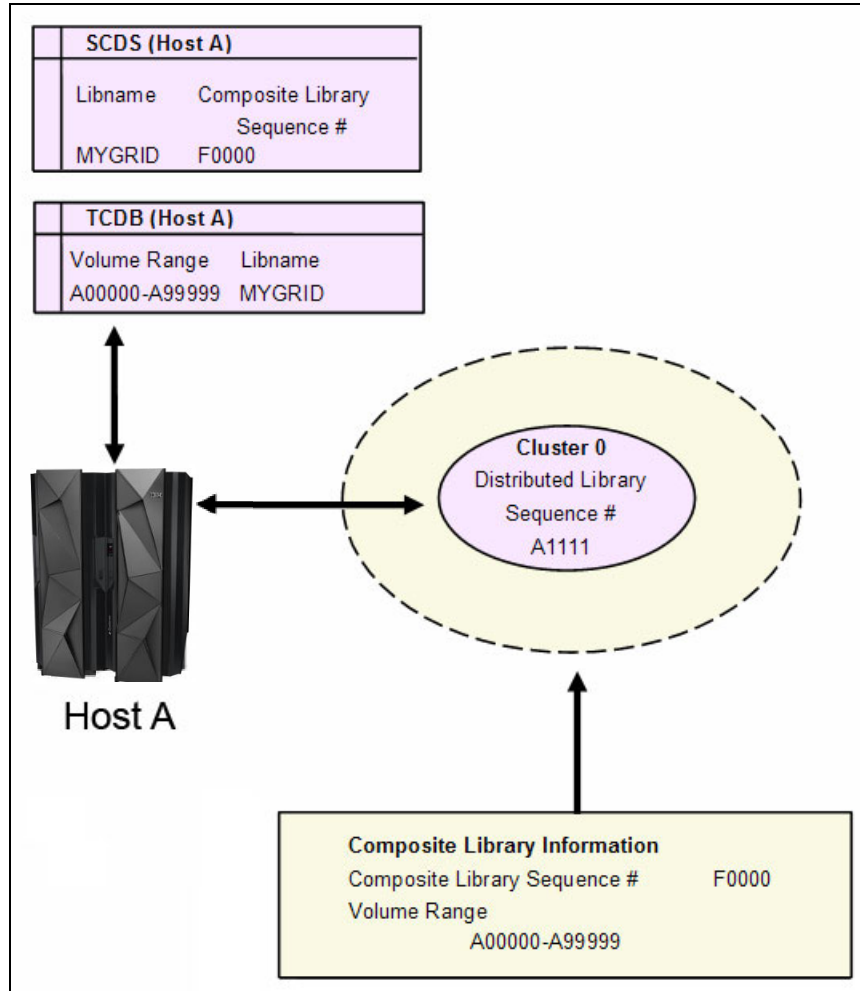


Figure 2-5 TS7700 stand-alone cluster configuration

Distributed library

Each cluster in a grid is a distributed library, which consists of a TS7700. In a TS7700T or TS7740, it is also attached to a physical TS3500 or TS4500 tape library. At the host, the distributed library is also defined to SMS. It is defined by using the existing ISMF windows, and has no tape devices defined. The virtual tape devices are defined to the composite library only.

A distributed library consists of the following cluster hardware components:

- ▶ A virtualization engine
- ▶ A TS7700 TVC
- ▶ A 3952-F05 frame or 3952-F06 frame
- ▶ Attachment to a physical library (TS7700T or TS7740)
- ▶ Several physical tape drives (TS7700T or TS7740)

Important: A composite library ID must be defined both for a multi-cluster grid and a stand-alone cluster. For a stand-alone cluster, 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.

The Library ID is used to tie the host's definition of the library to the actual hardware.

2.2.2 Tape Volume Cache

The TS7700 TVC is a disk buffer that receives all emulated tape write data, and provides all emulated tape read data.

The host operating system (OS) sees the TVC as virtual IBM 3490E Tape Drives, and the 3490 tape volumes are represented by storage space in a fault-tolerant disk subsystem. The host never writes directly to the physical tape drives attached to a TS7740 or TS7700T.

The following fault-tolerant TVC options are available. The TS7760 CSA/XSA cache are protected by the Dynamic Disk Pooling (DDP). For TS7740 configurations that use CC6, CC7, or CC8 technology, the TVC is protected with RAID 5. For all TS7720D, TS7720T, or TS7740 configurations that use CC9 technology, RAID 6 is used.

DDP on the TS7760 models provides not only a higher protection level, but also allow faster rebuild times. A DDP is built up from either one or two drawers. In a single drawer DDP, the data can be re-created when up to two disks in a DDP becomes unavailable. In a two-drawer DDP configuration, up to four disks can become unavailable and the data can still be re-created, but only two disks can be rebuilt at the same time.

Whether a DDP is built up from a single drawer or from two drawers depends on your configuration. Even numbers are always bounded to a two-drawer DDP. If an uneven number of drawers is installed, the last drawer in the frame is configured as a single drawer DDP.

The DDP does not use a global spare concept anymore, but provide free space on each of the 12 DDMs in a CSA/XSA drawer. In case of a DDM failure, the data is read from all remaining DDMs, and write to all remaining DDMs into the free space on the remaining DDMs. This procedure is called *reconstruction*.

If a second DDM fails during the reconstruction, the drawer pauses the first reconstruction, and starts a "Critical reconstruction." This process allows much faster rebuild times. After the Critical reconstruction is finished, the paused reconstruction will be started again.

As opposed to a RAID-protected system, the data will not be copied to the original DDM after the failing DDM has been replaced. Instead, newly arriving data is used to rebalance the usage of the DDMs. This behavior uses less internal resources and allows faster return to normal processing.

For older cache models, the RAID configurations provide continuous data availability to users. If up to one data disk (RAID 5) or up to two data disks (RAID 6) in a RAID group become unavailable, the user data can be re-created dynamically from the remaining disks by using parity data that is provided by the RAID implementation. The RAID groups contain global hot spare disks to take the place of a failed hard disk drive (HDD).

Using parity, the RAID controller rebuilds the data from the failed disk onto the hot spare as a background task. This process enables the TS7700 to continue working while the IBM SSR replaces the failed HDD in the TS7700 Cache Controller or Cache Drawer.

The TS7720T and the TS7760T support Cache Partitions. Virtual volumes in resident-only partition (CP0) are treated as one partition in TS7720. Virtual volumes in tape-attached partition (CP1 - CP7) are treated as one partition in TS7740. For a detailed description about Cache Partition, see 2.1.6, "Introduction of the TS7700T" on page 21.

2.2.3 Virtual volumes and logical volumes

Tape volumes that are created and accessed through the TS7700 virtual devices are referred to as *logical volumes* or *virtual volumes*. Either name can be used interchangeably. Logical volumes are objects that are in the TVC. They can optionally replicate to peer locations and also offload to back-end physical tape.

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

- ▶ Has a unique volume serial number (VOLSER) known to the host and to the TS7700.
- ▶ Is loaded and unloaded on a virtual device.
- ▶ Supports all tape write modes, including Tape Write Immediate mode.
- ▶ Contains all standard tape marks and data blocks.
- ▶ Supports an IBM, International Organization for Standardization (ISO), or American National Standards Institute (ANSI) standard label.
- ▶ Prior to R3.2, Non-Initialized tapes or scratch mounts required that the tape be written from beginning of tape (BOT) for the first write. Appends could then occur at any legal position.
- ▶ With R3.2 and later, Non-Initialized auto-labeled tapes allow the first write to occur at any position between BOT and just after the first tape mark after the volume label.
- ▶ The application is notified that the write operation is complete when the data is written to a buffer in vnode. The buffer is implicitly or explicitly synchronized with the TVC during operation. Tape Write Immediate mode suppresses write data buffering.
- ▶ 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 reaches one of the following limits:
 - 400 mebibytes (MiB) for an emulated cartridge system tape (CST)
 - 800 MiB for an emulated enhanced capacity cartridge system tape (ECCST) volume
 - 1000, 2000, 4000, 6000, or 25,000 MiB using the larger volume size options that are assigned by DC

The default logical volume sizes of 400 MiB or 800 MiB are defined at insert time. These volume sizes can be overwritten at every individual scratch mount, or any private mount where a write from BOT occurs, by using a DC construct option.

Virtual volumes can exist only in a TS7700. You can direct data to a virtual tape library by assigning a system-managed tape SG through the ACS routines. SMS passes DC, MC, SC, and SG names to the TS7700 as part of the mount operation. The TS7700 uses these constructs outboard to further manage the volume. This process uses the same policy management constructs defined through the ACS routines.

Beginning with TS7700 R2.0, a maximum of 2,000,000 virtual volumes per stand-alone cluster or multi-cluster grid was introduced. With a model V07/VEB server with R3.0 followed by model VEC server with R4.0, a maximum number of 4,000,000 virtual volumes per stand-alone cluster or multi-cluster grid are supported.

The default maximum number of supported logical volumes is still 1,000,000 per grid. Support for extra logical volumes can be added in increments of 200,000 volumes by using FC5270. Larger capacity volumes (beyond 400 MiB and 800 MiB) can be defined through DC and associated with CST (MEDIA1) or ECCST (MEDIA2) emulated media.

The VOLSERs for the logical volumes are defined through the MI when inserted. Virtual volumes go through the same cartridge entry processing as native cartridges inserted into a tape library that is attached directly to a z Systems host.

After virtual volumes are inserted through the MI, they are placed in the *insert* category and handled exactly like native cartridges. When the TS7700 is varied online to a host, or after an insert event occurs, the host operating system interacts by using the object access method (OAM) with the Library.

Depending on the definitions in the DEVSUPxx and EDGRMMxx parmlib members, the host operating system assigns newly inserted volumes to a particular scratch category. The host system requests a particular category when it needs scratch tapes, and the TS7700 knows which group of volumes to use to satisfy the scratch request.

Data compression is based on the IBMLZ1 algorithm within the FICON channel adapter in a TS7700. The actual host data that is stored on a virtual CST or ECCST volume can vary 1200 MiB - 75,000 MiB (assuming a 3:1 compression ratio).

2.2.4 Mounting a scratch virtual volume

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

Scratch volumes at the mounting cluster are chosen by using the following priority order:

1. All volumes in the source or alternative source category that are owned by the local cluster, not currently mounted, and do not have pending reconciliation changes against a peer cluster
2. All volumes in the source or alternative source category that are owned by any available cluster, not currently mounted, and do not have pending reconciliation changes against a peer cluster
3. All volumes in the source or alternative source category that are owned by any available cluster and not currently mounted
4. All volumes in the source or alternative source category that can be taken over from an unavailable cluster that has an explicit or implied takeover mode enabled

The first volumes that are chosen in the preceding steps are the volumes that have been in the source category the longest. Volume serials are also toggled between odd and even serials for each volume selection.

For all scratch mounts, the volume is temporarily initialized as though the volume was initialized by using the **EDGINERS** or **IEHINITT** program. The volume has an IBM-standard label that consists 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. In addition to choosing a volume, TVC selection processing is used to choose which TVC acts as the input/output (I/O) TVC, as described in 2.3.4, "I/O TVC selection" on page 63.

Important: In Release 3.0 or later of the TS7700, all categories that are defined as scratch inherit the Fast Ready attribute. There is no longer a need to use the MI to set the Fast Ready attribute to scratch categories. However, the MI is still needed to indicate which categories are scratch.

When the Fast Ready attribute is set or implied, no recall of content from physical tape is required in a TS7740 or TS7700T. No mechanical operation is required to mount a logical scratch volume. In addition, the volume's current consistency is ignored because a scratch mount requires a write from BOT.

The TS7700 with SAA function activated uses policy management with z/OS host software to direct scratch allocations to specific clusters within a multi-cluster grid.

2.2.5 Mounting a specific virtual volume

In a stand-alone environment, the mount is directed to the virtual drives of this cluster. In a grid environment, specific mounts are more advanced. See 2.3.12, "Mounting a specific virtual volume" on page 69.

In the stand-alone environment, the following scenarios are possible:

1. There is a valid copy in the TVC. In this case, the mount is signaled as complete and the host can access the data immediately.
2. There is no valid copy in the TVC. In this case, there are further options:
 - a. If it is a TS7760D, TS7720D, or TS7700T CP0, the mount fails.
 - b. If it is a TS7740 or TS7700T C1-CP7, and if it is on back-end physical tape, the virtual volume is recalled from a stacked volume. Mount completion is signaled to the host system only after the entire volume is available in the TVC.

The recalled 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 or the *Recalls Preferred to be Removed from Cache* override is enabled by using the **TS7700 Library Request** command.

If the mounted virtual volume was modified, the volume is again premigrated.

If modification of the virtual volume did not occur when it was mounted, the TS7740 or TS7700T 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 clusters in a grid configuration are not required if modifications were not made. If the primary or secondary pool location has changed, it is recognized now, and one or two new copies to tape are queued for premigration.

In a z/OS environment, to mount a specific volume in the TS7700, that volume must be in a private category within the library. The tape management system (TMS) prevents a scratch volume from being mounted in response to a specific mount request. Also, the TS7700 treats any specific mount that targets a volume that is assigned to a scratch category, which is also configured through the MI as scratch (Fast Ready), as a host scratch mount. *In Release 3.0 or later of TS7700, all scratch categories are Fast Ready.* If this occurs, the temporary tape header is created, and no recalls take place.

In this case, DFSMS Removable Media Manager (DFSMSrmm) or other TMS 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 are left intact, which accounts for categories that are incorrectly configured as scratch (Fast Ready) within the MI.

2.2.6 Logical WORM support and characteristics

The TS7700 supports the LWORM function through TS7700 software emulation. The host views the TS7700 as an LWORM-compliant library that contains WORM-compliant 3490E logical drives and media.

The LWORM implementation of the TS7700 emulates physical WORM tape drives and media. TS7700 provides the following functions:

- ▶ Provides an advanced function DC construct property that enables 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 worldwide identifier that is surfaced to host software during host software open and close processing, and then bound to the volume during the first write from BOT
- ▶ Generates and maintains a persistent Write-Mount Count for each LWORM volume, and keeps the value synchronized with host software
- ▶ Enables only appends to LWORM volumes by 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.

TS7700 reporting volumes (BVIR) cannot be written in LWORM format. For more information, refer to "Overview of the BVIR function" on page 700.

Clarification: Cohasset Associates, Inc. has assessed the LWORM capability of the TS7700. The conclusion is that the TS7700 meets all US Securities and Exchange Commission (SEC) requirements in Rule 17a-4(f), which expressly enables records to be retained on electronic storage media.

2.2.7 Virtual drives

From a host perspective, each TS7700 appears as 16 logical IBM 3490E tape CUs. With R3.2, up to 31 logical control units (LCUs) can be defined with the 496 drives. Each CU has 16 unique drives that are attached through FICON channels. Virtual tape drives and CUs are defined just like physical IBM 3490 systems through the hardware configuration definition (HCD). Defining a preferred path for the virtual drives gives you no benefit.

Each virtual drive has the following characteristics of physical 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 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
- ▶ Supports manual stand-alone mount processing for host initial program load (IPL) when initiated from the MI

For software transparency reasons, the functions of the 3490E integrated cartridge loader (ICL) are also included in the virtual drive's capability. All virtual drives indicate that they have an ICL. For scratch mounts, using the emulated ICL in the TS7700 to preinstall virtual cartridges is of no benefit.

With FC 5275, you can add 1 LCU with 16 drives up to the maximum of 496 logical drives per cluster.

Note: 8 Gigabit (Gb) FICON adapters (features #3438 or #3439) must be installed in a cluster before these additional devices can be defined. Existing configurations with 4 Gb FICON adapters do not support these additional devices.

RPQ 8B3643 must be requested and approved together with the first instance of FC 5275.

2.2.8 Selective Device Access Control

Due to the expanding capacity and throughput characteristics of the TS7700, there is an increased need for multiple system plexes or tenants that share a common TS7700 or TS7700 grid. Selective Device Access Control (SDAC) meets this need by enabling a secure method of *hard partitioning*. The primary intent of this function is to prevent one host logical partition (LPAR) or sysplex with an independent TMS from inadvertently modifying or removing data that is owned by another host.

This is valuable in a setup where you have a production system and a test system with different security settings on the hosts, and you want to separate the access to the grid in a more secure way. It can also be used in a multi-tenant service provider to prevent tenants from accessing each other's data, or when you have different z Systems operating systems that share the TS7700, such as z/OS, IBM z/VSE, IBM z/Transaction Processing Facility (IBM z/TPF), and IBM z/VM.

Hard partitioning is a way to give a fixed number of LCUs to a defined host group, and connect the units to a range of logical volumes that are dedicated to a particular host or hosts. SDAC is a useful function when multiple partitions have the following characteristics:

- ▶ Separate volume ranges
- ▶ Separate TMS
- ▶ Separate tape configuration database

SDAC enables you to define a subset of all of the logical devices per host (CUs in ranges of 16 devices based on the LIBPORT definitions in HCD). It enables exclusive control on host-initiated mounts, ejects, and attribute or category changes. The implementation of SDAC is described in Appendix I, "Case study for logical partitioning of a two-cluster grid" on page 937.

Implementing SDAC requires planning and orchestration with other system areas, to map the wanted access for the device ranges from individual servers or LPARs, and consolidate this information in a coherent input/output definition file (IODF) or HCD. From the TS7700 subsystem standpoint, SDAC definitions are set up using the TS7700 MI.

Important: SDAC is based on the availability of LIBPORT definitions or another equivalent way to define device ranges and administratively protect those assignments. Device partitions must be defined on 16 device boundaries to be compatible with SDAC.

2.2.9 Physical drives

The physical tape drives used by a TS7740 or TS7700T are installed in an IBM TS3500 or IBM TS4500 tape library. The physical tape drives are not addressable by any attached host system, and are controlled by the TS7740 or TS7700T. The TS7740 and TS7700T support TS1150, TS1140, TS1130, TS1120, and IBM 3592-J1A physical tape drives installed in an IBM TS3500 tape library. If an IBM TS4500 tape library is used, only TS1140 and TS1150 are supported.

Remember: Do not change the assignment of physical tape drives attached to a TS7740 or TS7700T in the IBM TS3500 IBM or IBM TS4500 Tape Library web interface. Consult your IBM SSR for configuration changes.

Before Release 3.3, all attached physical drives had to be homogeneous. With Release 3.3, support was added for the use of a mix between the TS1150 and one other tape drive generation. This is called heterogeneous tape drive support and is for migration purposes only. Although the TS1150 does not support JA and JB cartridges, it might be necessary to read the existing data with a tape drive from the previous generation, and then write the data with the TS1150 to a JC or JD cartridge. No new data can be placed on the existing JA and JB cartridges using the heterogeneous support. This is referred to as *sunset media*.

To support the heterogeneous tape drives, additional controls were introduced to handle the reclaim value for sunset media differently from the rest of the tape media. Also, two more SETTING ALERTS were introduced to allow the monitoring of the sunset drives.

For more information, see 7.1.5, “TS7700 tape library attachments, drives, and media” on page 248.

2.2.10 Stacked volume

Physical cartridges that are used by the TS7740 and TS7700T to store logical volumes are under the control of the TS7740 or TS7700T node. The physical cartridges 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.

Tip: Stacked volumes do not need to be initialized before inserting them into the TS3500 or TS4500. However, the internal VOL1 labels must match the external labels if they were previously initialized or used.

After the host closes and unloads a virtual volume, the storage management software inside the TS7740 or TS7700T schedules the virtual volume to be copied (also known as *premigration*) onto one or more physical tape cartridges. The TS7740 or TS7700T attempts to maintain a minimal amount of stacked volume to which virtual volumes are copied.

Therefore, mount activity is reduced because a minimal number of physical cartridges are mounted to service multiple virtual volume premigration requests that target the same physical volume pool. How many physical cartridges for premigration per pool can be mounted in parallel is defined within the MI as part of the pool property definitions. Virtual volumes are already compressed and are written in that compressed format to the stacked volume. This procedure maximizes the use of a cartridge’s storage capacity.

A logical volume that cannot fit in the currently filling stacked volume does 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.

Due to business reasons, it might be necessary to separate logical volumes from each other (selective dual write, multi-client environment, or encryption requirements). Therefore, you can influence the location of the data by using volume pooling. For more information, see “Using physical volume pools” on page 48.

Through the TS3500/ TS4500 web interface physical cartridge ranges should be assigned to the appropriate library partition associated with your TS7700. This enables them to become visible to the correct TS7700. The TS7700 MI must further be used to define which pool physical tapes are assigned to when initially inserted into the TS3500 or TS4500, which includes the common scratch pool. How physical tapes can move between pools for scratch management is also defined by using the MI.

2.2.11 Selective Dual Copy function

In a TS7740 or TS7700T, a logical volume and its internal data usually exist 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 one or more logical volumes that are contained on the damaged physical tape. The TS7700 provides a method to create redundant copies on independent physical tapes to help reduce the risk of such a loss.

With the Selective Dual Copy function, storage administrators can selectively create two copies of logical volumes within two pools of a TS7740 or TS7700T. The Selective Dual Copy function can be used with the Copy Export function to provide a secondary offsite physical copy for DR purposes. For more information about Copy Export, see 2.2.25, “Copy Export function” on page 53.

The second copy of the logical volume is created in a separate physical pool to ensure physical cartridge separation. Control of Dual Copy is through the MC construct (see “Management Classes window” on page 439). The second copy is created when the original volume is pre-migrated.

Important: When used for Copy Export, ensure that reclamation in the secondary physical volume pool is self-contained (the secondary volume pool reclaims onto itself) to keep secondary pool cartridges isolated from the others. Otherwise, Copy Export DR capabilities might be compromised.

The second copy that is created through the Selective Dual Copy function is only available when the primary volume cannot be recalled or is inaccessible. It cannot be accessed separately, and cannot be used if the primary volume is being used by another operation. The second copy provides a backup if the primary volume is damaged or inaccessible.

Selective Dual Copy is defined to the TS7740/TS7700T and has the following characteristics:

- ▶ The selective dual copy feature is enabled by the MC setting through the MI 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 (RUN) 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 or TS7700T.
- ▶ The secondary copy is created before the logical volume is migrated to the primary pool.

2.2.12 General TVC management in a stand-alone cluster

The TS7700 cluster manages the TVC cache. Through policy settings and **LI REQ** settings, you can influence the behavior of the way the cluster performs these actions. You can define which data to keep longer in the TVC, and which data is preferably removed from cache.

The following topics are described in the next sections:

- ▶ Rules for cache management
- ▶ Short introduction of how you control the contents of cache
- ▶ Description of how the TVC cache management mechanism works
- ▶ Description of which TVC cache management processes exist

Rules for Cache Management

Cache Management has the following rules:

- ▶ The TVC contents are managed by definitions in the SC.
- ▶ In a stand-alone TS7700D or TS7700T CP0, active data always remains in the cache.
- ▶ In a TS7740 or TS7700T CPx, if volumes are not in cache during a tape volume mount request, they are scheduled to be brought back into the disk cache from a physical tape device (recall).
- ▶ In a TS7740 configuration, if a modified virtual volume is closed and dismounted from the host, it is scheduled to be copied to a stacked volume (premigration).
- ▶ In a TS7700T CPx configuration, if a modified virtual volume is closed and dismounted from the host, it can be scheduled to be copied to a stacked volume (premigration) or kept in the delay premigration queue, depending on the SC definition. Virtual volumes in the delay premigration queue are only subject to an earlier premigration if the amount of data for this specific tape partition in the delay premigration queue is above the delay premigration threshold.
- ▶ In a TS7740 or TS7700T CPx, if the TVC runs out of space, the cache management removes or migrates previously premigrated volumes. Candidates for removal from cache are selected by using an LRU algorithm, and accept PG0/PG1 definitions.
- ▶ In addition, a TS7700T CPx partition can temporarily overflow into CP0 if CP0 space is available when the CPx partition has no migration candidates remaining.
- ▶ The TS7700 emulates a 3490E tape of a specific size that is chosen through the DC construct. However, the space that is used in the TVC is the number of bytes of data that is written to the virtual volume after compression and after a minimal amount of TS7700 metadata is introduced. When the virtual volume is written to the physical tape, it uses only the space that is occupied by the compressed data and resulting metadata.

How you control the content of the TVC (TS7700T and TS7740)

You control the content through the SC construct. Through the MI, you can define one or more SC names. If the selected cluster possesses a physical library, you can assign Preference Level 0 or 1. If the selected cluster does not possess a physical library, volumes in that cluster's cache display a Level 1 preference.

The following values are possible:

- Use IART** Volumes are removed according to the IBM TS7700s Initial Access Response Time (IART) assigned by the host during the volume's creation. The result is either Level 0 or Level 1.
- Level 0** Volumes are removed from the TVC as soon as they are copied to tape. This is called Preference Group 0 (PG0). This control is suitable for data that is unlikely to be read again.
- Level 1** Copied volumes remain in the TVC until more space is required, and then volumes are removed from disk cache in a least recently used order.

In a z/OS environment, the SC name that is assigned to a volume in the ACS routine is directly passed to the TS7700 and mapped to the predefined constructs. Figure 2-6 shows this process.

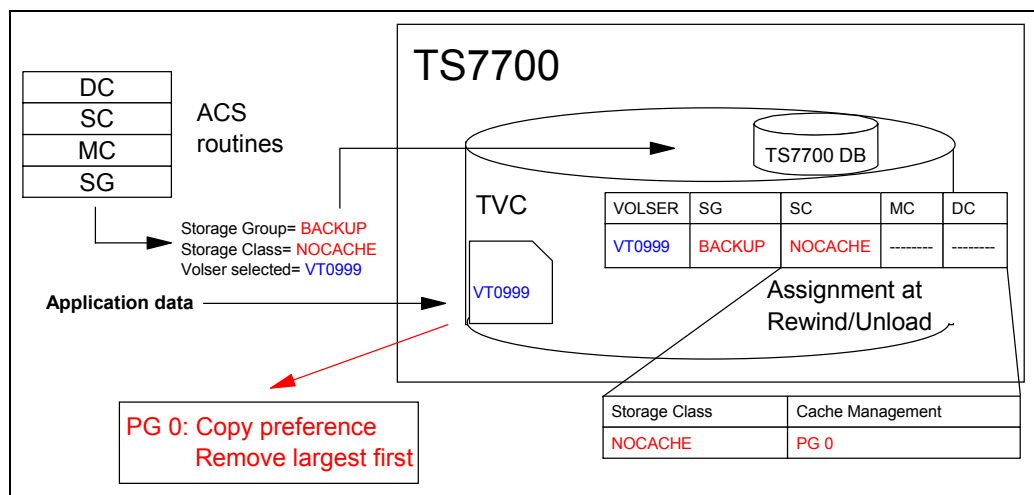


Figure 2-6 TS7740 TVC management through Storage Class

If the host passes a previously undefined SC name to the TS7700 during a scratch mount request, the TS7700 adds the name by using the definitions for the default SC.

Define SCs: Ensure that you predefine the SCs. The default SC might not support your needs.

For environments that are not z/OS (SMS) environments that use the MI, an SC can be assigned to a range of logical volumes during insert processing. The SC can also be updated to a range of volumes after they have been inserted through the MI.

To be compatible with the IART method of setting the preference level, the SC definition also enables a Use IART selection to be assigned. Even before Outboard Policy Management was made available for the previous generation VTS, you could assign a preference level to virtual volumes by using the IART attribute of the SC. The IART is an SC attribute that was originally added to specify the wanted response time (in seconds) for an object by using the OAM.

If you wanted a virtual volume to remain in cache, you assign an SC to the volume whose IART value is 99 seconds or less. Conversely, if you want to give a virtual volume preference to be out of cache, you assign an SC to the volume whose IART value was 100 seconds or more. Assuming that the Use IART selection is not specified, the TS7700 sets the preference level for the volume based on the Preference Level 0 or 1 of the SC assigned to the volume.

2.2.13 TVC Cache management in a TS7740 stand-alone cluster

As mentioned, virtual volumes with an SC with Preference Level 0 (PG0) are deleted from cache as soon as the logical volume is premigrated. If there are no more PG0 volumes that have been copied to physical volumes to remove, the TS7740 selects Preference Level 1 (Preference Group 1 or PG1) volumes. PG1 virtual volumes stay in the TVC for as long a time as possible.

When a volume is assigned Preference Level 1, the TS7740 adds it to the queue of volumes to be copied to physical tape after a 4-minute time delay, and after any volumes are assigned to Preference Level 0. The 4-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 first determines whether there are any PG0 volumes that can be removed. If not, the TS7740 selects PG1 volumes to remove based on an LRU algorithm. This process results in volumes that have been copied to physical tape, and have been in cache the longest without access, to be removed first.

Figure 2-7 shows cache usage with policy-based cache management.

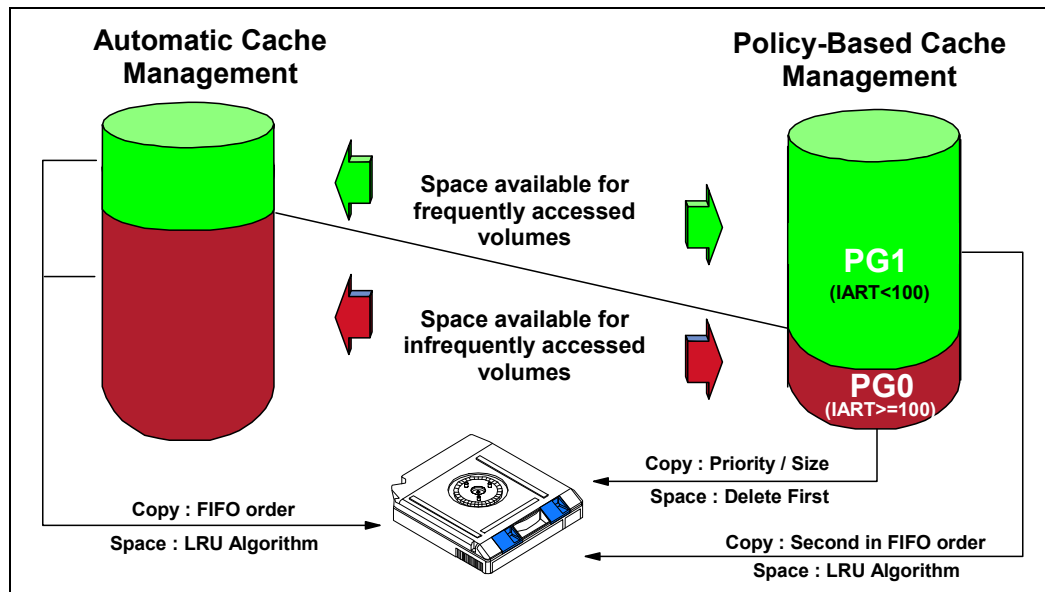


Figure 2-7 TS7740 cache usage with policy-based cache management

When a preference level is assigned to a volume, that assignment is persistent until the volume is reused for scratch and a new preference level is assigned. Or, if the policy is changed and a mount/dismount occurs, the new policy also takes effect.

Important: As of R2.1, all scratch volumes, independent of their preference group assignment, are favored for migration before selecting PG0 and PG1 candidates.

Recalled logical volumes are preferred for migration

Normally, a volume recalled into cache is managed as though it were newly created or modified, because it is in the TVC selected for I/O operations on the volume. A recalled volume displaces other volumes in cache, and moves to the end of the list of PG1 candidates to migrate due to how the LRU algorithm functions. The default behavior assumes any recall of a volume into the TVC might follow with additional host access.

However, there might be use cases where volumes recalled into cache are known to be accessed only once, and should be removed from disk cache as soon as they are read (for example, during a multi-volume data set restore). In this case, you wouldn't want the volumes to be kept in cache, because they require other more important cache resident data to be migrated.

Each TS7740 and TS7720T has an **LI REQ** setting that can determine how it handles recalled volumes. The **LI REQ SETTING RECLPG0** determines whether volumes that are recalled into cache are forced to PG0 or not. If forced to PG0, they are immediately migrated, freeing up space for other recalls without the need to migrate critical PG1 content.

Based on your current requirements, you can set or modify this control dynamically through the **LI REQ SETTING RECLPO** option:

- ▶ When **DISABLED**, which is the default, logical volumes that are recalled into cache are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7700.
- ▶ When **ENABLED**, logical volumes that are recalled into cache are managed as PG0 (preferable to be removed from cache). This control overrides the actions that are defined for the SC associated with the recalled volume.

2.2.14 About TVC cache management in a TS7700D and TS7700T CP0 stand-alone cluster

In a stand-alone TS7720 configuration, virtual volumes always remain in the TVC, because no physical tape drives are attached to the TS7700D. With a TS7700T configuration, contents in the CP0 partition are also not candidates for moving to physical tape. There is no autoremoval function in a stand-alone environment.

In a TS7700D stand-alone cluster, you can influence the TVC content only with the **Delete Expired and EJECT** setting. No further cache management is available. For a TS770T, the **LI REQ PARTRFSH** command can be used to move data between partitions.

If a TS7700D runs out of cache space or TS7700T runs out of CP0 space, warning messages and critical messages are shown. If the TS7700D enters the **Out of cache** condition, it moves to a read-only state. If a TS7720T CP0 partition becomes full, it becomes read-only regarding workloads that target CP0.

Important: Monitor your cache in a TS7700D stand-alone environment to avoid an **Out of Cache Resources** situation.

2.2.15 TVC Cache management in a TS7700T CPx stand-alone cluster

The TS7700T Tape partitions (CPx) have the same TVC cache controls as the TS7740. Two extra features exist on the TS7700T for cache management control when compared to the TS7740. These features are described in the following sections.

Multiple tape partitions

As mentioned earlier, you can specify 1 - 7 independent tape partitions. Each tape partition has its own independent cache management. The TVC management based on PG0/PG1 and LRU management is on a tape partition level, which means that data in one tape partition has no influence on the cache management in another tape partition.

Time-delayed premigration

You can use the delay premigration to delay the premigration of data. The delay premigration is controlled by two different components:

- ▶ In the partition itself, the amount of data that can be stored in the delay premigration queue is determined. The minimum is 500 gigabytes (GB). The maximum is the tape partition size minus 500 GB.
- ▶ In the SC, the delay premigration settings are defined. You can choose either *after Volume creation* or *Volume last accessed* as the used reference point. In addition, you define the time of the grace period. The time is specified in hours, 0 - 65535. A value of 0 means that no delay premigration time is set.

You can have multiple SCs, with different delay premigration definitions, pointing to the same tape partition. You can also have multiple SCs with different delay premigration definitions that point to different tape partitions.

Consider a situation where the amount of delayed premigration content, which has also not yet met its delay criteria, exceeds a partition's configured maximum delay-premigration limit. In this case, delayed content that has been present in the partition for the longest time is moved to the premigration queue proactively to maintain the configured limit. A too-small defined delay premigration size leads to the situation where data is always pushed out to tape too early, which can create excess back-end tape usage and its associated activity.

One important aspect of using delay premigration is that the content that is delayed for premigration is not added to the premigration queue until its delay criteria has been met. This means that if a large amount of delayed content meets its criteria at the same time, the premigration queue can rapidly increase in size. This rapid increase can result in unexpected host throttling.

Ensure that your FC5274 feature counts can accommodate these large increases in premigration activity. Alternatively, try to ensure that multiple workloads that are delayed for premigration do not reach their criteria at the same time.

Assume that you have three different tape partitions and a unique SC for each one. The following list describes the SC definitions:

- ▶ CP1: Delay premigration 12 hours after volume creation
- ▶ CP2: Delay premigration 6 hours after volume creation
- ▶ CP3: Delay premigration 3 hours after volume creation

In CP1 at 22:00, 6 TB are written every night. The 12-hour delay ensures that they are premigrated later in the day when there is a lower workload. To make the example simpler, we assume that no compression exists for all data.

In CP2 at 04:00, 2 TB are written. The six-hour delay makes them eligible for premigration also at 10:00 in the morning.

In CP3 at 07:00, 1 TB is written. The three-hour delay has them eligible for premigration at the same time as the other two workloads.

Therefore, all 9 TB of the workload is meeting its delay criteria at roughly the same time, producing a large increase in premigration activity. If the premigration queue size is not large enough, workloads into the TS7700T are throttled until the premigration process can reduce the queue size. Ensure that the number of FC 5274 features are suitable, or plan the delay times so that they do not all expire at the same time.

2.2.16 Expired virtual volumes and the Delete Expired function

To remain compatible with physical tape, logical volumes that are returned to scratch, or that are expired, retain all previously written content until they are reused or written from BOT. In a virtual tape environment, the retention of this scratched content can lead to any of the following situations:

- ▶ TVCs might fill up with large amounts of expired data.
- ▶ Stacked volumes might retain an excessive amount of expired data.
- ▶ Stacked volumes fill up with already expired data.

To help manage the expired content, the TS7700 supports a function referred to as *delete expire*. When enabling delete expire processing against a configured scratch category, you can set a grace period for expired volumes ranging 1 hour - 144 weeks (the default is 24 hours). If the volume has not already been reused when the delay period has passed, the volume is marked as a candidate for auto deletion or delete expire.

When deleted, its active space in TVC is freed. If it was also stacked to one or more physical tapes, that region of physical tape is marked inactive.

The start timer for delete expire processing is set when the volume is moved to a designated scratch category, or a category with the Fast Ready attribute set, which has defined a delete expire value. If the scratch category has no delete expire value, the timer is not set.

During the delete expire process, the start timer and the delete expire value are used to determine whether the logical volume is eligible for the delete expire processing. If so, the content is deleted immediately.

If the logical volume is reused during a scratch mount before the expiration delete time expires, the existing content is immediately deleted at the time of first write.

It does not matter whether the volume is in cache or on back-end tape; after the delete expire time passes, the volume is no longer accessible without IBM SSR assistance. The default behavior is to Delete Expire up to 1000 delete-expire candidates per hour. This value can be modified by using the **LI REQ** command.

For more information about expired volume management, see “Defining the logical volume expiration time” on page 551. The explicit movement of a volume out of the delete expired configured category can occur before the expiration of this volume.

Important: Disregarding the Delete Expired Volumes setting can lead to an out-of-cache state in a TS7700D. With a TS7740 or TS7700T, it can cause excessive tape usage. In an extreme condition, it can cause an out-of-physical scratch state.

The disadvantage of not having this option enabled is that scratched volumes needlessly use TVC and physical stacked volume resources, so they demand more TVC active space while also requiring more physical stacked volumes in a TS7740 or TS7700T. The time that 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 logical volumes to be moved from one stacked volume to another during reclamation.

Expire Hold settings

A volume that is expired or returned to scratch might be reused during a scratch mount before its delete expire grace period has passed. If retention of expired content is required, an extra Expire Hold setting can be enabled.

When Expire Hold is enabled as part of the delete expire settings, the expired or scratched volume is moved into a protected hold state in which it is not a candidate for scratch mounts. The volume is also not accessible from any host operation until the configured expire time grace period has passed. Starting with Release 2.1 of the TS7700, these held volumes can be moved back to a private category while still in a held state.

This additional option is made available to prevent any malicious or unintended overwriting of scratched data before the duration elapses. After the grace period expires, the volume is simultaneously removed from a held state and made a deletion candidate.

Remember: Volumes in the Expire Hold state are excluded from DFSMS OAM scratch counts, and are not candidates for TS7700 scratch mounts.

Delete Expired data that was previously stacked onto physical tape remains recoverable through an IBM services salvage process if the physical tape has not yet been reused, or if the secure erase process was not performed against it. Contact your IBM SSR if these services are required. Also, disabling reclamation as soon as any return to scratch mistake is made can help retain any content still present on physical tape.

Important: When Delete Expired is enabled for the first time against a scratch category, all volumes that are contained within that category are *not* candidates for delete expire processing. Only volumes that moved to the scratch category after the enablement of the Delete Expired are candidates for delete expire processing.

Changes to the Delete Expired values are effective to all logical volumes that are candidates for delete expire processing.

2.2.17 TVC management processes for TS7740 or TS7700T CPx

Two processes manage the TVC of the TS7740 and TS7700T in a stand-alone environment:

- ▶ Premigration Management (TS7740 and TS7700T CPx)

This process is always actively queuing newly created volumes for movement to back-end physical tape. When the TS7700 determines that minimal host activity is occurring, or if the **PMPRIOR** threshold has been crossed, it begins servicing the premigration queue by copying the logical volume content to one or more physical tapes.

When copied or stacked to physical tape, the volume is a candidate for migration or the deletion out of TVC. Content in a TS7700T CPx partition that is configured with a delay premigration time is not queued for premigration until the delay criteria is met, or until the maximum delay premigration threshold for the partition is exceeded.

If your TS7740/TS7700T exceeds its **PMPRIOR** threshold of content queued for premigration, the TS7740 or TS7700T likely enters the sustained mode of operation in which the speed of which the TS7740 or TS7700T can absorb new workload is at the speed of which it can be premigrated to physical tape. Thresholds that are associated with the priority of premigration and the sustained mode of operations are tunable by using the **LI REQ** commands.

For more information, see 10.1.3, “Host Console Request function” on page 608.

► Free-space Management (TS7740 and TS7700T CPx)

This process manages the amount of free space within the TVC of a TS7740 or TS7700T. When the premigration process completes, a volume is a candidate for deletion from disk cache, otherwise known as *migration*. Volumes that are selected for migration are chosen based on how much free space is needed, LRU algorithms, and configured policies, such as preference group.

In a TS7740 or in CPx partitions of a TS7700, the inability to free disk cache space through premigration and migration can lead to heavy host throttling. Volumes targeting the CP0 partition of a TS7720T are not susceptible to any throttling associated with moving content to physical tape.

2.2.18 TVC handling in outage situations

In a TS7740 environment, a “force paused” mode of the TS3500 or a resource shortage (out of physical scratch) leads to the situation where the stand-alone TS7740 does not accept any writes. That is true even when the TS7740 has free cache space. This function was introduced to ensure that no cache overflow occurs.

However, the cache sizes in a TS7700T are much bigger and depend on the installed configuration, so this behavior might not be appropriate. Therefore, in Release 3.3, a new command that is called the **LI REQ** command defines how a TS7700T behave in such a condition.

You can now specify whether a TS7700T CPx reacts the same as a TS7740 or accepts the incoming write in a stand-alone mode until the cache resources are exhausted.

2.2.19 Copy Consistency Point: Copy policy modes in a stand-alone cluster

In a stand-alone cluster, you cannot define any Copy Consistency Point.

2.2.20 TVC selection in a stand-alone cluster

Because there is only one TVC in a stand-alone cluster available, no TVC selection occurs.

2.2.21 TVC encryption

With R3.0, a TVC encryption feature was introduced.

TVC encryption is turned on for the whole disk cache. You cannot encrypt a disk cache partially. Therefore, all DDMs in all strings must be full disk encryption (FDE)-capable to enable the encryption. The disk cache encryption is supported for all TS7760 models with CSA, all TS7720 models with 3956-CS9 cache or higher, and for TS7740 with 3956-CC9.

Encryption can be enabled in the field at any time, and retroactively encrypts all existing content that is stored within the TVC. Because the encryption is done at the HDD level, encryption is not apparent to the TS7700 and has no effect on performance.

Starting with R3.0, only local key management is supported. Local key management is automated. There are no encryption keys (EKs) for the user to manage.

Release 3.3 now supports the usage of an external key manager. The following encryption key managers are supported:

- ▶ IBM Security Key Lifecycle Manager (formerly IBM Tivoli Key Lifecycle Manager)
- ▶ IBM Security Key Lifecycle Manager for z/OS

If you want to use an external key manager for both TVC and physical tape, you must use the same external key manager instance for both of them.

There are two differences between the usage of local or external key management:

- ▶ If you have no connection to the external key manager, TS7700 will not run. Therefore, you must plan carefully to have a primary and an alternative key manager that are reachable in a disaster situation.
- ▶ If a cluster that uses disk encryption with an external key manager is unjoined from a grid, the encryption must be disabled during this process. Otherwise, the TS7700 cannot be reused. Therefore, during the unjoin, the cluster is *secured erased*.

2.2.22 Physical volume pools

You can use the TS7740 and TS7700T to group volumes by pools when stacking to physical tape takes place.

The following list includes some examples of why physical volume pools are helpful:

- ▶ Data from separate customers on the same physical volume can compromise certain outsourcing contracts.
- ▶ Customers want to be able to “see, feel, and touch” their data by having only their data on dedicated media.
- ▶ Customers need separate pools for different environments, such as test, user acceptance test (UAT), and production.
- ▶ 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 reduce their media charges.
- ▶ Recall times depend on the media length. Small logical volumes on the tape cartridges (JA, JB, and JC) can take a longer time to recall than volumes on the economy cartridge (JJ or JK). Therefore, pooling by media type is also beneficial.
- ▶ Some workloads have a high expiration rate, which causes excessive reclamation. These workloads are better suited in their own pool of physical volumes.
- ▶ Protecting data through encryption can be set on a per pool basis, which enables you to encrypt all or some of your data when it is written to the back-end tapes.
- ▶ Migration from older tape media technology.
- ▶ Reclaimed data can be moved to a different target pool, which enables aged data to move to a specific subset of physical tapes.
- ▶ Second dedicated pool for key workloads to be Copy Exported.

There are benefits to using physical volume pools, so plan for the number of physical pools. See also “Relationship between reclamation and the number of physical pools” on page 52.

Using physical volume pools

Physical volume pool properties enable the administrator to define pools of stacked volumes within the TS7740/TS7700T. You can direct virtual volumes to these pools by using SMS constructs. There can be up to 32 general-purpose pools (01 - 32) and one common pool (00). A common scratch pool (Pool 00) is a reserved pool that contains only scratch stacked volumes for the other pools.

Each TS7740/TS7700T that is attached to an IBM TS4500 or IBM TS3500 tape library has its own set of pools.

Common scratch pool (Pool 00)

The *common scratch pool* is a pool that contains only scratch stacked volumes, and serves as a reserve pool. You can define a primary pool to borrow scratch stacked cartridges from the common scratch pool (Pool 00) if a scratch shortage occurs. This can be done either on a temporary or permanent basis.

Each pool can be defined to borrow single media type (for example, JA, JB, JC, JD), borrow mixed media, or have a first choice and a second choice. The borrowing options can be set by using the MI when you are defining stacked volume pool properties.

Remember: The common scratch pool must have at least three scratch cartridges available when one or more reports low scratch count warnings.

General-purpose pools (Pools 01 - 32)

There are 32 general-purpose pools available for each TS7740/TS7700T cluster. These pools can contain both empty and full or filling stacked volumes. All physical volumes in a TS7740/TS7700T cluster are distributed among available pools according to the physical volume range definitions in place. The distribution is also based on the pools' borrow and return attribute settings.

Those pools can have their properties tailored individually by the administrator for various purposes. When initially creating these pools, it is important to ensure that the correct borrowing properties are defined for each one. For more information, see "Stacked volume pool properties" on page 50.

By default, there is one pool, Pool 01, and the TS7740/TS7700T stores virtual volumes on any stacked volume available to it. This creates an intermix of logical volumes from differing sources, for example, an LPAR and applications on a physical cartridge.

The user cannot influence the physical location of the logical volume within a pool. Having all of the logical volumes in a single group of stacked volumes is not always optimal.

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

- ▶ Separate different clients or LPAR data from each other.
- ▶ Intermix or segregate media types.
- ▶ Map separate SGs to the same primary pools.
- ▶ Set up specific pools for Copy Export.
- ▶ Set up pool or pools for encryption.
- ▶ Set a reclamation threshold at the pool level.
- ▶ Set reclamation parameters for stacked volumes.
- ▶ Set up reclamation cascading from one pool to another.
- ▶ Set maximum number of devices to use concurrent premigration on pool base.
- ▶ Assign or eject stacked volumes from specific pools.

Physical pooling of stacked volumes is identified by defining a pool number, as shown in Figure 2-8.

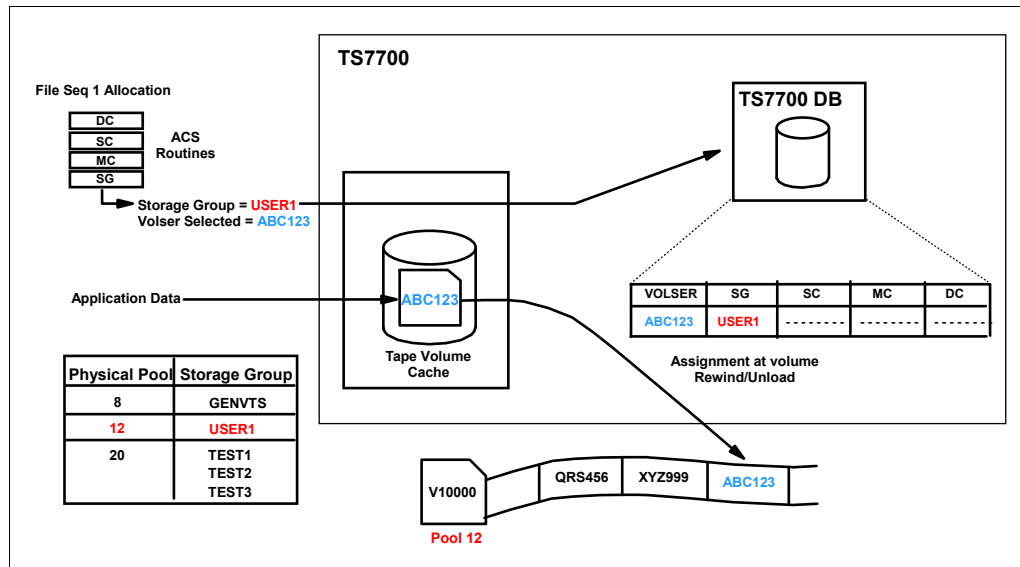


Figure 2-8 TS7740/TS7700T Logical volume allocation to specific physical volume pool flow

Through the MI, you can add an SG construct, and assign a primary storage pool to it. Stacked volumes are assigned directly to the defined storage pools. The pool assignments are stored in the TS7740/TS7700T database. During a scratch mount, a logical volume is assigned to a selected SG.

This SG is connected to a storage pool with assigned physical volumes. When a logical volume is copied to tape, it is written to a stacked volume that belongs to this storage pool. In addition, MC can be used to define a secondary pool when two copies on physical tape are required.

Physical 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. When also disabling borrow/return for that pool, this provides a method to have a specific range of volumes that are used exclusively by a specific pool.

Tip: Primary Pool 01 is the default private pool for TS7740/TS7700T stacked volumes.

Borrowing and returning: Out of physical stacked volume considerations

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/TS7700T can manage out-of-scratch and low scratch scenarios, which can occur within any TS7740/TS7700T from time to time.

You need at least two empty stacked volumes in the CSP to avoid any out of scratch condition. Empty pvols in other pools (regardless of the pool properties) are not considered. Ensure that non-borrowing active pools have at least two scratch volumes.

One physical pool with an out of stacked volume condition results in an out of stacked volume condition to the whole TSS740 / TS7700T cluster. Therefore, it is necessary to monitor all active pools.

Remember: Pools that have borrow/return enabled, and that contain no active data, eventually return all of the scratch volumes to the common scratch pool after 48 - 72 hours of inactivity.

Stacked volume pool properties

Logical volume pooling supports cartridge type selection. This can be used to create separate pools of 3592 tape cartridges with various capacities 128 GB - 10 TB, depending upon the type of media and tape drive technology used.

Lower capacity JJ, JK, or JL cartridges can be designated to a pool to provide consistently faster access to application data, such as hierarchical storage management (HSM) or Content Manager. Higher capacity JA, JB, JC, or JD cartridges that are assigned to a pool can address archival requirements, such as full volume dumps.

2.2.23 Logical and stacked volume management

Every time that a logical volume is modified (either by modification or by reuse of a scratch volume), the data from the previous use of this logical volume, which is on a stacked volume, becomes obsolete. The new virtual volume is placed in the cache and written to a stacked volume afterward (TS7740 or TS7700T). The previous copy on a stacked volume is invalidated, but it still uses up space on the physical tape. Auto delete of expired volumes or ejected volumes can also result in the creation of inactive space on physical tapes.

Virtual volume reconciliation

The reconciliation process periodically checks for active volume usage percentages, which remain on physical tape cartridges. This process automatically adjusts the active data values of the physical volumes, which are the primary attributes that are used for automatic physical volume reclamation processing.

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

- ▶ A host has assigned the logical volume to a scratch category. Later, the volume is 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. The category has a nonzero delete-expired data parameter value. The parameter value was exceeded, and the TS7740/TS7700T deleted the logical volume.
- ▶ A host has modified the contents of the volume. This can be a complete rewrite of the volume or an append to it. The new version of the logical volume is premigrated to a separate physical location and the older version is invalidated.
- ▶ The logical volume is ejected, in which case the version on physical tape is invalidated.
- ▶ The pool properties change during a mount/demount sequence and a new pool is chosen.

The TS7740/TS7700T tracks the amount of active data on a physical volume. During a premigration or reclamation, the TS7700 attempts to fill the targeted volume and mark it 100% active. Although the granularity of the percentage of full TS7740/TS7700T tracks is 1/10 of 1%, it rounds down, so even *1 byte* of inactive data drops the percentage to 99.9%. TS7740/TS7700T tracks 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 nonzero setting
- ▶ Comparing this time against the current time to determine whether the volume is eligible for reclamation

Physical volume reclamation

Physical volume reclamation consolidates active data and frees stacked volumes for return-to-scratch use. Reclamation is part of the internal management functions of a TS7740/TS7700T. The reclamation process is basically a tape-to-tape copy. The physical volume to be reclaimed is mounted to a physical drive, and the active logical volumes that are there are copied to another filling cartridge under control of the TS7740/TS7700T.

One reclamation task needs two physical tape drives to run. At the end of the reclaim, the source volume is empty, and it is returned to the specified reclamation pool as an empty (scratch) volume. The data that is being copied from the reclaimed physical volume does not go to the TVC. Instead, it is transferred directly from the source to the target tape cartridge. During the reclaim, the source volume is flagged to be in READ.ONLY mode.

Physical tape volumes become eligible for space reclamation when they cross the occupancy threshold level that is specified by the administrator in the home pool definitions where those tape volumes belong. This reclaim threshold is set for each pool individually according to the specific needs for that client, and is expressed in a percentage (%) of tape usage.

Volume reclamation can be concatenated with a Secure Data Erase for that volume, if required. This configuration causes the volume to be erased after the reclamation. For more information, see 2.2.24, “Secure Data Erase function” on page 52.

Consider *not* running reclamation during peak workload hours of the TS7740/TS7700T. This ensures that recalls and migrations are not delayed due to physical drive shortages. You must choose the best period for reclamation by considering the workload profile for that TS7740/TS7700T cluster, and inhibit reclamation during the busiest period for the system.

A physical volume that is being ejected from the library is also reclaimed in a similar way before it can be ejected. The active logical volumes that are contained in the cartridge are moved to another physical volume, according to the policies defined in the volume’s home pool, before the physical volume is ejected from the library.

An MI-initiated PVOL move also runs this reclamation process.

Reclamation can also be used to migrate older data from a pool to another while it is being reclaimed, but only by targeting a separate specific pool for reclamation.

With Release 3.3, it is now possible to deactivate the reclaim on a physical pool base by specifying a “0” value in the Reclaim Threshold.

With the introduction of heterogeneous tape drive support for migration purposes, the data from the old cartridges (for example, JA and JB) is reclaimed to the new media (for example, JC and JD). To support a faster migration, the reclaim values for the sunset media can be different from the reclaim values for the current tape media. To allow the reclaim for sunset media, at least 15 scratch cartridges from the newer tape media needs to be available. For more information “Physical Volume Pools” on page 410.

Relationship between reclamation and the number of physical pools

The reclaim process is done on a pool basis, and each reclamation process needs two drives. If you define too many pools, it can lead to a situation where the TS7740/TS7700T is incapable of processing the reclamation for all pools in an appropriate manner. Eventually, pools can run out of space (depending on the *borrow* definitions), or you need more stacked volumes than planned.

The number of physical pools, physical drives, stacked volumes in the pools, and the available time tables for reclaim schedules must be considered and balanced.

You can limit the number of reclaim tasks running concurrent with the **LI REQ** setting.

2.2.24 Secure Data Erase function

Another concern is the security of old data. The TS7740/TS7700T provides physical volume erasure on a physical volume pool basis controlled by an extra reclamation policy. When Secure Data Erase is enabled, a physical cartridge is not made available as a scratch cartridge until an erasure procedure is complete. The Secure Data Erase function supports the erasure of a physical volume as part of the reclamation process. The erasure is performed by running a long erase procedure against the media.

A Long Erase operation on a TS11xx drive is completed by writing a repeating pattern from the beginning to the end of the physical tape, making all data previously present inaccessible through traditional read operations. The key here is that it is not a fully random from beginning to end pattern, and it has only one pass. The erasure is writing a single random pattern repeatedly with one pass, which might not be as secure as a multi-pass fixed pattern method, as explained by the US Department of Defense (DoD).

Therefore, the logical volumes that are written on this stacked volume are no longer readable. As part of this *data erase* function, an extra reclaim policy is added. The policy specifies the number of days that a physical volume can contain invalid logical volume data before the physical volume becomes eligible to be reclaimed.

When a physical volume contains encrypted data, the TS7740/TS7700T is able to run a fast erase of the data by erasing the EKs on the cartridge. Basically, it erases only the portion of the tape where the key information is stored. This form of erasure is referred to as a *cryptographic erase*.

Without the key information, the rest of the tape cannot be read. This method significantly reduces the erasure time. 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 a read-only recovery (ROR).

If you use the eject stacked volume function, the data on the volume is not erased before ejecting. The control of expired data on an ejected volume is your 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 a volume is removed for recovery actions.

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

2.2.25 Copy Export function

One of the key reasons to use tape is for recovery of critical operations in a disaster. If you are using a grid configuration that is designed for DR purposes, the recovery time objectives (RTO) and recovery point objectives (RPO) can be measured in minutes. In case you do not require such low recovery times for all or a mixture of your workload, there is a function called *Copy Export* for the TS7740 and TS7700T.

The Copy Export function enables a copy of selected logical volumes that are written to secondary pools within the TS7740/TS7700T to be removed and taken offsite for DR purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. Because the physical volumes that are being exported are from a secondary physical pool, the primary logical volume remains accessible to the production host systems.

The following logical volumes are excluded from the export:

- ▶ Volumes that are mounted during any portion of the export process
- ▶ Volumes that are unable to create a valid primary or secondary pool copy
- ▶ Volumes that had not completed replication into the source TS7740 or TS7700T at the start of the export process

These volumes will be candidates in the next copy export request.

The Copy Export sets can be used to restore data at a location that has equal or newer tape technology and equal or newer TS7700 Licensed Internal Code. A TS7700T Copy Export set can be restored into both TS7740 and TS7700T. A TS7740 Copy Export set can also be restored into both TS7740 and TS7700T. However, some rules apply:

- ▶ TS7700T exported content that is restored to a TS7740 loses all knowledge of partitions.
- ▶ TS7700T to TS7700T retains all partition information.
- ▶ TS7740 exported content that is restored into a TS7700T has all content target the primary tape partition.

There is an offsite reclamation process against copy-exported stacked volumes. This process does not require the movement of physical cartridges. Rather, the logical volume is written newly to a copy-exported stacked volume, and the original copy exported stacked volume is marked invalid. For more information, see 12.1.3, “Reclaim process for Copy Export physical volumes” on page 765.

2.2.26 Encryption of physical tapes

The importance of data protection has become increasingly apparent with news reports of security breaches, loss, and theft of personal and financial information, and with government regulation. Encrypting the stacked cartridges 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:

- ▶ The encryption key manager
- ▶ The TS1150, TS1140, TS1130, and TS1120 encryption-enabled tape drives
- ▶ The TS7740/TS7700T

Encryption key manager

TS7700 can use one of the following encryption key managers:

- ▶ IBM Encryption Key Manager (EKM)
- ▶ IBM Security Key Lifecycle Manager (formerly IBM Tivoli Key Lifecycle Manager)
- ▶ IBM Security Key Lifecycle Manager for z/OS

This book uses the general term *key manager* for all three EK managers.

Important: *The EKM is no longer available and does not support the TS1140 and TS1150.* If you need encryption support for the TS1140 or higher, you must install either IBM Security Key Lifecycle Manager or IBM Security Key Lifecycle Manager for z/OS.

IBM Security Key Lifecycle Manager replaces Tivoli Key Lifecycle Manager.

The key manager is the central point from which all EK information is managed and served to the various subsystems. The key manager server communicates with the TS7740/TS7700T and tape libraries, CUs, and Open Systems device drivers. For more information, see 4.4.7, “Planning for tape encryption in a TS7740, TS7720T, and TS7760T” on page 174.

The TS1150, TS1140, TS1130, and TS1120 encryption-enabled tape drives

The IBM TS1150, TS1140, TS1130, and TS1120 tape drives provide hardware that performs the encryption without reducing the data transfer rate.

The TS7740/TS7700T

The TS7740/TS7700T provides the means to manage the use of encryption and the keys that are used on a storage pool basis. It also acts as a proxy between the tape drives and the key manager servers, by using redundant Ethernet to communicate with the key manager servers and FICONs to communicate with the drives. Encryption must be enabled in each of the tape drives.

Encryption on the TS7740/TS7700T is controlled on a storage pool basis. The SG 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/TS7700T.

The storage pools were originally created for management of physical media, and they have been enhanced to include encryption characteristics. Storage pool encryption parameters are configured through the TS7740/TS7700T MI under Physical Volume Pools.

For encryption support, all drives that are attached to the TS7740/TS7700T must be Encryption Capable, and encryption must be enabled. If TS7740/TS7700T 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 MI, one or more of the 32 pools can be enabled for encryption.

Each pool can be defined to use *specific EKs* or the *default EKs* defined at the key manager server:

- ▶ Specific EKs

Each pool that is defined in the TS7740/TS7700T can have its own unique EK. As part of enabling a pool for encryption, enter two key labels for the pool and an associated key mode. The two keys might or might not be the same. Two keys are required by the key 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.

The MI 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 MI, you provide IP addresses for a primary and an optional secondary key manager.

► **Default EKs**

The TS7740/TS7700T encryption 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/TS7700T. 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 by using the SG construct name. All data that is directed to a pool that is enabled for encryption is encrypted when they are premigrated to the physical stacked volumes, or reclaimed to the stacked volume during the reclamation process. The SG construct name is bound to a logical volume when it is mounted as a scratch volume.

Through the MI, the SG 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/TS7700T determines whether a new physical volume needs to be mounted. If a new cartridge is required, the TS7740/TS7700T directs the drive to use encryption during the mount process.

The TS7740/TS7700T also provides the drive with the key labels specified for that pool. When the first write data is received by the drive, a connection is made to a key manager and the key that is needed to perform the encryption is obtained. Physical scratch volumes are encrypted with the keys in effect at the time of first write to BOT.

Any partially filled physical volumes continue to use the encryption settings in effect at the time that the tape was initially written from BOT. The encryption settings are static until the volumes are reclaimed and rewritten again from BOT.

Figure 2-9 illustrates that the method for communicating with a key manager is through the same Ethernet interface that is used to connect the TS7740/TS7700T to your network for access to the MI.

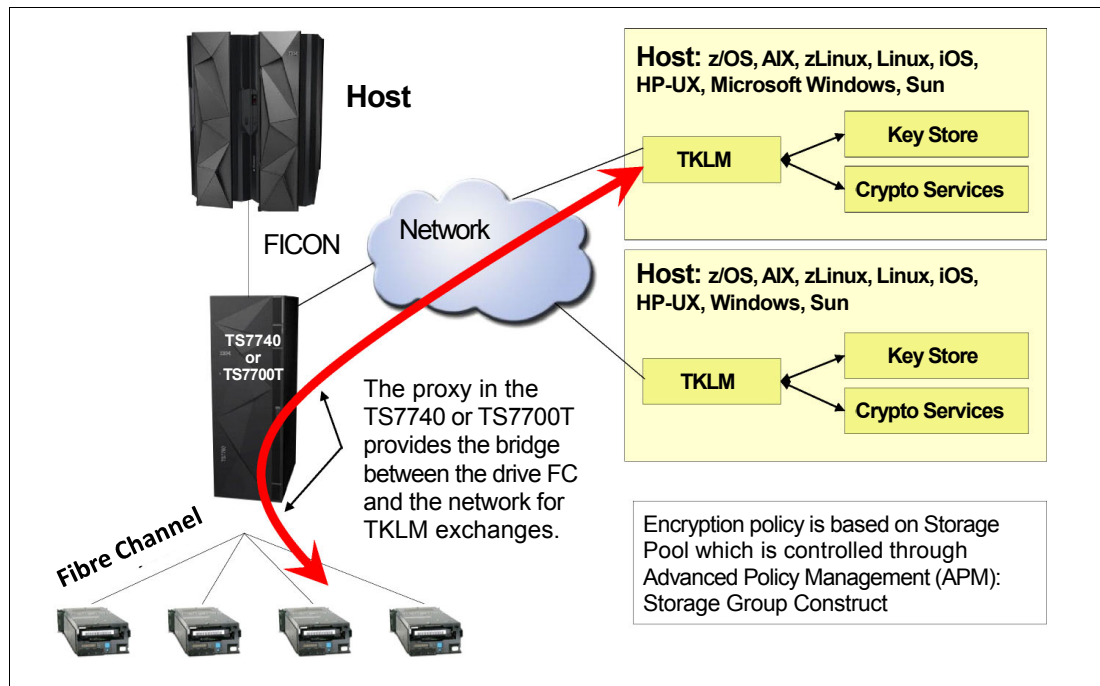


Figure 2-9 TS7740/TS7700T encryption

The request for an EK is directed to the IP address of the primary key manager. Responses are passed through the TS7740/TS7700T 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 TS11x0 drive completes 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 that is enabled for encryption, either as part of a recall or reclamation operation, the TS7740/TS7700T uses the key manager to obtain the necessary information to decrypt the data.

The affinity of the logical volume to a specific EK, or the default key, can be used as part of the search criteria through the TS7700 MI.

Remember: If you want to use external key management for both cache and physical tapes, you must use the same external key manager instance.

2.2.27 User Management: Roles and profiles

The TS7700 offers you internal user management, but also external user management through LDAP support.

You can use this user management to specify independent User IDs. Each User ID is assigned a role. The role identifies the access rights for this user. You can use this method to restrict the access to specific tasks.

In R3.2, a new *read only* role was introduced. Users who are assigned to this role can view information only about the MI, but cannot change any information.

You should consider restricting access to specific items. Especially the Tape Partition management and the access to the LIBRARY REQUEST should be considered carefully.

2.2.28 Security identification by using Lightweight Directory Access Protocol

Previous implementations are based on Tivoli System Storage Productivity Center to authenticate users to a client's Lightweight Directory Access Protocol (LDAP) server. Beginning with Release 3.0 of Licensed Internal Code (LIC), both the TS7700 clusters and TS3000 System Console (TS3000 TSSC) have native support for an LDAP server (currently, only Microsoft Active Directory (MSAD) is supported).

Starting with R3.0, when LDAP is enabled, the TS7700 MI is controlled by the LDAP server. Also, the local actions that are run by the IBM SSR are secured by the LDAP server. All IBM standard users can no longer access the system without a valid LDAP user ID and password. You must have a valid account in the LDAP server, and the roles that are assigned to your user, to be able to communicate with the TS7700.

If your LDAP server is not available, you are not able to interact with TS7700 (not with IBM SSR or an operator).

Important: Create at least one external authentication policy for IBM SSRs before a service event.

With R3.2, IBM RACF® can now be used to control the access. That means that all users are defined to RACF and, in case of an access, the password is verified on the RACF database. Roles and profiles still must be maintained because the RACF database runs only the password authentication.

In Release 3.2, a change was introduced to allow specific access without the usage of LDAP (IBM SSR and second-level dial-in support).

2.2.29 Service preparation mode

This function is available only in a multi-cluster grid.

2.2.30 Service mode

This function is available only in a multi-cluster grid.

2.3 Multi-cluster grid configurations: Components, functions, and features

Multi-cluster grids are combinations of two clusters, three clusters, four clusters, five clusters, or six clusters that work together as one logical entity. TS7700D, TS7700T, and TS7740 can be combined as a *hybrid grid*, but you can also form a grid just from TS7700D, TS7700T, or TS7740 clusters. The configuration that is suitable for you depends on your requirements.

To enable multiple clusters to work together as a multi-cluster grid, some hardware configurations must be provided. Also, logical considerations need to be planned and implemented. The following topics are described in this section:

- ▶ The base rules that apply in a multi-cluster grid
- ▶ Required grid hardware
- ▶ Implementation concepts for the grid
- ▶ Components and features that are used in a grid

Figure 2-10 shows a four-cluster hybrid grid. The configuration consists of two TS7720s and two TS7740s. More examples are available in 2.4, “Grid configuration examples” on page 95.

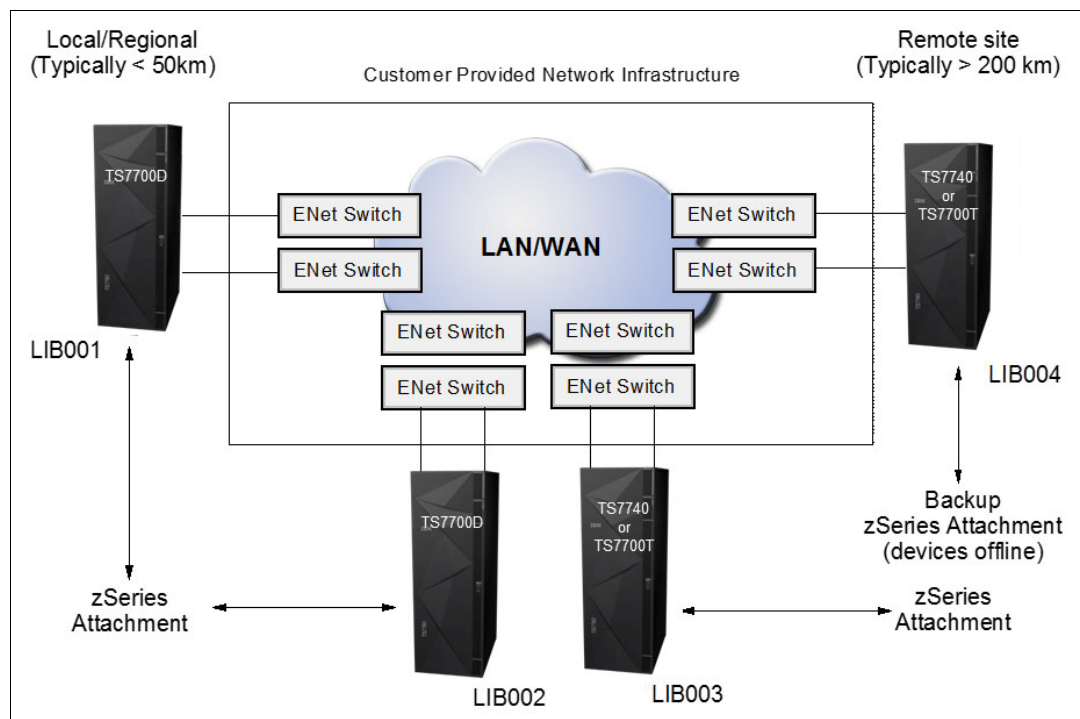


Figure 2-10 TS7700D 4-cluster grid

2.3.1 Rules in a multi-cluster grid

In a multi-cluster grid, some general rules apply:

- ▶ A grid configuration looks like a single tape library and tape drives to the hosts.
- ▶ It is a composite library with underlying distributed libraries.
- ▶ Up to six clusters can form a grid.
- ▶ Data integrity is accomplished by the concept of volume ownership.
- ▶ All TS7700 models can coexist in a grid. If only a disk and a tape-attached model are combined, that configuration is called a *hybrid grid*.

- ▶ If one cluster is not available, the grid still continues to work.
- ▶ Clusters can be grouped into cluster families.
- ▶ Mounts, both scratch (Fast Ready) and private (non-Fast Ready), can be satisfied from any cluster in the grid, which is controlled by your implementation.

Remember: Five-cluster grid and six-cluster grid configurations are available with an RPQ.

In a multi-cluster grid, some rules for virtual and logical volumes apply:

- ▶ You can store a logical volume or virtual volume in the following ways:
 - Single instance in only one cluster in a grid.
 - Multiple instances (two, three, four, five, or six) in different clusters in the grid, up to the number of clusters in the grid.
 - Each TS7740/TS7700T cluster in the grid can store dual copies on physical tape. Each copy is a valid source for the virtual or logical volume.
 - Selective dual copy is still a valid option in a TS7740/TS7700T. (In an extreme case, you can end up with 12 instances of the same data spread out on six different clusters using selective dual copy.)
- ▶ You control the number of instances, and the method of how the instances are generated through different copy policies.

In a multi-cluster grid, the following rules for access to the virtual and logical volumes apply:

- ▶ A logical volume can be accessed from any virtual device in the system.
- ▶ Any logical volume (replicated or not) is accessible from any other cluster in the grid.
- ▶ Each distributed library has access to any logical volumes within the composite library.

Note: You can still restrict access to clusters by using host techniques (for example, HCD).

With this flexibility, the TS7700 grid provides many options for business continuance and data integrity, meeting requirements for a minimal configuration up to the most demanding advanced configurations.

2.3.2 Required grid hardware

To combine single clusters into a grid, several requirements must be met:

- ▶ Each of the TS7700 must have the Grid Enablement feature installed.
- ▶ Each of the TS7700 engines must be connected to all other clusters in the grid through the *grid network*. Each cluster can have two or four links to the grid network.

Grid enablement

FC4015 must be installed on all clusters in the grid.

Grid network

A grid network is the client-supplied TCP/IP infrastructure that interconnects the TS7700 grid. Each cluster has two Ethernet adapters that are connected to the TCP/IP infrastructure. The single-port 10 gigabits per second (Gbps) long-wave optical fiber adapter is supported. This configuration accounts for two or four grid links, depending on the cluster configuration. See 7.1.1, “Common components for the TS7700 models” on page 230.

Earlier TS7740 might still have the single-port adapters for the copper connections and SW 1 Gbps connections. A miscellaneous equipment specification (MES) is available to upgrade the single port to dual-port adapters.

Dynamic Grid Load Balancing

Dynamic Grid Load Balancing is an algorithm that is used within the TS7700. It continually monitors and records the rate at which data is processed by each network link. Whenever a new task starts, the algorithm uses the stored information to identify the link that can most quickly complete the data transfer. The algorithm also identifies degraded link performance, and sends a warning message to the host.

Remote mount automatic IP failover

If a grid link fails during a remote mount, the Remote Mount IP Link Failover function attempts to reestablish the connection through an alternative link. During a failover, up to three extra links are attempted. If all configured link connections fail, the remote mount fails, resulting in a host job failure or a Synchronous mode copy break. When a remote mount or a Synchronous copy is in use and a TCP/IP link failure occurs, this intelligent failover function recovers by using an alternative link. The following restrictions apply:

- ▶ Each cluster in the grid must operate by using a Licensed Internal Code level of 8.21.0.xx or later.
- ▶ At least two grid connections must exist between clusters in the grid (either two or four 1 Gbps grid links or two or four 10 Gbps grid links).

Internet Protocol Security for grid links

When you are running the TS7700 R3.1 level of LIC, the TS7760 and the 3957-V07 and 3957-VEB models support Internet Protocol Security (IPSec) on the grid links. Use IPSec capabilities Only if required by the nature of your business.

Tip: Enabling grid encryption significantly affects the replication performance of the TS7700 grid.

Date and Time coordination

The TS7700 cluster tracks time in relation to Coordinated Universal Time. Statistics are also reported in relation to Coordinated Universal Time. All nodes in the grid subsystem coordinate their time with one another. There is no need for an external time source, such as an NTP server, even in a grid with large distances between the clusters.

However, if the grid is not connected to an external time source, the time that is presented from the grid (VEHSTATS and so on) might not show the same time as your LPARs, which can lead to some confusion during problem determination or for reporting, because the different time stamps do not match.

Therefore, the preferred method to keep nodes synchronized is by using a Network Time Protocol (NTP) server. The NTP server can be a part of the grid wide area network (WAN) infrastructure, your intranet, or a public server on the internet (Figure 2-11 on page 61).

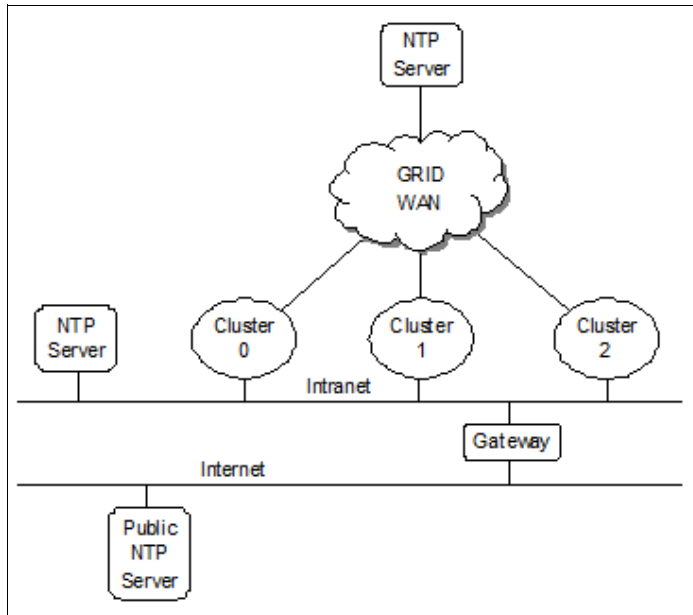


Figure 2-11 Time coordination with NTP servers

The NTP server address is configured into the system vital product data (VPD) on a system-wide scope. Therefore, all nodes access the same NTP server. All clusters in a grid need to be able to communicate with the same NTP server that is defined in VPD. In the absence of an NTP server, all nodes coordinate time with Node 0 or the lowest cluster index designation. The lowest index designation is Cluster 0, if Cluster 0 is available. If not, it uses the next available cluster.

2.3.3 Data integrity by volume ownership

In a multi-cluster grid, only one cluster at a time can modify volume data or attributes. To manage this, the concept of ownership was introduced.

Ownership

Any logical volume, or any copies of it, can be accessed by a host from any virtual device that is 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 rules*. At any point in time, a logical volume is owned by only one cluster. The owning cluster controls access to the data and the attributes of the volume.

Remember: The volume ownership protects the volume from being accessed or modified by multiple clusters simultaneously.

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 must obtain ownership first. When required, the TS7700 node transfers the ownership of the logical volume as part of mount processing. This action ensures that the cluster with the virtual device that is associated with the mount has ownership.

If the TS7700 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 not apparent to the host.

If a TS7700 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 fails unless more direction is given.

Ownership can also be transferred manually by an LI REQ,OTCNTL for special purposes. For more information, see the *IBM TS7700 Series z/OS Host Command Line Request User's Guide* on the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

If a cluster is not reachable, clusters do not automatically assume or take ownership of a logical volume without being directed. This can either be done manually, or can be automated with the Autonomic Ownership Takeover Manager (AOTM). Service outages have *implied* ownership takeover. The manual ownership possibility is presented on the MI when the grid determines that a cluster cannot be reached, and AOTM is either not installed or the failure cannot be managed by AOTM.

To support the concept of ownership, it was necessary to introduce tokens.

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. Tokens have the following characteristics:

- ▶ Every logical volume has a corresponding token.
- ▶ The grid component manages updates to the tokens.
- ▶ Tokens are maintained in an IBM DB2® database that is 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 that are generated with the Bulk Volume Information Retrieval (BVIR) facility.

Tokens are part of the architecture of the TS7700. Even in a stand-alone cluster, they exist and are used in the same way as they are used in the grid configuration (with only one cluster running the updates and keeping the database). In a grid configuration, all members in the grid have the information for all tokens (also known as logical volumes) within the composite library mirrored in each cluster. Token information is updated real time at all clusters in a grid.

Ownership takeovers

In some situations, the ownership of the volumes might not be transferable, such as when there is a cluster outage. Without the AOTM, you need to take over manually. The following options are available:

- ▶ Read-only Ownership Takeover

When Read-only Ownership Takeover is enabled for a failed cluster, ownership of a volume is taken from the failed TS7700 Cluster. Only read access to the volume is allowed through the other TS7700 clusters in the grid. After ownership for a volume has been taken in this mode, any operation that attempts to modify data on that volume or change its attributes fails. The mode for the failed cluster remains in place until another mode is selected or the failed cluster is restored.

- ▶ Write Ownership Takeover (WOT)

When WOT is enabled for a failed cluster, ownership of a volume is taken from the failed TS7700 Cluster. Full access is allowed through the requesting TS7700 Cluster in the grid, and all other available TS7700 clusters in the grid.

The automatic ownership takeover method that is used during a service outage is identical to WOT, but without the need for a person or AOTM to initiate it. The mode for the failed cluster remains in place until another mode is selected or the failed cluster has been restored.

Scratch mounts continue to prefer volumes that are owned by the available clusters. Only after all available candidates have been exhausted does it take over a scratch volume from the unavailable cluster.

You can set the level of ownership takeover, Read-only or Write, through the TS7700 MI.

In the service preparation mode of a TS7700 cluster, ownership takeover is automatically enabled, making it possible for the remaining clusters to gracefully take over volumes with full read and write access. The mode for the cluster in service remains in place until it is taken out of service mode.

Important: You cannot set a cluster in service preparation after it has already failed.

For more information about an automatic takeover, see 2.3.34, “Autonomic Ownership Takeover Manager” on page 90.

2.3.4 I/O TVC selection

All vnodes in a grid have direct access to all logical volumes in the grid. The cluster that is selected for the mount is not necessarily the cluster that is chosen for I/O TVC selection. All I/O operations that are associated with the virtual tape drive are routed to and from its vnode to the I/O TVC.

When a TVC that is different from the local TVC at the actual mount point is chosen, this is called a *remote mount*. The TVC is then accessed by the grid network. You have several ways to influence the TVC selection.

During the logical volume mount process, the best TVC for your requirements is selected, based on the following considerations:

- ▶ Availability of the cluster
- ▶ Copy Consistency policies and settings
- ▶ Scratch allocation assistance (SAA) for scratch mount processing
- ▶ DAA for specific mounts
- ▶ Override settings
- ▶ Cluster family definitions

2.3.5 Copy Consistency Points

In a multi-cluster grid configuration, several policies and settings can be used to influence the location of data copies and when the copies are run.

Consistency point management is controlled through the MC storage construct. Using the MI, you can create MCs and define where copies are placed and when they are synchronized relative to the host job that created them. Depending on your business needs for more than one copy of a logical volume, multiple MCs, each with a separate set of definitions, can be created.

The following key questions help to determine copy management in the TS7700:

- ▶ Where do you want your copies to be placed?
- ▶ When do you want your copies to become consistent with the originating data?
- ▶ Do you want logical volume copy mode retained across all grid mount points?

For different business reasons, data can be synchronously created in two places, copied immediately, or copied asynchronously. Immediate and asynchronous copies are pulled and not pushed within a grid configuration. The cluster that acts as the mount cluster informs the appropriate clusters that copies are required and the method they need to use. It is then the responsibility of the target clusters to choose an optimum source and pull the data into its disk cache.

There are currently five available consistency point settings:

- Sync** As data is written to the volume, it is compressed and then simultaneously written or duplexed to two TS7700 locations. The mount point cluster is not required to be one of the two locations. Memory buffering is used to improve the performance of writing to two locations. Any pending data that is buffered in memory is hardened to persistent storage at both locations only when an implicit or explicit sync operation occurs. This provides a zero RPO at tape sync point granularity.
- Tape workloads in z Systems environments already assume sync point hardening through explicit sync requests or during close processing, enabling this mode of replication to be performance-friendly in a tape workload environment. When sync is used, two clusters must be defined as sync points. All other clusters can be any of the remaining consistency point options, enabling more copies to be made.
- RUN** The copy occurs as part of the Rewind Unload (RUN) operation, and completes before the RUN operation at the host finishes. This mode is comparable to the immediate copy mode of the PtP VTS.
- Deferred** The copy occurs after the rewind unload operation at the host. This mode is comparable to the Deferred copy mode of the PtP VTS. This is also called *Asynchronous* replication.
- Time Delayed** The copy occurs only after a specified time (1 hour - 379 days). If the data expires before the Time Delayed setting is reached, no copy is produced at all. For Time Delayed, you can specify *after creation* or *after access* in the MC.
- No Copy** No copy is 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 can be different on each cluster 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 is accepted, unless Retain Copy mode was turned on at the MC.

For more information, see 2.3.5, “Copy Consistency Points” on page 63.

Remember: The mount point (allocated virtual device) and the actual TVC used might be in different clusters. The Copy Consistency Policy is one of the major parameters that are used to control the TVC.

2.3.6 Cluster family concept

In earlier releases, copy consistency points were the primary rules that were used to determine I/O TVC selection and how replication occurred. When two or more clusters are in proximity to each other, these behaviors were not always ideal. For example, remote clusters could be used for I/O TVC selection versus adjacent clusters, or copies to remote clusters could pass data across distant links more than once.

The concept of *families* was introduced to help with the I/O TVC selection process, and to help make distant replication more efficient. For example, two clusters are at one site, and the other two are at a remote site. When the two remote clusters need a copy of the data, cluster families enforce that only one copy of the data is sent across the long grid link.

Also, when a cluster determines where to source a volume, it gives 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 geographical proximity to optimize the use of grid bandwidth. Family members are given higher weight when determining which cluster to prefer for TVC selection.

Figure 2-12 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 prefers a cluster within its family. The example uses Copy Consistency Points of Run, Run, Deferred, Deferred [R,R,D,D].

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 2-12 illustrates the concept of cooperative replication.

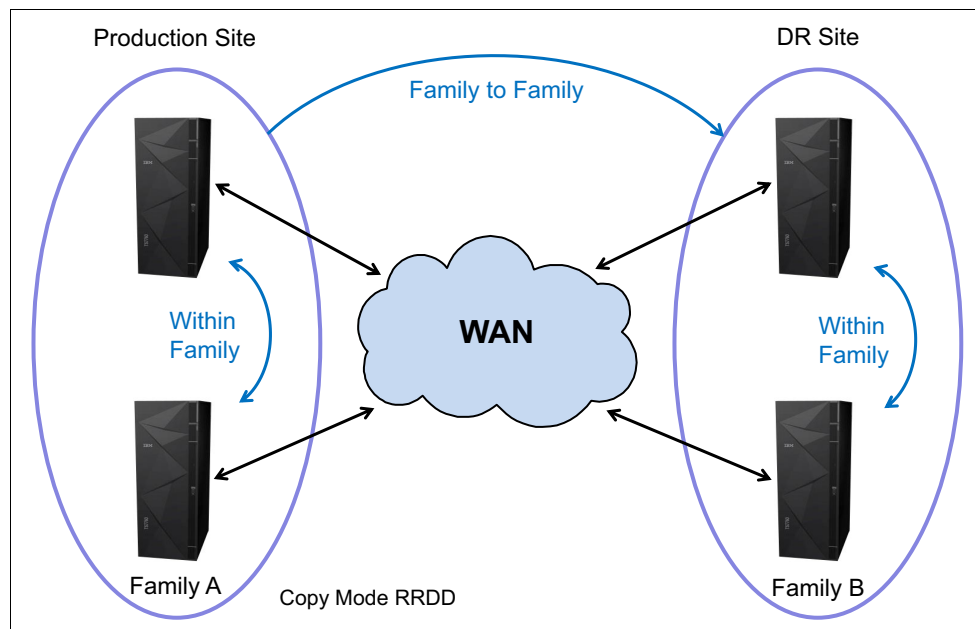


Figure 2-12 Cluster families

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 rather than 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 separate volume, this process makes a consistent copy of all volumes to the family quicker. With cooperative replication, a family prefers retrieving a new volume that the family does not have a copy of yet, over copying a volume within a family. With fewer than 20 (or the number of configured replication) tasks, copies must be sourced from outside of the family, and the family begins to replicate among itself.

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 that need a copy have received their copy. With families, a cluster's responsibility to keep the volume in cache is released after all families that need a copy have it. This process enables PG0 volumes in the source cluster to be removed from cache sooner.

Another benefit is the improved TVC selection in cluster families. For cluster families already using cooperative replication, the TVC algorithm favors using a family member as a copy source. Clusters within the same family are favored by the TVC algorithm for remote (cross) mounts. This favoritism assumes that all other conditions are equal for all the grid members.

For more information about cluster families, see *IBM Virtualization Engine TS7700 Series Best Practices -TS7700 Hybrid Grid Usage*, found at the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>

2.3.7 Override settings concept

With the prior generation of PtP VTS, several optional override settings influenced how an individual VTC selected a VTS to run the I/O operations for a mounted tape volume. In the existing VTS, the override settings were only available to an IBM SSR. With the TS7700, you define and set the optional override settings that influence the selection of the I/O TVC and replication responses by using the MI.

Note: Remember, Synchronous mode copy is not subject to copy policy override settings.

TS7700 overrides I/O TVC selection and replication response

The settings are specific to a cluster, which means that each cluster can have separate settings, if wanted. The settings take effect for any mount requests received after the settings were saved. All mounts, independent of which MC is used, use the same override settings. 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 prefers the mount point cluster as the I/O TVC for scratch mounts if it is available and contains a valid copy consistency definition other than No Copy.

- ▶ Prefer Local Cache for non-Fast Ready Mount Requests

This override prefers the mount point cluster as the I/O TVC for private mounts if it is available, contains a valid copy consistency definition other than No Copy, and contains a valid copy of the volume. If the local valid copy is only on physical tape, a recall occurs versus using a remote cache resident copy.

- ▶ Force Local TVC to have a copy of the data

The default behavior of the TS7700 is to make only a copy of the data based on the definitions of the MC associated with the volume mounted, and to select an I/O TVC that was defined to have a copy and a valid Copy Consistency Point defined. If the mount vnode is associated with a cluster for which the specified MC 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 RUN clusters are preferred. This overrides the specified MC with a Copy Consistency Point of RUN for the local cluster independent of its currently configured Copy Consistency Point. Furthermore, it requires that the local cluster is always chosen as the I/O TVC. If the mount type is private (non-Fast Ready), and a consistent copy is unavailable in the local TVC, a copy is run to the local TVC before mount completion. The copy source can be any participating TS7700 in the grid.

In a TS7740/TS7700T, the logical volume might 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 fails, even if remote TVC choices are still available when this override is enabled.

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

- ▶ Copy Count Override

This override limits the number of RUN consistency points in a multi-cluster grid that must be consistent before the surfacing device end to a RUN command. Only Copy Consistency Points of RUN are counted. For example, in a three-cluster grid, if the MC specifies Copy Consistency Points of RUN, RUN, RUN, and the override is set to two, initial status or device end is presented after at least two clusters that are configured with a RUN consistency point are consistent.

This includes the original I/O TVC if that site is also configured with a RUN consistency point. The third RUN consistency point is changed to a Deferred copy after at least two of the three RUN consistency points are consistent. The third site that has its Copy Consistency Point changed to Deferred is called the *floating deferred site*. A floating deferred site has not completed its copy when the Copy Count value is reached.

- ▶ Ignore cache preference groups for copy priority

If this option is selected, copy operations ignore the cache preference group when determining the priority of volumes that are copied to other clusters. When not set, preference group 0 volumes are preferred to enable the source cluster, which retains the volume in cache for replication purposes, to migrate the volume as quickly as possible. When set, the priority is in first-in first-out (FIFO) order.

Overrides for Geographically Dispersed Parallel Sysplex

The default behavior of the TS7700 is to follow the MC definitions and configuration characteristics to provide the best overall job performance. In certain IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®) use cases, all I/O must be local to the mount vnode. There can be other requirements, such as DR testing, where all I/O must go only to the local TVC to ensure that the correct copy policies are implemented and that data is available where required.

In these GDPS use cases, you must set the Force Local TVC override to ensure that the local TVC is selected for all I/O. This setting includes the following options:

- ▶ Prefer Local for Fast Ready Mounts
- ▶ Prefer Local for non-Fast Ready Mounts
- ▶ Force Local TVC to have a copy of the data

Consideration: Do *not* use the Copy Count Override in a GDPS environment.

2.3.8 Host view of a multi-cluster grid and Library IDs

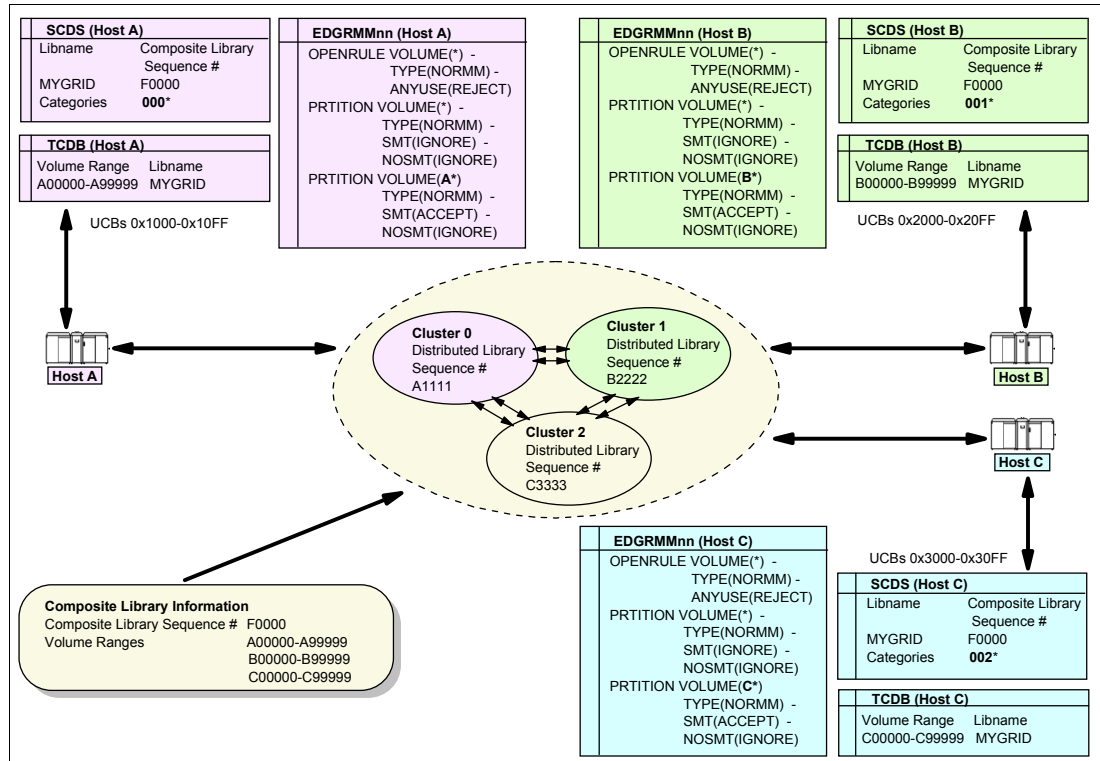
In addition to the stand-alone cluster, the grid is represented by only one composite library to the host. But each of the multiple TS7700s must have a unique distributed library defined. It is necessary to enable the host to differentiate between the entire grid versus each cluster within the grid. This differentiation is required for messages and certain commands that target the grid or clusters within the grid.

Composite library

The *composite library* is the logical image of all clusters in a multi-cluster grid, and is presented to the host as a single library. The host sees a logical tape library with up to 96 CUs in a standard six cluster grid, or up to 186 CUs if all six clusters have been upgraded to support 496 drives.

The virtual tape devices are defined for the composite library only.

Figure 2-13 illustrates the host view of a three-cluster grid configuration.



Distributed library

Each cluster in a grid is a distributed library, which consists of a TS7700. In a TS7740/TS7700T, it is also attached to a physical tape library. Each distributed library can have up to 31 3490E tape controllers per cluster. Each controller has 16 IBM 3490E tape drives, and is attached through up to four FICON channel attachments per cluster. However, the virtual drives and the virtual volumes are associated with the composite library.

There is no difference from a stand-alone definition.

2.3.9 Tape Volume Cache

In general, the same rules apply as for stand-alone clusters.

However, in a multi-cluster grid, the different TVCs from all clusters are potential candidates for containing logical volumes. The group of TVCs can act as one composite TVC to your storage cloud, which can influence the following areas:

- ▶ TVC management
- ▶ Out of cache resources conditions
- ▶ Selection of I/O cache

For more information, see 2.3.20, “General TVC management in multi-cluster grids” on page 74 and 2.3.25, “Copy Consistency Point: Copy policy modes in a multi-cluster grid” on page 80.

2.3.10 Virtual volumes and logical volumes

There is no difference between multi-cluster grids and stand-alone cluster.

Remember: Starting with V07/VEB servers and R3.0, the maximum number of supported virtual volumes is 4,000,000 virtual volumes per stand-alone cluster or multi-cluster grid. The default maximum number of supported logical volumes is still 1,000,000 per grid. Support for extra logical volumes can be added in increments of 200,000 volumes by using FC5270.

Important: All clusters in a grid must have the same quantity of installed instances of FC5270 configured. If you have configured a different number of FC5270s in clusters that are combined to a grid, the cluster with the lowest number of virtual volumes constrains all of the other clusters. Only this number of virtual volumes is then available in the grid.

2.3.11 Mounting a scratch virtual volume

In addition to the stand-alone capabilities, you can use SAA. This function widens the standard DAA support (for specific allocations) to scratch allocations, and enables you to direct a scratch mount to a set of specific candidate clusters. For more information, see “Scratch allocation assistance” on page 72.

2.3.12 Mounting a specific virtual volume

A mount for a specific volume can be sent to any device within any cluster in a grid configuration. With no additional assistance, the mount uses the TVC I/O selection process to locate a valid version of the volume.

The following scenarios are possible:

- ▶ There is a valid copy in the TVC of the cluster where the mount is placed. In this case, the mount is signaled as complete and the host can access the data immediately.
- ▶ There is no valid copy in the TVC of the cluster where the mount is placed. In this case, there are further options:
 - Another cluster has a valid copy already in cache. The virtual volume is read over the grid link from the remote cluster, which is called a *remote mount*. No physical mount occurs. In this case, the mount is signaled as complete and the host can access the data immediately. However, the data is accessed through the grid network from a different cluster.
 - No clusters have a copy in disk cache. In this case, a TS7740 or TS7700T CP1 - CP7 is chosen to recall the volume from physical tape to disk cache. Mount completion is signaled to the host system only after the entire volume is available in the TVC.
 - No copy of the logical volume can be determined in an active cluster, in cache, or on a stacked volume. The mount fails. Clusters in service preparation mode or in service mode are considered inactive.

To optimize your environment, DAA can be used. See “Device allocation assistance” on page 72.

If the virtual volume was modified during the mount operation, it is premigrated to back-end tape (if present), and has all copy policies acknowledged. The virtual volume is transferred to all defined consistency points. If you do not specify the Retain Copy Mode, the copy policies from the mount cluster are chosen at each close process.

If modification of the virtual volume did not occur when it was mounted, the TS7740/TS7720T 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 clusters are not required if modifications were not made.

The exception is if the Retain Copy policy is not set, and the MC at the mounting cluster has different consistency points defined compared to the volume’s previous mount. If the consistency points are different, the volume inherits the new consistency points and creates more copies within the grid, if needed. Existing copies are not removed if already present. Remove any non-required copies by using the **LIBRARY REQUEST REMOVE** command.

2.3.13 Logical WORM support and characteristics

There is no difference between LWORM in multicluster and stand-alone cluster environments.

2.3.14 Virtual drives

From a technical perspective, there is no difference between virtual drives in a multi-cluster grid versus a stand-alone cluster. Each cluster has 256 drives per default. See Table 2-1.

Table 2-1 Number of maximum virtual drives in a multi-cluster grid

Cluster type	Number of maximum virtual drives
Stand-alone cluster	256
Dual-cluster grid/Two-cluster grid	512
Three-cluster grid	768

Cluster type	Number of maximum virtual drives
Four-cluster grid	1024
Five-cluster grid	1280
Six-cluster grid	1536

With the new FC 5274, you can add one LCU with 16 drives up to the maximum of 496 logical drives per cluster. This results in the following maximum numbers of virtual drives. See Table 2-2.

Table 2-2 Number of maximum virtual drives in a multi-cluster grid with FC 5275 installed

Cluster type	Number of maximum virtual drives
Stand-alone cluster	496
Dual-cluster grid/Two-cluster grid	992
Three-cluster grid	1488
Four-cluster grid	1984
Five-cluster grid	2480
Six-cluster grid	2976

To support this number of virtual drives, specific authorized program analysis reports (APARs) are needed to install the appropriate program temporary fixes (PTFs) for the Preventive Service Planning (PSP) bucket.

2.3.15 Allocation assistance

Scratch and private allocations in a z/OS environment can be more efficient or more selective using the allocation assistance functions incorporated into the TS7700 and z/OS software. DAA is used to help specific allocations choose clusters in a grid that provides the most efficient path to the volume data.

DAA is enabled, by default, in all TS7700 clusters. If random allocation is preferred, it can be disabled by using the **LIBRARY REQUEST** command for each cluster. If DAA is disabled for the cluster, DAA is disabled for all attached hosts.

SAA was introduced in TS7700 R2.0, and is used to help direct new allocations to specific clusters within a multi-cluster grid. With SAA, clients identify which clusters are eligible for the scratch allocation and only those clusters are considered for the allocation request. SAA is tied to policy management, and can be tuned uniquely per defined MC.

SAA is disabled, by default, and must be enabled by using the **LIBRARY REQUEST** command before any SAA MC definition changes take effect. Also, the allocation assistance features might not be compatible with Automatic Allocation managers based on offline devices. Verify the compatibility before you introduce either DAA or SAA.

Important: Support for the allocation assistance functions (DAA and SAA) was first added to the job entry subsystem 2 (JES2) environment. Starting with z/OS V2R1, DAA and SAA are also available to JES3.

Device allocation assistance

DAA enables the host to query the TS7700 to determine which clusters are preferred for a private (specific) mount request before the actual mount is requested. DAA returns to the host a ranked list of clusters (the preferred cluster is listed first) where the mount must be run.

The selection algorithm orders the clusters in the following sequence:

1. Those clusters with the highest Copy Consistency Point
2. Those clusters that have the volume already in cache
3. Those clusters in the same cluster family
4. Those clusters that have a valid copy on tape
5. Those clusters without a valid copy

If the mount is directed to a cluster without a valid copy, a remote mount can be the result. Therefore, in special cases, even if DAA is enabled, remote mounts and recalls can still occur.

Later, host processing attempts to allocate a device from the first cluster that is returned in the list. If an online non-active device is not available within that cluster, it moves to the next cluster in the list and tries again until a device is chosen. This process enables the host to direct the mount request to the cluster that results in the fastest mount, which is 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 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. With DAA, using the copy policy overrides to *Prefer local TVC for Fast Ready mounts* provides the best overall performance. Configurations that include the TS7760 and TS7720 deep cache dramatically increase their cache hit ratio.

Without DAA, configuring the cache management of replicated data as PG1 (prefer to be kept in cache with an LRU algorithm) is the best way to improve private (non-Fast Ready) mount performance by minimizing cross-cluster mounts. However, this performance gain includes 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 might be required.

DAA (JES2) requires updates in host software (APAR OA24966 for z/OS V1R8, V1R9, and V1R10). DAA functions are included in z/OS V1R11 and later. DAA (JES3) is available starting with z/OS V2R1.

Scratch allocation assistance

With the grid configuration, using TS7760, TS7720, and TS7740 clusters is becoming more popular. There is a growing need for a method to enable z/OS to favor particular clusters over others for a workload. For example, OAM or DFSMSHsm Migration Level 2 (ML2) migration might favor a TS7760 or TS7720 with its deep cache versus an archive workload that favors a TS7740 within the same grid configuration.

SAA functions extend the capabilities of DAA to the scratch mount requests. SAA filters the list of clusters in a grid to return to the host a smaller list of candidate clusters that are designated as scratch mount candidates. By identifying a subset of clusters in the grid as sole candidates for scratch mounts, SAA optimizes scratch mounts to a TS7700 grid.

Figure 2-14 shows the process of scratch allocation.

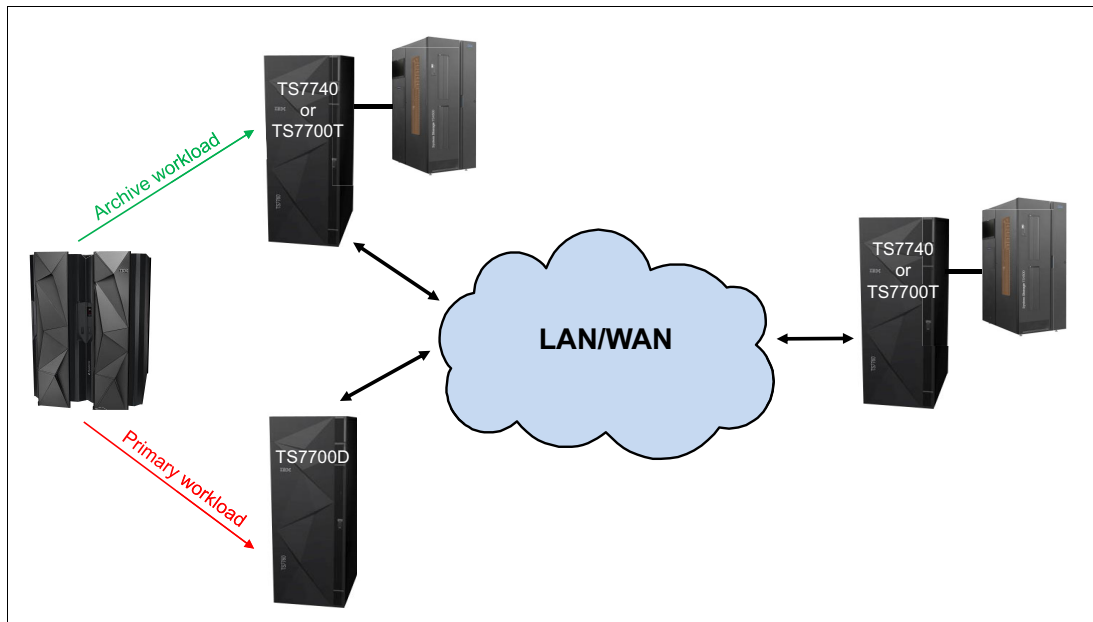


Figure 2-14 Scratch allocation direction to preferred cluster

A cluster is designated as a candidate for scratch mounts by using the Scratch Mount Candidate option on the MC construct, which is accessible from the TS7700 MI. Only those clusters that are specified through the assigned MC are considered for the scratch mount request.

When queried by the host that is preparing to issue a scratch mount, the TS7700 considers the candidate list that is associated with the MC, and considers cluster availability. The TS7700 then returns to the host a filtered, but unordered, list of candidate clusters suitable for the scratch mount operation.

The z/OS allocation process then randomly chooses a device from among those candidate clusters to receive the scratch mount. If all candidate clusters are unavailable or in service, all clusters within the grid become candidates. In addition, if the filtered list returns clusters that have no devices that are configured within z/OS, all clusters in the grid become candidates.

Be aware that SAA (and therefore this behavior) influences only the mount selection of the logical volume. If in the management class the unavailable cluster is defined as the only cluster where the data should be written to (TVC selection), the mount will be processed. However, the job is still unable to run because the selected TVC is unavailable. You will see CBR4000I and CBR4171I messages, and get a CBR4196D for a reply.

If either of the following events occurs, the mount enters the mount recovery process and does not use non-candidate cluster devices:

- ▶ All devices in the selected cluster are busy.
- ▶ Too few or no devices in the selected cluster are online.

You can use a new **LIBRARY REQUEST** option to enable or disable globally the function across the entire multi-cluster grid. Only when this option is enabled does the z/OS software run the additional routines that are needed to obtain the candidate list of mount clusters from a certain composite library. This function is disabled by default.

All clusters in the multi-cluster grid must be at release 2.0 level before SAA is operational. A supporting z/OS APAR OA32957 is required to use SAA in a JES2 environment of z/OS. Any z/OS environment with earlier code can exist, but it continues to function in the traditional way in relation to scratch allocations. SAA is also supported in a JES3 environment, starting with z/OS V2R1.

2.3.16 Selective Device Access Control

There is no difference between SDAC in multicluster and stand-alone cluster environments. However, configure SDAC so that each plex gets a portion of a cluster's devices in a multicluster configuration to achieve HA.

2.3.17 Physical drives

In a multi-cluster grid, each TS7740/TS7700T can have different drives, media types, and Licensed Internal Code levels. The TS7740/TS7700T that is used to restore the export data for merging or DR purposes must have compatible drive hardware and equal or later Licensed Internal Code than the source TS7740/TS7700T. Ensure that if you use Copy Export that the restore TS7740/TS7700T has compatible hardware and a compatible Licensed Internal Code level.

2.3.18 Stacked volume

There is no difference between stacked volume in multicluster and stand-alone environments.

2.3.19 Selective Dual Copy function

The Selective Dual Copy function is used often in stand-alone clusters. However, you can also use it in a multi-cluster grid. There is no difference in its usage in a multicluster and a stand-alone environment.

2.3.20 General TVC management in multi-cluster grids

In multicluster configurations, the TS7700 cache resources are accessible by all participating clusters in the grid. The architecture enables 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.

To use this effective cache size, you need to manage the cache content. This is done by copy policies (how many copies of the logical volume need to be provided in the grid) and the cache management and removal policy (which data to keep preferably in the TVC). If you define your copy and removal policies in a way that every cluster maintains a copy of every logical volume, the effective cache size is no larger than a single cluster.

Therefore, You can configure your grid to take advantage of removal policies and a subset of consistency points to have a much larger effective capacity without losing availability or redundancy. Any logical volume that is stacked in physical tape can be recalled into TVC, making them available to any cluster in the grid.

Replication order

Volumes that are written to an I/O TVC that is configured for PG0 have priority, based on the peer TS7700 replication priority. Therefore, copy queues within TS7700 clusters handle volumes with I/O TVC PG0 assignments before volumes configured as PG1 within the I/O TVC. This behavior is designed to enable those volumes that are marked as PG0 to be flushed from cache as quickly as possible, and not left resident for replication purposes.

This behavior overrides a pure FIFO-ordered queue. There is a new setting in the MI under Copy Policy Override, *Ignore cache Preference Groups for copy priority*, to disable this function. When selected, it causes all PG0 and PG1 volumes to be treated in FIFO order.

Tip: These settings in the Copy Policy Override window override default TS7700 behavior, and can be different for every cluster in a grid.

Treatment of data that is not yet replicated to other clusters

Logical volumes that need to be replicated to one or more peer clusters are retained in disk cache regardless of their preference group assignments. This enables peer clusters to complete the replication process without requiring a recall. After the copy completes, the assigned preference group then takes effect. For example, those assigned as preference group 0 are then immediately migrated.

If replication is not completing and the retention backlog becomes too large, the original preference groups are recognized, enabling data that is not yet replicated to be migrated to tape. These volumes likely need to be recalled into disk cache later for replication to complete. The migration of not yet replicated data might be expected when replication is not completing due to an extended outage within the grid.

2.3.21 Expired virtual volumes and the Delete Expired function

The Delete Expired function is based on the time that a volume enters the scratch category. Each cluster in a multi-cluster grid uses the same time to determine whether a volume becomes a candidate, but each cluster independently chooses from the candidate list when it deletes data. Therefore, all clusters do not necessarily delete-expire a single volume at the same time. Instead, a volume that expires is eventually deleted on all clusters within the same day.

2.3.22 TVC management for TS7740 and TS7700T CPx in a multi-cluster grid

In addition to the TVC management features from a stand-alone cluster, you can decide the following information in a multi-cluster grid:

- ▶ How copies from other clusters are treated in the cache
- ▶ How recalls are treated in the cache

Copy files preferred to reside in cache for local clusters- COPYFSC

Normally, all caches in a multi-cluster grid are managed as one composite cache. This configuration increases the likelihood that a needed volume is 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 be in the cache on that cluster. The copy that is made on the other clusters is preferred to be removed from cache.

For example, in a two-cluster grid, consider that you set up a Copy Consistency Point policy of RUN, RUN, and that the host has access to all virtual devices in the grid. After that, the

selection of virtual devices that are 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) are preferred to be in cache, and the copies are preferred to be removed from cache. The result is that each TVC is filled with unique newly created or modified volumes, roughly doubling the effective amount of cache available to host operations.

This behavior is controlled by the **LI REQ SETTING CACHE COPYFSC** option. When this option is disabled (default), logical volumes that are copied into cache from a Peer TS7700 are managed as PG0 (prefer to be removed from cache).

Copy files preferred to reside in cache for remote clusters: COPYFSC

For a multi-cluster grid that is used for DR consideration, particularly when the local clusters are used for all I/O (remote virtual devices varied offline), the default cache management method might not be wanted. If 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. Using the default setting would result in the situation, that the cache of the DR cluster is nearly empty, because all incoming logical volumes are copies and treated as PG0.

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 that are copied into the cache from a peer TS7700 are managed as PG0 (preferred to be removed from cache).
- ▶ When on, logical volumes that are copied into the cache from a peer TS7700 are managed by using the actions that are defined for the SC construct associated with the volume, as defined at the TS7700 receiving the copy.

Note: COPYFSC is a cluster-wide control. All incoming copies to that specific cluster are treated in the same way. All clusters in the grid can have different settings.

Recalls preferred for cache removal

There is no difference in a stand-alone cluster environment.

2.3.23 TVC management for TS7760 or TS7720 in a multi-cluster grid

Compared to the possibilities of TVC management from a TS7760 or TS7720 stand-alone cluster, a multi-cluster grid with TS7760/TS7720 has several options of cache management. The following options are true for TS7700D TS7700T CP0.

TS7760 and TS7720 Enhanced Removal Policies

The TS7720 Enhanced Volume Removal Policy provides tuning capabilities in grid configurations where one or more TS7760 and TS7720s are present. The tuning capabilities increase the flexibility of the subsystem effective cache in responding to changes in the host workload.

Because the TS7700D has a maximum capacity (the size of its TVC), after this cache fills, the Volume Removal Policy enables logical volumes to be automatically removed from this TS7700D TVC while a copy is retained within one or more peer clusters in the grid. When coupled with copy policies, TS7700D Enhanced Removal Policies provide various automatic data migration functions between the TS7700 clusters within the grid. This is also true for a TS7700T CP0.

In addition, when the automatic removal is run, it implies an override to the current Copy Consistency Policy in place, resulting in a lowered number of consistency points compared with the original configuration defined by the user.

When the automatic removal starts, all volumes in scratch categories are removed first, because these volumes are assumed to be unnecessary. To account for any mistake where private volumes are returned to scratch, these volumes must meet the same copy count criteria in a grid as the private volumes. The pinning option and minimum duration time criteria described next are ignored for scratch (Fast Ready) volumes.

To ensure that data will always be in a TS7700D or TS7700T CP0, or be there for at least a minimal amount of time, a *volume retention time* can be associated with each removal policy. This *volume retention time* (in hours) enables volumes to remain in a TS7720 TVC for a certain time before the volume becomes a candidate for removal. The time varies 0 - 65,536 hours. A volume retention time of zero assumes no minimal requirement.

In addition to the volume retention time, three policies are available for each volume in a TS7700D or TS7700T CP0:

► Pinned

The copy of the volume is never removed from this cluster. There is no volume retention time applicable, and is implied as infinite. After a pinned volume is moved to scratch, it becomes a priority candidate for removal similar to the next two policies. This policy must be used cautiously, to prevent cache overruns.

► Prefer Remove: When Space is Needed Group 0 (LRU)

The copy of a private volume is removed if the following conditions exist:

- An appropriate number of copies exist on peer clusters.
- The pinning duration (in number of hours) has elapsed since the last access.
- The available free space on the cluster has fallen below the removal threshold.

The order in which volumes are removed under this policy is based on their LRU access times. Volumes in Group 0 are removed before the removal of volumes in Group 1, except for any volumes in scratch categories, which are always removed first. Archive and backup data can be a good candidate for this removal group, because it is not likely accessed after it is written.

► Prefer Keep: When Space is needed Group 1 (LRU)

The copy of a private volume is removed if the following conditions exist:

- An appropriate number of copies exist on peer clusters.
- The pinning duration (in number of hours) has elapsed since the last access.
- The available free space on the cluster has fallen below removal threshold.
- Volumes with the Prefer Remove (LRU Group 0) policy have been exhausted.

The order in which volumes are removed under this policy is based on their LRU access times. Volumes in Group 0 are removed before the removal of volumes in Group 1, except for any volumes in scratch categories, which are always removed first.

Prefer Remove and Prefer Keep policies are similar to cache preference groups PG0 and PG1, except that removal treats both groups as LRU versus using their volume size. In addition to these policies, volumes that are assigned to a scratch category, and that were not previously delete-expired, are also removed from cache when the free space on a cluster falls below a threshold. Scratch category volumes, regardless of their removal policies, are always removed before any other removal candidates in descending volume size order.

Volume retention time is also ignored for scratch volumes. Only if the removal of scratch volumes does not satisfy the removal requirements are PG0 and PG1 candidates analyzed

for removal. If an appropriate number of volume copies exist elsewhere, scratch removal can occur. If one or more peer copies cannot be validated, the scratch volume is not removed.

Figure 2-15 shows a representation of the TS7720 cache removal priority.

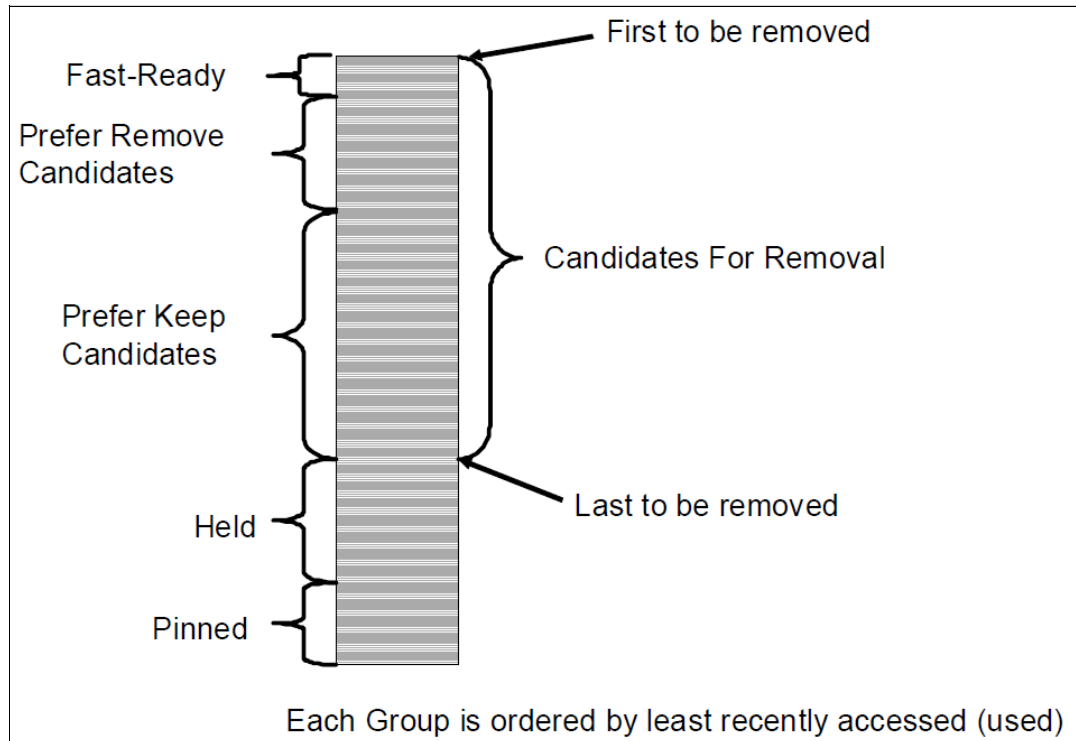


Figure 2-15 TS7720 cache removal priority

Host command-line query capabilities are supported that help override automatic removal behaviors and disable automatic removal within a TS7700D cluster, or for the CP0 in a TS7700T. For more information, see the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* on Techdocs. It is available at the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

The following host console requests are related:

```
LVOL {VOLSER} REMOVE
LVOL {VOLSER} REMOVE PROMOTE
LVOL {VOLSER} PREFER
SETTING CACHE REMOVE {DISABLE|ENABLE}
```

Delayed Replication in R3.1 changed the auto-removal algorithm so that removal of volumes where one or more delayed replication consistency points exist can take place only after those delayed replications have completed. If families are defined, only delayed consistency points within the same family must have completed.

This restriction prevents the removal of the only copy of a group before the delayed replications can complete. If no candidates are available for removal, any delayed replication tasks that have not had their grace period elapse replicate early, enabling candidates for removal to be created.

TS7700D and TS7700T CP0 cache full mount redirection

If the Enhanced Volume Removal Policies were not defined correctly or are disabled, a TS7700D TVC can become full. That is also true for a TS7700T CP0. Before becoming full, a warning message appears. Eventually, the disk cache becomes full and the library enters the Out of Cache Resources state. For multi-cluster grid configurations where more clusters are present, an Out of Cache Resources event causes mount redirection, so that an alternative TVC can be chosen.

During this degraded state, if a private volume is requested to the affected cluster, all TVC candidates are considered, even when the mount point cluster is in the Out of Cache Resources state. The grid function chooses an alternative TS7700 cluster with a valid consistency point and, if you have a TS7700D or TS7700T CP0, available cache space.

Scratch mounts that involve a TVC candidate that is Out of Cache Resources fail only if no other TS7700 cluster is eligible to be a TVC candidate. Private mounts are only directed to a TVC in an Out of Cache Resources state if there is no other eligible (TVC) candidate. When all TVCs within the grid are in the Out of Cache Resources state, private mounts are mounted with read-only access.

When all TVC candidates are either in the Paused, Out of Physical Scratch Resource, or Out of Cache Resources state, the mount process enters a queued state. The mount remains queued until the host issues a dismount command, or one of the distributed libraries exits the unwanted state. This behavior can be influenced by a new `LI REQ,distlib,SETTING,PHYSLIB` command.

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

Tip: Ensure that Removal Policies, Copy Consistency Policies, and threshold levels are applied to avoid an out-of-cache-resources situation.

Temporary removal threshold

This process is used in a TS7700 Tape Virtualization multi-cluster grid where automatic removal is enabled and a service outage is expected. Because automatic removal requires validation that one or more copies exist elsewhere within the grid, a cluster outage can prevent a successful check that leads to disk cache full conditions.

A temporary removal threshold is used to free enough cache space of the TS7700D or TS7700T CP0 cache in advance so that it does not fill up while another TS7700 cluster is in service. This temporary threshold is typically used when there are plans of taking down one TS7700 cluster for a considerable amount of time.

The process is run on the TS7700 MI.

In addition, the temporary removal threshold can also be used to free up space before a disaster recovery test with Flash Copy. During the disaster recovery test, no autoremoval or delete expire process is allowed. Therefore, you should use the temporary removal threshold to ensure that enough free space for the usual productions and the additional flash copies is available in the clusters in the DR family.

2.3.24 TVC management processes in a multi-cluster grid

The TVC management processes are the same as for stand-alone clusters. In addition to the already explained premigration management and free-space management, two further processes exist:

- ▶ Copy management (TS7740 and TS7700T CP1 - CP7)

This process applies only to a multi-cluster grid configuration, and becomes effective when the amount of non-replicated data retained in the TVC reaches a predefined threshold. It applies in particular to Deferred Copy mode, and when started reduces the incoming host data rate independently of premigration or free-space management. The purpose of this process is to prevent logical volumes from being migrated to physical tape before being copied to one or more other TS7700 clusters.

This is done to avoid a possible recall operation from being initiated by remote clusters in the grid. Only when replication target clusters are known to be unavailable, or when the amount of retained data to be copied becomes excessive, is this retained data migrated ahead of the copy process, which might lead to a future recall to complete the copy. This process is also called *copy throttling*.

- ▶ Copy time management

This process applies to multi-cluster grid configurations where the 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. If limiting the host input helps, the TS7700 enables the job to succeed before the MIH timer expires.

If limiting the host input does not help, the job changes to Deferred mode, and an alert is posted to the host console that the TS7700 has entered the Immediate-deferred state. You can modify this setting through the Host Console Request function to customize the level of throttling that is applied to the host input when this condition is detected. Because Synchronous mode copy is treated as Host I/O to the remote cluster, this is not applicable to Synchronous copies.

2.3.25 Copy Consistency Point: Copy policy modes in a multi-cluster grid

In a TS7700 Grid, you might want multiple copies of a virtual volume on separate clusters. You might also want to specify when the copies are created relative to the job that has written to a virtual volume, and that must be unique for each cluster.

Copy management is controlled through the MC storage construct. Using the MI, you can create MCs, and define where copies exist and when they are synchronized relative to the host job that created them.

When a TS7700 is included in a multi-cluster grid configuration, the MC definition window lists each cluster by its distributed library name, and enables a copy policy for each. For example, assume that three clusters are in a grid:

- ▶ LIBRARY1
- ▶ LIBRARY2
- ▶ LIBRARY3

A portion of the MC definition window includes the cluster name and enables a Copy Consistency Point to be specified for each cluster. If a copy is to exist on a cluster's TVC, you indicate a Copy Consistency Point. If you do not want a cluster to have a copy of the data, you specify the No Copy option.

As described in 2.3.5, “Copy Consistency Points” on page 63, you can either define Sync, Run, Deferred, Time Delayed, or No Copy.

Note: The default MC is *deferred* at all configured clusters, including the local. The default settings are applied whenever a new construct is defined through the MI, or to a **mount** command where MC was not previously defined.

Synchronous mode copy

To enable the synchronous mode copy (SMC), create an MC that specifies exactly two specific grid clusters with the Sync S mode.

Data is written into one TVC and simultaneously written to the secondary cluster as opposed to a RUN or DEFERRED copy where the data is not written to the cache in the I/O TVC and then read again from the cache to produce the copy. Instead, the data is written directly with a remote mount to the synchronous mode copy cluster. One or both locations can be remote.

All remote writes use memory buffering to get the most effective throughput across the grid links. Only when implicit or explicit sync operations occur does all data at both locations get flushed to persistent disk cache, providing a zero RPO of all data up to that point on tape. Mainframe tape operations do not require that each tape block is synchronized, enabling improved performance by only hardening data at critical sync points.

Applications that use data set-style stacking, and migrations, are the expected use cases for SMC. But also, any application that requires a zero RPO at sync point granularity can benefit from the Synchronous mode copy feature.

Important: The Synchronous mode copy takes precedence over any Copy Override settings.

Meeting the zero RPO objective can be a flexible requirement for certain applications and users. Therefore, a series of extra options are provided if the zero RPO cannot be achieved. For more information, see the *IBM TS7700 Series Best Practices - Synchronous Mode Copy* white paper that is available at the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102098>

Several new options are available with the synchronous mode copy. These options are described in the following sections.

Synchronous Deferred On Write Failure option

The default behavior of SMC is to fail a write operation if both clusters with the S copy mode are not available or become unavailable during the write operations.

Enable this option to enable update operations to continue to any valid consistency point in the grid. If there is a write failure, the failed S locations are set to a state of synchronous-deferred. After the volume is closed, any synchronous-deferred locations are updated to an equivalent consistency point through asynchronous replication. If the Synchronous Deferred On Write Failure option is not selected, and a write failure occurs at either of the S locations, host operations fail.

During allocation, an R or D site is chosen as the primary consistency point only when both S locations are unavailable.

Whenever a Synchronous copy enters a synchronous-deferred state, the composite library enters a Degraded state. This can be prevented by using the **LI REQ DEFDEG** option.

Open Both Copies On Private Mount option

Enable this option to open both previously written S locations when a private mount occurs. If one or both S locations are on back-end tape, the tape copies are first recalled into disk cache within those locations. The Open Both Copies On Private Mount option is useful for applications that require synchronous updates during appends. Private mounts can be affected by cache misses when this option is used. Consider these other circumstances:

- ▶ If a private mount on both locations is successfully opened, all read operations use the primary location. If any read fails, the host read also fails, and no failover to the secondary source occurs unless a z/OS dynamic device reconfiguration (DDR) swap is initiated.
- ▶ If a write operation occurs, both locations receive write data, and they must synchronize it to TVC disk during each implicit or explicit synchronization command.
- ▶ If either location fails to synchronize, the host job either fails or enters the synchronous-deferred state, depending on whether the Synchronous Deferred On Write Failure option is enabled.

Open Both Copies On z/OS implied Private Mount option

Enable this option to use the **DISP=xxxx** from the JCL to identify if a volume can be read only, or can be modified:

- ▶ If **DISP=OLD** is specified, the TS7700 assumes that only a read occurred, and opens only a single copy on a private mount.
- ▶ If **DISP=SHR** is specified, z/OS converts that to a **DISP=OLD**, because a tape does not support **DISP=SHR**. Then the mount is treated as coded with described with **DISP=OLD**.
- ▶ If **DISP=MOD** is specified, the TS7700 assumes that an append occurs, and opens the copy on both sides.

Some applications open the virtual volume with a **DISP=OLD** parameter, and still append the volume. In this case, the append is successful, and a synchronous-deferred copy is produced.

Tip: With the introduction of the new z/OS implied update option, we advise you to use this option for DFSMSshm or equivalent products.

Rewind Unload (RUN)

If a Copy Consistency Point of *RUN* is defined for a cluster in the MC that is assigned to the volume, a consistent copy of the data must exist in that cluster's TVC before command completion is indicated for the **Rewind/Unload** command.

If multiple clusters have a Copy Consistency Point of *RUN*, all of their associated TVCs must have a copy of the data before command completion is indicated for the **Rewind/Unload** command. These copies are produced in parallel. Options are available to override this requirement for performance tuning purposes.

Deferred

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.

Time Delayed Replication Policy in R3.1 or later

In the TS7700, all types of data can be stored. Some of this data usually has a short lifetime, and is replaced with other (more current) data. This is true for daily database backups, logs, and daily produced reports, such as generation data groups (GDGs), but also for other data. However, in certain conditions, this data is not replaced with more current content. Therefore, the actual logical volumes need to be treated as archive data (for example, GDGs).

Without Time Delayed Replication Policy, there was only the choice between *replication* or *nocopy* to the target clusters. Replication to the TS7700 (especially to TS7740 or TS7700T CP1 - CP7 in a multicluster hybrid grid) led to the situation that resources were being used for data that was soon to expire. No replication of this data to the TS7700 meant that if this data must be treated as archive, no additional copy could be created automatically.

Therefore, customers normally chose the *always copy* option and accepted the processor burden of replicating data that might soon expire. With the Time Delayed Replication Policy, you can now specify when the replication is done. A deferred copy will be made to all T sites after X hours have passed since volume creation or last access. The process to identify newly created T volumes runs every 15 minutes. You can specify only one T time for all Time replication target clusters in the MC. You can specify 1 - 65,535 hours.

Data already expired is still copied to the target clusters in these circumstances:

- ▶ The TMS has not yet returned the volume to scratch.
- ▶ The logical volume is scratch but not reused, and the scratch category has no expire delete definition.
- ▶ The logical volume is scratch but not reused, and the scratch category has an expire delete setting, which has not yet been reached for this specific logical volume. Ensure that the expire delete setting for the scratch category and the Time Replication time combination fit together.

Using the Time Delayed policy, the automatic removal in the TS7700D or TS7700T CP0 can be influenced. The following rules apply:

- ▶ In a grid without cluster families, all T copies need to be processed before an automatic removal can occur on any TS7700D or TS7700T CP0 in the grid.
- ▶ If cluster families are defined, all T copies in the family must be processed before auto removal to any TS7700D or TS7700T CP0 in the cluster family can occur. However, a logical volume in a TS7700D or TS7700T CP0 can be removed even if all T copies in a different family have not been processed.

A TS7700D or TS7700T CP0 might run out of removal candidates, and the only candidates in sight are those delayed replications that have not yet had their time expire. In this case, the TS7700D or TS7700T CP0 detects this condition and triggers a subset of time-delayed copies to replicate early to create removal candidates. The TS7700D and TS7700T CP0 prioritizes these copies as fast as it can replicate them. To avoid this situation, configure delay times to be early enough to provide enough removal candidates to complete production workloads.

No Copy

No copy to this cluster is performed.

For examples of how Copy Consistency Policies work in different configurations, see 2.4, “Grid configuration examples” on page 95.

A mixture of Copy Consistency Points can be defined for an MC, enabling each cluster to have a unique consistency point.

Tip: The Copy Consistency Point is considered for both scratch and specific mounts.

Management Class locality to a cluster

MCs for the TS7700 are created at the MI associated with a cluster. The same MC can be defined differently at each cluster, and there are valid reasons for doing so. For example, one of the functions that are controlled through MC is to have a logical volume copied to two separate 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 MI associated with the cluster where you want the second copy, specify a secondary pool when defining the MC. For the MC definition on the other clusters, 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 take to a DR site.

Important: During mount processing, the Copy Consistency Point information that is used for a volume is taken from the MC definition for the cluster with which the mount vnode is associated.

Define the Copy Consistency Point definitions of an MC to be the same on each cluster to avoid confusion about where copies are. You can devise a scenario in which you define separate Copy Consistency Points for the same MC 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 differs, depending on which cluster a mount is processed.

In these scenarios, use the Retain Copy mode option. When the Retain Copy mode is enabled against the currently defined MC, the previously assigned copy modes are retained independently of the current MC definition.

Retain Copy mode across grid mount points

Retain Copy mode is an optional setting in the management class where a volume's existing Copy Consistency Points are used rather than applying the Copy Consistency Points that are 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 wanted.

Figure 2-16 shows a four-cluster grid where Cluster 0 replicates to Cluster 2, and Cluster 1 replicates to Cluster 3. The wanted 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 that is written to Cluster 0. On systems where DAA is supported, DAA is used to determine which cluster is the best cluster from which to request the mount. DAA asks the grid from which cluster to allocate a virtual drive. The host then attempts to allocate a device from the best cluster (Cluster 0).




Figure 2-16 Four-cluster grid with DAA

Remember: DAA support for JES3 was added in z/OS V2R1.

On systems where DAA is not supported, 50% of the time, the host allocates to the cluster that does not have a copy in its cache. When the alternative cluster is chosen, the existing copies remain present, and more copies are made to the new Copy Consistency Points defined in the management class, resulting in more copies. If host allocation selects the cluster that does not have the volume in cache, one or two extra copies are created on Cluster 1 and Cluster 3 because the Copy Consistency Points indicate that the copies need to be made to Cluster 1 and Cluster 3.

For a read operation, four copies remain. For a write append, three copies are created. This process is illustrated in Figure 2-17.

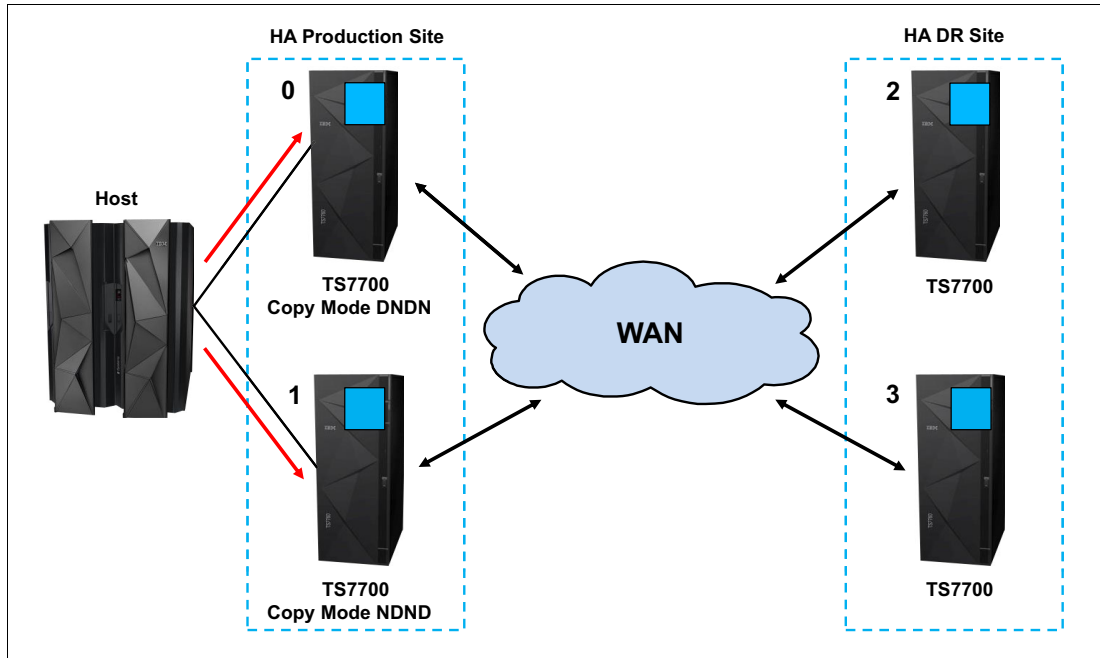


Figure 2-17 Four-cluster grid without DAA, Retain Copy mode disabled

With the Retain Copy mode option set, the original Copy Consistency Points of a volume are used rather than applying the management class with the corresponding Copy Consistency Points 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 Copy Consistency Points of the original mount are used. In this case, the result is that Cluster 0 and Cluster 2 have the copies, and Cluster 1 and Cluster 3 do not. This is shown in Figure 2-18.




Figure 2-18 Four-cluster grid without DAA, 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 exists only in Cluster 0 and Cluster 2, the mount to Cluster 1 results in a total of three or four copies. This applies to JES2 and JES3 without Retain Copy mode enabled (Figure 2-19).

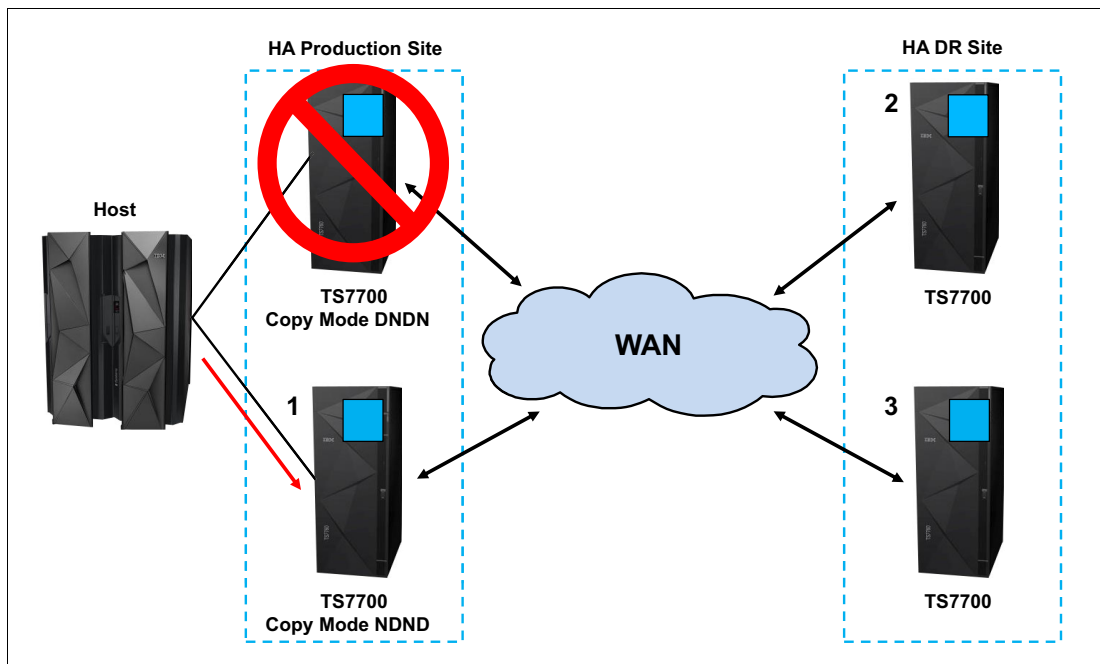


Figure 2-19 Four-cluster grid, one production cluster down, Retain Copy mode disabled

Figure 2-20 shows that the Retain Copy mode is enabled, and one of the production clusters is down. In the scenario where the cluster that contains the volume to be mounted is down, the host allocates to a device on the other cluster, in this case Cluster 1. A cross-cluster mount that uses the Cluster 2 cache occurs, and the original two copies remain. If the volume that is appended to it is changed on Cluster 2 only, Cluster 0 gets a copy of the altered volume when it rejoins the grid. Currently, only one valid copy is available in the grid.

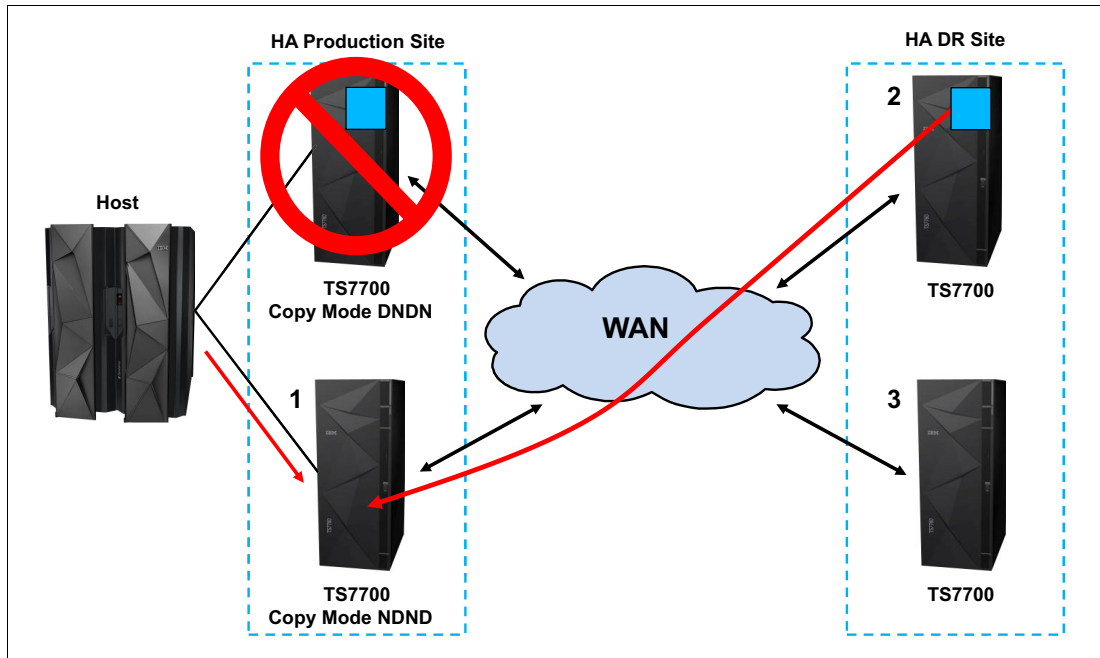


Figure 2-20 Four-cluster grid, one production cluster down, Retain Copy mode enabled

For more information, see the *IBM Virtualization Engine TS7700 Series Best Practices - TS7700 Hybrid Grid Usage* white paper at the Techdocs website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>

2.3.26 TVC (I/O) selection in a multi-cluster grid

The TVC associated with one of the clusters in the grid is selected as the I/O TVC for a specific tape mount request during mount request. The vnode is referred to as the *mount vnode*.

The TS7700 filters based on the following elements:

- ▶ Cluster availability (offline cluster, cluster in service prep, or degraded are cleared)
- ▶ Mount type:
 - Scratch. Clear all TS7700Ds and TS7700T CP0s with out of cache conditions, and remove no copy clusters.
 - Private. Clear cluster without a valid copy.
- ▶ Preferences regarding the consistency point, override policies, and families

With these three elements, an obvious favorite can be considered. If not, further filtering occurs where choices are ranked by certain performance criteria:

- ▶ Cache residency
- ▶ Recall times

- ▶ Network latency
- ▶ Host workload

The list is ordered favoring the clusters that are thought to provide the optimal performance.

With Release 3.3, two new **LI REQ** parameter settings are introduced that influence the TVC selection. You can use the **SETTING2,PHYSLIB** parameter to determine how a shortage or unavailability condition is treated in a TS7700T.

In addition, you can use the **LI REQ** parameter **LOWRANK** to give a cluster a lower ranking in the TVC selection. This parameter can be used under special conditions before you enter Service Mode. This parameter influences the TVC selection for Host I/O and the copy and mount behavior. In addition, it is a persistent setting, and can be set on every cluster independently. To avoid a negative impact to your data availability, set **LOWRANK** to the default after the maintenance is done.

For more information, see *IBM TS7700 Series z/OS Host Command Line Request User's Guide*, found on the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

2.3.27 TVC handling in an unavailability condition

In a stand-alone environment with Release 3.3, you can define that TS7700T CPx partitions react as they did in prior releases and not accept further Host I/O. You can also ignore the unavailability condition and let further Host I/O processing proceed.

In a grid, you can define that the cluster is treated as degraded, which means that this cluster has a lower priority in the TVC selection. However, all TVC selection criteria are acknowledged, and if no other cluster can fulfill the selection criteria, the degraded cluster is chosen as the TVC cluster.

In addition, you can specify whether this cluster pull s copies from other clusters.

2.3.28 Remote (cross) cluster mounts

A remote (also known as *cross*) cluster mount is created when the I/O TVC selected is not in the cluster that owns the allocated virtual device. The logical volume is accessed through the grid network by using TCP/IP. Each I/O request from the host results 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 that is transferred depends on many factors, one of which is the data compression ratio provided by the host FICON adapters. To minimize grid bandwidth requirements, only compressed data that is used or provided by the host is transferred across the network. Read-ahead and write buffering is also used to get the maximum from the remote cluster mount.

2.3.29 TVC encryption

From a technical point of view, TVC encryption is a cluster feature. Each cluster can be treated differently from the others in the multi-cluster grid. There is no difference from a stand-alone cluster.

2.3.30 Logical and stacked volume management

There is no real difference from a stand-alone environment. Each cluster is a separate entity. You can define different stacked volume pools with different rules on each distributed library.

2.3.31 Secure Data Erase

There is no difference from a stand-alone cluster.

2.3.32 Copy Export

In general, the Copy Export feature has the same functions as a stand-alone cluster. However, there are further considerations:

- ▶ The Copy Export function is supported on all configurations of the TS7740/TS7700T, including grid configurations. In a grid configuration, each TS7740/TS7700T is considered a separate source TS7740/TS7700T.

Only the physical volumes that are exported from a source TS7740/TS7700T can be used for the recovery of a source TS7740/TS7700T. Physical volumes from more than one source TS7740/TS7700T in a grid configuration cannot be combined for recovery use.

Important: Ensure that scheduled Copy Export operations are always run from the same cluster for the same recovery set. Other clusters in the grid can also initiate independent copy export operations if their exported tapes are kept independent and used for an independent restore. Exports from two different clusters in the same grid cannot be joined.

- ▶ Recovery that is run by the client is only to a stand-alone cluster configuration. After recovery, the Grid MES offering can be applied to re-create a grid configuration.
- ▶ When a Copy Export operation is initiated, only the following logical volumes are considered for export:
 - They are assigned to the secondary pool specified in the Export List File Volume.
 - They are also on a physical volume of the pool or in the cache of the TS7700 running the export operation.

For a Grid configuration, if a logical volume is to be copied to the TS7700 that will run the Copy Export operation, but that copy has not yet completed when the export is initiated, it is not included in the current export operation. Ensure that all logical volumes that need to be included have completed replication to the cluster where the export process is run.

- ▶ A service from IBM is available to merge a Copy Export set in an existing grid. Talk to your IBM SSR.

2.3.33 Encryption of physical tapes

There is no difference to a stand-alone cluster.

2.3.34 Autonomic Ownership Takeover Manager

AOTM is an optional function by which, after a TS7700 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 uses the TS3000 TSSC associated with each TS7700 in a grid to provide an alternative path to check the status of a peer TS7700. Therefore, every TS7700 in a grid must be connected to a TSSC. To take advantage of AOTM, you must provide an IP communication path between the TS3000 TSSCs at the cluster sites. *Ideally, the AOTM function uses an independent network between locations*, but this is not a requirement.

With AOTM, the user-configured takeover mode is enabled if normal communication between the clusters is disrupted, and the cluster that is running the takeover can verify that the other cluster has failed or is otherwise not operational. For more information, see 9.2.11, “The Service icon” on page 496.

When a cluster loses communication with another peer cluster, it prompts the attached local TS3000 to communicate with the remote failing cluster’s TS3000 to confirm that the remote TS7700 is down. If it is verified that the remote cluster is down, the user-configured takeover mode is automatically enabled. If it cannot validate the failure, or if the system consoles cannot communicate with each other, AOTM does not enable a takeover mode. In this scenario, ownership takeover mode can be enabled only by an operator through the MI.

Without AOTM, an operator must determine whether one of the TS7700 clusters has failed, and then enable one of the ownership takeover modes. This process is required to access the logical volumes that are owned by the failed cluster. It is important that WOT be enabled only when a cluster has failed, and not when there is only a problem with communication between the TS7700 clusters.

If ownership takeover is enabled in the read/write mode against a network-inaccessible cluster, and the inaccessible cluster is in fact handling host activity, volumes can be modified at both locations. This results in conflicting volume versions. When the Read Ownership Takeover (ROT) is enabled rather than read/write mode, the original owning cluster can continue to modify the volume, where peers have read-only access to an earlier version.

Therefore, manually enabling ownership takeover when only network issues are present should be limited to only those scenarios where host activity is not occurring to the inaccessible cluster. If two conflicting versions are created, the condition is detected when communications are resolved, and the volumes with conflicting versions are moved into an error state. When in this error state, the MI can be used to choose which version is most current.

Even if AOTM is not enabled, configure it to provide protection from a manual takeover mode being selected when the cluster is functional. This additional TS3000 TSSC path is used to determine whether an unavailable cluster is still operational or not. This path is used to prevent the user from forcing a cluster online when it must not be, or enabling a takeover mode that can result in dual volume use.

2.3.35 Selective Write Protect for disaster recovery testing

This function enables clients to emulate disaster recovery events by running test jobs at a DR location within a TS7700 grid configuration, and enabling volumes only within specific categories to be manipulated by the test application. This configuration prevents any changes to production-written data. Up to 32 categories can be identified and set to be included or excluded from Write Protect Mode by using the Category Write Protect Property table.

When a cluster is write protect-enabled, all volumes that are protected cannot be modified or have their category or storage construct names modified. As in the TS7700 write protect setting, the option is at the cluster scope and configured through the MI. Settings are persistent.

Also, the new function enables any volume that is assigned to one of the categories that are contained within the configured list to be excluded from the general cluster's write protect state. The volumes that are assigned to the excluded categories can be written to or have their attributes modified. In addition, those scratch categories that are not excluded can optionally have their Fast Ready characteristics ignored, including Delete Expire and hold processing. This enables the DR test to mount volumes as private that the production environment has since returned to scratch (they are accessed as read-only).

One exception to the write protect is those volumes in the *insert* category. To enable a volume to be moved from the insert category to a write protect-excluded category, the source category of insert cannot be write-protected. Therefore, the insert category is always a member of the excluded categories.

Be sure that you have enough scratch space when Expire Hold processing is enabled to prevent the reuse of production scratched volumes when you are planning for a DR test. Suspending the volumes' Return-to-Scratch processing during the DR test is also advisable.

Because selective write protect is a cluster-wide function, separated DR drills can be conducted simultaneously within one multi-cluster grid, if each cluster has its own independent client-configured settings.

For more information, see "Selective write protect for disaster recovery testing" on page 205.

2.3.36 FlashCopy for disaster recovery testing R3.1

This new function builds upon the TS7700s ability to provide DR testing capabilities by introducing FlashCopy consistency points within a DR location. A DR test host can use this DR family to run a DR test, while production continues on the remaining clusters of the grid.

For the DR host, the FlashCopy function provides data on a time consistent basis (Time zero). The production data continues to replicate during the entire test. The same volumes can be mounted at both sites at the same time, even with different data. To differentiate between read-only production data at time zero and fully read/write-enabled content that is created by the DR host, the selective write protect features must be used.

All access to write-protected volumes involves a snapshot from the time zero FlashCopy. Any production volumes that are not yet replicated to the DR location at the time of the snapshot cannot be accessed by the DR host, which mimics a true disaster.

Through selective write protect, a DR host can create new content to segregated volume ranges. There are 32 write exclusion categories now supported, versus the previous 16. Write protected media categories cannot be changed (by the DR host) while the Write Protection mode is enabled. This is true not only for the data, but also for the status of the volumes.

Therefore, it is not possible (by the DR host) to set production volumes from scratch to private or vice versa. When the DR site has just TS7700Ds, the flash that is initiated during the DR test is across all TS7700Ds in the DR-Family. As production returns logical volumes to scratch, deletes them, or reuses them, the DR site holds on to the old version in the flash. Therefore, return to scratch processing can now run at the production side during a test, and there is no need to defer it or use expire hold.

The TS7740 can be a part of a DR Family, but it has no FlashCopy capability itself. Therefore, the TS7740 can be used only for remote mounts from the TS7720 or TS7760. The devices of the TS7740 must not be used for mounting purposes. Enablement is done by configuring DR Families by using **LI REQ** and Write Protect or Flash by using the **LI REQ** (Library Request command) against all clusters in a DR Family.

For more information about FlashCopy setup, see Chapter 9, “Operation” on page 319. For DR testing examples, see Chapter 13, “Disaster Recovery Testing” on page 787.

The following items are extra notes for R3.1 FlashCopy for DR Testing:

- ▶ Only TS7700 Grid configurations where all clusters are running R3.1 or later, and at least one TS7720 or TS7760 cluster exists, are supported.
- ▶ Disk cache snapshot occurs to one or more TS7720 and TS7760 clusters in a DR family within seconds. TS7740 clusters do not support snapshot.
- ▶ All logical volumes in a TS7700T CP0 partition, and all logical volumes from CPx kept in cache, are part of the DR-Flash.
- ▶ If a TS7740 cluster is present within a DR family, an option is available enabling the TS7740 live copy to be accessed if it completed replication before time zero of the DR test. Although the initial snapshot itself does not require any extra space in cache, this might apply if the TS7720 or TS7760 has its live copy removed for some reason.
- ▶ Volumes in the TS7720T that are stored in CPx partitions, and that are already migrated to physical tape, are not part of the DR-Flash. They can still be accessed if the LIVECOPY Option is enabled and the logical volume was created before time zero.
- ▶ TS7720 clusters within the DR location should be increased in size to accommodate the delta space retained during the test:
 - Any volume that was deleted in production is not deleted in DR.
 - Any volume that is reused in production results in two DR copies (old at time zero and new).
- ▶ Automatic removal is disabled within TS7720 clusters during DR test, requiring a pre-removal to be completed before testing.
- ▶ **LI REQ** DR Family settings can be completed in advance, enabling a single **LI REQ** command to be run to initiate the flash and start DR testing.
- ▶ DR access introduces its own independent ownership, and enables DR read-only volumes to be mounted in parallel to the production-equivalent volumes.

The following terminology is used for FlashCopy for DR Testing:

Live copy	<p>A real-time instance of a virtual tape within a grid that can be modified and replicated to peer clusters.</p> <p>This is the live instance of a volume in a cluster that is the most current true version of the volume. It is altered by a production host, or as the content created during a DR test.</p>
FlashCopy	<p>A snapshot of a live copy at time zero. The content in the FlashCopy is fixed, and does not change even if the original copy is modified. A FlashCopy might not exist if a live volume was not present at time zero. In addition, a FlashCopy does not imply consistency, because the live copy might have been an obsolete or incomplete replication at time zero.</p>
DR family	<p>A set of TS7700 clusters (most likely those at the DR site) that serve the purpose of disaster recovery. One to five clusters can be assigned to a DR family.</p> <p>The DR family is used to determine which clusters are affected by a flash request or write-protect request by using the LI REQ (Library Request command).</p>

Write Protect Mode	When Write Protect Mode is enabled on a cluster, host commands fail if they are sent to logical devices in that cluster and attempt to modify a volume's data or attributes. The FlashCopy is created on a cluster when it is in the write protect mode only. Also, only write-protected virtual tapes are flashed. Virtual tapes that are assigned to the excluded categories are not flashed.
Time zero	The time when the FlashCopy is taken within a DR family. The time zero mimics the time when a real disaster happens. Customers can establish the time zero by using LI REQ (Library Request command).

2.3.37 Service preparation mode

The transition of a cluster into service mode is called *service prep*. Service prep enables a cluster to be gracefully and temporarily removed as an active member of the grid. The remaining sites can acquire ownership of the volumes while the site is away from the grid. If a volume owned by the service cluster is not accessed during the outage, ownership is retained by the original cluster. Operations that target the distributed library that is entering service are completed by the site going into service before the move to service completes.

Other distributed libraries within the composite library remain available. The host device addresses that are associated with the site in service send Device State Change alerts to the host, enabling those logical devices that are associated with the service preparation cluster to enter the *pending offline* state.

If a cluster enters service prep, the following copy actions are processed:

- ▶ All copies in flight (running currently), regardless of whether they are going to or from the cluster, are finished.
- ▶ No copies from other clusters to the cluster that is entering the service mode are started.
- ▶ All logical volumes that have not been copied yet, and that need to have at least one copy outside the cluster that is entering the service mode, are copied to at least one other cluster in the grid. This is true for all copies except those that are Time Delayed.
- ▶ For time delayed copies, all data that should be copied in the next 8 hours to the target clusters is copied. All other data is not copied, even if the data is only in the cluster that is entering the service mode.

When service prep completes and the cluster enters service mode, nodes at the site in service mode remain online. However, the nodes are prevented from communicating with other sites. This stoppage enables service personnel to run maintenance tasks on the site's nodes, run hardware diagnostics, and so on, without affecting other sites.

Only one service prep can occur within a composite library at a time. If a second service prep is attempted at the same time, it fails. You should put only one cluster in the service mode at the same point in time.

A site in service prep automatically cancels and reverts to an ONLINE state if any ONLINE peer in the grid experiences an unexpected outage. The last ONLINE cluster in a multicluster configuration cannot enter the service prep state. This restriction includes a stand-alone cluster. Service prep can be canceled by using the MI, or by the IBM SSR at the end of the maintenance procedure. Canceling service prep returns the subsystem to a normal state.

Important: We advise you not to place multiple clusters at the same time in service or service preparation. If you must put multiple clusters in service at once, wait for a cluster to be in final service mode before you start the service preparation for the next cluster.

If you use SAA, you might consider disabling SAA for the duration of the maintenance. This is necessary if you usually offline the drives to the z/OS systems before you enter the service preparation mode. In this case, the cluster is not yet identified as *in service*, but no devices are online to the z/OS. That would cause the job to go into device allocation recovery, if this were the only SAA candidate.

If you have multiple SAA candidates defined, you still might consider disabling SAA. This would be necessary if otherwise the amount of SAA selectable devices are not sufficient to run all of the jobs concurrently.

After SAA is enabled again, you should restart all attached OAM address spaces to ensure that the changed SAA state is recognized by the attached z/OS. If you do not restart the OAM address spaces, the system may react as if SAA is still disabled.

2.3.38 Service mode

After a cluster completes service prep and enters service mode, it remains in this state. The cluster must be explicitly taken out of service mode by the operator or the IBM SSR.

In smaller grid configurations, put only a single cluster into service at a time to retain the redundancy of the grid. This is only a suggestion, and does not prevent the action from taking place, if necessary.

If it is necessary to put multiple clusters in service mode, it is mandatory to bring them back to normal state together. In this situation a cluster cannot come back online if another cluster is still in service mode. Using the MI, you need to select each cluster independently and select **Return to normal mode**. The clusters wait until all clusters in service mode are brought back to “normal mode” before they exit the service mode.

Tip: Ensure that you can log on to the MIs of the clusters directly. A direct logon is possible, but you cannot navigate to or from other clusters in the grid when the cluster is in service mode.

2.4 Grid configuration examples

Several grid configuration examples are provided. These examples describe the requirements for high availability (HA) and DR planning.

2.4.1 Homogeneous versus hybrid grid configuration

Homogeneous configurations contain either only TS7700Ds or TS7700Ts, or only TS7740. If you have an intermix of disk-only and tape-attached models, it is a *hybrid configuration*. Consider the following information when you choose whether a TS7720, TS7740, or a mixture of the two types is appropriate.

Requirement: Fast read response times and many reads

When your environment needs to process many reads in a certain amount of time, or it needs fast response times, the TS7700Ds or TS7700Ts CP0 is the best choice. The TS7740 is susceptible to disk cache misses, resulting in a recall, making the TS7740 not optimal for workloads that need the highest cache hit read percentages.

Although TS7760 disk-only configurations can store over 1.3 PB of post-compressed content in disk cache, your capacity needs might be far too large, especially when a large portion of your workload does not demand the highest read hit percentage. This is when the introduction of a TS7700T makes sense.

Requirement: No physical tape or dark site

Some clients are looking to completely eliminate physical tape from one or more data center locations. The TS7700Ds or a hybrid configuration supports these requirements. The complete elimination of physical tape might not be the ideal configuration, because the benefits of both physical tape and deep disk cache can be achieved with hybrid configurations.

Requirement: Big data

The TS7740/TS7700T is attached to an IBM TS3500 or IBM TS4500 tape library, and can store multiple PB of data while still supporting writes at disk speeds and read hit ratios up to 90% for many workloads. Depending on the size of your tape library (the number of library frames and the capacity of the tape cartridges that are being used), you can store more than 175 PB of data without compression.

Requirement: Offsite vault of data for DR purposes with Copy Export

Some clients require an extra copy on physical tape, require a physical tape to be stored in a vault, or depend on the export of physical tape for their DR needs. For these accounts, the TS7740/TS7700T is ideal.

Requirement: Workload movement with Copy Export

In specific use cases, the data that is associated with one or more workloads must be moved from one grid configuration to another without the use of TCP/IP. Physical tape and TS7740/TS7700T Copy Export with merge (available as a service offering) provide this capability.

2.4.2 Planning for high availability or disaster recovery in limited distances

In many HA configurations, two TS7700 clusters are located within *metro distance* of each other. They are in one of the following situations:

- ▶ The same data center within the same room
- ▶ The same data center, in different rooms
- ▶ Separated data centers, on a campus
- ▶ Separated data centers, at a distance in the same metropolitan area

These clusters are connected through a local area network (LAN). If one of them becomes unavailable because it failed, is being serviced, or is being updated, data can be accessed through the other TS7700 Cluster until the unavailable cluster is available. The assumption is that continued access to data is critical, and no single point of failure, repair, or upgrade can affect the availability of data.

For these configurations, the multi-cluster grid can act as both an HA and DR configuration that assumes that all host and disk operations can recover at the metro distant location. However, metro distances might not be ideal for DR, because some disasters can affect an entire metro region. In this situation, a third location is ideal.

Configuring for high availability or metro distance

As part of planning a TS7700 Grid configuration to implement this solution, consider the following information:

- ▶ Plan for the virtual device addresses in both clusters to be configured to the local hosts.
- ▶ Plan a redundant FICON attachment of both sites (an extender that is longer than 10 kilometers (km), equivalent to 6.2 miles, for the FICON connections is suggested).
- ▶ Determine the appropriate Copy Consistency Points. For the workloads that require the highest recovery point objective (RPO), use Sync, or use RUN. For those workloads that are less critical, use deferred replication.
- ▶ Design and code the DFSMS ACS routines that point to a TS7700 MC with the appropriate Copy Consistency Point definitions.
- ▶ Ensure that the AOTM is configured for an automated logical volume ownership takeover in case a cluster becomes unexpectedly unavailable within the grid configuration. Alternatively, prepare written instructions for the operators that describe how to perform the ownership takeover manually, if needed. See 2.3.34, “Autonomic Ownership Takeover Manager” on page 90.

2.4.3 Disaster recovery capabilities in a remote data center

A mechanical problem or human error event can make the local site’s TS7700 Cluster unavailable. Therefore, one or more grid members can be introduced, separated by larger distances, to provide business continuance or DR functions.

Depending on the distance to your DR data center, consider connecting your grid members in the DR location to the host in the local site.

No FICON attachment of the remote grid members

In this case, the only connection between the local site and the DR site is the grid network. There is no host connectivity between the local hosts and the DR site’s TS7700.

FICON attachment of the remote grid members

For distances longer than 10 km (6.2 miles), you need to introduce dense wavelength division multiplexing (DWDM) or channel extension equipment. Depending on the distance (latency), there might be a difference in read or write performance compared to the virtual devices on the local TS7700 Cluster:

- ▶ The distance separating the sites can affect performance.
- ▶ If the local TS7700 Cluster becomes unavailable, use this remote access to continue your operations by using a remote TS7700 Cluster.
- ▶ If performance differences are a concern, consider using only the virtual device addresses in a remote TS7700 Cluster when the local TS7700 is unavailable. If these differences are an important consideration you need to provide operator procedures to take over ownership *and* to vary the virtual devices in a remote TS7700 from online to offline.

As part of planning a TS7700 grid configuration to implement this solution, consider the following information:

- ▶ Plan for the necessary WAN 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 SYNC or RUN, because any delays in copy time that are caused by bandwidth limitations can result in an elongation of job run times.

If you have limited bandwidth available between sites, copy critical data with a consistency point of SYNC or RUN, with the rest of the data using the *Deferred Copy Consistency Point*. Consider introducing cluster families only for three or more cluster grids.

- ▶ Depending on the distance, the latency might not support the use of RUN or SYNC at all.
- ▶ Under certain circumstances, you might consider the implementation of an IBM SAN42B-R SAN Extension Switch to gain higher throughput over large distances.
- ▶ Plan for host connectivity at your DR site with sufficient resources to run your critical workloads.
- ▶ Design and code the DFSMS ACS routines that point to the appropriate TS7700 MC constructs to control the data that gets copied, and by which Copy Consistency Point.
- ▶ Prepare procedures that your operators run when the local site becomes unusable. The procedures include several tasks, such as bringing up the DR host, varying the virtual drives online, and placing the DR TS7700 Cluster in one of the ownership takeover modes. Even if you have AOTM configured, prepare the procedure for a manual takeover.

2.4.4 Configuration examples

The various examples in this section are installed in the field, depending on the requirements of the clients. In all of these examples, you can also replace the TS7740 with a TS7720T, depending on the customer requirements.

Example 1: Two-cluster grid

With a two-cluster grid, you can configure the grid for DR, HA, or both.

This example is a two-site scenario where the sites are separated by a 10 km (6.2 miles) distance. Although the customer needs big data, and read processes are limited, two TS7760Ts were installed, one in each site. Because of the limited distance, both clusters are FICON-attached to each host.

The client chooses to use Copy Export to store a third copy of the data in an offsite vault (Figure 2-21).

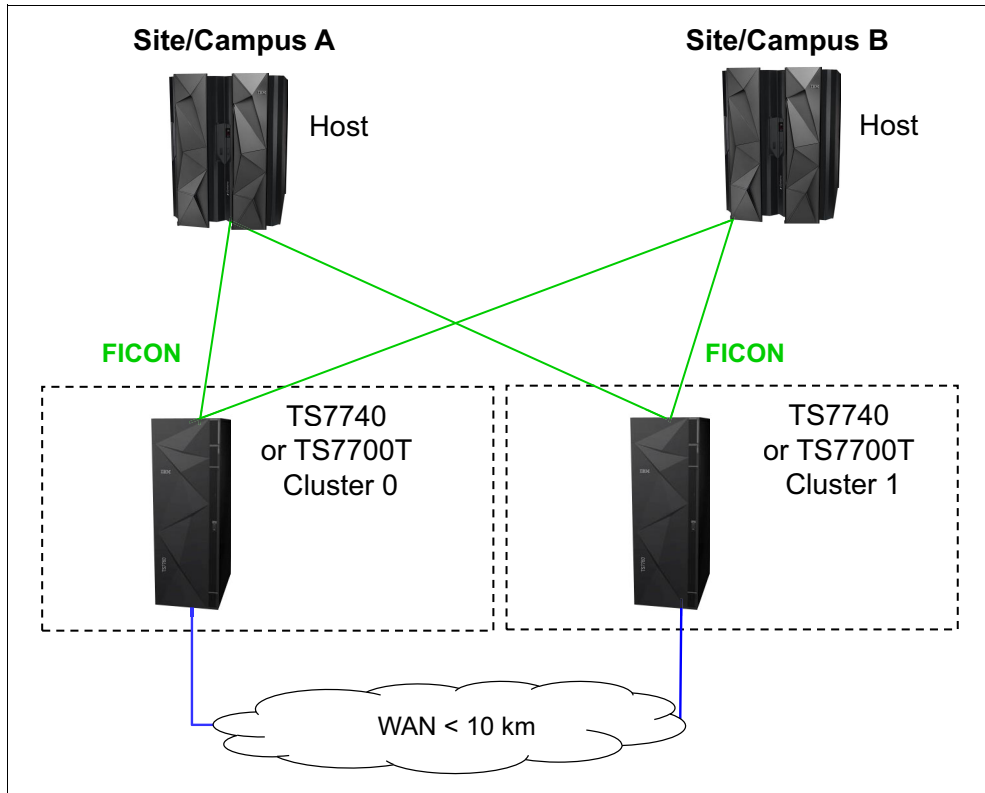


Figure 2-21 Two-cluster grid

Example 2: Three-cluster grid in two locations

In this example (Figure 2-22), one of the data center locations has several departments. The grid and the hosts are spread across the different departments. For DR purposes, the client introduced a remote site, where the third TS7740/TS7720T is installed.

The client runs many OAM and HSM workloads, so the large cache of the TS7760 provides the necessary bandwidth and response times. Also, the client wanted to have a third copy on a physical tape, which is provided by the TS776T in the remote location.

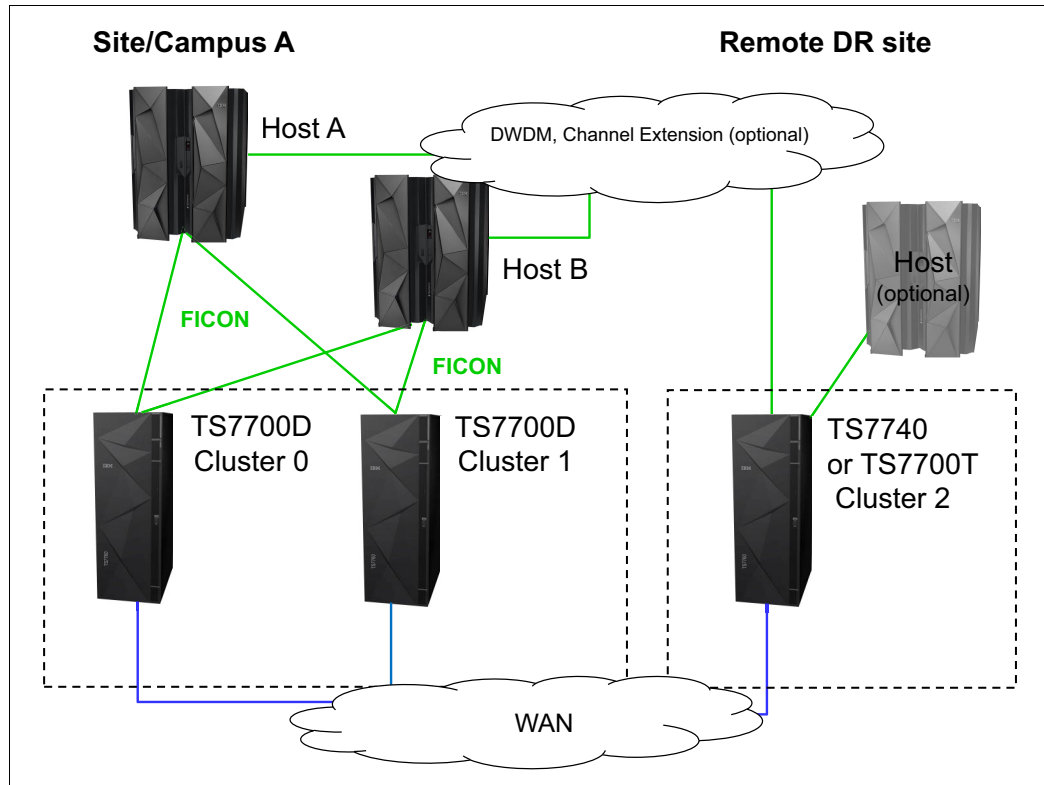


Figure 2-22 Three-cluster grid in two locations

Example 3: Three-cluster grid in three locations

This example is the same as configuration example 2. However, in this case, the two TS7760s and the attached hosts are spread across two data centers that are at a distance further than 10 km (6.2 miles). Again, the third location is a data-only store, where an existing TS7740 was used. See Figure 2-23.

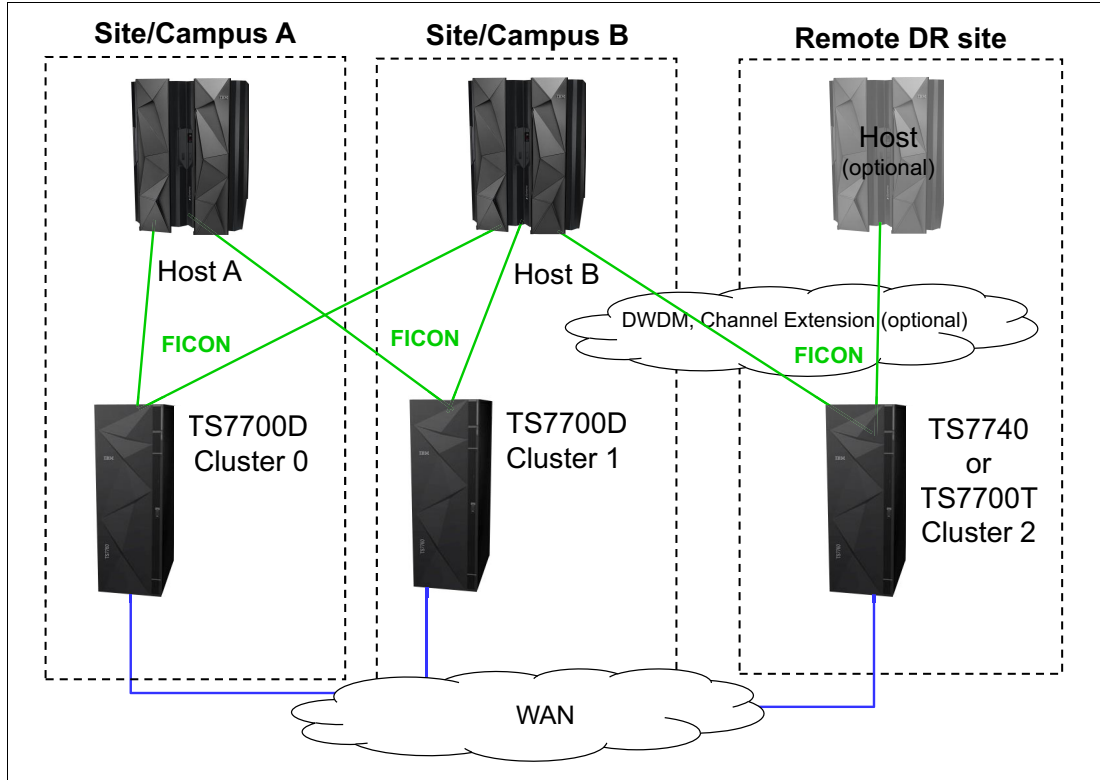


Figure 2-23 Three-cluster grid in three locations

Example 4: Four-cluster grid in three locations

The setup in Figure 2-24 shows the configuration after a merge of existing grids. Before the merge, the grids were only spread across 10 km (6.2 miles). The client's requirements changed. The client needed a third copy in a data center at a longer distance.

By merging environments, the client can address the requirements for DR and still use the existing environment.




Figure 2-24 Four-cluster grid in three locations



IBM TS7700 usage considerations

This chapter provides a general overview of the necessary information and considerations to choose the best configuration for your needs.

The IBM TS7700 offers a great variety of configuration options, features, functions, and implementation parameters. Many of the options serve different purposes, and their interactions with other options can affect how they contribute to achieving your business goals.

For some environments, these features and functions are mandatory, whereas for other environments, they only raise the level of complexity. There is no “one configuration and implementation fits all needs” solution. Therefore, you need a plan to build an environment to meet your requirements.

This chapter summarizes what to consider during the planning phase, especially for introducing new features. It also provides some general suggestions about considerations for the day-to-day operation of a TS7700 environment.

This chapter includes the following sections:

- ▶ Introduction
- ▶ Gather your business requirements
- ▶ Features and functions for all TS7700 models
- ▶ Features and functions available only for the TS7700T
- ▶ Operation aspects: Monitoring and alerting
- ▶ Choosing a migration method

3.1 Introduction

Since the first days of tape usage, the world has changed dramatically. The amount of data that is stored has increased, as have the sources of data and data legal requirements. The technical data management possibilities have grown dramatically.

In each new release of the IBM Virtual Tape Server (VTS), IBM has delivered new features to support your most demanding client needs. Consider that although some of these functions are needed in your environment, others are not:

- ▶ Some features are independent from all others, others are not.
- ▶ Certain features have a strong effect on the behavior of your environment, for example, performance or data availability.
- ▶ Specific features influence your setup of the environment.
- ▶ Some features can be overruled by Override settings.

Therefore, although these functions might be necessary to support different client requirements, they might not be required in all use cases. In fact, they might only add complexity to the solution, or they might result in unexpected behaviors. Therefore, understanding the available features, and how they complement, interact, and affect each other, helps you plan an optimal and successful implementation.

3.1.1 A short look at history

At first, data was measured in megabytes (MB). Terabytes (TB) of data were hosted only by a few mainframe clients. You always knew the location of your data. When you mounted a tape (by operator or by robot), you could be sure that your data was written directly to that specific tape. The ownership of physical tapes was clear. If you wanted to have two tapes, you needed duplex writing from the host. If you wanted to relocate specific data to a different storage location, you moved that data to a specific tape.

Your batch planners ensured that, if multifile was used, they belonged to the same application, and that the same rules (duplexing and moving) applied. Sharing resources between multiple different logical partitions (LPARs), multiple IBM z Systems operating systems, or even multiple clients was mostly not wanted or needed. You were always sure where your data on tape was located. Another important aspect was that users did not expect fast read response time for data on tape.

3.1.2 Challenges of today's businesses

The amount of data is increasing tremendously. Legal retention requirements, compliance needs, and increased redundancy for business continuance has driven much of this growth. In addition, other drivers are new sources of data:

- ▶ Email.
- ▶ Social networks and their global implications.
- ▶ Web shops record not only your actual shopping, but also your interest and buying patterns. Based on this information, you get personalized information per email or other communication paths.
- ▶ Electronic retention of paper documents.
- ▶ Digital media.

This large amount of data must all be stored and protected, and this data must be quickly and readily accessible. With more primary workloads ending on tape, response time requirements have become more demanding.

Due to cost pressures, businesses are enforcing a Tier-Technology environment. Older or not often-used data must be on less expensive storage, whereas highly accessed data must stay on primary storage, which enables fast access. Applications such as Content Manager, Hierarchical Storage Manager (HSM), or output archiver are rule-based, and are able to shift data from one storage tier to another, which can include tape. If you are considering using such applications, the tier concept needs to be planned carefully.

3.1.3 Challenges of technology progress

With advanced technology, there are challenges, such as having many new options to meet many client needs. For example, the TS7700 has many options regarding where data can be located, and where and how it must be replicated. Investing some time in choosing the correct set of rules helps you meet your requirements.

Also, the TS7700 itself decides which workloads must be prioritized. Depending on the cluster availability in the grid, actual workload, or other storage conditions, the copy queues might be delayed. In addition, the TS7700 automates many decisions to provide the most value. This dynamic behavior can sometimes result in unexpected behaviors or delays. Understanding how your environment behaves, and where your data is stored at any point in time, is key to having a successful implementation, including the following considerations:

- ▶ During a mount, a remote Tape Volume Cache (TVC) was chosen over a local TVC.
- ▶ Copies are intentionally delayed due to configuration parameters, yet they were expected to complete sooner.
- ▶ Copy Export sets do not include all of the expected content because the export was initiated from a cluster that was not configured to receive a replica of all the content.

A reaction might be to configure your environment to define synchronous and immediate copies to all locations, or to set all overrides. This likely increases the configuration capacity and bandwidth needs, which can introduce negative results. Planning and working with your IBM account team so that the optimal settings can be configured helps eliminate any unexpected behaviors.

Other features, such as scratch allocation assistance (SAA) and device allocation assistance (DAA), might affect your methodology of drive allocation, whereas some customizing parameters must always be used if you are a Geographically Dispersed Parallel Sysplex (GDPS) user.

So, it is essential for you to understand these mechanisms to choose the best configuration and customize your environment. You need to understand the interactions and dependencies to plan for a successful implementation.

Important: There is no “one solution that fits all requirements.” Do not introduce complexity when it is not required. Allow IBM to help you look at your data profile and requirements so that the best solution can be implemented for you.

3.2 Gather your business requirements

There are several different types of business requirements that you need to consider.

3.2.1 Requirement types

Consider the following lists as a starting point.

Requirements from the data owners, application administrators, and the applications

The following items should be considered when you gather data and application requirements:

- ▶ How important is the data? Consider multiple copies, Copy Consistency Points, retention requirements, and business recovery time expectations.
- ▶ How often will the data be accessed, and what retrieval times are expected? Consider sizing and performance.
- ▶ How will the application react if the tape environment is not available? Consider high availability (HA) and disaster recovery (DR) planning and copy consistency.
- ▶ How will the application react if specific data is not available? Consider HA and DR planning and copy consistency.
- ▶ How much storage for the data is needed? Factor in any future growth.
- ▶ What are the performance expectations during an outage or disaster event?

It can be difficult to get all of the required information from the owners of the data and the owners of the applications to best manage the data. Using service level agreement (SLA) requirements and an analysis of your existing tape environment can help with the process.

Requirements from the IT department

The following items should be considered when you gather information technology (IT) requirements:

- ▶ Support of the general IT strategy (data center strategy and DR site support)
- ▶ Sharing of a TS7700 environment between multiple LPARs or sysplexes (definition of logical pools, physical pools, and performance)
- ▶ Sharing of a TS7700 in a multi-tenancy environment (logical pools, physical pools, Selective Device Access Control (SDAC), export and migration capabilities, and performance)
- ▶ Support of zAutomation concepts (monitoring and validation)
- ▶ Environmental requirements (cooling and space)
- ▶ Financial requirements
- ▶ Multiple platforms required (z Systems operating systems)
- ▶ Monitoring and automation capabilities to identify issues and degradations
- ▶ Floor space requirements
- ▶ Network infrastructure
- ▶ Power availability

Depending on your overall IT strategy, application requirements and data owner requirements can be used to select an appropriate TS7700 configuration. If you have multiple data centers, spread your clusters across the data centers, and ensure that copies of the data are in each data center.

If your approach is that each data center can host the total workload, plan your environment accordingly. Consider the possible outage scenarios, and verify whether any potential degradations for certain configurations can be tolerated by the business until the full equipment is available again.

In a two-cluster environment, there is always a trade-off between availability and a nonzero point of recovery. Assume that data protection is the highest priority within your workload. During a planned outage, new or modified workloads do not have a redundant copy that is generated, which might be unacceptable. Putting new allocations on hold during this period might be optimal. If availability is rated higher, you might want to take the risk of a single copy during an outage so that operations can continue.

However, more advanced TS7700 configurations can be implemented that enable both availability and data protection to be equally important, for example, a four cluster grid. Consider what type of data you store in your TS7700 environment. Depending on your type of data, you have multiple configuration choices. This section starts with a general view before looking closer at the specific types of data.

3.2.2 Environment: Source of data

Depending on the method of creating data, you might have different requirements. Assume that you have all four types of systems to create data:

- ▶ Sandbox system: Used to verify new operating and subsystem versions
- ▶ Development system: Used to develop new applications
- ▶ User Acceptance Test (UAT) system: Used for integration and performance testing
- ▶ Production system

Consider the following guidelines:

- ▶ Data from a sandbox system (regardless of whether it is backup or active data) might not need multiple copies because you can re-create the information from other sources (new installation, and so on).
- ▶ Application data from a development system might not need multiple copies in different storage pools or data centers because the data can be re-created from production systems.
- ▶ Application code from a development system likely needs multiple copies because that data might not be re-created from elsewhere.
- ▶ If physical tape is present, have UAT-created content migrate to physical tape so that precious disk cache space is not used.
- ▶ Not all production or backup workloads that target the TS7700 might be replicated. Perhaps, you have those workloads managed differently for DR needs, or you do not need that workload in a DR event. These non-replicated workloads can optionally be Copy Exported as a DR alternative if replication is not feasible.

Data from your sandbox, test, UAT, or production system might share the tape environment, but it can be treated differently. That is important for sizing, upgrades, and performance considerations as well.

Note: Plan your environments and the general rules for different types of environments. Understand the amount of data that these environments host today.

3.2.3 Backup data, active data, and archive data

In general, data from different applications has different requirements for your tape environment. Your tape processing environment can be all virtual, all physical, or a combination of the two.

Backup data

The data on tape is only a backup. Under normal conditions, it will not be accessed again. It might be accessed again only if there are problems, such as direct access storage device (DASD) hardware problems, logical database failures, and site outages.

Expiration

The expiration is mostly a short or medium time frame.

Availability requirements

If tape environment is not available for a short time, the application workload can still run without any effect. When the solution is unavailable, the backup to tape cannot be processed.

Retrieval requirements

Physical tape recall can normally be tolerated, or at least for previous generations of the backup.

Multiple copies

Depending on your overall environment, a single copy (not in the same place as the primary data) might be acceptable, perhaps on physical tape. However, physical media might fail or a storage solution or its site might experience an outage. Therefore, one or more copies are likely needed. These copies might exist on more media within the same location or ideally at a distance from the initial copy.

If you use multiple copies, a Copy Consistency Point of *Deferred* might suffice, depending on your requirements.

Active data on tape

The data is stored only on tape. This data is not also somewhere in DASD. If the data needs to be accessed, it is read from the tape environment.

Expiration

The expiration depends on your application.

Availability requirements

When the tape environment is not available, your original workload might be severely affected.

Retrieval requirements

Physical tape recalls might not be tolerated, depending on your data source (sandbox, test, or production) or the type of application. Older, less-accessed active data might tolerate physical tape recalls.

Multiple copies

Although tape is the primary source, a single copy is not suggested. Even a media failure can result in data loss. Multiple copies should be stored in different locations to be prepared for a data center loss or outage. In a stand-alone environment, dual copies on physical tape are suggested.

Depending on the recovery point objectives (RPO) of the data, choose an appropriate Consistency Point Policy. For example, synchronous mode replication is a good choice for these workloads because it can achieve a “zero point RPO at sync point” granularity.

Especially for DFSMSHsm ML2 and OAM objects, use the synchronous mode copy.

Archive data on tape

Archive data on tape is also active data. However, archive data is stored for a long time. Expiration dates for 10 - 30 years to satisfy regulatory requirements are common. Sometimes, logical Write Once Read Many (LWORM) data is required.

Expiration

The expiration depends on your application, but it is usually many years.

Availability requirements

Archive data is seldom accessed for read. If the tape environment is not available, your original workload might still be affected because you cannot write new archive data.

Retrieval requirements

Physical tape recalls might be tolerated.

Multiple copies

Although the tape is the primary source, a single copy is not suggested. Even a media failure results in data loss. Store multiple copies in different locations to be prepared for a data center loss. In a standalone environment, dual copies on physical tape are suggested.

Depending on the criticality of the data, choose an appropriate Copy Consistency Point Policy.

Archive data sometimes must be kept for 10 - 30 years. During such long time periods, the technology progresses, and data migration to newer technology might need to take place. If your archive data is on physical tapes in a TS7740/TS7700T, you must also consider the life span of physical tape cartridges. Some vendors suggest that you replace their cartridges every five years, other vendors, such as IBM, offer tape cartridges that have longer lifetimes.

If you are using a TS7740/TS7700T and you store archive data in the same storage pools with normal data, there is a slight chance that, due to the reclaim process, the number of stacked volumes that contain only archive data will increase. In this case, these cartridges might not be used (either for cartridge reads or reclaim processing) for a longer time. Media failures might not be detected. If you have more than one copy of the data, the data can still be accessed. However, you have no direct control over where this data is stored on the stacked volume, and the same condition might occur in other clusters, too.

Therefore, consider storing data with such long expiration dates on a specific stacked volume pool. Then, you can plan regular migrations (even in a 5 - 10-year algorithm) to another stacked volume pool. You might also decide to store this data in the common data pool.

3.2.4 IBM DB2 archive log handling

With IBM DB2, you have many choices about how to handle your DB2 archive logs. You can put both of them to DASD and maybe rely on a later migration to tape through DFSMSHsm or an equivalent application. You can write one archive log to DASD and another one to tape. Alternatively, you can put them both directly to tape.

Depending on your choice, the tape environment is more or less critical to your DB2 application. This depends also on the number of active DB2 logs that you define in your DB2 environment. In some environments, due to peak workload, logs are switched every two minutes. If all DB2 active logs are used and they cannot be archived to tape, DB2 stops processing.

Scenario

You have a four-cluster grid, spread over two sites. A TS7760D and a TS7760T are at each site. You store one DB2 archive log directly on tape and the other archive log on disk. Your requirement is to have two copies on tape:

- ▶ Using the TS7760 can improve your recovery (no recalls from physical tape needed).
- ▶ Having a consistency point of R, N, R, N provides two copies, which are stored in both TS7760s. If one TS7760 is available, DB2 archive logs can be stored to tape. However, if one TS7760 is not available, you have only one copy of the data. In a DR situation where one of the sites is not usable for a long time, you might want to change your policies to replicate this workload to the local TS7760T as well.
- ▶ If the TS7760D enters the Out of cache resources state, new data and replications to that cluster are put on hold. To avoid this situation, consider having this workload also target the TS7760T and enable the Automatic Removal policy to free space in the TS7760D. Until the Out of cache resources state is resolved, you might have fewer copies than expected within the grid.
- ▶ If one TS7760D is not available, all mounts must be run on the other TS7760D cluster.
- ▶ In the unlikely event that both TS7760Ds are not reachable, DB2 stops working as soon as all DB2 logs on the disk are used.
- ▶ Having a consistency point of R, N, R, D provides you with three copies, which are stored in both TS7760Ds and in the TS7760T of the second location. That exceeds your original requirement, but in an outage of any component, you still have two copies. In a loss of the primary site, you do not need to change your DB2 settings because two copies are still written. In an Out of Cache resources condition, the TS7760D can remove the data from cache because there is still an available copy in the TS7760T.

Note: Any application with the same behavior can be treated similarly.

3.2.5 DFSMSHsm Migration Level 2

Several products are available on the z Systems platform for hierarchical storage management (HSM). IBM Data Facility Storage Management Subsystem Hierarchical Storage Manager (DFSMSHsm) provides different functions. DFSMSHsm migrates active data from disk pools to ML2 tape in which the only copies of these data sets are on tape.

To achieve fast recall times to DASD, you should consider storing the data in a TS7760D or a TS7760T CP0 at least for a certain period. With time-delay copies to additional tape-attached (TS7760T, TS7720T, or TS7740) clusters in a grid, you can ensure that the data is kept first in

the TS7760D, and later copied to a cluster with tape attachment. Auto removal processing (if enabled) can then remove the content from the TS7760 as it ages.

Ideally, DFSMSHsm ML2 workloads should be created with synchronous mode copy to ensure that a data set is copied to a second cluster before the DFSMSHsm migration processes the next data set. The DFSMSHsm application marks the data set candidates for deletion in DASD. With z/OS 2.1, MIGRATION SUBTASKING enables DFSMSHsm to offload more than one data set at a time, so it can do batches of data sets per sync point.

Using TS7700 replication mechanisms rather than DFSMSHsm local duplexing can save input/output (I/O) bandwidth, improve performance, reduce the number of logical volumes that are used, and also reduces the complexity of bringing up operations at a secondary location.

Other vendor applications might support similar processing. Contact your vendor for more information.

Tip: To gain an extra level of data protection, run ML2 migration only after a DFSMSHsm backup runs.

3.2.6 Object access method: Object processing

You can use the object access method (OAM) to store and transition object data in a storage hierarchy. Objects can be on disk (in DB2 tables, the IBM z Systems file system (zFS), or Network File System (NFS) mountable file systems), optical, and tape storage devices. You can choose how long an object is stored on disk before it is migrated to tape. If the object is moved to tape, this is active data.

Users accessing the data on tape (in particular the TS7700T or TS7740) might have to wait for their document until it is read from physical media. The TS7700D or the TS7700T CP0 is traditionally a better option for such a workload given the disk cache residency can be much longer and even indefinite.

For OAM primary objects, use Synchronous mode copy on two clusters and depending on your requirements, additional immediate or deferred copies elsewhere if needed.

With OAM, you can also have up to two backup copies of your object data. Backup copies of your data (managed by OAM) are in addition to any replicated copies of your primary data that are managed by the TS7700. Determine the copy policies for your primary data and any additional OAM backup copies that might be needed. The backups that are maintained by OAM are only used if the primary object is not accessible. The backup copies can be on physical or virtual tape.

3.2.7 Batch processing: Active data

If you create data in the batch process, which is not stored on disk, it is also considered active data. The access requirements of these data types can determine whether the data should be placed on a TS7700D or TS7700T CP0, a TS7700T, CPx, a TS7740, or a mix of them. For example, active data that needs quick access is ideal for a TS7700D or the resident partition of a TS7700T.

Depending on your environment, you also can place this data on a tape partition of a TS7700T and ensure that the data is kept in cache. This can be done either by delaying premigration or defining the size of the tape partition to keep all of this data in cache.

Rarely accessed data that does not demand quick access times can be put on a TS7700T tape partition with PG0 and a not-delayed migration, or on the TS7700T in a second cluster.

Data that becomes less important with time can also use the TS7700D or TS7700T CP0 auto-removal policies to benefit from both technologies.

Assume that you have the same configuration as the DB2 archive log example:

- ▶ With a Consistency Copy Point policy of [N,R,N,R], your data is stored only on the TS7700T CPx or TS7740s (fast access is not critical).
- ▶ With a Consistency Copy Point policy of [R,N,R,N], your data is stored only on the TS7700Ds (fast access is critical).
- ▶ With a Consistency Copy Point policy of [R,D,R,D], your copy is on the TS7700Ds first and then also on the TS7700Ts, enabling the older data to age off the TS7700Ds by using the auto-removal policy.

3.2.8 Data type and cache control

You can differentiate the type of data held on a TS7700, as shown in Table 3-1.

Table 3-1 Type of data

Type of Data	Application Examples	Fits best on	Suitable cache Control
Data needs a 100% cache hit	OAM objects (primary data), HSM ML2	TS7700 Disk-Only TS7700T CP0	Pinned / PG1 CP0/ Pinned/ PG1
Data that benefits from a longer period in disk cache	Depending on the user requirements. OAM objects (primary data), HSM ML2	TS7700D with autoremoval TS7700T CPx TS7740 if duration in disk cache can be minimal	PG1 PG1 PG1
Data that is needed for a specific time in cache, but then should be kept on tape	DB2 log files (depending on your requirements), Active Batch data	TS7700T TS7740	CPx / PG1 delay premigration PG1
Data with limited likelihood to be read, only cache pass through	Backups, Dumps	TS7700T TS7740	CPx / PG0 PG0

3.3 Features and functions for all TS7700 models

Based on the gathered requirements, you can now decide which features and functions you want to use in your TS7700 environment.

3.3.1 Stand alone versus grid environments

Consider a stand-alone cluster in the following conditions:

- ▶ You do not need a high availability or an electronic DR solution.
- ▶ You can handle the effect to your application in a cluster outage.
- ▶ In a data center loss, a data loss is tolerable or a recovery from Copy Export tapes is feasible (time and DR site).
- ▶ You can plan outages for Licensed Internal Code loads or upgrade reasons.

If you cannot tolerate any of these items, consider implementing a grid environment.

3.3.2 Sharing a TS7700

Sharing TS7700 resources is supported in most use cases. Whether the environment includes different applications within a common sysplex, independent sysplexes, or z Systems operating systems, the TS7700 can be configured to provide shared access. The TS7700 can also be shared between multiple tenants.

Because the TS7700 is policy-managed, each independent workload can be treated differently depending on how the data is managed within the TS7700. For example, different workloads can be on different tape partitions in a TS7700T and use independent physical volume pools within a TS7740 or TS7700T. Alternatively, different workloads can use different replication requirements.

All applications within a Parallel Sysplex can use the same logical device ranges and logical volume pools, simplifying sharing resources. When independent sysplexes are involved, device ranges and volume ranges are normally independent, but are still allowed to share the disk cache and physical tape resources.

Of all the sharing use cases, most share the FICON channels into the TS7700. Although the channels can also be physically partitioned, it is not necessary because each FICON channel has access to all device and volume ranges within the TS7700.

However, there are still considerations:

- ▶ The TVC is used in common in a TS7700D, a TS7700T CP0, or a TS7740. You cannot define a physical limit to the amount of space a client is using in the TVC. However, through policy management, you can use preference groups differently in these models or the removal policies can be configured differently, giving more TVC priority to some workloads over others. In a TS7720T, you can specify multiple tape partitions to enable that each tenant get its own dedicated disk cache residency.
- ▶ Define the scratch categories that the different systems use. The scratch categories are specified in the DEVSUPxx parmlib member.
- ▶ Decide which VOLSER ranges the different systems use. This is typically handled through the tape management system (TMS). For DFSMSrmm, this is handled through their PARTITION and OPENRULE parameters.
- ▶ Another main item to consider is how the drives are managed across the different systems, and which systems share which drives. This is typically handled through a tape device sharing product.

- ▶ Storage management subsystem (SMS) constructs and constructs on the TS7700 must match. If not, new constructs in SMS lead to new constructs in the TS7700 that are created with default parameters. To avoid the uncontrolled buildup of constructs in the TS7700, SMS should be controlled by a single department.
- ▶ SMS constructs used by different workloads need to use unique names when the TS7700 behavior is expected to be different. This enables each unique application's behavior to be tuned within the TS7700. If the behavior is common across all shared workloads, the same construct names can be used.
- ▶ Ensure that the single defined set of constructs within the TS7700 are configured with a behavior that is acceptable to all users. If not, different constructs must be used for those customers.
- ▶ Control of the TS7700 Management Interfaces (MIs), TS3500 GUI, and TS4500 GUI must be allowed only to a single department that controls the entire environment. Control must not be given to a single customer.
- ▶ Review the IBM RACF statements for the **Devserv** and **Library** commands on all LPARs. These commands must be protected. In a multiple-client environment, the use of Library commands must be restricted.

When independent sysplexes are involved, the device ranges and corresponding volume ranges can be further protected from cross-sysplex access through the SDAC feature.

When device partitioning is used, consider assigning the same number of devices per cluster per sysplex in a grid configuration so that the availability for a given sysplex is equal across all connected clusters.

Override policies set in the TS7700 apply to the whole environment and cannot be enabled or disabled by an LPAR or client.

For more considerations, see the *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409.

Note: Some parameters can be updated by the **Library Request** command. This command changes the cluster behavior. This is not only valid for the LPAR where the command was run, but for all LPARs that use this cluster.

Ensure that only authorized personnel can use the **Library Request** command.

If you share a library for multiple customers, establish regular performance and resource usage monitoring. See 3.4, "Features and functions available only for the TS7700T" on page 119.

Note: With APAR OA49373 (z/OS V2R1 and above), the individual IBM MVS™ LIBRARY command functions (EJECT, REQUEST, DISPDRV, and so on) can be protected using a security product such as RACF. This APAR adds security product resource-names for each of the LIBRARY functions.

3.3.3 Tape Volume Cache selection

Depending on your Copy Consistency Policies, the cluster where the virtual tape mount occurred is not necessarily the cluster that is selected as the TVC. When a TVC other than the local TVC is chosen, this is referred to as a *remote mount*. Plan the Copy Consistency Policy so that you are aware where your data is at any point in time.

TVC selection is also influenced by some LI REQ parameters. For more information about the parameters **LOWRANK** and **SETTINGS,PHYSLIB**, see the *Library Request Command* white paper, found at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

TVC selection might also influence the Copy Export. See 12.1, “Copy Export overview and considerations” on page 758.

3.3.4 Copy Consistency policy

Define the consistency policy for each Management Class (MC). For more information, see 2.3.5, “Copy Consistency Points” on page 63.

The following list describes several general considerations:

- ▶ When a cluster is assigned a policy of N, this cluster is not the target of a replication activity:
 - This cluster cannot be chosen as the TVC (it can be chosen as the mount point).
 - If only N clusters are available, any mount that uses that MC fails.
 - If **Force local copy override** was selected, this makes the local mount an R if previously an N.
- ▶ A consistency point of [D, D, D, D] means that the selected TVC is treated as RUN, and the additional copies are created asynchronously. For a scratch selection, the mount point cluster is normally chosen as the TVC, although it is not required. Copy Override settings can be used to prefer that it also acts as the TVC.
- ▶ A consistency point of [D, R, D, D] means that Cluster 1 is preferred as the TVC, even if the other cluster is chosen as the mount point. Therefore, the ‘R’ location is preferred, which can result in a remote mount when the mount point is not the same as the ‘R’ location. This can be done intentionally to create a remote version as the initial instance of a volume.

If you do not care which TVC is chosen and you prefer a balanced grid, use [D, D, D, D].

With the new Time Delayed Replication policy, you can now decide that certain data is only copied to other clusters after a specified time. This policy is designed for data that usually has a short lifecycle, and is replaced shortly with more current data, such as backups and generation data groups (GDGs). In addition, this policy can also be used for data with an unknown retention time, and where the data should be copied only to another cluster when this data is still valid after the given time. Time Delayed Replication policy is targeted especially for multi-cluster grids (3 or more).

With the usage of the consistency copy policies before R3.1, this data was either never replicated or always replicated to the specified clusters. In case the target was a TS7740, the data was copied to the backend tape and can result in excessive reclamation when expired early. Now you can specify that this type of data is only replicated if the data is still valid and the specified time (after creation or last access) has expired. That might reduce replication traffic, and the backend activities in TS7740s.

However, plan to have at least two copies for redundancy purposes, such as on a local TS7700D/TS7700T and a remote TS7700D/TS7700T.

3.3.5 Synchronous mode copy

Synchronous mode copy creates a copy of the data whenever an explicit or implicit sync point is written from an application. This enables a much more granular copy than all other consistency points, such as Run or Deferred.

This consistency point is ideal for applications that move primary data to tape, such as DFSMSshm or OAM Object Support, which can remove the primary instance in DASD after issuing an explicit sync point.

Therefore, you should use Synchronous mode copy for this type of applications.

The synchronous mode copy offers three options for how to handle private mounts:

- ▶ Always open both instances on private mount.
- ▶ Open only one instance on private mount.
- ▶ Open both instances on z/OS implied update.

Plan the usage of this option. Dual open is necessary for workloads that can append to existing tapes. When only reads are taking place, the dual open can introduce unnecessary resource use, especially when one of the instances requires a recall from a physical tape. Using the dual open z/OS implied update helps reduce any resource use to only those mounts where an update is likely to take place.

In addition, synchronous mode copy provides an option to determine its behavior when both instances cannot be kept in sync. One option is to move to the synch-deferred state. Another option is to fail future write operations. Depending on your requirements, determine whether continued processing is more important than creating synchronous redundancy of the workload. For more information, see 2.3.5, “Copy Consistency Points” on page 63.

3.3.6 Override policies

Override policies overrule the explicit definitions of Copy policies.

Note: Synchronous mode is not subject to override policies.

The Override policies are cluster-based. They cannot be influenced by the attached hosts or policies. With Override policies, you can help influence the behavior on how the TS7000 cluster chooses a TVC selection during the mount operation, and whether a copy needs to be present in that cluster (for example, favoring the local mount point cluster).

Copy Count Override enables the client to define for this cluster that at least two or more copies exist at RUN time, but the client does not care which clusters have a copy. If you use Copy Count Override, the grid configuration and available bandwidth between locations likely determines which RUN copies meet the count criteria. Therefore, the limited numbers of copies can be within the closest locations versus at longer distances. Remember this if you use this override.

3.3.7 Cluster family

Cluster families can be introduced to help with TVC selection or replication activity. You might want to use them for the following conditions:

- ▶ You have an independent group or groups of clusters that serve a common purpose within a larger grid.

- ▶ You have one or more groups of clusters with limited bandwidth between the groups and other clusters in the grid.

Cluster families provide two essential features:

- ▶ During mounts, clusters within the same family as the mount point cluster are preferred for TVC selection.
- ▶ During replication, groups of clusters in a family cooperate and distribute the replication workload inbound to its family, which provides the best use of the limited network outside of the family.

Therefore, grid configurations with three or more clusters can benefit from cluster families.

3.3.8 Logical Volume Delete Expire Processing versus previous implementations

When a system TMS returns a logical volume to a SCRATCH category during housekeeping processing, the TS7700 is aware that the volume is not in a SCRATCH pool. The default behavior of the TS7700 is to retain the content on the virtual volume and its used capacity within the TS7700 until the logical volume is reused or ejected. Delete expire provides a means for the TS7700 to automatically delete the contents after a period of times has passed.

A scratch category can have a defined expiration time, enabling the volume contents for those volumes that are returned to scratch to be automatically deleted after a grace period passes. The grace period can be configured from 1 hour to many years. Volumes in the scratch category are then either expired with time or reused, whichever comes first.

If physical tape is present, the space on the physical tape that is used by the deleted or reused logical volume is marked inactive. Only after the physical volume is later reclaimed or marked full inactive is the tape and all inactive space reused. After the volume is deleted or reused, content that was previously present is no longer accessible.

An inadvertent return to scratch might result in loss of data, so a longer expiration grace period is suggested to enable any *return to scratch* mistakes to be corrected within your host environment. To prevent reuse during this grace period, enable the additional hold option to prevent such reuse. This provides a window of time where a host-initiated mistake can be corrected, enabling the volume to be moved back to a private category while retaining the previously written content.

3.3.9 Encryption

Depending on your legal requirements and your type of business, data encryption might be mandatory.

Consider the following information:

- ▶ If you use the Copy Export feature and encrypt the export pool, you must ensure that you can decrypt the tapes in the restore location:
 - You need to have access to an external key manager that has the appropriate keys available.
 - The same or compatible drives that can read the exported media format must be available.
- ▶ TVC encryption for data at rest in disk cache can be enabled only against the entire cache repository.

- ▶ Both physical tape and TVC encryption can be enabled at any time. After TVC encryption is enabled, it cannot be disabled without a rebuild of the disk cache repository. If you use an external key manager for physical tape and TVC encryption, the same external key manager instance must be used.
- ▶ Disk-based encryption can be enabled in the field retroactively on all Encryption Capable hardware. Therefore, enabling encryption can occur after the hardware has been configured and used.

3.3.10 z/OS Allocation assistance

Allocation assistance is a function that is built into z/OS and the TS7700 that enables both private and scratch mounts to be more efficient when they choose a device within a grid configuration where the same sysplex is connected to two or more clusters in an active-active configuration.

Remember: Support for the allocation assistance functions (DAA and SAA) was initially only supported for the job entry subsystem 2 (JES2) environment. With z/OS V2R1, JES3 is also supported.

If you use the allocation assistance, the device allocation routine in z/OS is influenced by information from the grid environment. Several aspects are used to find the best mount point in a grid for this mount. For more information, see 2.3.15, “Allocation assistance” on page 71.

Depending on your configuration, your job execution scheduler, and any automatic allocation managers you might use, the allocation assist function might provide value to your environment.

If you use any dynamic tape manager, such as the IBM Automatic Tape Allocation Manager, plan the introduction of SAA and DAA carefully. Some dynamic tape managers manage devices in an offline state. Because allocation assist functions assume online devices, issues can surface.

Therefore, consider keeping some drives always online to a specific host, and leave only a subset of drives to the dynamic allocation manager. Alternatively, discontinue working with a dynamic tape allocation manager.

Automatic tape switching (ATSSTAR), which is included with z/OS, works with online devices, and is compatible with DAA and SAA.

3.3.11 25 GB logical volumes

The TS7700 has traditionally supported 400 megabyte (MB), 800 MB, 1 gigabyte (GB), 2 GB, 4 GB, and 6 GB logical volumes. As of R3.2, 25 GB logical volumes are also supported. Using 25 GB logical volumes can have several advantages:

- ▶ Fewer virtual volumes to insert and have managed by your TMS.
- ▶ Migration from other tape libraries can be easier.
- ▶ Large multi-volume workloads, such as a large database backup, can be stored with fewer logical tapes.

Here are some considerations if you choose to use 25 GB logical volumes:

- ▶ 25 GB logical volumes that use RUN copy consistency points are viewed as Deferred consistency points.

- ▶ If present only on physical tape, the entire volume must be recalled into disk cache before completing the logical mount.
- ▶ Appending data to larger volumes requires a full replication to peers, and can result in larger inactive spaces on physical tape.
- ▶ Many jobs running to 25 GB logical volumes can create a large increase in disk cache content, which can result in non-optimal performance.

To avoid any performance effect, you should review your installation before you use the 25 GB volumes.

3.4 Features and functions available only for the TS7700T

With the introduction of tape support behind the TS7700T, additional features that are unique to the TS7700T are now available:

- ▶ Multiple tape-managed partitions
- ▶ Delay premigration to physical tape

Having multiple tape partitions enables you to define how much disk cache is used by a workload or group of workloads. Through partitioning, a workload's TVC residency footprint can be fixed when compared to other independent workloads. Therefore, independent of how much content other partition workloads create, their activity does not alter the residency footprint of the partition of interest.

In addition, delay premigration was introduced to help manage the movement of data to physical tape. By using policies that can delay premigration of specific workloads from one to many hours, only content that has not yet expired when the delay period passes ends up on tape. This creates a solution where the aged or archive component of a workload is the only content that moves to tape. Until then, the data is only resident in disk cache.

When the data expires from a host perspective while it is still in cache, it is not premigrated or migrated to a tape. That reduces your back-end activities (migrate and reclaim).

3.5 Operation aspects: Monitoring and alerting

To ensure that your TS7700 environment works as expected, and to be notified of any potential issues or trends, two different topics should be reviewed:

- ▶ Message handling:
 - Check for the introduction of new messages into your automation and alerting tool.
 - Use automation to trap on alerts of interest that are surfaced to the hosts.
- ▶ Regularly scheduled performance monitoring:
 - Gather long-term statistics through tools, such as VEHSTATS and BVIR, to retain data for trending.
 - Analyze any changes in the workload profile or behavior of the grid environment to ensure that the overall configuration operates as expected, and to determine whether changes should be made.

In addition, optional checks might be useful, especially after complex migrations or changes in your environment.

3.5.1 Message handling

With each new feature or Licensed Internal Code release, new messages might be introduced. Usually, they are described in the PTF description or mentioned in the messages and codes books. Identify all new messages for the TS7700 (usually CBRxxxx) and review them. The main message is the CBR3750 message, which contains many submessages. Evaluate the meanings to understand how they relate to your business.

For a complete list of all possible CBR3750 submessages, see the *IBM Virtualization Engine TS7700 Series Operator Informational Messages* white paper:

<https://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101689>

Identify the appropriate action that an operator or your automation tool must run. Introduce the new messages in your automation tool (with the appropriate action) or alert the message for human intervention.

3.5.2 Regularly scheduled performance monitoring

Regularly scheduled performance monitoring enables you to complete the following tasks:

- ▶ See trends in your workload profile so that you can tune your environment before any issues arise.
- ▶ Store historical information of your environment for trends and performance analysis.

The TS7700 keeps performance data for the last 90 days. If more than 90 days is required, running tools periodically to collect the information and store it is required. Then, set up regular Bulk Volume Information Retrieval (BVIR) runs and keep the data. Check this data on a periodic basis to see the usage trends, especially for shortage conditions.

3.5.3 Optional checks

Especially after major changes to the environments, you should consider running extra checks.

Verifying your data redundancy

The TS7700 in a grid configuration provides both high availability and disaster recovery. Both require one or more replicas of content within the grid. The BVIR Copy Audit provides a method to verify that replicas of all volumes exist at specific clusters or specific groups of clusters. The audit can run in a way that assumes that all volumes should replicate, and also has methods to verify replicas only based on assigned copy policies.

Consider running copy audits after major changes in the environment, such as joins, merges and before the removal of one or more clusters. You can also run the Copy Audit periodically as a method to audit your expected business continuance requirements.

Checking the SMS environment

Make sure that all distributed library names within a grid are configured within z/OS, even if they are not connected to the specific z/OS host.

Checking the settings environment

To check the settings environment and ensure that all parameters are correct, run the **LIBRARY REQUEST** command.

3.6 Choosing a migration method

To introduce new technology, sometimes data migration is needed because a hardware upgrade itself is not sufficient. In addition, you might need a data migration method to support a data center move. In general, there are two different methodologies:

- ▶ Host-based migration
- ▶ TS7700 internal data migration

TS7700 Release 3.3 introduces a new data migration method that is called *Grid to Grid Migration* (GGM), which is offered as a service from IBM.

The following section provides an overview of the different migration techniques.

3.6.1 Host-based migration

Host-based migration means that the data is read by the host through the FICON channels from the tape and written into the new tape environment, which has some consequences:

1. The logical volume number changes because the data is transferred by the host from one logical volume to another one.
2. Without manipulation of the Tape Management Catalog (TMC), you lose the origin creation date, job, owner, expiration date, and so on. Therefore, copy tools are often used to keep the origin information.
3. The data on the “old” logical volumes must be deleted manually.

The biggest advantage of this migration is that it is technology- and vendor-independent. However, it is resource-intensive (human effort and processor resources) and manual actions are error-prone.

3.6.2 TS7700 internal data migration

With the introduction of Release 3.3, there are two different data-migration methods provided by the TS7700 technology:

- ▶ Join and Copy Refresh Processing
- ▶ The GGM tool

Still other possibilities, for example Host Tape Copy, exist.

Join and Copy Refresh processing

If you want to move to a new data center, or do a technical refresh, use this method to migrate the data to a new cluster without using host-based migration. To do so, complete the following steps:

1. Join a new cluster.
2. Change the MC contents to allow copies to the new cluster.
3. Use the **LI REQ** parameter with the **CopyRefresh** parameter from the host for each logical volume, to produce a new copy of the data in the new cluster.

While the command is submitted from a host, the data is copied internally through the gridlinks. There is no Host I/O through the FICON adapters, and all data in the TCDB and tape management remain unchanged.

This method can be used only if the data migration is inside a grid. Inside a grid, it is a fast and proven copy method. In addition, the BVIR **AUDIT** parameter provides an easy method to ensure that all data is copied.

Grid to Grid migration tool

The GGM tool is a service offering from IBM. You can use it to copy logical volumes from one grid to another grid while both grids have a separated grid network. After the GGM is set up by an IBM Service Support Representative (IBM SSR), the data from the logical volumes is transferred from one grid to the other grid through the existing IP addresses for the gridlinks. Much like Join and Copy Refresh processing, there is no host I/O with the FICON adapters.

The GGM tool should be considered if the following situations are true:

- ▶ There are already six clusters installed in the grid.
- ▶ The Join and Copy Refresh processing cannot be used (there are floor space requirements, microcode restrictions, or other considerations).
- ▶ Source and Target grid belongs are maintained by different providers.

The GGM tool also provides several different options, such as how the new data (new device categories) and the old data (keep or delete the data in the source grid) is treated.

To access the data in the new grid, TCDB and the TMC must be changed. These changes are the responsibility of the customer, and must be processed manually.

The GGM is controlled by the **LI REQ** command, and reporting is provided by additional BVIR reports. A summary of this command can be found in the following white paper:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5328>

In addition, several supporting tools to create the necessary input control statements, and the necessary TCDB entry changes and TMC entry changes, are provided at the IBM Tape Tool website:

<ftp://public.dhe.ibm.com/storage/tapetool>

For more information refer to Chapter 8., "Migration" on page 283, or ask your local IBM SSR.

3.6.3 Tape drive technology behind a TS7700

Before Release 3.3, all tape drives that were attached to a TS7700 had to be homogeneous. There was no intermixing allowed.

With Release 3.3, you now can mix the TS1150 with *only one* older drive technology. This intermix is for migration purposes because a TS1150 cannot read content from JA and JB cartridges.

The following considerations apply:

- ▶ The "old" technology is used only for reads. You cannot write data on the legacy cartridge tape media by using the older drive technology.
- ▶ The maximum of 16 back-end drives must be divided by two tape technologies. Plan ahead to have enough tape drives in the older technology for recalls, and maybe for reclaim. But, have enough TS1150 tape drives to allow premigration, recalls, and reclaim for newly written data.

- ▶ Use the **LI REQ** to define the values for the alerts for missing physical drives for both technologies and the TS1150.
- ▶ Run **VEHSTATS** to understand the physical drive behavior.

With reclamation, the data from the discontinued media is moved to the new data. If you do not want that situation to occur, modify the “Sunset Media Reclaim Threshold Percentage (%)” for the specific physical pool on the MI to 0, and 0 reclaim runs for the discontinued media inside that pool.



Preinstallation planning and sizing

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

This chapter includes the following sections:

- ▶ Hardware installation and infrastructure planning
- ▶ Planning for a grid operation
- ▶ Planning for software implementation
- ▶ Planning for logical and physical volumes
- ▶ Tape analysis and sizing the TS7700

Remember: For this chapter, the term *tape library* refers to the IBM TS3500 and TS4500 tape libraries.

4.1 Hardware installation and infrastructure planning

This section describes planning information that is related to your TS7700. The topics that are covered include system requirements and infrastructure requirements. Figure 4-1 shows an example of the connections and infrastructure resources that might be used for a TS7700 grid configuration with two separate data centers.

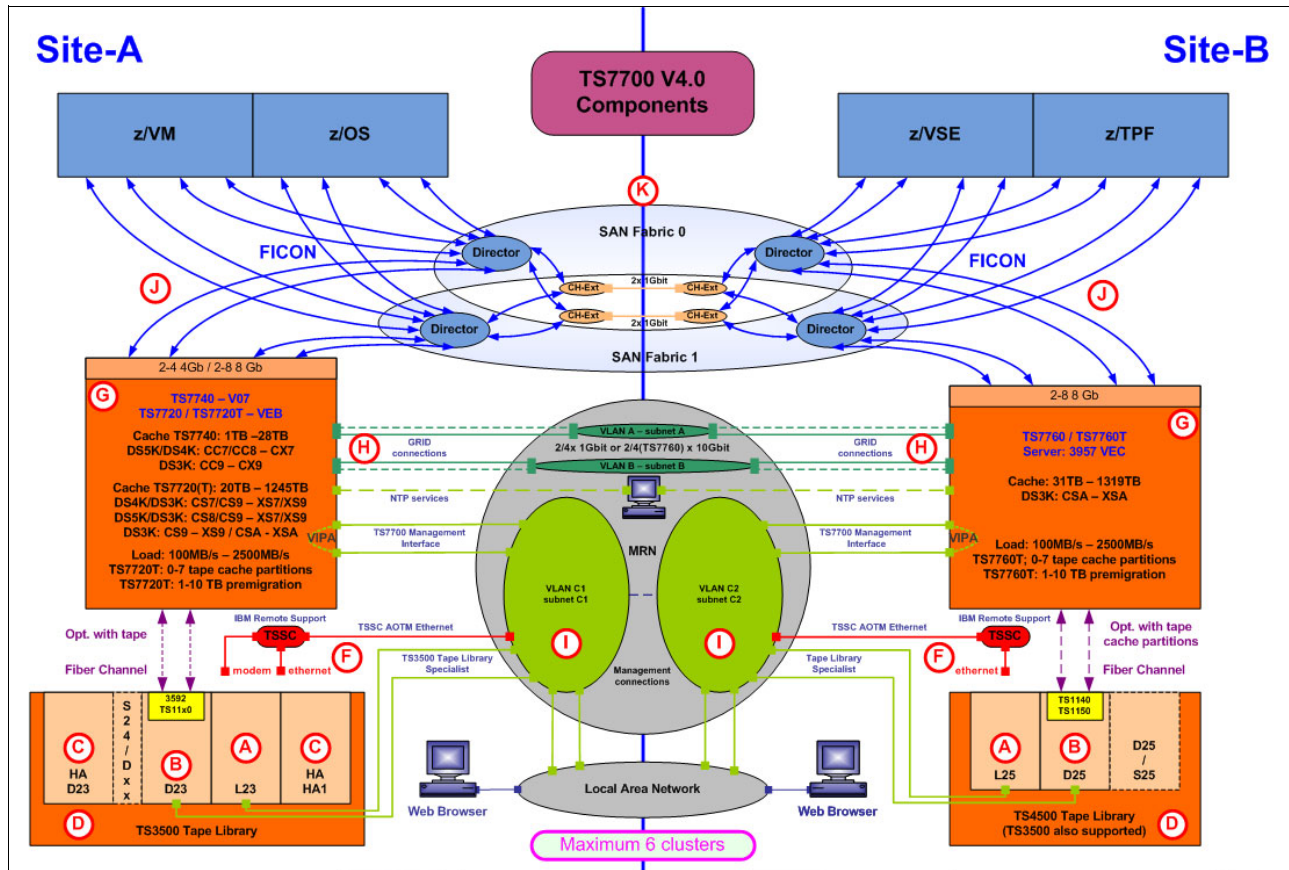


Figure 4-1 TS7700 grid configuration example

The letters in Figure 4-1 refer to the following items:

- ▶ A: TS7740/TS7720T/TS7760T – 3584-L23 library control frame
TS7760T – 3584-L25 library control frame
- ▶ B: TS7740/TS7720T/TS7760T – 3584-D23 frames with 3592, TS1120, TS1130, TS1140, or TS1150 drives, 3584-S24 HD storage frames
TS7760T – 3584-D25 frames with TS1140 or TS115 drives, 3584-S25 storage frames
- ▶ C: TS7740/TS7720T/TS7760T – 3584-HA1 frame and 3584-D23/HA frame (optional)
- ▶ D: TS7740/RS7720T/TS7760T – 3592 Advanced media type JA/JB/JC/JD and 3592 Advanced Economy media type JJ/JK/JL data cartridges for the data repository
TS7760T – 3592 Advanced media type JA/JB/JC/JD and 3592 Advanced Economy media type JJ/JK/JL data cartridges for the data repository
- ▶ F: Total Storage System Console (TSSC) for IBM Service Support Representatives (IBM SSRs) and Autonomic Ownership Takeover Manager (AOTM)
- ▶ G: TS7700

- ▶ H: TS7740/TS7720 – two or four 1 Gb Ethernet (copper or SW fibre) or two 10 Gb Ethernet for Grid communication.
TS7760 – two or four 1 Gb Ethernet (copper or SW fibre) or two or four 10 Gb Ethernet for Grid communication
- ▶ I: Ethernet connections for Management Interfaces (MIs)
- ▶ J: TS7740/TS7720 – FICON connections for host workload, two - four 4 Gb or two - eight 8 Gb
TS7760 - FICON connections for host workload, two - eight 8 Gb
- ▶ K: FICON fabric infrastructure with extension technology when appropriate

4.1.1 System requirements

Ensure that your facility meets the system requirements for the TS7700 when you plan for installation. System requirements for installation include requirements for power, cooling, floor leveling, loading, distribution, clearance, environmental conditions, and acoustics.

For a detailed listing of system requirements, see IBM TS7700 R4.0 IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_system_requirements.html

IBM 3952 Tape Frame specifications

The 3952 Tape Frame F05 houses the components of the TS7720 and TS7740, while 3952 F06 houses the components of the TS7760. Table 4-1 lists the dimensions of the frame that encloses the TS7700.

Table 4-1 Physical characteristics of a maximally configured 3952 Tape Frame

Characteristic	3952 F05	3952 F06 Value
Height	1804 mm (71.0 in.)	1930.4 mm (76 in.)
Width	644 mm (25.4 in.)	635 mm (25 in.)
Depth	1098 mm (43.23 in.)	1409.7 mm (55.5 in.)
Weight	270 kg (595.25 lb.) empty 669.1 kg (1475 lb.) maximally configured	746 kg (1645 lb.) maximally configured
Power	240 Vac, 15 amp (single phase)	240 Vac, 15 amp (single phase)
Unit height	36 U	40 U

Environmental operating requirements

Your facility must meet specified temperature and humidity requirements before you install the TS7700. Table 4-2 shows the preferred environmental conditions for the TS7700.

Table 4-2 Environmental specifications

Condition	Air temperature	Altitude	Relative humidity ^a	Wet bulb temperature
Operating (low altitude)	10°C - 32°C (50°F - 89.6°F)	Up to 5000 ft. above mean sea level (AMSL)	20% - 80%	23°C (73°F)
Operating (high altitude)	10°C - 28°C (50°F - 82.4°F)	5001 ft. AMSL - 7000 ft. AMSL	20% - 80%	23°C (73°F)
Preferred operating range ^b	20°C - 25°C (68°F - 77°F)	Up to 7000 ft. AMSL	40% - 55%	N/A
Power off	10°C - 43°C (50°F - 109°F)	N/A	8% - 80%	27°C (80°F)
Storage	1°C - 60°C (33.8°F - 140°F)	N/A	5% - 80%	29°C (84°F)
Shipping	-40°C - 60°C (-40°F - 140°F)	N/A	5% - 100%	29°C (84°F)

a. Non-condensing

b. Although the TS7700 can operate outside this range, it is advised that you adhere to the preferred operating range.

Power considerations

Your facility must have ample power to meet the input voltage requirements for the TS7700.

The standard 3952 Tape Frame includes one internal power distribution unit. However, feature code 1903, Dual AC power, is required to provide two power distribution units to support the high availability (HA) characteristics of the TS7700. The 3952 Tape Expansion Frame has two power distribution units and requires two power feeds.

TS7720 Base Frame power requirements

Your facility must have ample power to meet the input voltage requirements for the TS7720 Base Frame. Table 4-3 displays the maximum input power for a fully configured TS7720 Base Frame.

Table 4-3 TS7720 Base Frame maximum input power requirements

Power requirement	Value
Voltage	200 - 240 V AC (single phase)
Frequency	50 - 60 Hz (+/- 3 Hz)
Current	20 A
Inrush current	250 A
Power (W)	3,140 W
Input power required	4.0 kVA (single phase)
Thermal units	11.0 KBtu/hr, 2.76 kcal/hr

TS7720 Storage Expansion Frame power requirements

Your facility must have ample power to meet the input voltage requirements for the TS7720 Storage Expansion Frame. Table 4-4 displays the maximum input power for a fully configured TS7720 Expansion Frame.

Table 4-4 TS7720 Storage Expansion Frame maximum input power requirements

Power requirement	Value
Voltage	200 - 240 V AC (single phase)
Frequency	50 - 60 Hz (+/- 3 Hz)
Current	20 A
Inrush current	250 A
Power (W)	3,460 W
Input power required	4.0 kVA (single phase)
Thermal units	11.8 KBtu/hr, 2.96 kcal/hr

TS7740 Base Frame power requirements

Your facility must have ample power to meet the input voltage requirements for the TS7740 Base Frame. Table 4-5 displays the maximum input power for a fully configured TS7740 Base Frame.

Table 4-5 TS7740 Base Frame maximum input power requirements

Power requirement	Value
Voltage	200 - 240 V AC (single phase)
Frequency	50 - 60 Hz (+/- 3 Hz)
Current	20 A
Inrush current	250 A
Power (W)	1786 W
Input power required	4.0 kVA (single phase)
Thermal units	6.05 kBtu/hr, 1.52 kcal/hr

TS7760 Base Frame power requirements

Your facility must have ample power to meet the input voltage requirements for the TS7760 Base Frame. Table 4-6 displays the maximum input power for a fully configured TS7760.

Table 4-6 TS7760 Base Frame maximum input power requirements

Power requirement	Value
Voltage	200 - 240 V AC (single phase)
Frequency	50 - 60 Hz (+/- 3 Hz)
Current	24 amp
Inrush current	250 amp
Power (W)	3280 watts
Input power required	4.8 kVa (single phase)
Thermal units	11.5 kBtu/hr, 2.9 kcal/hr

TS7760 Storage Expansion Frame power requirements

Your facility must have ample power to meet the input voltage requirements for the TS7760 Storage Expansion Frame. Table 4-7 displays the maximum input power for a fully configured TS7760.

Table 4-7 TS7760 Storage Expansion Frame maximum input power requirements

Power requirement	Value
Voltage	200 - 240 V AC (single phase)
Frequency	50 - 60 Hz (+/- 3 Hz)
Current	24 amp
Leakage current	13.5 ma
Inrush current	250 amp

Power requirement	Value
Power (W)	3200 watts
Input power required	4.8 kVa (single phase)
Thermal units	11.2 kBtu/hr, 2.9 kcal/hr

Tape drives and media support (TS7740,TS7720T, and TS7760T)

The TS7740, TS7720T, and TS7760T support the 3592 Tape Cartridge (JA), 3592 Expanded Capacity Cartridge (JB), 3592 Advanced Tape Cartridge (JC), 3532 Advanced Data Tape Cartridge (JD), 3592 Economy Tape Cartridge (JJ), 3592 Economy Advanced Tape Cartridge (JK) media, and 3592 Economy Tape Cartridge (JL).

The TS7740, TS7720T, and TS7760T support the 3592 Extended Tape Cartridge (JB) media and require TS1120 model E05 Tape Drives in E05 mode, TS1130 Model E06/EU6 tape drives, TS1140 Model E07 or EH7 tape drives. Alternatively, they require a heterogeneous setup involving the TS1150 Model E08 or EH8 tape drives and either of the TS1120 Model E05, TS1130 Model E06/EU6, or TS1140 Model E07 or EH7 tape drives, depending on the library generation.

In a TS3500 tape library, all tape drives and media are supported. In a TS4500 tape library, only TS1140 and TS1150 with the corresponding media is supported.

The TS7740, TS7720T, and TS7760T tape encryption require that all the backend drives be encryption capable. TS1130 Model E06/EU6, TS1140 Model E07 or EH7, and TS1150 Model E08 or EH8 drives are encryption capable. TS1120 Model E05 Tape Drives in E05 mode and are encryption capable, with either FC 9592 from the factory, or FC 5592 as a field upgrade.

Support for the fourth generation of the 3592 drive family is included in TS7700 Release 2.0 PGA1. At this code level, the TS1140 tape drive that is attached to a TS7740 and TS7720T cannot read JA or JJ media. Ensure that data from all JA and JJ media has been migrated to JB media before you replace older-generation tape drives with TS1140 drives. Starting with Release 2.1 PGA0 the reading of JA and JJ media by the TS1140 drive is supported. The client can choose to keep the data on the JA/JJ media or can plan to migrate the data to newer generations of media.

Heterogeneous Tape Drive Support

Starting with release R3.3 the TS7740 and TS7720 Tape attach supports Heterogeneous Tape Drives. Heterogeneous Tape Drives are also supported by TS7760T and release R4.0.

TS1150 tape drives can be intermixed with some other drive types. The E08 drives are supported in a limited heterogeneous configuration.

The new media types JD and JL are supported by the TS7700. Up to 10 TB of data can be written to the JD cartridge in the 3593 E08 tape drive recording format. Up to 2 TB of data can be written to the JL cartridge in the 3592 E08 tape drive recording format. The 3592 E08 tape drive also supports writing to prior generation media types JK and JC. When the 3592 E08 recording format is used starting at the beginning of tape, up to 7 TB of data can be written to a JC cartridge and up to 900 GB of data can be written to a JK cartridge.

The 3592 E08 tape drive does not support reading or writing to media type JJ, JA, and JB cartridges. The TS7700 does not support any type of Write Once Read Many (WORM) media.

Important: Not all cartridge media types and media formats are supported by all 3592 tape drive models. For the media, format, and drive model compatibility to see which tape drive model is required for a certain capability, see Table 4-8.

Table 4-8 Supported 3592 read/write formats

3592 Tape Drive	EFMT1 512 tracks, 8 R/W channels	EFMT2 896 tracks, 16 R/W channels	EFMT3 1152 tracks, 16 R/W channels	EFMT4 664 tracks (JB/JX) 2176 tracks (JC/JK), 32 R/W channels	EFMT5 4608 tracks (JC/JK) 5120 tracks (JD/JL), 32 R/W channels
Model J1A	Read/write	Not supported	Not supported	Not supported	Not supported
Model E05	Read/write ^a	Read/write	Not supported	Not supported	Not supported
Model E06/EU6	Read	Read/write	Read/write	Not supported	Not supported
Model E07/EH&	Read ^b	Read ^b	Read/write ^c	Read/write	Not supported
Model E08/EH8	Not supported	Not supported	Not supported	Read/write	Read/write

a. Model E05 can read and write EFMT1 operating in native or J1A emulation mode.

b. Model E07/EH7 can read JA and JJ cartridge types only with a tape drive firmware level of D3I3_5CD or higher.

c. Cartridge type JB only.

Table 4-9 summarizes the tape drive models, capabilities, and supported media by tape drive model.

Table 4-9 3592 Tape Drive models and characteristics versus supported media and capacity

3592 drive type	Supported media type	Encryption support	Capacity	Data rate
TS1150 Tape Drive (3592-E08/EH8 Tape Drive)	JC JD JK JL	Yes	7 TB (JC native) 10.0 TB (JD native) 900 GB (JK native) 2 TB (JL native) 10.0 TB (maximum all)	360 MBps
TS1140 Tape Drive (3592-E07/EH7 Tape Drive)	JB JC JK Media read only: JA JJ	Yes (IBM SecurityKey Lifecycle Manager or IBM SecurityKey Lifecycle Manager for z/OS only)	1.6 TB (JB native) 4.0 TB (JC native) 500 GB (JK native) 4.0 TB (maximum all)	250 MBps
TS1130 Tape Drive (3592-EU6 or 3592-E06 Tape Drive)	JA JB JJ	Yes	640 GB (JA native) 1.0 TB (JB native) 128 GB (JJ native) 1.0 TB (maximum all)	160 MBps

3592 drive type	Supported media type	Encryption support	Capacity	Data rate
TS1120 Tape Drive (3592-E05 Tape Drive)	JA JB JJ	Yes	500 GB (JA native) 700 GB (JB native) 100 GB (JJ native) 700 GB (maximum all)	100 MBps
3592-J1A	JA JJ	No	300 GB (JA native) 60 GB (JJ native) 300 GB (maximum all)	40 MBps
Notes: <ul style="list-style-type: none"> ▶ To use tape encryption, all drives that are associated with the TS7740, TS7720T, or TS7760T must be Encryption Capable and encryption-enabled. ▶ Encryption is not supported on 3592 J1A tape drives. 				

The *media type* is the format of the data cartridge. The media type of a cartridge is shown by the last two characters on standard bar code labels. The following media types are supported:

▶ **JA: An Enterprise Tape Cartridge (ETC)**

A JA cartridge can be used in native mode in a 3592-J1A drive or a 3592-E05 Tape Drive operating in either native mode or J1A emulation mode. The native capacity of a JA tape cartridge that is used in a 3592-J1A drive or a 3592-E05 Tape Drive in J1A emulation mode is 300 GB, equivalent to 279.39 gibibytes (GiB). The native capacity of a JA tape cartridge that is used in a 3592-E05 Tape Drive in native mode is 500 GB (465.6 GiB). The native capacity of a JA tape cartridge that is used in a 3592-E06 drive in native mode is 640 GB (596.04 GiB).

▶ **JB: An Enterprise Extended-Length Tape Cartridge (ETCL)**

Use of JB tape cartridges is supported only with TS1140 Tape Drives, TS1130 Tape Drives, and TS1120 Tape Drives operating in native capacity mode. When used with TS1140 Tape Drives, JB media that contains data that is written in native E05 mode is only supported for *read-only* operations.

After this data is reclaimed or expired, the cartridge can be written from the beginning of the tape in the new E07 format. If previously written in the E06 format, appends are supported by the TS1140 drive.

The native capacity of a JB tape cartridge that is used in a 3592-E05 drive is 700 GB (651.93 GiB). When used in a 3592-E06 drive, the JB tape cartridge native capacity is 1000 GB (931.32 GiB). When used within a Copy Export pool, a JB tape cartridge can be written in the E06 format with a TS1140 drive, enabling Copy Export restores to occur with TS1130 hardware. The native capacity of JB media that are used in a 3592-E07 tape drive in native mode is 1600 GB (1490.12 GiB).

▶ **JC: Advanced Type C Data (ATCD)**

This media type is supported for use with TS1150 and TS1140 tape drives. The native capacity of JC media that is used in a 3592-E07 drive is 4 TB (3.64 TiB) and in a 3592-E08 drive is 7 TB (6.52 TiB).

▶ **JD: Advanced Type D Data (ATDD)**

This media type is supported for use only with TS1150 tape drives.

- ▶ JJ: An Enterprise Economy Tape Cartridge (EETC)
A JJ cartridge can be used in native mode in a 3592-J1A drive or a 3592-E05 Tape Drive operating in either native mode or J1A emulation mode. The native capacity of a JJ tape cartridge that is used in a 3592-J1A drive or 3592-E05 Tape Drive in J1A emulation mode is 60 GB (58.88 GiB). The native capacity of a JJ tape cartridge that is used in a 3592-E05 Tape Drive in native mode is 100 GB (93.13 GiB). A JJ cartridge can be used in native mode in a 3592-J1A drive or a 3592-E05 Tape Drive operating in either native mode or J1A emulation mode.
- ▶ JK: Advanced Type K Economy (ATKE)
This media type is supported for use only with TS1150 and TS1140 tape drives.
- ▶ JL: Advanced Type L Economy (ATLE)
This media type is supported for use only with TS1150 tape drives.

The following media identifiers are used for diagnostic and cleaning cartridges:

- ▶ CE: Customer Engineer diagnostic cartridge for use only by IBM SSRs. The VOLSER for this cartridge is CE xxxJA, where a space occurs immediately after CE and xxx is three numerals.
- ▶ CLN: Cleaning cartridge. The VOLSER for this cartridge is CLN xxxJA, where a space occurs immediately after CLN and xxx is three numerals.

Planning for a TS7740, TS7720T, or TS7760T tape drive model change

Important: WORM cartridges, including JW, JR, JX, JY, and JZ, are *not supported*. Capacity scaling of 3592 tape media is also *not supported* by TS7740, TS7720T, and TS7760T.

When you change the model of the 3592 tape drives of an existing TS7740, TS7720T or TS7760T, the change must be in the later version direction, from an older 3592 tape drive model to a newer 3592 tape drive model.

3592 E08 drives can be mixed with one other previous generation tape drive through heterogeneous tape drive support, which allows a smooth migration of existing TS7700 tape drives with older tape drives to TS1150 tape drives.

For more information, see 7.2.6, “Upgrading drive models in an existing TS7740 or TS7700T” on page 260.

4.1.2 TS7700 specific limitations

Consider the following restrictions when you perform your TS7700 preinstallation and planning:

- ▶ Release 4.0 is only supported on models 3957-V07, 3957-VEB, and 3957-VEC.
- ▶ Release 3.3 is only supported on models 3957-V07 and 3957-VEB.
- ▶ TS1120 Tape Drives set in static emulation mode are not supported by the TS7740, TS7720T, and TS7760T. Static emulation mode forces the 3592-E05 to operate as a 3592-J1A drive.
- ▶ The maximum FICON cable distance for a direct connection between a TS7700 and host processor using short wavelength attachments at the 4 Gbps speed is up to 150 meters on 50 micron fiber cable, and up to 55 meters using 62.5 micron fiber.

At 8 Gbps speed, the short wave total cable length cannot exceed the following measurements:

- One hundred and fifty meters using 50 micron OM3 (2000 MHz*km) Aqua blue colored fiber
- Fifty meters using 50 micron OM2 (500 MHz*km) Orange colored fiber
- Twenty-one meters using 62.5 micron OM1 (200 MHz*km) Orange colored fiber
- ▶ Long wavelength attachments, both 4 Gb and 8 Gb provide a direct link of up to 10 km between the TS7700 and host processor on 9-micron fiber.
- ▶ Short and long wavelength attachments provide for up to 100 km between the TS7700 and host processor using appropriate fiber switches, and up to 250 km with DWDMs. Support is not provided through more than one dynamic switch.

For more information, see the *FICON Planning and Implementation Guide*, SG24-6497, for details about FICON connectivity:

<http://www.redbooks.ibm.com/abstracts/sg246497.html>

- ▶ The maximum length of the Cat 5e or Cat 6 cable between the grid Ethernet adapters in the TS7700 and the customer's switches or routers is 100 meters.
- ▶ The TS7700 does not support capacity scaling of 3592 tape media.
- ▶ The TS7700 does not support physical WORM tape media.
- ▶ The TS7700 does not support 3590 tape drives or 3590 tape media.
- ▶ The TS3500/TS4500 and TS7700 must be within 100 feet of the TSSC.
- ▶ The 3592 back-end tape drives for a TS7740, TS7720T, or TS7760T cluster must be installed in a TS3500 or TS4500 tape library. Connections to 3494 tape libraries are no longer supported (since the R2.0 machine code).
- ▶ The TS7740 and TS7720T support only 4 Gb or 8 Gb fiber switches for connection to the back-end drives.
- ▶ Clusters running Release 4.0 machine code can be joined only in a grid with clusters that are running either Release 2.1 or later. Release 4.0 supports up to three different code levels within the same grid. This situation can happen when grids are composed of clusters V06/VEA intermixed with clusters V07/VEB/VEC within the same grid. TS7700 cluster models V06/VEA are not compatible with LIC R4.0, and must stay at Release 2.1 or Release 3.0. V07/VEB/VEC clusters can be upgraded to Release 4.0.

Note: Existing TS7700 (3957-V06 with 3956-CC7 or 3956-CC8, 3957-VEA, 3957-V07, or 3957-VEB) can be upgraded to Release 3.0. To upgrade to Release 3.0, the existing cluster must be at least at 8.20.x.x (R2.0) level or later. Upgrade from 8.7.x.x (R1.7) level to Release 3.0 is only supported by RPQ.

3957-V06 with 3956-CC6 is not supported by Release 3.0.

- ▶ For this reason, during the code upgrade process, one grid can have clusters that are simultaneously running three different levels of code. Support for three different levels of code is available on a short-term basis (days or a few weeks), which should be long enough to complete the Licensed Internal Code upgrade in all clusters in a grid. The support for two different levels of code in a grid enables an indefinite coexistence of V06/VEA and V07/VEB/VEC clusters within the same grid.
- ▶ Because one new cluster can be joined in an existing grid with clusters that are running up to two different code levels, the joining cluster must join to a target cluster at the higher of the two code levels. Merging of clusters with mixed code levels is *not* supported.

The grid-wide functions available to a multi-cluster grid are limited by the lowest code level present in that grid.

4.1.3 TCP/IP configuration considerations

The Transmission Control Protocol/Internet Protocol (TCP/IP) configuration considerations and local area network/wide area network (LAN/WAN) requirements for the TS7700 are described in the following sections. Single and multi-cluster grid configurations are covered.

Figure 4-2 shows you the different networks and connections that are used by the TS7700 and associated components. This two-cluster TS7740/TS7720T/TS7760T grid shows the TS3500 and TS4500 tape library connections (not present in a TS7720D and TS7760D configuration).

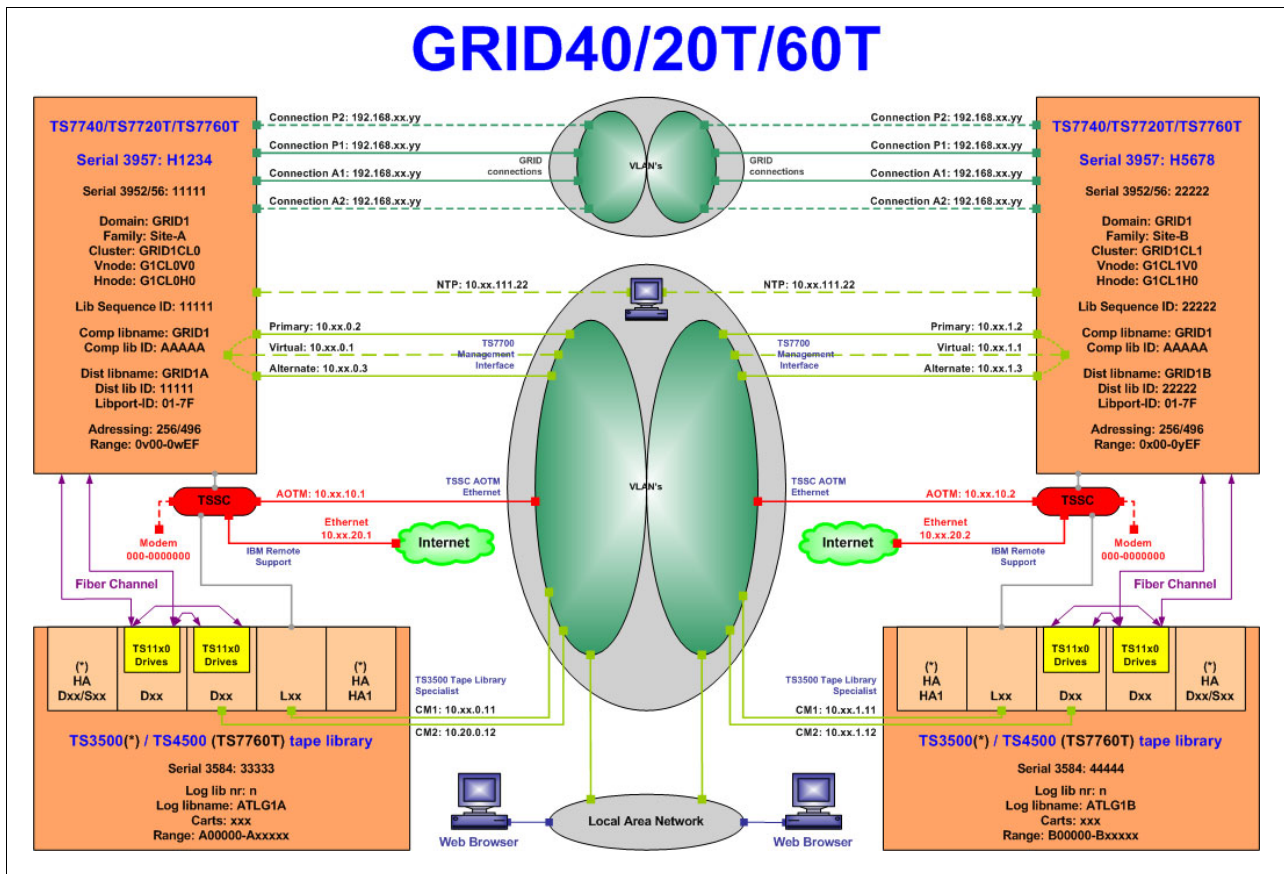


Figure 4-2 TCP/IP connections and networks

TS7700 grid interconnect LAN/WAN requirements

The LAN/WAN requirements for the TS7700 cross-site grid Internet Protocol network infrastructure are described in this section.

The TS7700 grid IP network infrastructure must be in place before the grid is activated so that the clusters can communicate with one another as soon as they are online. Two or four 1-GbE or 10-GbE connections must be in place before grid installation and activation.

An Ethernet extender or other extending equipment can be used to complete extended distance Ethernet connections.

Extended grid Ethernet connections can be any of the following connections:

- ▶ 1 Gb copper 10/100/1000 Base-TX

This adapter conforms to the Institute of Electrical and Electronics Engineers (IEEE) 802.3ab 1000Base-T standard, which defines gigabit Ethernet operation over distances up to 100 meters by using four pairs of CAT6 copper cabling.

- ▶ 1 Gb optical SW

This SX adapter has an LC Duplex connector that attaches to 50-micron (μ) or 62.5- μ multimode fiber cable. It is a standard SW, 850-nanometer (nm) adapter that conforms to the IEEE 802.3z standards. This adapter supports distances of 2 - 260 meters for 62.5- μ multimode fiber (MMF) and 2 - 550 meters for 50.0- μ MMF.

- ▶ 10 Gb optical LW

This 10 Gb grid optical LW connection provides a single-port, 10-Gbps Ethernet LW adapter for grid communication between TS7700 tape systems. This adapter has an LC Duplex connector for attaching 9- μ , single-mode fiber cable. This is a standard LW (1310 nm) adapter that conforms to the IEEE 802.3ae standards. It supports distances up to 10 kilometers (km), equivalent to 6.21 miles.

The default configuration for a TS7700 server from manufacturing (3957-VEB, 3957-V07, or 3957-VEC) is two dual-ported PCIe 1-GbE adapters. You can use FC 1035, 10 Gb grid optical LW connection to add support for two 10-Gb optical LW Ethernet adapters.

If the TS7700 server is a 3957-V07, 3957-VEB or 3957-VEC, two instances of either FC 1036 (1 Gb grid dual port copper connection) or FC 1037 (1 Gb dual port optical SW connection) must be installed. You can use FC 1038 to activate the second port on 10 Gb optical LW Ethernet adapter to support four 10 Gb grid links in the 3957-VEC.

Clusters that are configured by using four 10-Gb, two 10-Gb, four 1-Gb, or two 1-Gb clusters, can be interconnected within the same TS7700 grid, although the explicit same port-to-port communications still apply.

Important: Identify, order, and install any new equipment to fulfill grid installation and activation requirements. The connectivity and performance of the Ethernet connections must be tested before grid activation. You must ensure that the installation and testing of this network infrastructure is complete before grid activation.

To avoid performance issues, the network infrastructure should *not* add packet metadata (increase its size) to the default 1500-byte maximum transmission unit (MTU), such as with an encryption device or extender device.

The network between the TS7700 tape drives must have sufficient bandwidth to account for the total replication traffic. If you are sharing network switches among multiple TS7700 paths or with other devices, the total bandwidth of the network must be sufficient to account for all of the network traffic.

Consideration: Jumbo Frames are not supported.

The TS7700 uses TCP/IP for moving data between each cluster. Bandwidth is a key factor that affects throughput for the TS7700. The following key factors can also affect throughput:

- ▶ Latency between the TS7700 clusters
- ▶ Network efficiency (packet loss, packet sequencing, and bit error rates)
- ▶ Network switch capabilities

- ▶ Flow control to pace the data from the TS7700 tape drives
- ▶ Inter-switch link capabilities (flow control, buffering, and performance)

The TS7700 clusters attempt to drive the grid network links at the full speed that is allowed by the adapter (1 Gbps or 10 Gbps rate), which might exceed the network infrastructure capabilities. The TS7700 supports the IP flow control frames so that the network paces the level at which the TS7700 attempts to drive the network. The preferred performance is achieved when the TS7700 can match the capabilities of the underlying grid network, resulting in fewer dropped packets.

Remember: When the system exceeds the grid network capabilities, packets are lost. This causes TCP to stop, resync, and resend data, resulting in a less efficient use of the network. Flow control helps to reduce this behavior. 1-Gb and 10-Gb clusters can be within the same grid, but compatible network hardware must be used to convert the signals because 10 Gb cannot negotiate down to 1 Gb.

Note: It is advised to enable flow control in both directions to avoid grid link performance issues.

To maximize throughput, ensure that the underlying grid network meets these requirements:

- ▶ Has sufficient bandwidth to account for all network traffic that is expected to be driven through the system to eliminate network contention.
- ▶ Can support the flow control between the TS7700 clusters and the switches, which enables the switch to pace the TS7700 to the WAN capability. Flow control between the switches is also a potential factor to ensure that the switches can pace their rates to one another. The performance of the switch should be capable of handling the data rates that are expected from all of the network traffic.

Latency can be defined as *the time interval elapsed between a stimulus and a response*. In the network world, latency can be understood as how much time it takes for a data package to travel from one point to another in a network infrastructure. This delay is introduced by some factors, such as the electronic circuitry used in processing the data signals, or plainly by the universal physics constant, the speed of light. Considering the current speed of data processing, this is the most important element for an extended distance topology.

In short, latency between the sites is the primary factor. However, packet loss due to bit error rates or insufficient network capabilities can cause TCP to resend data, which multiplies the effect of the latency.

The TS7700 uses clients LAN/WAN to replicate logical volumes, access logical volumes remotely, and run cross-site messaging. The LAN/WAN must have adequate bandwidth to deliver the throughput necessary for your data storage requirements.

The cross-site grid network is 1 GbE with either copper (RJ-45) or SW fiber (single-ported or dual-ported) links. For copper networks, CAT5E or CAT6 Ethernet cabling can be used, but CAT6 cabling is preferable to achieve the highest throughput. Alternatively, two or four 10-Gb LW fiber Ethernet links can be provided. Internet Protocol Security (IPSec) is now supported on grid links to support encryption.

Important: To avoid any network conflicts, the following subnets must *not* be used for LAN/WAN IP addresses, for MI primary, secondary, or virtual IP addresses:

- ▶ 192.168.251.xxx
- ▶ 192.168.250.xxx
- ▶ 172.31.1.xxx

For TS7700 clusters configured in a grid, the following extra assignments must be made for the grid WAN adapters. For each adapter port, you must supply the following information:

- ▶ A TCP/IP address
- ▶ A gateway IP address
- ▶ A subnet mask

Tip: In a TS7700 multi-cluster grid environment, you must supply two or four IP addresses per cluster for the physical links that are required by the TS7700 for grid cross-site replication.

The TS7700 provides up to four independent 1 Gb copper (RJ-45) or SW fiber Ethernet links for grid network connectivity, or up to four 10 Gb LW links. To be protected from a single point of failure that can disrupt all WAN operating paths to or from a node, connect each link through an independent WAN interconnection.

Note: It is a strongly preferred practice that the primary and alternative grid interfaces exist on separate subnets. Plan different subnets for each grid interface. If the grid interfaces are directly connected (without using Ethernet switches), you must use separate subnets.

Local IP addresses for Management Interface access

You must provide three TCP/IP addresses on the same subnet. Two of these addresses are assigned to physical links, and the third is a virtual IP address that is used to connect to the TS7700 MI.

Use the third IP address to access a TS7700. It automatically routes between the two addresses that are assigned to physical links. The virtual IP address enables access to the TS7700 MI by using redundant paths, without the need to specify IP addresses manually for each of the paths. If one path is unavailable, the virtual IP address automatically connects through the remaining path.

You must provide one gateway IP address and one subnet mask address.

Important: All three provided IP addresses are assigned to one TS7700 cluster for MI access.

Each cluster in the grid must be configured in the same manner as explained previously, with three TCP/IP addresses providing redundant paths between the local intranet and cluster. Two of these addresses are assigned to physical links, and the third address provides a virtual IP address to connect to the MI in this specific TS7700. Customer access to the cluster should require only the use of the virtual IP, and should function if one of the underlying physical IP links is operational.

Connecting to the Management Interface

This section describes how to connect to the IBM TS7700 MI. Table 4-10 lists the supported browsers.

Table 4-10 Supported browsers

Browser	Version supported	Version tested
Microsoft Edge	25.x	25.0
Internet Explorer	9, 10, or 11	11
Mozilla Firefox	24.0, 24.x ESR, 31.0, 31.x ESR, 38.0, or 38.x ESR	38.0 ESR
Google Chrome	39.x or 42.x	42.0

Perform the following steps to connect to the interface:

1. In the address bar of a supported web browser, enter `http://` followed by the *virtual IP* entered during installation, followed by `/Console`. The virtual IP is one of three IP addresses given during installation. The complete URL takes this form:
`http://<virtual IP address>/Console`
2. Press Enter on your keyboard or **Go** on your web browser.
The web browser redirects to `http://<virtual IP address>/<cluster ID>`, which is associated with the virtual IP address. If you bookmark this link and the cluster ID changes, you must update your bookmark before the bookmark resolves correctly. Alternatively, you can bookmark the more general URL, `http://<virtual IP address>/Console`, which does not require an update after a cluster ID change.
3. The login page for the MI loads. The default login name is `admin` and the default password is `admin`.

For the list of required TCP/IP port assignments, see Table 4-11 on page 142.

The MI in each cluster can access all other clusters in the grid through the grid links. From the local cluster menu, select a remote cluster. The MI goes automatically to the selected cluster through the grid link. Alternatively, you can point the browser to the IP address of the target cluster that you want.

This function is handled automatically by each cluster's MI in the background. Figure 4-3 shows a sample setup for a two-cluster grid.




Figure 4-3 TS7700 Management Interface access from a remote cluster

IPv6 support

Starting with TS7700 Licensed Internal Code R3.0, the 3957-V07 and 3957-VEB support IPv6.

Note: The TS7700 grid link interface does not support IPv6.

All network interfaces that support monitoring and management functions are now able to support IPv4 or IPv6:

- ▶ MI
- ▶ Key manager server:
 - IBM Security Key Lifecycle Manager
 - IBM Security Key Lifecycle Manager for z/OS
- ▶ Simple Network Management Protocol (SNMP) servers
- ▶ Lightweight Directory Access Protocol (LDAP) server
- ▶ Network Time Protocol (NTP) server

Important: All of these client interfaces must be *either* IPv4 or IPv6 for each cluster. Mixing IPv4 and IPv6 is *not* supported within a single cluster. For grid configurations, *each* cluster can be either all IPv4 or IPv6 unless an NTP server is used, in which case *all* clusters within the grid must be all one or the other.

For implementation details, see “Enabling IPv6” on page 564.

IPSec support for the grid links

Support for IPSec on the grid links was introduced with TS7700 R3.0 level for the 3957-V07 and 3957-VEB models. Use IPSec capabilities only if they are required by the nature of your business. Grid encryption might cause a considerable slowdown in all grid link traffic, such as in the following situations:

- ▶ Synchronous, immediate, or deferred copies
- ▶ Remote read or write

For implementation details, see “Enabling IPSec” on page 565.

TSSC Network IP addresses

The TS3000 Total Storage System Console (TSSC) uses an internal isolated network that is known as the TSSC network. All separate elements in the TS7700 tape subsystem connect to this network and are configured in the TSSC by the IBM SSR.

Each component of your TS7700 tape subsystem that is connected to the TSSC uses at least one Ethernet port in the TSSC Ethernet hub. For example, a TS7700 cluster needs two connections (one from the primary switch and other from the alternative switch). If your cluster is a TS7740, TS7720T, or TS7760T, you need a third port for the TS3500 or TS4500 tape library. Depending on the size of your environment, you might need to order a console expansion for your TSSC. For more information, see FC2704 in the IBM TS7700 R4.0 IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7700_feature_codes_all.html

Generally, there should be at least one TSSC available per location in proximity of the tape devices, such as TS7700 clusters and TS3500 tape libraries. Apart from the internal TSSC network, the TSSC can also have another two Ethernet physical connections:

- ▶ External Network Interface
- ▶ Grid Network Interface

Those two Ethernet adapters are used by advanced functions, such as AOTM, LDAP, Assist On-site (AOS), and Call Home (not using a modem). If you plan to use them, provide one or two Ethernet connections and the corresponding IP addresses for the TSSC. The ports in the table must be opened in the firewall for the interface links to work properly. Table 4-11 shows the network port requirements for the TSSC.

Table 4-11 TSSC TCP/IP port requirement

TSSC interface link	TCP/IP port	Role
TSSC External	80	Call Home
	443	
	53	Advise to remain open for the domain name server
	Internet Control Message Protocol (ICMP)	
TSSC Grid	80	Autonomic Ownership Takeover Mode (AOTM)
	22	
	443	
	9666	
	ICMP	

Network switches and TCP/IP port requirements

The network switch and TCP/IP port requirements for the WAN of a TS7700 in the grid configuration are shown in Table 4-12.

Clarification: These requirements apply only to the LAN/WAN infrastructure. The TS7700 internal network is managed and controlled by internal code.

Table 4-12 Infrastructure grid WAN TCP/IP port assignments

Link	TCP/IP port	Role
TS7700 MI	ICMP	Dead gateway detection
	123 ^a	NTP uses the User Datagram Protocol (UDP) time server
	443	Access the TS7700 MI (HTTPS)
	80	Access the remote MI when clusters are operating at different code levels (HTTP)
	1443	Encryption key (EK) server, Secure Sockets Layer (SSL)
	3801	EK server (TCP/IP)
T7700 GRID	ICMP	Check cluster health
	9	Discard port for speed measurement between grid clusters
	80	Access the remote MI when clusters are operating at different code levels
	123 ^a	NTP time server
	1415/1416	IBM WebSphere® message queues
	443	Access the TS7700 MI
	350	TS7700 file replication, Remote Mount, and Sync Mode Copy (distributed library file transfer)
	20	For use by IBM Support
	21	For use by IBM support
	500	IPSec Key Exchange (TCP and UDP): Must remain open when grid encryption is enabled.
	8500	IPSec Key Exchange (TCP and UDP): Must remain open when grid encryption is enabled.

a. Port 123 is used for grid link time synchronization within clusters, not for an external time server.

4.1.4 Factors that affect performance at a distance

Fibre Channel distances depend on many factors:

- ▶ Type of laser used: Long wavelength or short wavelength
- ▶ 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
- ▶ Link extenders

Native SW Fibre Channel transmitters have a maximum distance of 500 m with 50-micron diameter, multi-mode, optical fiber (at 1 or 2 Gbps). Although 62.5-micron, multimode fiber can be used, the larger core diameter has a greater dB loss and maximum distances are shortened to 300 or 150 meters. Native LW Fibre Channel transmitters have a maximum distance of 10 km (6.2 miles) when used with 9-micron diameter single-mode optical fiber. See the Table 4-13 on page 146 for a comparative table.

Link extenders provide a signal boost that can potentially extend distances to up to about 100 km (62 miles). These link extenders act as a large, 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 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 sends 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 must 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 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 has been to allocate one buffer credit for every 2 km (1.24 miles) 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. Generally, to maintain acceptable performance, one buffer credit is required for every 2 km (1.24 miles) distance covered. See *IBM SAN Survival Guide*, SG24-6143, for more information.

4.1.5 Host attachments

The TS7700 attaches to z Systems hosts through the FICON adapters on the host, either FICON LW or SW, at speeds of 2, 4, 8, or 16 Gbps. 1 Gbps is no longer supported by 8 Gb FICON Adapters:

- ▶ ESCON channel attachment is not supported.
- ▶ FICON channel extension and DWDM connection are supported.
- ▶ FICON directors and director cascading are supported.

Note: Considerations for host FICON connections:

- ▶ z Systems 8 Gbps FICON supports any TS7700 FICON (4 Gbps and 8 Gbps) direct attached.
- ▶ IBM z13@ 16 Gbps FICON supports only TS7700 8 Gbps FICON direct-attached.
- ▶ z13 16 Gbps FICON supports TS7700 4 Gbps FICON if FICON Director provides proper speed conversion.

Host attachment supported distances

When directly attaching to the host, the TS7700 can be installed at a distance of up to 10 km (6.2 miles) from the host. With FICON switches, also called *FICON Directors* or *Dense Wave Division Multiplexers (DWDMs)*, the TS7700 can be installed at extended distances from the host.

Figure 4-4 shows a sample diagram that includes the DWDM and FICON Directors specifications. For more information, see “FICON Director support” on page 146.

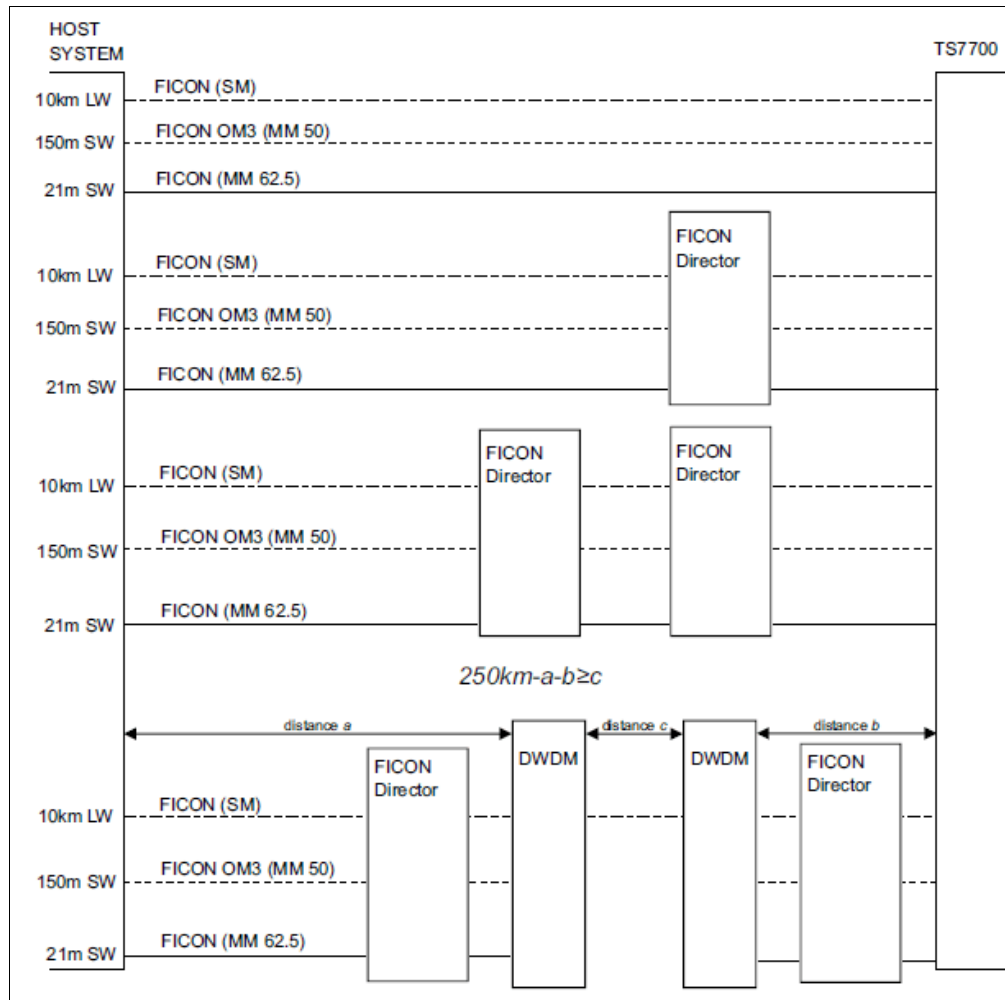


Figure 4-4 The z Systems host attachment to the TS7700 (at speed of 8 Gbps)

The maximum distances vary depending on the cable type and on the speed and type of optical transducer. There are three basic types of optical cable fiber:

- ▶ The orange colored cables are SW, multimode **OM2** type cables.
- ▶ The aqua colored multimode cables are **OM3, OM4** type and are laser-optimized.
- ▶ The yellow colored LW cables are single mode. The connection speed in Gbps determines the distance that is allowed.

Table 4-13 shows the relationship between connection speed and distance by cable type.

Table 4-13 Connection speed and distance by cable type

Cable type	Speed	Distance
OM2	1 Gbps	500 m (1640 ft.)
OM3	1 Gbps	500 m (1640 ft.)
OM2	2 Gbps	300 m (900 ft.)
OM3	2 Gbps	500 m (1640 ft.)
OM2	4 Gbps	150 m (492 ft.)
OM3	4 Gbps	270 m (886 ft.)
OM3	8 Gbps	150 m (492 ft.)
OM4	8 Gbps	190 m (623 ft.)

Figure 4-4 on page 145 shows the supported distances using different fiber cables for single-mode long wave laser and multimode short wave laser.

These attachments used the following abbreviations:

- ▶ SM: Single Mode fiber
- ▶ LW: Long Wave Laser
- ▶ MM: Multimode Fiber
- ▶ SW: Short Wave Laser

The TS7700 supports z Systems servers by using 8 Gb IBM FICON at distances up to 250 km (155 miles) by using dense wavelength division multiplexing (DWDM) in combination with switches, or more extended distances by using supported channel extension products.

Distances greater than 30 km require DWDM in combination with qualified switches or directors with adequate random access memory (RAM) buffer online cards. An *adequate* RAM buffer is defined as capable of reaching distances of 100 - 250 km.

Note: Long wave cables attach only to long wave adapters and short wave cables attach only to short wave adapters. There is no intermixing.

FICON Director support

All FICON Directors are supported for single and multi-cluster grid configurations with 1 Gbps, 2 Gbps, 4 Gbps, or 8 Gbps links. The components auto-negotiate to the highest speed allowed. 8 Gbps links cannot negotiate down to 1 Gbps.

You cannot mix different vendors, such as Brocade (formerly McData, CNT, and InRange) and CISCO, but you can 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 URL:

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>

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 performing any special channel or control unit (CU) 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 CU by presenting command responses and channel enablement (CE)/device end (DE) status ahead of the controller, and emulating the channel when running the pre-acknowledged write operations to the real remote tape device. Therefore, data is accepted early and forwarded to the remote device to maintain a full pipe throughout the write channel program.

The supported channel extenders between the z Systems host and the TS7700 are in the same matrix as the FICON switch support under the following URL (see the FICON Channel Extenders section):

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133>

Cascaded switches

The following list summarizes the general configuration rules for configurations with cascaded switches:

- ▶ Director Switch ID

This is defined in the setup menu.

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.

The following allowable Director Switch ID ranges have been established by the manufacturer:

- McDATA range: x'61' - x'7F'
- CNT/Inrange range: x'01' - x'EF'
- Brocade range: x'01' - 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 have been established by the manufacturer:

- McDATA range: x'61' - x'7F'
- CNT/Inrange range: x'01' - x'EF'
- Brocade range: x'01' - x'EF'

► Director Ports

The Port Address might not be the same as the Port Number. The Port Number identifies the physical location of the port, and the Port Address is used to route packets.

The Inboard Director Port is the port to which the CPU is connected. The Outboard Director Port is the port to which the CU 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 by using the available setup menu in the equipment. Figure 4-5 shows an example of host connection that uses DWDM and cascaded switches.

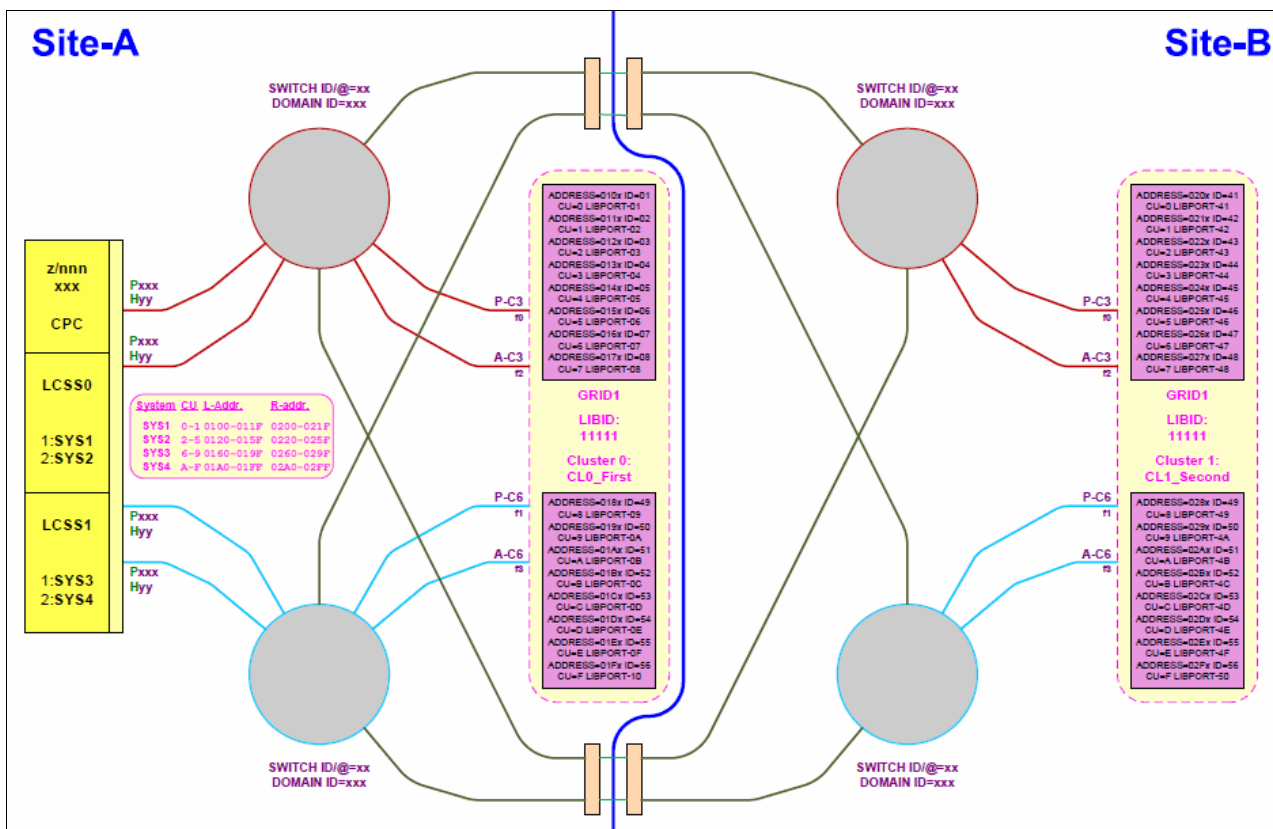


Figure 4-5 Host connectivity that uses DWDM and cascaded switches

4.1.6 Planning for LDAP for user authentication in your TS7700 subsystem

Depending on the security requirements in place, the user of the TS7700 can choose to have all of the TS7700 users' authentications controlled and authorized centrally by an LDAP server.

Important: Enabling LDAP requires that *all* users must authenticate with the LDAP server. All interfaces to the TS7700, such as MI, remote connections, and even the local serial port, are blocked. The TS7700 might be inaccessible if the LDAP server is unreachable.

The previous implementation relied on System Storage Productivity Center to authenticate users to a client's LDAP server. Beginning with Release 3.0 of LIC, both the TS7700 clusters and the TSSC have native support for the LDAP server (currently, only Microsoft Active Directory (MSAD) is supported).

Tip: System Storage Productivity Center continues to be a valid approach for LDAP.

Enabling authentication through an LDAP server means that all personnel with access to the TS7700 subsystem, such as computer operators, storage administrators, system programmers, and IBM SSRs (local or remote), must have a valid account in the LDAP server, along with the roles assigned to each user. The role-based access control (RBAC) is also supported. If the LDAP server is down or unreachable, it can render a TS7700 inaccessible from the outside.

Important: Create at least one external authentication policy for IBM SSRs before a service event.

When LDAP is enabled, the TS7700 MI is controlled by the LDAP server. Record the Direct LDAP policy name, user name, and password that you created for IBM SSRs and keep this information easily available in case you need it.

Note: Service access requires the IBM SSR to authenticate through the normal service login and then to authenticate again by using the IBM SSR Direct LDAP policy.

For more information about how to configure LDAP availability, see "Defining security settings" on page 566.

4.1.7 Cluster time coordination

All nodes in the entire 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 with a Network Time Protocol (NTP) Server. The NTP server can be a part of the Grid and WAN infrastructure, it can be a part of a customer intranet, or it can be a public server on the internet.

Figure 4-6 shows the NTP server configuration in grid.

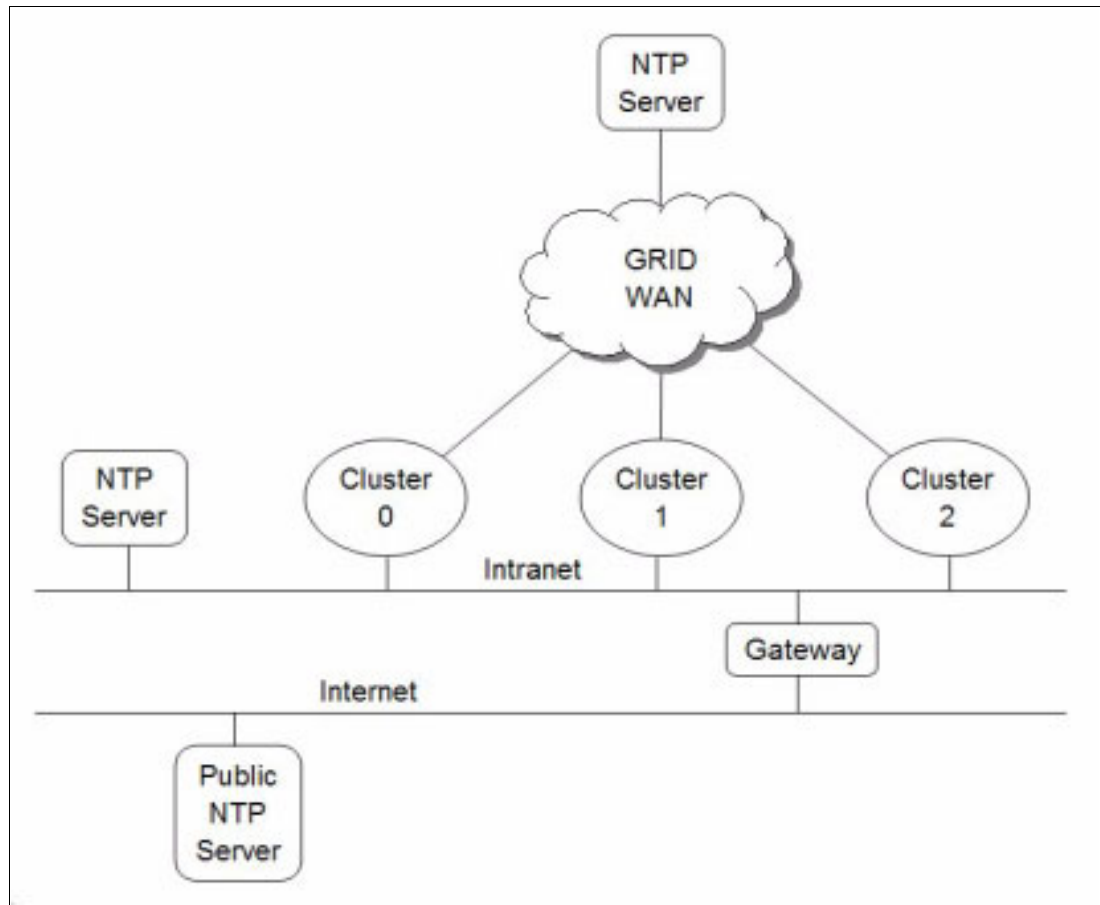


Figure 4-6 NTP server configuration

The NTP server address is configured into system VPD on a system-wide scope, so that all clusters access the same NTP server. All of the clusters in a grid need to be able to communicate with the same NTP server that is defined in VPD. In the absence of an NTP server, all nodes coordinate time with Cluster 0. In the absence of cluster 0, cluster 1 is used.

4.2 Planning for a grid operation

The TS7700 grid provides configuration flexibility to meet various requirements. Those requirements depend on both your business and your applications. This section specifically addresses planning a two-cluster grid configuration to meet HA needs. However, the configuration easily converts to a three-cluster grid configuration with two production clusters of HA and disaster recovery (DR). The third cluster is strictly a DR site.

4.2.1 Autonomic Ownership Takeover Manager (AOTM) Considerations

The Autonomic Ownership Takeover Manager (AOTM) is an optional function that was introduced in TS7700 R1.2. Following a TS7700 cluster failure, AOTM automatically enables one of the methods for ownership takeover without operator intervention, improving the availability of the TS7700. It uses the TS3000 System Console associated with each TS7700 to provide an alternate path to check the status of a peer TS7700.

Without AOTM, an operator must determine if one of the TS7700 clusters has failed, and then enable one of the ownership takeover modes. This is required to access the logical volumes that are owned by the failed cluster. It is very important that write ownership takeover be enabled only when a cluster has failed, and not when there is a problem only with communication between the TS7700 clusters.

If it is enabled and the cluster in question continues to operate, data might be modified independently on other clusters, resulting in a corruption of the data. Although there is no data corruption issue with the read ownership takeover mode, it is possible that the remaining clusters might not have the latest version of the logical volume and present previous data.

Even if AOTM isn't enabled, it is advised that it be configured. This provides protection from a manual takeover mode being selected when the other cluster is still functional.

With AOTM, one of the takeover modes is enabled if normal communication between the clusters is disrupted and the cluster to perform takeover can verify that the other cluster has failed or is otherwise not operating. If a TS7700 suspects that the cluster that owns a volume it needs has failed, it asks the TS3000 System Console to which it is attached to query the System Console attached to the suspected failed cluster.

If the remote system console can validate that its TS7700 has failed, it replies back and the requesting TS7700 enters the default ownership takeover mode. If it cannot validate the failure, or if the system consoles cannot communicate, an ownership takeover mode can only be enabled by an operator.

To take advantage of AOTM, the customer should provide IP communication paths between the TS3000 System Consoles at the cluster sites. For AOTM to function properly, it should not share the same paths as the Grid interconnection between the TS7700s.

Note: When the TSSC code level is Version 5.3.7 or higher, the AOTM and Call Home IP addresses can be on the same subnet. However, earlier levels of TSSC code require the AOTM and Call Home IP addresses to be on different subnets. It is advised to use different subnets for those interfaces.

IBM service enables or disables AOTM, and also sets the default ownership takeover mode that is to be enabled.

4.2.2 Defining grid copy mode control

When upgrading a stand-alone cluster to a grid, FC4015, Grid Enablement must be installed on all clusters in the grid. Also, you must set up the Copy Consistency Points in the Management Class (MC) definitions on all clusters in the new grid. The data consistency point is defined in the MC's construct definition through the MI. You can perform this task only for an existing grid system.

In a stand-alone cluster configuration, you can choose between three consistency points per cluster:

- ▶ No Copy (NC): No copy 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 (DEF): A replication of the modified logical volume is made to this cluster after the volume had been unloaded.

- ▶ **Synchronous Copy:** Provides tape copy capabilities up to synchronous-level granularity across two clusters within a multi-cluster grid configuration. For more information, see “Synchronous mode copy” on page 81.
- ▶ You might consider the new option that is introduced in Release 3.1, the *Time Delayed Replication* policy, which enables better control of what data needs to be replicated to the TS7740, TS7720T, or TS7760T in a mixed grid. If a large portion of the data that is written to tape expires quickly in your environment, Time Delayed Replication makes it possible to delay the copies to a remote cluster for later than the average Lifecycle of your data.

Most of the data then expires before the time set for the delayed copies runs out, avoiding the processor burden introduced by the replication of archive or short retention data, and later the additional reclamation activity on the TS7740, TS7720T, or TS7760T cluster. Time delay can be set from 1 hour to 65,535 hours.

For more information, see the following links for details about this subject:

- ▶ *IBM TS7700 Series Best Practices - TS7700 Hybrid Grid Usage:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>
- ▶ *IBM TS7700 Series Best Practices - Copy Consistency Points:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101230>
- ▶ *IBM TS7700 Series Best Practices - Synchronous Mode Copy:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102098>

Define Copy Policy Override settings

With the TS7700, 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 each cluster in a multi-cluster grid configuration, which means that each cluster can have different settings, tailored to meet your requirements. The settings take effect for any mount requests received after you save the changes. 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**
A scratch (Fast Ready) mount selects a local copy if a cluster Copy Consistency Point is not specified as No Copy in the MC for the mount. The cluster is not required to have a valid copy of the data.
- ▶ **Prefer local cache for private (non-Fast Ready) mount requests**
This override causes the local cluster to satisfy the mount request if 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, the default cluster selection criteria applies.

Important: The Synchronous mode copy feature takes precedence over any Copy Override settings.

- ▶ **Force volumes that are mounted on this cluster to be copied to the local cache**
For a private (non-Fast Ready) mount, this override causes a copy to be created on the local cluster as part of mount processing. For a scratch (Fast Ready) mount, this setting overrides the specified MC with a Copy Consistency Point of Rewind-Unload for the cluster. This does not change the definition of the MC, but serves to influence the Replication policy.

- ▶ Enable fewer RUN consistent copies before reporting RUN command complete
If selected, the value that is entered for Number of required RUN consistent copies, including the source copy, is used to determine the number of copies to override before the RUN operation reports as complete. If this option is not selected, the MC definitions are used explicitly. Therefore, the number of RUN copies can be from one to the number of clusters in the grid.
- ▶ Ignore cache preference groups for copy priority
If this option is selected, copy operations ignore the cache preference group when determining the priority of volumes that are copied to other clusters.

Consideration: 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 cluster is preferred for all I/O operations. If the TS7700 cluster of the GDPS primary site fails, you must complete the following recovery actions:

1. Vary on virtual devices from a remote TS7700 cluster from the primary site of the GDPS host.
2. Manually start, through the TS7700 MI, a read/write Ownership Takeover (WOT), unless AOTM already has transferred ownership.

4.2.3 Defining scratch mount candidates

Scratch allocation assistance (SAA) is an extension of the device allocation assistance (DAA) function for scratch mount requests. SAA filters the list of clusters in a grid to return to the host a smaller list of candidate clusters that are designated as scratch mount candidates.

If you have a grid with two or more clusters, you can define scratch mount candidates. For example, in a hybrid configuration, the scratch allocation assist (SAA) function can be used to direct certain scratch allocations (workloads) to one or more TS7700Ds or cache partition (CP0) of a TS7700Ts for fast access, while other workloads can be directed to TS7740s or the cache partition (CPx) of TS7760Ts for archival purposes.

Clusters not included in the list of scratch mount candidates are not used for scratch mounts at the associated MC unless those clusters are the only clusters that are known to be available and configured to the host. If you have enabled SAA, but not selected any cluster as SAA candidates in the management class, all clusters are treated as SAA candidates.

Understand that SAA only influences the mount behavior of the grid. While other clusters can be selected as mount point if the original SAA clusters are not available or not configured to the host, they will not be considered for the TVC selection. If all clusters specified in the management class as target are not available, the mount may be processed, but the job will hang afterwards.

Before SAA is visible or operational, the following prerequisites must be true:

- ▶ All clusters in the grid have a Licensed Internal Code level of 8.20.0.xx. The necessary z/OS host software support is installed. For more information, see authorized program analysis report (APAR) OA32957 in the Tech docs Library link in “Related publications” on page 959.
- ▶ The z/OS environment uses job entry subsystem 2 (JES2) or JES3 (starting with z/OS V2R1).

Tip: DAA and SAA support for JES3 was added in z/OS V2R1. Before this release (and to avoid JES3 job abends), if the composite library was being shared between JES2 and JES3, the MCs being used for JES3 did not enable SAA through the Scratch Mount Candidate option on the MCs assigned to JES3 jobs.

- ▶ The SAA function must be enabled in the grid using the **LI REQ SETTING**.

4.2.4 Retain Copy mode

Retain Copy mode is an optional setting where a volume's existing Copy Consistency Points are accepted rather than 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 from being created in the grid than wanted. This is important in a grid with three or more clusters that has two or more clusters online to a host.

4.2.5 Defining cluster families

If you have a grid with three or more clusters, you can define *cluster families*.

This function introduces a concept of grouping clusters together into families. Using cluster families, you can define a common purpose or role to a subset of clusters within a grid configuration. The role that is assigned, for example, production or archive, is used by the TS7700 Licensed Internal Code to make improved decisions for tasks, such as replication and TVC selection. For example, clusters in a common family are favored for TVC selection, or replication can source volumes from other clusters within its family before using clusters outside of its family.

4.2.6 TS7720 and TS7760 cache thresholds and removal policies

These thresholds determine the state of the cache as it relates to remaining free space.

Cache thresholds for a TS7720 or TS7760 cluster

There are three thresholds that define the capacity of CP0 in a TS7720T or TS7760T and the active cache capacity in a TS7720D or TS7760D. These thresholds determine the state of the cache as it relates to remaining free space.

The following list describes the three thresholds in ascending order of occurrence:

- ▶ Automatic Removal

The policy removes the oldest logical volumes from the TS7720 or TS7760 cache if a consistent copy exists elsewhere in the grid. This state occurs when the cache is 3 TB below the out-of-cache-resources threshold. In the automatic removal state, the TS7720 or TS7760 automatically removes volumes from the disk-only cache to prevent the cache from reaching its maximum capacity.

This state is identical to the limited-free-cache-space-warning state unless the Temporary Removal Threshold is enabled. You can also lower the removal threshold in the LI REQ. The default is 4 TB.

To perform removal operations in a TS7720T or TS7760T, the size of CP0 must be at least 10 TB:

- You can disable automatic removal within any specific TS7720 cluster by using the following **LIBRARY REQUEST** command:

```
LIBRARY REQUEST,library-name,CACHE,REMOVE,{ENABLE|DISABLE}
```

- The default automatic removal threshold can be changed from the command line by using the following library request command:

```
LIBRARY REQUEST,library-name,CACHE,REMVTHR,{VALUE}
```

Note: The automatic removal function was introduced in R1.6, whereas the library request support was introduced in R1.7.

Automatic removal is temporarily disabled while disaster recovery write protect is enabled on a disk-only cluster so that a DR test can access all production host-written volumes. When the write protect state is lifted, automatic removal returns to normal operation.

► Limited free cache space warning

This state occurs when there is less than 3 TB of free space that is left in the cache. After the cache passes this threshold and enters the limited-free-cache-space-warning state, write operations can use only an extra 2 TB before the out-of-cache-resources state is encountered. When a TS7720 or TS7760 enters the limited-free-cache-space-warning state, it remains in this state until the amount of free space in the cache exceeds 3.5 TB.

The following messages can be displayed on the MI during the limited-free-cache-space-warning state:

- HYDME0996W
- HYDME1200W

For more information about these messages, see the TS7700 IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_removal_policies.html

Clarification: Host writes to the TS7720 or TS7760 and inbound copies continue during this state.

► Out of cache resources

This state occurs when there is less than 1 TB of free space that is left in the cache. After the cache passes this threshold and enters the out-of-cache-resources state, it remains in this state until the amount of free space in the cache exceeds 3.5 TB. When a TS7720 or TS7760 is in the out-of-cache-resources state, volumes on that cluster become read-only and one or more out-of-cache-resources messages are displayed on the MI. The following messages can display:

- HYDME0997W
- HYDME1133W
- HYDME1201W

For more information about these messages, see the TS7700 IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_removal_policies.html

Clarification: Although host allocations are not aware of a TS7720 or TS7760 in the out of cache resource state, the TS7700 grid avoids using a TS7720 or TS7760 in this state as a valid TVC candidate. New host allocations sent to a TS7720 or TS7760 in this state choose a remote TVC instead.

If all valid clusters are in this state or unable to accept mounts, the host allocations fail. Read mounts can choose the TS7720 or TS7760 in this state, but modify and write operations fail. Copies inbound to this TS7720 or TS7760 are queued as Deferred until the TS7720 or TS7760 exits this state.

Table 4-14 displays the start and stop thresholds for each of the active cache capacity states that are defined.

Table 4-14 Active cache capacity state thresholds

State	Enter state (free space available)	Exit state (free space available)	Host message displayed
Automatic removal	< 4 TB	> 4.5 TB	CBR3750I when automatic removal begins
Limited free cache space warning (CP0 for a TS7720T)	< 3 TB	> 3.5 TB or 15% of the size of CP0, whichever is less	CBR3792E upon entering state CBR3793I upon exiting state
Out of cache resources (CP0 for a TS7720T)	< 1 TB	> 3.5 TB or 5% of the size of CP0, whichever is less	CBR3794A upon entering state CBR3795I upon exiting state
Temporary removal ^a	< (X = 1 TB) ^b	> (X + 1.5 TB) ^b	Console message

a. When enabled

b. Where X is the value set by the TVC window on the specific cluster

Volume removal policies in a grid configuration

Removal policies determine when virtual volumes are removed from the cache of a TS7720 or TS7760 cluster in a grid configuration. These policies provide more control over the removal of content from a TS7720 or TS7760 cache as the active data reaches full capacity. To perform removal operations in a TS7720T or TS7760T cluster, the size of CP0 must be at least 10 TB.

To ensure that data is always in a TS7720 or TS7760, or is in for at least a minimal amount of time, a volume copy retention time must be associated with each removal policy. This volume retention time in hours enables volumes to remain in a TS7720 or TS7760 TVC for at least x hours before it becomes a candidate for removal, where x is 0 - 65,536. A volume retention time of zero assumes no minimal requirement.

In addition to pin time, three policies are available for each volume within a TS7720D or TS7760D and for CP0 within a TS7720T or TS7760T. For more information, see Chapter 2, "Architecture, components, and functional characteristics" on page 15.

Removal threshold

The default, or permanent, removal threshold is used to prevent a cache overrun condition in a TS7720 or TS7760 cluster that is configured as part of a grid. By default, it is a 4 TB (3 TB fixed plus 1 TB) value that, when taken with the amount of used cache, defines the upper size limit for a TS7720 or TS7760 cache, or for a TS7720T or TS7760T CP0.

Above this threshold, virtual volumes are removed from a TS7720 or TS7760 cache.

Note: Logical volumes are only removed if there is another consistent copy within the grid.

Logical volumes are removed from a TS7720 or TS7760 cache in this order:

1. Volumes in scratch categories
2. Private volumes that are least recently used by using the enhanced removal policy definitions

After removal begins, the TS7720 or TS7760 continues to remove logical volumes until the stop threshold is met. The stop threshold is a value that is the removal threshold minus 500 GB.

A particular logical volume cannot be removed from a TS7720 or TS7760 cache until the TS7720 or TS7760 verifies that a consistent copy exists on a peer cluster. If a peer cluster is not available, or a volume copy has not yet completed, the logical volume is not a candidate for removal until the appropriate number of copies can be verified later. Time delayed replication can alter the removal behavior.

Tip: This field is only visible if the selected cluster is a TS7720 or TS7760 in a grid configuration.

Temporary removal threshold

The temporary removal threshold lowers the default removal threshold to a value lower than the stop threshold in anticipation of a service mode event, or before a DR test where FlashCopy for DR testing is used.

Logical volumes might need to be removed before one or more clusters enter service mode. When a cluster in the grid enters service mode, remaining clusters can lose their ability to make or validate volume copies, preventing the removal of enough logical volumes. This scenario can quickly lead to the TS7720 or TS7760 cache reaching its maximum capacity.

The lower threshold creates more free cache space, which enables the TS7720 or TS7760 to accept any host requests or copies during the service outage without reaching its maximum cache capacity.

The temporary removal threshold value must be greater than or equal to (\geq) the expected amount of compressed host workload that is written, copied, or both to the TS7720 or TS7760 during the service outage. The default temporary removal threshold is 4 TB, which provides 5 TB (4 TB plus 1 TB) of existing free space. You can lower the threshold to any value from 2 TB to full capacity minus 2 TB.

All TS7720 or TS7760 clusters in the grid that remain available automatically lower their removal thresholds to the temporary removal threshold value that is defined for each one. Each TS7720 or TS7760 cluster can use a different temporary removal threshold. The default temporary removal threshold value is 4 TB or 1 TB more data than the default removal threshold of 3 TB. Each TS7720 or TS7760 cluster uses its defined value until the originating cluster in the grid enters service mode or the temporary removal process is canceled. The cluster that is initiating the temporary removal process does not lower its own removal threshold during this process.

4.2.7 Data management settings (TS7740/TS7700T CPx in a multi-cluster grid)

The following settings for the TS7700 are optional. Your IBM SSR configures these settings during the installation of the TS7700, or later through the TS7700 System Management Interface Tool (SMIT) menu. The following data management settings can be selected in SMIT:

- ▶ Copy Files Preferred to Reside in Cache
- ▶ Recalls Preferred for Cache Removal

Copy files preferred to reside in cache

Normally, the TVCs in both TS7700 tape drives in a multi-cluster grid are managed as one TVC to increase the likelihood that a needed volume is in cache. By default, the volume on the TS7700 that is selected for I/O operations is preferred to stay in cache on that TS7700. The copy that is made on the other TS7700 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 TS7700 running in a dual production multi-cluster grid configuration, both TS7700 tape drives are selected as the I/O TVCs, and have the original volumes (newly created or modified) preferred in cache. The copies to the other TS7700 are preferred to be removed from cache. Therefore, each TS7700 TVC is filled with unique, newly created, or modified volumes, roughly doubling the amount of cache seen by the host.

For a TS7700 running in a multi-cluster grid configuration that is used for business continuance, particularly when all I/O is preferred to the local TVC, this default management method might not be wanted. If 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 needed is to have the most recent copy volumes remain in the cache, not being preferred out of cache.

Based on your requirements, your IBM SSR can set or modify this control through the TS7700 SMIT menu for the remote TS7700:

- ▶ The default is set to off.
- ▶ When off, copy files are managed as PG0 volumes (prefer largest files out of cache first).
- ▶ When set to on, copy files are managed based on the Storage Class (SC) construct definition at the copy target cluster.

Recalls preferred for cache removal

Normally, a volume recalled into cache is managed as though it were newly created or modified because it is in the TS7700 that is selected for I/O operations on the volume. A recalled volume displaces other volumes in cache.

If the remote TS7700 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, so some number of recalls is 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 later in the recovery process.

After a restore completes from a recalled volume, that volume is no longer needed. These volumes must be removed from the cache after they have been accessed so that they minimally displace other volumes in the cache.

Based on your requirements, the IBM SSR can set or modify this control through the TS7700 SMIT menu of the remote TS7700:

- ▶ When off, which is the default, recalls are managed as PG1 volumes (LRU).
- ▶ When on, recalls are managed as PG0 volumes (prefer out of cache first by largest size).

This control is independent of and not affected by cache management controlled through the SC storage management subsystem (SMS) construct. SC cache management affects only how the volume is managed in the I/O TVC.

High availability means being able to provide continuous access to logical volumes through planned and unplanned outages with as little user effect or intervention as possible. It does not mean that all potential for user effect or action is eliminated. The following guidelines relate to establishing a grid configuration for HA:

- ▶ The production systems, which are the sysplexes and logical partitions (LPARs), have FICON channel connectivity to both clusters in the grid. The IBM Data Facility Storage Management Subsystem (DFSMS) library definitions and input/output definition file (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 a planned or unplanned outage to enable production to continue.

- ▶ For the workload placed on the grid configuration, when using only one of the clusters, performance throughput needs to be sufficient to meet service level agreements (SLAs). Assume that both clusters are normally used by the production systems (the virtual devices in both clusters are varied online to production). In the case where one of the clusters is unavailable, the available performance capacity of the grid configuration can be reduced by up to one half.
- ▶ For all data that is critical for high availability, consider using an MC whose Copy Consistency Point definition has both clusters with a Copy Consistency Point of RUN (immediate copy) or SYNC (sync mode copy). Therefore, each cluster has a copy of the data when the following conditions occur:
 - The volume is closed and unloaded from the source cluster for immediate copy.
 - Both clusters have copies that are written at the same time with Synchronous mode copy.
- ▶ The following types of applications can benefit from Synchronous mode copy (SMC):
 - DFSMS Hierarchical Storage Manager (DFSMSHsm)
 - DFSMS Data Facility Product (DFSMSdfp) OAM Object Support
 - Other applications that use data set-style stacking
 - Any host application that requires zero recovery point objective (RPO) at sync point granularity

The copy is updated at the same time as the original volume, keeping both instances of this logical volume synchronized at the record level. See Chapter 2, “Architecture, components, and functional characteristics” on page 15 for a detailed description.

- ▶ The distance of grid links between the clusters might influence the grid link performance. Job execution times that use Synchronous or Immediate mode might be affected by this

factor. Low-latency directors, switches, or DWDMs might help to optimize the network performance. Avoid network quality of service (QoS) or other network sharing methods because they can introduce packet loss, which directly reduces the effective replication bandwidth between the clusters.

- ▶ To improve performance and take advantage of cached versions of logical volumes, do not configure the Prefer Local Cluster for private mounts and Force Local Copy Override settings in either cluster. This setting is suggested for homogeneous TS7720D or TS7760D grids. See 11.21, “Virtual Device Allocation in z/OS with JES2” on page 741.
- ▶ To minimize operator actions when a failure has occurred in one of the clusters, which makes it unavailable, set up the AOTM to automatically place the remaining cluster in at least the Read Ownership Takeover (ROT) mode. Use read/WOT mode if you want to modify existing tapes, or if you think that your scratch pool might not be large enough without using those scratch volumes that are owned by the downed cluster.

If AOTM is not used, or it cannot positively determine whether a cluster has failed, an operator must determine whether a cluster has failed and, through the MI on the remaining cluster, manually select one of the ownership takeover modes.

- ▶ If multiple grid configurations are available for use by the same production systems, you can optionally remove the grid that experienced an outage from the Storage Group (SG) for scratch allocations. This directs all scratch allocations to fully functional grids while still enabling reads to access the degraded grid. This approach might be used if the degraded grid cannot fully complete the required replication requirements. Use this approach only for read access.

By following these guidelines, the TS7700 grid configuration supports the availability and performance goals of your workloads by minimizing the effect of the following outages:

- ▶ Planned outages in a grid configuration, such as Licensed Internal Code 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. For the logical volumes with an Immediate or Synchronous Copy policy effective, all jobs that completed before the outage 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 reissued 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).

If it is necessary, access existing data to complete the job. For more details about AOTM, see 2.3.34, “Autonomic Ownership Takeover Manager” on page 90. For jobs that were writing data, the written data is not accessible and the job must start again.

Important: Scratch categories and Data Classes (DCs) definitions are defined at the system level. Therefore, if you modify them in one cluster, it applies to all clusters in that grid.

4.3 Planning for software implementation

This section provides information for planning tasks that are related to host configuration and software requirements for use with the TS7700.

4.3.1 Host configuration definition

Library names, Library IDs, and port IDs are used to define the TS7700 to the host at the hardware, operating system, and SMS levels. Some of these identifiers are also used by the IBM SSR in the hardware configuration phase of installation.

On the host side, definitions must be made in HCD and in the SMS. For an example, see Table 4-15, and create a similar one during your planning phase. It is used in later steps. The Library ID must contain only hexadecimal characters (0 - 9 and A - F).

Table 4-15 Sample of library names and IDs in a four-cluster grid implementation

TS7700 virtual library names	SMS name ^a	LIBRARY-ID	Defined in HCD	Defined in SMS
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. The composite library name is linked to the composite library ID when defining the tape library to DFSMS, as shown in Figure 6-6 on page 224. In the same manner, the distributed library name is linked to the distributed library ID, as shown in Figure 6-9 on page 225. Use names that are similar to those listed in Table 4-15.

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 the configuration parameters that are used by the IBM SSR during the installation of the TS7700. These names are not defined to the TS7700 hardware. However, to make administration easier, associate the LIBRARY-IDs with the SMS library names through the nickname setting in the TS7700 MI.

Remember: Match the distributed and composite library names that are entered at the host with the nicknames that are defined at the TS7700 MI. Although they do not have to be the same, this guideline simplifies the management of the subsystem.

LIBRARY-ID and LIBPORT-ID

LIBRARY-ID and LIBPORT-ID are z/OS HCD parameters that enable 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 enables the logical tape devices to be varied online (when they later become available to the system) without reactivating the IODF.

Tip: Specify the LIBRARY-ID and LIBPORT-ID in your HCD/IOCP definitions, even in a stand-alone configuration. This configuration reduces the likelihood of having to reactivate the IODF when the library is not available at IPL, and provides enhanced error recovery in certain cases. It might also eliminate the need to have an IPL when you change your I/O configuration. In a multicluster configuration, LIBRARY-ID and LIBPORT-ID must be specified in HCD, as shown in Table 4-15.

Distributed library ID

During installation planning, each cluster is assigned a unique, five-digit hexadecimal number (that is, the sequence number). This number 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 can start with the letter D.

In addition to the letter D, you can use the last four digits of the hardware serial number if it consists only of hexadecimal characters. For each distributed library ID, it is the last four digits of the TS7700 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 following examples can be the distributed library IDs of a four-cluster grid configuration:

Cluster 0	DA01A
Cluster 1	DA01B
Cluster 2	DA01C
Cluster 3	DA01D

The composite library ID for this four-cluster grid can then be CA010.

Important: Whether you are using your own or IBM nomenclature, the important point is that the subsystem identification must be clear. Because the identifier that appears in all system messages is the SMS library name, it is 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 configuration during hardware installation. All TS7700 tape drives participating in a grid have the same composite library ID. In the 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 identify uniquely each composite library in a meaningful way. The sequence number must match the LIBRARY-ID that is used in the HCD library definitions and the LIBRARY-ID that is listed in the Interactive Storage Management Facility (ISMF) Tape Library definition windows.

Remember: In all configurations, each LIBRARY-ID, whether distributed or composite, must be unique.

LIBPORT-ID

Each logical control unit (LCU), or 16-device group, must present a unique subsystem identification to the z Systems host. This ID is a 1-byte field that uniquely identifies each LCU within the cluster, and is called the *LIBPORT-ID*. The value of this ID cannot be 0.

Table 4-16 shows the definitions of the LIBPORT-IDs in a multi-cluster grid. For Cluster 0, 256 devices is 01 - 10 and 496 devices is 01 - 1F. LIBPORT-ID is always one more than CUADD.

Table 4-16 Subsystem identification definitions

Cluster	Logical CU (hex)	LIBPORT-ID (hex)
0	0 - 1E	X'01'-X'1F'
1	0 - 1E	X'41'-X'5F'
2	0 - 1E	X'81'-X'9F'
3	0 - 1E	X'C1'-X'DF'
4	0 - 1E	X'21'-X'3F'
5	0 - 1E	X'61'-X'7F'

Virtual tape drives

The TS7700 presents a tape drive image of a 3490 C2A, identical to the IBM Virtual Tape Server (VTS) and peer-to-peer (PTP) subsystems. Command sets, responses to inquiries, and accepted parameters match the defined functional specifications of a 3490E drive. Depending on the machine model and installed features, this collection can contain up to 31 LCUs and 496 virtual drives. Virtual drives are organized in groups of 16 drive addresses under a single LCU address.

4.3.2 Software requirements

The TS7700 is supported at z/OS V2R1 or later (earlier release level support must be done through the RPQ process). For more information about the support that is provided for the specified releases of the TS7700, see the following APARs:

- ▶ APAR OA32957 for Release 2.0.
- ▶ APAR OA37267 for Release 2.1.
- ▶ No additional host software support is provided for Release 3.0.
- ▶ APAR OA40572 for Release 3.1.
- ▶ APAR OA44351 is advised for Release 3.2.
- ▶ APAR OA47487 is advised for Release 3.3.
- ▶ APAR OA49373 is advised for Release 4.0.

In general, install the host software support. See the VTS, PTP, and 3957 Preventive Service Planning (PSP) topics on the IBM Support and Downloads web page (ibm.com/support) for the current information about Software Maintenance.

4.3.3 System-managed storage tape environments

System-managed tape enables you to manage tape volumes and tape libraries according to a set of policies that determine the 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 (DC, SC, MC, and SG) that reflect your installation's policies for the data on that volume.

The ACS routines are started for every new allocation. Tape allocations are passed to the OAM, which uses its Library Control System (LCS) component to communicate with the Integrated Library Manager.

The SC ACS routine determines whether a request is SMS-managed. If no SC 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 SG routine assigns the request to an SG. The assigned SG determines which LPARs in the tape library are used. Through the SG construct, you can direct logical volumes to specific tape libraries.

In addition to defining new SMS classes in z/OS, the new SMS classes must be defined in the TS7700 through the MI. This way, the name is created in the list and the default parameters are assigned to it. Figure 4-7 shows the default MC in the first line and another MC defined as described in the second line.

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"/>	IIINNNNN	0	The default Management Class	No	Rewind Unload (RUN)	Rewind Unload (RUN)

Figure 4-7 Default construct

4.3.4 Sharing and partitioning considerations

This section includes the following topics:

- ▶ Tape management system and OAM
- ▶ Partitioning the physical media (TS7740, TS7720T, or TS7760T) between multiple hosts

Tape management system and OAM

Your tape management system (TMS) enables the management of removable media across systems. The TMS manages your tape volumes and protects the data sets on those volumes. It handles expiring data and scratch tapes according to policies that you define.

Data Facility System Managed Storage Removable Media Manager (DFSMSrmm) is one such TMS that is included as a component of z/OS. The placement and access to the disk that contains the DFSMSrmm control data set (CDS) determines whether a standard or client/server subsystem (RMMplex) should be used. If all z/OS hosts have access to a shared disk, an RMMplex is not necessary.

Review the Redbooks publication *DFSMSrmm Primer*, SG24-5983 and the *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874 for further information about which RMM subsystem is correct for your environment.

The OAM is a component of DFSMSdfp that is included with z/OS as part of the storage management system (SMS). Along with your TMS, OAM uses the concepts of system-managed storage to manage, maintain, and verify tape volumes and tape libraries

within a tape storage environment. OAM uses the tape configuration database (TCDB), which consists of one or more volume catalogs, to manage volume and library entries.

If tape libraries are shared among hosts, they must all have access to a single TCDB on shared disk, and they can share the DEVSUPxx parmlib member. If the libraries are to be partitioned, each set of sharing systems must have its own TCDB. Each such TCDBplex must have a unique DEVSUPxx parmlib member that specifies library manager categories for each scratch media type, error, and private volumes.

Planning what categories are used by which hosts is an important consideration that needs to be addressed before the installation of any tape libraries. For more information about OAM implementation and category selection, see *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

Partitioning the physical media (TS7740, TS7720T, or TS7760T) between multiple hosts

The virtual drives and virtual volumes in a TS7700 can be partitioned just like physical drives and real volumes. Any virtual volume can go to any physical stacked volume when you use a TS7740, TS7720T, or TS7760T. The TS7700 places no restrictions on the use and management of those resources. When you use a TS7740, TS7720T, or TS7760T, you can partition your stacked media in up to 32 separate pools by assigning an SG to a defined range of stacked volumes before insertion time.

4.3.5 Sharing the TS7700 by multiple hosts

Each multi-cluster grid or stand-alone grid has its own library sequence number, which is used to define the logical library to the host. Each logical library that is identified as a composite library looks like a separate library to the host. A TS7700 can be shared by multiple z/OS, VM, VSE, and TPF systems.

Sharing can be achieved in two ways:

- ▶ By logically dividing the TS7700 into separate partitions (*partitioning*)
- ▶ By enabling all attached systems to sequentially access all physical and logical volumes (*sharing*)

Sharing of an IBM automated tape library (ATL) means that all attached hosts have the same access to all volumes in the tape library. To achieve this sharing, you need to share the host CDSs, that is, the TMS inventory and the integrated catalog facility (ICF) catalog information, among the attached hosts.

Additionally, you need to have the same categories defined in the DEVSUPxx member on all hosts. In a non-SMS environment, all systems must share the ICF catalog that contains the Basic Tape Library Support (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 per media type is available.

4.3.6 Partitioning the TS7700 between multiple hosts

Partitioning is the solution if you need to dedicate the use of volume ranges to certain systems or complexes, or separate host platforms. Dividing one or more libraries into logical libraries is the easiest way to enable different hosts to access them. Each host or complex owns its own set of drives, volumes, and DEVSUPxx scratch categories that another system or complex cannot access. Each system knows only about its part of the library. Partitioning is also appropriate for the attachment to a z/OS LPAR for testing.

This partition is implemented through values that are updated in the DEVSUPxx category definitions. Until now, to modify a category value you needed to change the DEVSUPxx member and restart the system. A new command, **DS QLIB, CATS**, enables you to query and modify these category values without an initial program load (IPL). However, this command must be used with great care because a discrepancy in this area causes scratch mounts to fail.

Partitioning the TS7700 with Selective Device Access Control

SDAC enables exclusive access to one or more volume serial number (VOLSER) ranges by only certain LCUs or subsystem IDs within a composite library for host-initiated mounts, ejects, and changes to attributes or categories.

You can use SDAC to configure hard partitions at the LIBPORT-ID level for independent host LPARs or system complexes. Hard partitioning prevents a host LPAR or system complex with an independent tape management configuration from inadvertently modifying or removing data that is owned by another host. It also prevents applications and users on one system from accessing active data on volumes that are owned by another system.

SDAC is enabled by using FC 5271, Selective Device Access Control. This feature license key must be installed on all clusters in the grid before SDAC is enabled. You can specify one or more LIBPORT-IDs per SDAC group. Each access group is given a name and assigned mutually exclusive VOLSER ranges. Use the Library Port Access Groups window on the TS7700 MI to create and configure Library Port Access Groups for use with SDAC.

Access control is imposed as soon as a VOLSER range is defined. As a result, selective device protection applies retroactively to pre-existing data. A case study about sharing and partitioning the TS7700 is in Appendix I, "Case study for logical partitioning of a two-cluster grid" on page 937.

4.3.7 Logical path considerations

The TS7700 attaches to the host system or systems through two or four FICON adapters. For the 8 Gb FICON adapter, each channel that is connected to the FICON adapter port supports 512 logical paths (4 Gb FICON adapter continues to support 256 paths per port). A four FICON (8 Gb adapter) configuration with dual ports enabled results in a total of 4,096 logical paths per TS7700:

Four adapters x 2 ports x 512 paths_per_port=4,096 total paths

To calculate the number of logical paths that are required in an installation, use the following formula:

Number of logical paths per FICON channel = number of LPARs x number of CUs

This formula assumes that all LPARs access all CUs in the TS7700 with all channel paths. For example, if one LPAR has 16 CUs defined, you are using 16 logical paths of the 512 logical paths available on each FICON adapter port.

The *FICON Planning and Implementation Guide*, SG24-6497, covers the planning and implementation of FICON channels, and describes operating in FICON native (Fibre Channel (FC)) mode. It also describes the FICON and FC architectures, terminology, and supported topologies.

Define one tape CU in the HCD dialog for every 16 virtual devices. Up to eight channel paths can be defined to each CU. A logical path might be thought of as a three-element entity:

- ▶ A host port
- ▶ A TS7700 port
- ▶ A logical CU in the TS7700

Remember: A reduction in the number of physical paths reduces the throughput capability of the TS7700 and the total number of available logical paths per cluster. A reduction in CUs reduces the number of virtual devices available to that specific host.

4.4 Planning for logical and physical volumes

As part of your planning process, you need to determine the number of virtual and stacked physical volumes that are required for your workload. The topics in this section provide information to help you determine the total number of virtual volumes that are required, suggestions about the volume serial number (VOLSER) ranges to define, and the number of physical volumes required.

The VOLSER of the virtual and physical volumes must be unique throughout a system-managed storage complex (SMSplex) and throughout all storage hierarchies, such as DASD, tape, and optical storage media. To minimize the risk of misidentifying a volume, the VOLSER should be unique throughout the grid and across different clusters in different TS3500 or TS4500 tape libraries.

4.4.1 Volume serial numbering

Before you define logical volumes to the TS7700, consider the total number of logical volumes that are required, the volume serial ranges to define, and the number of volumes within each range. The VOLSERs for logical volumes and physical volumes in any cluster within the same grid must be unique.

The VOLSERs must be unique throughout an SMSplex and throughout all storage hierarchies. It must also be unique across LPARs connected to the grid. Have independent plexes use unique ranges in case volumes ever need to be shared. In addition, future merges of grids require that their volume ranges be unique.

Tip: Try not to insert an excessive amount of scratch that isn't likely to be used over a few months duration given that it can add processor burden to allocations, especially when expire with hold is enabled. Volumes can always be inserted later if scratch counts become low.

When you insert volumes, you do that by providing starting and ending volume serial number range values.

The TS7700 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 add logical volumes with VOLSERs of ABC000, ABC001, ABC002...ABC998, and ABC999 into the inventory of the TS7700. You might find it helpful to plan for growth by reserving multiple ranges for each TS7700 that you expect to install.

If you have multiple partitions, it is better to plan which ranges will be used in which partitions, for example, A* for the first sysplex and B* for the second sysplex. If you need more than one range, you can select A* and B* for the first sysplex, C* and D* for the second sysplex, and so on.

4.4.2 Logical volumes

Determine the number of logical volumes that are required to handle the workload that you are planning for the TS7700. The default number of logical volumes that is supported is 1,000,000. You can add support for more logical volumes in 200,000 volume increments (FC5270), up to a total of 4,000,000 logical volumes.

Tip: For 3957-V06/VEA, the limit is 2,000,000 logical volumes.

The TS7700 supports logical WORM (LWORM) volumes. Consider the size of your logical volumes, the number of scratch volumes you need per day, the time that is required for return-to-scratch processing, how often scratch processing is run, and whether you need to define LWORM volumes.

Size of logical volumes

The TS7700 supports logical volumes with maximum sizes of 400, 800, 1000, 2000, 4000, 6000, and 25,000 mebibytes (MiB), 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 6000 MiB logical volume is 18,000 MiB.

Depending on the logical volume sizes that you choose, you might see the number of volumes that are 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 6000 MiB logical volumes, you need more TS7700 logical volumes to store the data, which is stored as multivolume data sets.

The 400 MiB cartridge storage tape (CST)-emulated cartridges or 800 MiB with emulated enhanced capacity cartridge system tape (ECCST)-emulated cartridges are the two types you can specify when adding volumes to the TS7700. You can use these sizes directly, or use policy management to override them to provide for the 1000, 2000, 4000, 6000, or 25,000 MiB sizes.

A logical volume size can be set by VOLSER, and can change dynamically by using the DFSMS DC storage construct when a job requires a scratch volume or writes from beginning of tape (BOT). The amount of data that is copied to the stacked cartridge is only the amount of data that was written to a logical volume. The choice between all available logical volume sizes does not affect the real space that is used in either the TS7700 cache or the stacked volume.

In general, unless you have a special need for CST emulation (400 MiB), specify the ECCST media type when you insert volumes in the TS7700.

In planning for the number of logical volumes that is needed, first determine the number of private volumes that make up the current workload that 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, accounting for the compressibility of your data. If you do not know the average ratio, use the conservative value of 2:1.

If you choose to use only the 800 MiB volume size, the total number that is needed might increase depending on whether current volumes that contain more than 800 MiB compressed need to expand to a multivolume set. Take that into account for planning the number of logical volumes required. Consider using smaller volumes for applications such as DFSMSHsm and larger volumes for backup and full volume memory dumps.

If you plan to use 25,000 MiB logical volumes, a maximum size of 25,000 MiB for logical volumes is allowed without any restriction if all clusters in a grid operate at R3.2 or higher level of Licensed Internal Code. Otherwise, the 25,000 MiB is not supported if one or more TS7740 clusters are present in the grid, and at least one cluster in the grid operates at an Licensed Internal Code level earlier than R3.2.

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 must 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 therefore 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 2x) and by the frequency with which you want to perform returns to scratch processing.

The following formula is suggested to calculate the number of logical volumes needed:

$$Vv = Cv + Tr + (Sc)(Si + 1)$$

The formula contains the following values:

- Vv** Total number of logical volumes needed
- Cv** Number of logical volumes that is 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), set to equal the number of scratch volumes that are used per night
- Sc** Number of empty scratch volumes that are used per night, rounded up to the nearest 500
- Si** Number of days between scratch processing (return-to-scratch) by the TMS

For example, assuming the current volume requirements (that use all the available volume sizes), that use 2500 scratch volumes per night, and running return-to-scratch processing every other day, you need to plan on the following number of logical volumes in the TS7700:

$$75,000 \text{ (current, rounded up)} + 2,500 + 2,500 (1+1) = 82,500 \text{ logical volumes}$$

If you plan to use the expired-hold option, take the maximum planned hold period into account when calculating the **Si** value in the previous formula.

If you define more volumes than you need, you can always eject the additional volumes. Unused logical volumes do not use space, but excessive scratch counts in the 100,000+ might add processor burden to scratch allocation processing.

The default number of logical volumes that is supported by the TS7700 is 1,000,000. You can add support for more logical volumes in 200,000 volume increments, up to a total of 4,000,000 logical volumes. This is the maximum number either in a stand-alone or grid configuration.

To make this upgrade, see how to use FC 5270 in the Increased logical volumes bullet in 7.2.1, “TS7700 concurrent system component upgrades” on page 251.

Consideration: Up to 10,000 logical volumes can be inserted at one time. Attempting to insert over 10,000 logical volumes at one time returns an error.

Number of scratch volumes needed

As you run your daily production workload, you need enough logical volumes in SCRATCH status to support the data that is written to the TS7700. This can be hundreds or thousands of volumes, depending on your workload.

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 to change the status of those volumes from private to scratch.

The amount of time the process takes depends on the type of TMS being employed, how busy the TS7700 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 TMSs return-to-scratch process when you migrate to or install a multicluster configuration solution.

If the number of logical volumes that is used daily is small (fewer than a few thousand), you might choose to run return-to-scratch processing only 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 4 hours, enough time exists during first shift to run the process twice if problems are encountered during the first attempt. Unless there are specific reasons, run return-to-scratch processing one time per day.

With z/OS V1.9 or later, return-to-scratch in DFSMSrmm has been enhanced to speed up this process. To reduce the time that is required for housekeeping, it is now possible to run several return-to-scratch processes in parallel. For more information about the enhanced return-to-scratch process, see the *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874.

Tip: The expire-hold option might delay the time that the scratch volume becomes usable again, depending on the defined hold period.

Preferred migration of scratch volumes

TS7740 clusters operating at Licensed Internal Code level 8.21.0.xx or later, TS7720T, and TS7760T use the preferred migration of scratch volumes enhancement, which migrates scratch volumes before migrating non-scratch volumes. This enhancement modifies the least recently used (LRU) algorithm to ensure that more critical data remains in cache for a longer period.

Under this preferred migration, hierarchical storage management (HSM) first migrates all volumes in a scratch category according to size (largest first). Only when all volumes (PG0 or PG1) in a scratch category have been migrated and the PG1 threshold is still unrelieved does HSM operate on private PG1 volumes in LRU order.

Note: You must define all scratch categories before using the preferred migration enhancement.

4.4.3 Logical WORM

TS7700 supports the LWORM function through TS7700 software emulation. The LWORM enhancement can duplicate most of the 3592 WORM behavior. The host views the TS7700 as an LWORM-compliant library that contains WORM-compliant 3490E logical drives and media. Similar to volume emulated capacity, the LWORM capability is dynamically selected through DATACLASS.

TS7700 reporting volumes (BVIR) cannot be written in LWORM format. For more information, refer to “Overview of the BVIR function” on page 700.

4.4.4 Physical volumes for TS7740, TS7720T, and TS7760T

This section describes the number of physical volumes that are required to accommodate the workload you are planning for the TS7740, TS7720T, and TS7760T. To determine the number of physical volumes that are required to accommodate your workload, consider the following information:

- ▶ Amount of data that is stored for a given host workload
- ▶ Average compression ratio that is achieved per workload
- ▶ Average utilization rate of filling physical volumes
- ▶ Scratch physical volumes

Amount of data that is stored for a given host workload

The amount of data that is stored per workload can be extracted from your Tape Management System, such as RMM, or from TS7700 by using VEHSTATS.

Average compression ratio that is achieved per workload

The data that a host writes to a virtual volume might be compressible. The space that is required on a physical volume is calculated after the effect of compression. If you do not know the average number for your data, assume a conservative 2:1 ratio.

Average utilization rate of filling physical volumes

The average utilization rate of filling physical volumes can be calculated from the Reclaim Threshold Percentage. This is 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 is performed on the stacked volume. The valid range of possible values is 0 - 95%; 35% is the default value. Therefore, the utilization rate of filling physical volumes should range 35% - 100%. The average utilization rate of filling physical volumes can be calculated as $(35+100)/2 = 67.5\%$.

Scratch physical volumes

Answer the following questions to determine the number of scratch physical volumes you need for each pool.

Is the E08 or E07 drive installed?

- ▶ If yes
 - Is borrow/return sharing enabled?
 - Yes: 15 volumes in the common scratch pool
 - No: 15 volumes in each dedicated pool
- ▶ If no
 - Is borrow/return sharing enabled?

- If yes: 50 volumes in the common scratch pool
- If no: 50 volumes in each dedicated pool

If the number of scratch physical volumes in your system is fewer than these thresholds, the following situations can occur:

- ▶ Reclamation of sunset media does not occur
- ▶ Reclamation runs more frequently

You can have less than 15 or 50 volumes in your pool if these conditions are acceptable. Keep in mind that you need at least two scratch physical volumes to avoid an out of scratch state.

The following is a suggested formula to calculate the number of physical volumes needed:

- ▶ For each workload, calculate the number of physical volumes needed:

$$P_x = (D_a/C_r)/(P_c \times U_t/100)$$

- ▶ Next, add in physical scratch counts and the P_x results from all known workloads:

$$P_v = P_s + P_1 + P_2 + P_3 + \dots$$

The formula contains the following values:

P_v	Total number of physical volumes needed
D_a	Total amount of data that is returned from your Tape Management System or VEHSTATS per workload
C_r	Compression Ratio per workload (Use $C_r=1$ when D_a represents previously compressed data)
U_t	Average utilization rate of filling physical volumes
P_c	Capacity of a physical volume in TB
P_x	Resulting number of physical volumes that are needed for a particular workload x
P_s	Number of physical volumes in common scratch pool

Using the suggested formula and the assumptions, plan to use the following number of physical volumes in your TS7700:

- ▶ Example 1 by using the following assumptions:

- $D_a = 100$ TB
 - $C_r = 2$
 - $U_t = 67.5\%$
 - $P_c = 10$ TB (capacity of a JD volume)
- $$P_1 = (100/2)/(10 \times 67.5/100) = 8 \text{ physical volumes}$$

- ▶ Example 2 by using the following assumptions:

- $D_a = 150$ TB
 - $C_r = 2$
 - $U_t = 67.5\%$
 - $P_c = 7$ TB (capacity of a JC volume in 3592-E08 format)
- $$P_2 = (150/2)/(7 \times 67.5/100) = 16 \text{ physical volumes}$$

If the number of physical volumes in the common scratch pool is $P_s = 15$, you would need to plan on the following number of physical volumes in the TS7740, the TS7720T, or the TS7760T:

$$P_v = P_s + P_1 + P_2 = 15 + 8 + 16 = 39 \text{ physical volumes}$$

If you need dual copied virtual volumes in a single cluster, you need to double the number of physical volumes for that workload. If a workload uses dedicated pools with the borrow/return sharing disabled, then each workload must have its own dedicated additional scratch count versus the shared P_s count.

If you are planning to use the Copy Export function, plan for enough physical volumes for the Copy Export function and enough 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 2,000 physical volumes in the Copy Export state. This number includes offsite 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 MI to a maximum value of 10,000. After your Copy Export operations reach a steady state, approximately the same number of physical volumes is being returned to the library for reclamation as there are those being sent offsite as new members of the Copy Export set of volumes.

4.4.5 Data compression

When writing data to a virtual volume, the host compression definition is accepted. Compression is turned on or off by the JCL parameter **DCB=TRTCH=COMP** (or **NOCOMP**), the DC parameter **COMPACTION=YES|NO**, or the **COMPACT=YES|NO** definition in the DEVSUPxx PARMLIB member. The **TRTCH** parameter overrides the DC definition, and both override the PARMLIB definition.

Important: To achieve the optimum throughput, verify your definitions to ensure that you specified compression for data that is written to the TS7700.

4.4.6 Secure Data Erase function

Expired data on a physical volume remains readable until the volume is overwritten with new data. Some clients prefer to delete the content of a reclaimed stacked cartridge, due to security or business requirements.

TS7740, TS7720T, and TS7760T implement the Secure Data Erasure on a pool basis. With the Secure Data Erase function, all reclaimed physical volumes in that pool are erased by writing a random pattern across the whole tape before reuse. If a physical volume has encrypted data, the erasure is accomplished by deleting EKs on the volume, rendering the data unrecoverable on this cartridge. A physical cartridge is not available as a scratch cartridge if its data is not erased.

Consideration: If you choose this erase function and you are not using tape encryption, TS7740, TS7720T, or TS7760T need time to erase every physical tape. Therefore, the TS7740, TS7720T, or TS7760T need more time and more back-end drive activity every day to complete reclamation and erase the reclaimed cartridges afterward. With tape encryption, the Secure Data Erase function is relatively fast.

The Secure Data Erase function also monitors the age of expired data on a physical volume and compares it with the limit set by the user in the policy settings. Whenever the age exceeds the limit that is defined in the pool settings, the Secure Data Erase function forces a reclaim and subsequent erasure of the volume.

In a heterogeneous drive configuration, older generations of tape drives are used for read-only operation. However, the Secure Data Erase function uses older generations of tape drives to erase older media (discontinued media) that cannot be written by 3592-E08 tape drives.

Note: In a homogeneous drive configuration with 3592-E07 drives or a heterogeneous drive configuration with 3592-E08 and E07 drives, JA/JJ media cannot be erased because 3592-E07 drives cannot write to JA/JJ media. If JA/JJ media exists in a pool where Secure Data Erase is enabled with 3592-E07 drives installed, the following text message is shown:

```
AL5019 The erase of physical volume xxxxxx is skipped due to the functional
limitation of the installed physical tape drives.
```

For more information about the Secure Data Erase function, see 2.2.24, “Secure Data Erase function” on page 52 and “Defining physical volume pools in the TS7700T” on page 531.

Encryption and Secure Data Erasure

If a physical volume is encrypted, the TS7700 does not perform a physical overwrite of the data. The EK is shredded, rendering the encrypted data unrecoverable.

When compared to the normal or long erasure operation, EK shredding is much faster. Normal erasure is always used for non-encrypted tapes, and EK shredding is the default that is used for encrypted tapes. The first time an encrypted tape is erased, a normal erasure is performed, followed by an EK shredding. A TS7700 can be configured to perform a normal erasure with every data operation, but this function must be configured by an IBM SSR.

4.4.7 Planning for tape encryption in a TS7740, TS7720T, and TS7760T

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 the physical tapes that are used by a TS7740, TS720T, and TS7760T helps control the risks of unauthorized data access without excessive security management burden or subsystem performance issues.

Encryption on the TS7740, TS7720T, and TS7760T is controlled on a storage pool basis. SG and MC DFSMS constructs specified for logical tape volumes determine which physical volume pools are used for the primary and backup (if used) copies of the logical volumes. The storage pools, originally created for the management of physical media, have been enhanced to include encryption characteristics.

The tape encryption solution in a TS7740, TS7720T, and TS7760T consists of several components:

- ▶ The TS7740, TS7720T, and TS7760T tape encryption solution uses either the IBM Security Key Lifecycle Manager (SKLM) or the IBM Security Key Lifecycle Manager for z/OS (ISKLM) as a central point from which all EK information is managed and served to the various subsystems.
- ▶ The TS1120, TS1130, TS1140, or TS1150 encryption-enabled tape drives are the other fundamental piece of TS7740, TS7720T, and TS7760T tape encryption, providing hardware that runs the cryptography function without reducing the data-transfer rate.

- ▶ The TS7740, TS7720T, or TS7760T provides the means to manage the use of encryption and the keys that are used on a storage-pool basis. It also acts as a proxy between the tape drives and the IBM Security Key Lifecycle Manager (SKLM) or IBM Security Key Lifecycle Manager for z/OS (ISKLM) by using Ethernet to communicate with the SKLM or ISKLM (or in-band through FICONs) to the tape drives. Encryption support is enabled with FC9900.

Rather than user-provided key labels per pool, the TS7740, TS7720T, and TS7760T can also support the use of default keys per pool. After a pool is defined to use the default key, the management of encryption parameters is run at the key manager. The tape encryption function in a TS7740, TS7720T, or TS7760T does not require any host software updates because the TS7740, TS7720T, or TS7760T controls all aspects of the encryption solution.

Although the feature for encryption support is client-installable, check with your IBM SSR for the prerequisites and related settings before you enable encryption on your TS7740, TS7720T, or TS7760T.

Tip: Pool encryption settings are *disabled* by default.

Encryption key managers

The encryption key managers must be installed, configured, and operational before you install the encryption feature on the TS7740, TS7720T, or TS7760T.

Note: The IBM Encryption Key Manager is not supported for use with TS1140 3592 E07 and TS1150 3592 E08 tape drives. Either the IBM Security Key Lifecycle Manager (SKLM) or the IBM Security Key Lifecycle Manager for z/OS (ISKLM) is required.

You also need to create the certificates and keys that you plan to use for encrypting your back-end tape cartridges.

Although it is possible to operate with a single key manager, configure two key managers for redundancy. Each key manager needs to have all of the required keys in its respective keystore. Each key manager must have independent power and network connections to maximize the chances that at least one of them is reachable from the TS7740, TS7720T, and TS7760T when needed.

If the TS7740, TS7720T, and TS7760T cannot contact either key manager when required, you might temporarily lose access to migrated logical volumes. You also cannot move logical volumes in encryption-enabled storage pools out of cache.

IBM Security Key Lifecycle Manager (SKLM)

You can use IBM Security Key Lifecycle Manager (SKLM) (formerly called IBM Tivoli Key Lifecycle Manager) to create, back up, and manage the lifecycle of keys and certificates that an enterprise uses. You can manage encryption of symmetric keys, asymmetric key pairs, and certificates. IBM Security Key Lifecycle Manager also provides a graphical user interface (GUI), command-line interface (CLI), and REST interface to manage keys and certificates.

IBM Security Key Lifecycle Manager waits for and responds to key generation or key retrieval requests that arrive through TCP/IP communication. This communication can be from a tape library, tape controller, tape subsystem, device drive, or tape drive.

Additional information can be obtained from the IBM Security Key Lifecycle Manager website:

http://www.ibm.com/support/knowledgecenter/SSWPVP_2.5.0/com.ibm.sk1m.doc_2.5/welcome.htm

Information that is related to Tivoli Key Lifecycle Manager can be found on the following website:

http://www.ibm.com/support/knowledgecenter/SSWPVP_1.0.0.3/com.ibm.tklm.doc/welcome.htm

IBM Security Key Lifecycle Manager for z/OS (ISKLM)

The IBM Security Key Lifecycle Manager for z/OS (ISKLM) has been available since April 2011. ISKLM removes the dependency on IBM System Services Runtime Environment for z/OS and DB2, which creates a simpler migration from IBM Encryption Key Manager.

ISKLM manages EKs for storage, simplifying deployment and maintaining availability to data at rest natively on the z Systems mainframe environment. It simplifies key management and compliance reporting for the privacy of data and compliance with security regulations.

Additional information can be obtained from the IBM Security Key Lifecycle Manager for z/OS website:

<http://www.ibm.com/software/tivoli/products/security-key-lifecycle-mgr-z>

Encryption capable tape drives

Data is encrypted on the back-end tape drives, so the TS7740, TS7720T, and TS7760T must be equipped with Encryption Capable tape drives, such as these tape drives:

- ▶ TS1120 3592 E05 (Encryption Capable) tape drives. Must be running in 3592 E05 native mode. TS1120 tape drives with FC5592 or FC9592 are Encryption Capable.
- ▶ TS1130 3592 E06 tape drives.
- ▶ TS1140 3592 E07 tape drives.
- ▶ TS1150 3592 E08 tape drives.

The TS7740, TS7720T, and TS7760T must not be configured to force the TS1120 drives into J1A mode. This setting can be changed only by your IBM SSR. If you need to update the Licensed Internal Code level, be sure that the IBM SSR checks and changes this setting, if needed.

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* tape media and decrypting information being read *from* tape media. It must be available in the network, and the TS7740, TS7720T, and TS7760T must be configured with the IP addresses and TCP/IP ports to find the encryption key managers in the network.

For a comprehensive TS7740, TS7720T, and TS7760T encryption implementation plan, see “Implementing TS7700 Tape Encryption” in *IBM System Storage Tape Encryption Solutions*, SG24-7320.

4.4.8 Planning for cache disk encryption in the TS7700

Release 3.0 introduced full disk encryption (FDE) on the TVC. The TS7700 cache models, 3956-CC9/CS9, 3956-CS9XS9, and 3956-CSA/XSA, support FDE. FDE uses the Advanced Encryption Standard (AES) 256-bit data encryption to protect the data at the hard disk drive (HDD) level. Cache performance is not affected because each HDD has its own encryption engine, which matches the drive's maximum port speed.

FDE encryption uses two keys to protect HDDs:

- ▶ The data encryption key: Generated by the drive and never leaves the drive. It is stored in an encrypted form within the drive and runs symmetric encryption and decryption of data at full disk speed.
- ▶ The lock key or security key: A 32-byte random number that authenticates the drive with the CC9/CS9/CSA Cache Controller by using asymmetric encryption for authentication. When the FDE drive is *secure-enabled*, it must authenticate with the CC9/CS9/CSA cache controller or it does not return any data and remains locked. One security key is generated for all FDE drives that are attached to the CC9/CS9/CSA cache controller and CX9/XS9/CSA cache expansion drawers.

Authentication occurs after the FDE disk has powered on, where it will be in a locked state. If encryption was never enabled (the lock key is not initially established between CC9/CS9/CSA cache controller and the disk), the disk is considered unlocked with access unlimited just like a non-FDE drive.

The following feature codes are required to enable FDE:

- ▶ Feature Code 7404: Required on all 3956-CC9, 3956-CX9, 3956-CS9, 3956-XS9, 3956-CSA, and 3956-XSA cache drawers
- ▶ Feature Code 7730: Required on the 3952-F05 base frame for a TS7740
- ▶ Feature Code 7331: Required on the 3952-F05 base frame for a TS7720
- ▶ Feature Code 7332: Required on the 3952-F05 expansion frame
- ▶ Feature Code 7333: Required on the 3952-F06 base frame for a TS7760
- ▶ Feature Code 7334: Required on the 3952-F06 expansion frame for a TS7760

Disk-based encryption is activated with the purchase and installation of Feature Code 5272: Enable Disk Encryption, which is installable on the TS7720-VEB, TS7740-V07, and TS7760-VEC (Encryption Capable Frames, as listed in the previous required feature code list).

Key management for FDE does not use the IBM Security Key Lifecycle Manager (SKLM), or IBM Security Key Lifecycle Manager for z/OS (ISKLM). Instead, the key management is handled by the disk controller, either the 3956-CC9, 3956-CS9, or 3956-CSA. There are no keys to manage by the user, because all management is done internally by the cache controllers.

Disk-based encryption FDE is enabled on all HDDs that are in the cache subsystem (partial encryption is not supported). It is an “all or nothing” proposition. All HDDs, disk cache controllers, and drawers must be Encryption Capable as a prerequisite for FDE enablement. FDE is enabled at a cluster TVC level, so you can have clusters with TVC encrypted along with clusters with TVC that are not encrypted as members of the same grid.

When disk-based encryption is enabled on a system already in use, all previously written user data is encrypted retroactively, without a performance penalty. After disk-based encryption is enabled, it cannot be disabled again.

External key management

You can manage the EK for the disk drive modules (DDMs) externally since Release 3.3. For external key management of encryption, the encryption must be enabled onsite by an IBM SSR. The EK server, IBM Security Key Lifecycle Manager (SKLM) or IBM Security Key Lifecycle Manager for z/OS (ISKLM) is installed and configured in the network.

4.5 Tape analysis and sizing the TS7700

This section documents the process of using various tools to analyze current tape environments, and to size the TS7700 to meet specific requirements. It also shows you how to access a tools library that offers many jobs to analyze the current environment, and a procedure to unload specific System Management Facility (SMF) records for a comprehensive sizing with *BatchMagic*, which must be done by an IBM SSR or IBM Business Partner.

4.5.1 IBM tape tools

Most of the IBM tape tools are available to you, but some, such as *BatchMagic*, are only available to IBM personnel and IBM Business Partners. You can 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 in the IBMTOOLS.EXE file, which is a self-extracting compressed file that, after it has been downloaded to your PC, can expand into four separate files:

- ▶ IBMJCL.XMI: Job control language (JCL) for current tape analysis tools
- ▶ IBMJCL.XMI: Parameters that are 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 you in tuning your current tape environment by identifying the factors that influence the overall performance of the TS7700. An example of factors is bad block sizes, that is, smaller than 16 KB, and low compression ratios, both of which can affect performance in a negative way.

SMF record types

System Management Facilities (SMF) is a component of the mainframe z/OS that provides a standardized method for writing out records of activity to a data set. The volume and variety of information in the SMF records enable installations to produce many types of analysis reports and summary reports.

By keeping historical SMF data and studying its trends, an installation can evaluate changes in the configuration, workload, or job scheduling procedures. Similarly, an installation can use SMF data to determine wasted system resources because of problems, such as inefficient operational procedures or programming conventions.

The examples that are shown in Table 4-17 show the types of reports that can be created from SMF data. View the examples primarily as suggestions to assist you in planning SMF reports.

Table 4-17 SMF input records

Record type	Record description
04	Step End
05	Job End
14	End-of-volume (EOV) or CLOSE when open for reading. Called "open for input" in reports.

Record type	Record description
15	EOV or CLOSE when open for writing. Called "open for output" in reports.
21 ^a	Volume dismount
30 ^b	Address Space Record (Contains subtypes 04, 05, 34, 35, and others)
34	Step End (Time Sharing Option, called TSO)
35	Job End (TSO)

a. Type 21 records exist only for tape data.

b. Record type 30 (subtypes 4 and 5) is a shell record that contains the same information that is in record types 04, 05, 34, and 35. If a type 30 record has the same data as type 04, 05, 34, and 35 records in the input data set, use the data from the type 30 record and ignore the other records.

Tape compression analysis for TS7700

By analyzing the miscellaneous data records (MDRs) from the SYS1.LOGREC data set or the EREP history file, you can see how well current tape volumes are compressing.

The following job stream was created to help analyze these records. See the installation procedure in the member \$\$INDEX file:

- ▶ EREPMDR: JCL to extract MDR records from the EREP history file
- ▶ TAPECOMP: A program that reads either SYS1.LOGREC or the EREP history file and produces reports on the current compression ratios and MB transferred per hour

The SMF 21 records record 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

TAPEWISE

The 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)

MDR analysis for bad TS7700 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 and causing extended job run times.

The following job stream was 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, and then gathers the job name and data set name from a TMS copy

Data collection and extraction

To size the TS7700 correctly, the current workload must be analyzed. The SMF records that are required to run the analysis are record types 14, 15, and 21.

Collect the stated SMF records for all z/OS systems that share the current tape configuration and might have data that is migrated to the TS7700. The data that is collected must 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 their collection.

The following steps are shown in Figure 4-8:

1. The TMS data and SMF data collection use FORMCATS and SORTSMF. Select only the required tape processing-related SMF records and the TMS catalog information.
2. The files that are created are compressed by the BMPACKT and BMPACKS procedures.
3. Download the packed files (compressed file format) to your PC and send them by email to your IBM SSR.

Figure 4-8 shows the unload process.

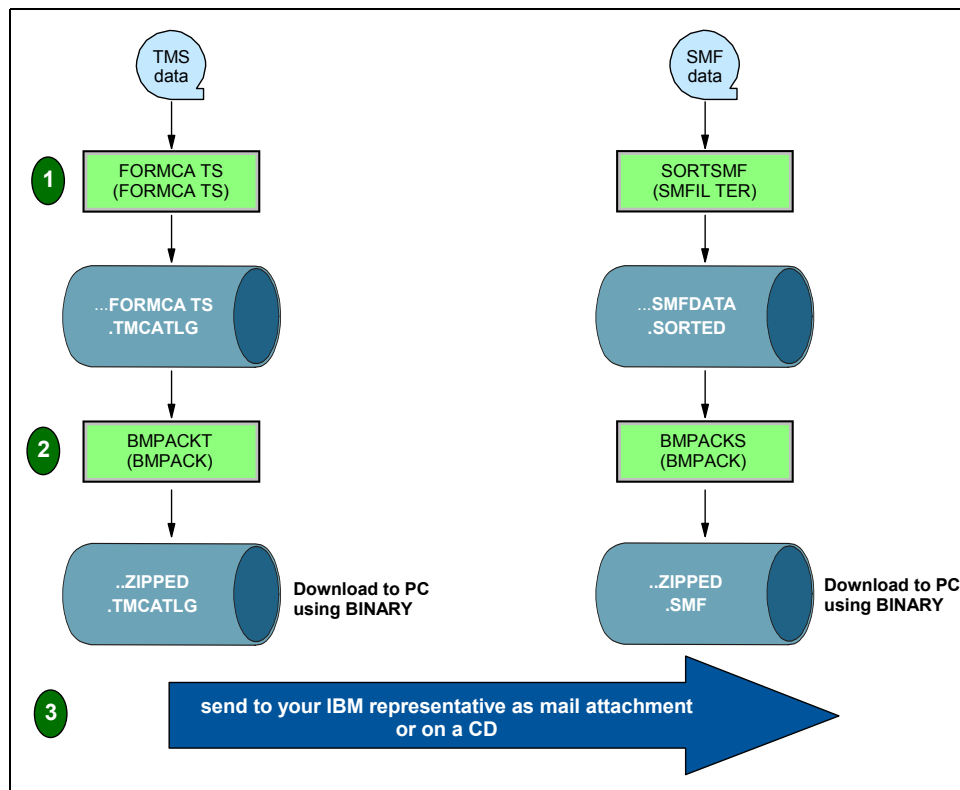


Figure 4-8 Unload process for TMS and SMF data

In addition to the extract file, the following information is useful for sizing the TS7700:

- ▶ Number of volumes in current tape library
This number includes all the tapes (located within automated libraries, on shelves, and offsite). If the unloaded Tape Management Catalog (TMC) data is provided, there is no need to collect the number of volumes.
- ▶ Criteria for identifying volumes
Because volumes are transferred offsite to be used as backup, their identification is important. Identifiers, such as high-level qualifiers (HLQs), program names, or job names, must be documented for easier reference.
- ▶ Number and type of tape CUs installed
This information provides a good understanding of the current configuration and helps identify the reasons for any apparent workload bottlenecks.
- ▶ Number and type of tape devices installed
This information, similar to the number and type of tape CUs installed, helps identify the reasons for any apparent workload bottlenecks.
- ▶ Number and type of host channels that are attached to tape subsystems
This information also helps you identify the reasons for any apparent workload bottlenecks.

4.5.2 BatchMagic

The BatchMagic tool provides a comprehensive view of the current tape environment and predictive modeling of workloads and technologies. The general methodology behind this tool involves analyzing SMF type 14, 15, 21, and 30 records, and data extracted from the TMS. 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 usage that is stored in the TMS catalog.

When you run BatchMagic, the tool extracts data, groups data into workloads, and then targets workloads to individual or multiple IBM tape technologies. BatchMagic examines the TMS catalogs and estimates cartridges that are required with new technology, and it models the operation of a TS7700 and 3592 drives (for TS7740, TS7720T, or TS7760T) and estimates the required resources.

The reports from BatchMagic give you a clear understanding of your current tape activities. They make projections for a TS7700 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.

4.5.3 Workload considerations

The TS7700 appears to the host systems as sixteen 3490E subsystems with a total of 496 virtual devices that are attached per cluster. Any data that can be on a 3480, 3490, 3590, or 3592, prior generations of VTS systems, or cartridges from other vendors, can be on the TS7700. However, processing characteristics of workloads differ, so some data is more suited for the TS7700 than other data.

This section highlights several important considerations when you are deciding what workload to place in the TS7700:

▶ **Throughput**

The TS7700 has a finite bandwidth capability, as does any other device that is attached to a host system. With 8 Gb FICON channels and large disk cache repositories that operate at disk speeds, most workloads are ideal for targeting a TS7700.

▶ **Drive concurrency**

Each TS7700 appears to the host operating system as up to the maximum of 496 3490E logical drives. If there are periods during the day when your tape processing jobs are limited by drive availability, the TS7700 might help considerably in the area of processing.

The TS7720 and TS7760 enable access to multiple logical volumes directly from cache, at disk speed.

The design of the TS7740/TS7700T enables access to multiple logical volumes on the same stacked physical volume because access to the logical volumes is solely through the TS7740/TS7700T TVC. If there is access that is needed 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 more information about scratch and specific allocation considerations in a TS7700 TVC, see the “Load Balancing Considerations” section in *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867, and 11.21, “Virtual Device Allocation in z/OS with JES2” on page 741.

▶ **Cartridge capacity usage**

A key benefit of the TS7740, TS7720T, and TS7760T is its ability to use fully the capacity of the 3592 cartridges independent of the data set sizes that are written, and to manage that capacity effectively without host or user involvement. A logical volume can contain up to 25,000 MiB of data (75,000 MiB, assuming a data compressibility of 3:1) by using the extended logical volume sizes.

The size of a logical volume is only the amount of data that is written by the host. Therefore, even if an application writes only 20 MB to a 6000 MiB volume, only the 20 MB is kept in the TS7700 cache, or on a TS7740, TS7720T, and TS7760T, a managed physical volume.

▶ **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 by using the TS7700 because the data is cached in the TS7700 cache, which effectively removes the rewind time, the robotics time, and load or thread times for the mount.

Figure 4-9 on page 183 shows an example effect that a TS7700 can have on a job and drive assignment as compared to a native drive. The figure is an out-of-scale freehand drawing. 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 4-9 on page 183, notice that the TS7700 cache hit timing does not include robotics, load, or thread time at the beginning of the timeline, and no rewind or unload time at the end of it.

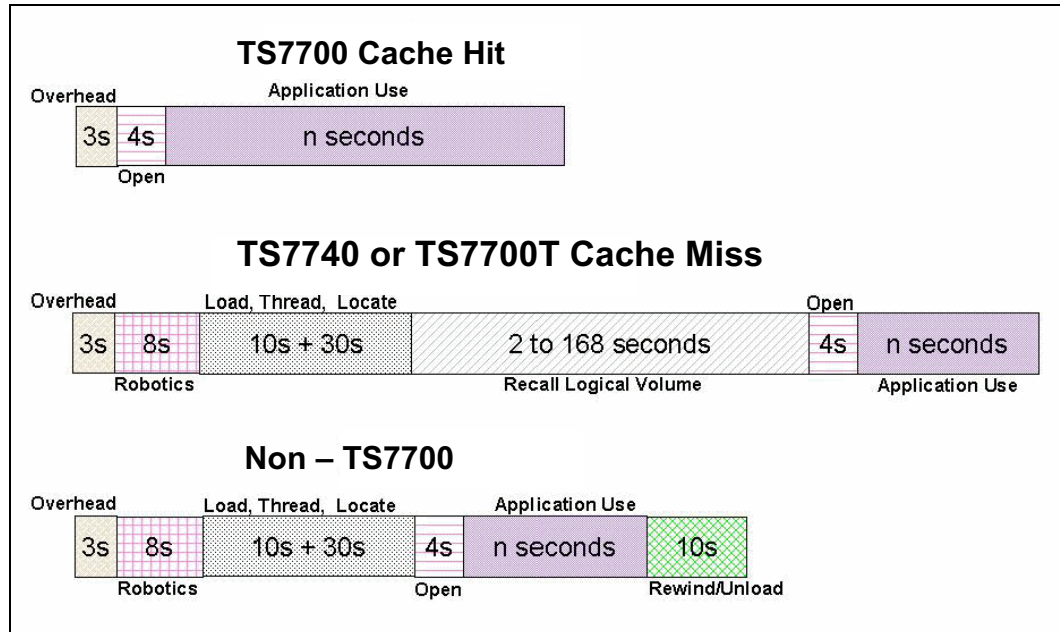


Figure 4-9 Tape processing time comparison example (not to scale)

In this example, the TS7700 cache hit results in 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 that is processing the tape.
- It frees up a drive earlier, so the next job that needs 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, TS7720T and TS7760T TVC, 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 affect the overall throughput of the TS7740, TS7720T, and TS7760T.

Remember: The TS7720 and TS7760 resident-only partition (CP0) features a large disk cache and no back-end tape drives. These characteristics result in a fairly consistent throughput at peak performance most of the time, operating with 100% of cache hits.

During normal operation of a TS7700 grid configuration, logical volume mount requests can be satisfied from the local TVC or a remote TVC. TS7700 algorithms can evaluate the mount request and determine the most effective way to satisfy the request from within the TS7700 grid.

If the *local* TVC does not have a current copy of the logical volume and a remote TVC does, the TS7700 can satisfy the mount request through the grid by accessing the volume in the TVC of a *remote* TS7700. The result is that in a multicluster configuration, the grid combines the TS7700 TVCs to produce a larger effective cache size for logical mount request.

Notes:

- ▶ The term *local* means the TS7700 cluster that is running the logical mount to the host.
- ▶ The term *remote* means any other TS7700 that is participating in the same grid as the local cluster.
- ▶ The acronym TVC means tape volume cache.

▶ Scratch mount times

When a program issues a scratch mount to write data, the TS7700 completes the mount request without having to recall the logical volume into the cache. With the TS7720D or TS7760D, all mounts are cache hit mounts. For workloads that create many tapes, this significantly reduces volume processing times and improves batch window efficiencies.

The effect of using the scratch category on the TVC improves mount performance in the TS7740, TS7720T, and TS7760T because no physical mount is required. The performance for scratch mounts is the same as for TVC read hits. The comparison between the time that is taken to process a mount request on a subsystem with cache to a subsystem without cache is made in Figure 4-9 on page 183.

▶ Disaster recovery

The TS7700 grid configuration is a perfect integrated solution for disaster recovery data. The TS7700 clusters in a multi-cluster grid can be separated over long distances and interconnected by using a TCP/IP infrastructure to provide for automatic data replication.

Data that is written to a local TS7700 is accessible at the remote TS7700 as though it were created there. Flexible replication policies make it easy to tailor the replication of the data to your business needs.

The Copy Export function provides another disaster recovery (DR) method. The copy-exported physical volumes can be used in an empty TS7700 to recover from a disaster or merged into an existing TS7700 grid. See 2.3.32, "Copy Export" on page 90 for more details.

▶ Multifile volumes

Stack multiple files onto volumes by using JCL constructs, or by using other methods, to better use cartridge capacity. Automatic use of physical cartridge capacity is one of the primary attributes of the TS7740, TS7720T, and TS7760T. Therefore, in many cases, manual stacking of data sets onto volumes is no longer required. If there is planning for a new application that uses JCL to stack data sets onto a volume, the TS7740, TS7720T, or TS7760T makes this JCL step unnecessary.

Multifile volumes that are moved to the TS7740, TS7720T, and TS7760T can also work without changing the stacking. However, the TS7740, TS7720T, and TS7760T recalls the complete logical volume to the TS7740, TS7720T, and TS7760T cache if the volume is not in cache, rather than moving each file as you access it.

Therefore, in certain cases, a possible advantage is to enable the TS7740, TS7720T, and TS7760T to do the stacking automatically. It can save not only manual management processor burden, but also in certain cases, host processor cycles, host channel bandwidth, direct access storage device (DASD) space, or a combination of all of these items.

- ▶ Interchange or offsite storage

As currently delivered, the TS7740, TS7720T, and TS7760T does not support the 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 offsite storage of data for the purposes of DR, or to merge into an existing TS7700 grid. See 2.3.32, “Copy Export” on page 90 for more details.

With the wide range of capabilities that the TS7700 provides, unless the data sets are large or require interchange, the TS7700 is likely a suitable place to store data.

4.5.4 Education and training

There is plenty of information in IBM Redbooks publications, operator manuals, IBM Knowledge Centers, and other places about the IBM TS7700 and IBM TS3500/TS4500 tape library. The amount of education and training your staff requires on the TS7700 depends on several factors:

- ▶ If you are using a TS7740, TS7720T, or TS7760T, are you installing the TS7740, TS7720T, or TS7760T in an existing TS3500/TS4500 tape library environment?
- ▶ Are both the TS7740, TS7720T, or TS7760T and the library new to your site?
- ▶ Are you installing the TS7700 into an existing composite library?
- ▶ Is the tape library or the TS7700 shared among multiple host systems?
- ▶ Do you have existing tape drives at your site?
- ▶ Are you installing the TS7720D or TS7760D solution?
- ▶ Are you migrating from existing B10/B20 hardware to a TS7740, TS7720T, or TS7760T?

A new TS7740, TS7720T, or TS7760T sharing an existing TS3500 or TS4500

When the TS7740, TS7720T, or TS7760T is installed and shares an existing TS3500 or TS4500 tape library, the amount of training that is needed for the operational staff, system programmers, and storage administrators is minimal. They are already familiar with the tape library operation, so the area to be covered with the operation staff must focus on the TS7740, TS7720T, or TS7760T management interface (MI).

Also, you must cover how the TS7740, TS7720T, or TS7760T relates to the TS3500 or TS4500, which helps operational personnel understand the tape drives that belong to the TS7740, TS7720T, or TS7760T, and which logical library and assigned cartridge ranges are dedicated to the TS7740, TS7720T, or TS7760T.

The operational staff must be able to identify an operator intervention, and perform the necessary actions to resolve it. They must be able to perform basic operations, such as inserting new volumes in the TS7740, TS7720T, or TS7760T, or ejecting a stacked cartridge by using the MI.

Storage administrators and system programmers need to be familiar with the operational aspects of the equipment and the following information:

- ▶ Be able to understand the advanced functions and settings, and how they affect the overall performance of the subsystem (TS7740, TS7720, TS7760 or grid)
- ▶ Software choices, takeover decision, and library request commands: How to use them and how they affect the subsystem
- ▶ Disaster recovery considerations

Storage administrators and system programmers need to also receive the same training as the operations staff, in addition to the following information:

- ▶ Software choices and how they affect the TS7700
- ▶ Disaster recovery considerations

There is extensive information about all of these topics in Chapter 2, “Architecture, components, and functional characteristics” on page 15, Chapter 6, “IBM TS7700 implementation” on page 213, and Chapter 9, “Operation” on page 319. More related information is in *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

4.5.5 Implementation services

A range of services is available to assist with the TS7700. IBM can deliver end-to-end storage services to help you throughout all phases of the IT lifecycle:

- ▶ 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

Helps with the collection of information that is required for tape analysis, analysis of your current environment, and the design of the automated tape library (ATL) environment, including coding and testing of customized DFSMS ACS routines.

- ▶ Implementation:

- TS7700 implementation provides technical consultation, software planning, and assistance and operator education to clients that are implementing an IBM TS7700.
- 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 or traditional automated tape library.
- ATL implementation provides technical consultation, software planning assistance, and operational education to clients that are implementing an ATL.
- Tape Copy Service runs copying of data on existing media into an ATL. This service is run after an Automated Library, TS7700, or grid implementation.

- ▶ Support

Support Line provides access to technical support professionals who are experts in all IBM tape products.

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 more information about storage services and IBM Global Services, contact your IBM marketing representative, or see the following website:

<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.

Planning steps checklist

This section lists the steps to be revised and run from initial planning up to the complete installation or migration. The list spans different competencies, such as hardware, software, educational, and performance monitoring activities.

Table 4-18 can help you when you plan the preinstallation and sizing of the TS7700. Use the table as a checklist for the main tasks that are needed to complete the TS7700 installation.

Table 4-18 Main checklist

Task	Reference
Initial meeting	N/A
Physical planning	4.1, "Hardware installation and infrastructure planning" on page 126 and your IBM SSR
Host connectivity	4.1.5, "Host attachments" on page 144
Hardware installation	Specific hardware manuals and your IBM SSR
IP connectivity	4.1.3, "TCP/IP configuration considerations" on page 136
HCD	6.4, "Hardware configuration definition" on page 216
Maintenance check (PSP)	Preventive Service Planning buckets
SMS	6.3, "Setting up the TS7700" on page 216
OAM	<i>z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries, SC23-6867</i>
Removable Media Management (RMM)	<i>z/OS DFSMSrmm Implementation and Customization Guide, SC23-6874</i>
TS7700 customization	Chapter 9, "Operation" on page 319
Setting up the BVIR	11.3.1, "TS7700 components and task distribution" on page 640
Specialist training	N/A
DR implications	Chapter 12, "Copy Export" on page 757
Functional/performance test	Chapter 11, "Performance and monitoring" on page 635
Cutover to production	N/A
Postinstallation tasks (if any)	11.3.1, "TS7700 components and task distribution" on page 640
Data migration (if required)	8.1, "Migration to a TS7700" on page 284



Disaster recovery

This chapter describes the use of the TS7700 in disaster recovery (DR).

This chapter includes the following sections:

- ▶ TS7700 disaster recovery principles
- ▶ Failover scenarios
- ▶ Planning for disaster recovery
- ▶ High availability and disaster recovery configurations
- ▶ Disaster recovery testing basics
- ▶ A real disaster
- ▶ Geographically Dispersed Parallel Sysplex for z/OS

5.1 TS7700 disaster recovery principles

To understand the DR capabilities of the TS7700 grid, the following topics are described:

- ▶ Data availability in the grid
- ▶ Deferred Copy Queue
- ▶ Volume ownership

5.1.1 Data availability

The fundamental function of the TS7700 is that all logical volumes are accessible through any of the virtual device addresses on the clusters in the grid configuration. If a copy of the logical volume is not available at that TS7700 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 cluster in the grid, the volume is accessed through the Tape Volume Cache (TVC) at the TS7700 cluster that has the available copy. If a recall is required to place the logical volume in the TVC on the other TS7700 cluster, it is done as part of the mount operation.

Whether a copy is available at another TS7700 cluster in a multi-cluster grid depends on the Copy Consistency Policy that was assigned to the logical volume when it was written. The Copy Consistency Policy is set through the Management Class (MC) storage construct. It specifies whether and when a copy of the data is made between the TS7700 clusters in the grid configuration. The following Copy Consistency Policies can be assigned:

- ▶ Synchronous Copy (Synch): Data that is written to the cluster is compressed and simultaneously written to another specified cluster.
- ▶ Rewind Unload (RUN): Data that is created on one cluster is copied to the other cluster as part of successful RUN command processing.
- ▶ Deferred Copy (Deferred): Data that is created on one cluster is copied to the specified clusters after successful RUN command processing.
- ▶ No Copy (None): Data that is created on one cluster is not copied to the other cluster.

Consider when the data is available on the cluster at the DR site. With Synchronous Copy, the data is written to a secondary cluster. If the primary site is unavailable, the volume can be accessed on the cluster that specified Synch. With RUN, unless the Copy Count Override is enabled, any cluster with Run specified has a copy of the volume available. With None, no copy is written in this cluster. With Deferred, a copy is available later, so it might be available at the cluster that specified Deferred.

When you enable Copy Count Override, it is possible to limit the number of RUN consistency points that are required before the application is given back device end, which can result in fewer copies of the data that is available than your copy policies specify.

The Volume Removal policy for hybrid grid configurations is available in any grid configuration that contains at least one TS7720 or TS7720T cluster and should be considered as well. The TS7720 Disk-Only solution has a maximum storage capacity that is the size of its TVC, and TS7720T CP0 works like TS7720. Therefore, after the cache fills, this policy enables logical volumes to be removed automatically from cache while a copy is retained within one or more peer clusters in the grid. If the cache is filling up, it is possible that fewer copies of the volume exist in the grid than is expected based on the copy policy alone.

5.1.2 Deferred Copy Queue

Besides a copy policy of No Copy, a Deferred Copy policy has the least impact to the applications that are running on the host. Immediately after the volume is closed, device end is passed back to the application and a copy is then queued to be made later. These copies are put on the Deferred Copy Queue.

With the standard settings, host application I/O always has a higher priority than the Deferred Copy Queue. It is normally expected that the configuration and capacity of the grid is such that the entire queue has the copies completed each day; otherwise, the incoming copies cause the Deferred Copy Queue to grow continually and the RPO might not be fulfilled.

When a cluster becomes unavailable due to broken grid links, error, or disaster, the incoming copy queue might not be complete, and the data might not be available on other clusters in the grid. You can use BVIR to analyze the incoming copy queue, but the possibility exists that volumes are not available. For backups, this might be acceptable, but for primary data, it might be preferable to use a Synch copy policy rather than Deferred.

5.1.3 Volume ownership

If a logical volume is written on one of the clusters in the grid configuration and copied to the other cluster, the copy can be accessed through the other cluster. This is subject to the so-called *volume ownership*.

At any time, a logical volume is owned by a single cluster. The owning cluster has control over access to the volume and changes to the attributes that are 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 cluster for that virtual device must have ownership of the volume to be mounted, or must obtain the ownership from the cluster that owns it. If the 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 to the operation of the TS7700.

However, if a cluster that owns a volume is unable to respond to requests from other clusters, the operation against that volume fails, unless more direction is given. 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 clusters, resulting in both clusters thinking that they have ownership of the volume. If more than one cluster has ownership of a volume, that might 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 cluster fails, is known to be unavailable (for example, a power fault in the IT center), or must be serviced, its ownership of logical volumes is transferred to the other cluster through one of the following modes.

These modes are set through the Management Interface (MI):

- ▶ **Read Ownership Takeover (ROT):** When ROT is enabled for a failed cluster, ownership of a volume is allowed to be taken from a cluster that has failed. Only read access to the volume is allowed through the other cluster in the grid. After ownership for a volume is taken in this mode, any operation that attempts 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 is restored.
- ▶ **Write Ownership Takeover (WOT):** When 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 cluster in the grid. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster is restored.
- ▶ **Service prep/service mode:** When a cluster is placed in service preparation mode or is in service mode, ownership of its volumes is allowed to be taken by the other cluster. Full access is allowed. The mode for the cluster in service remains in place until it is 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 clusters is attached to a TS3000 System Console (TSSC) and there is a communication path that is provided between the TSSCs. AOTM is enabled and defined by the IBM Service Support Representative (IBM SSR). If the clusters are near each other, 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 (Gb) links. AOTM must have a different network to be able to detect that a missing cluster is down, and that the problem is not caused by a failure in the grid gigabit wide area network (WAN) links.

When enabled by the IBM SSR, suppose that a cluster cannot obtain ownership from the other cluster because it does not get a response to an ownership request. In this case, a check is made through the TSSCs to determine whether the owning cluster is inoperable, or if the communication paths to it are not functioning. If the TSSCs determine that the owning cluster is inoperable, they enable either read or WOT, depending on what was set by the IBM SSR.

AOTM enables an ownership takeover mode after a grace period, and can be configured only by an IBM SSR. Therefore, jobs can intermediately fail with an option to try again until the AOTM enables the configured takeover mode. The grace period is set to 20 minutes, by default. The grace period starts when a cluster detects that another remote cluster 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* might be unavailable or a communication issue might 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* determined that library *libname* is unavailable. The Read-Only ownership takeover mode has been enabled.

A failure of a cluster causes the jobs that use its virtual device addresses to end abnormally (abend). To rerun the jobs, host connectivity to the virtual device addresses in the other cluster must be enabled (if it is not already), and an appropriate ownership takeover mode selected. If the other cluster has a valid copy of a logical volume, the jobs can be tried again.

If a logical volume is being accessed in a remote cache through the Ethernet link and that link fails, the job accessing that volume also fails. If the failed job is attempted again, the TS7700 uses another Ethernet link. If all links fail, access to any data in a remote cache is not possible.

5.2 Failover scenarios

As part of a total systems design, you must develop business continuity procedures to instruct information technology (IT) personnel in the actions that they need to take in a failure. Test those procedures either during the initial installation of the system or at another time.

The scenarios are described in detail in the *IBM Virtualization Engine TS7700 Series Grid Failover Scenarios* white paper, which was written to assist IBM specialists and clients in developing such testing plans. The white paper is available at the following web address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100831>

The white paper documents a series of TS7700 Grid failover test scenarios for z/OS that were run in an IBM laboratory environment. Simulations of single failures of all major components and communication links, and some multiple failures, are run.

5.3 Planning for disaster recovery

Although you 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 DR plan as it relates to a TS7700.

Many aspects of DR planning must be considered:

- ▶ Consider DR site connectivity input/output definition file (IODF).
- ▶ How critical is the data in the TS7700?
- ▶ 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 runs them?
- ▶ How will you test your procedures?

5.3.1 Disaster recovery site connectivity IODF considerations

If your production hosts have FICON connectivity to the TS7700 clusters at your DR site, you might consider including those virtual device addresses in your production IODF. Having those devices configured and offline to your production hosts makes it easier if there is a TS7700 failure that requires FICON access to the DR clusters, which is distance-dependent and might not be appropriate for all configurations.

To switch over to the DR clusters, a simple vary online of the DR devices is all that is needed by the production hosts to enable their usage. Another alternative is to have a separate IODF ready with the addition of the DR devices. However, that requires an IODF activation on the production hosts.

5.3.2 Grid configuration

With the TS7700, two types of configurations can be installed:

- ▶ Stand-alone cluster
- ▶ Multi-cluster grid

With a stand-alone system, a single cluster is installed. If the site at which that system is installed is destroyed, the data that is associated with the TS7700 might be lost unless COPY EXPORT was used and the tapes were removed from the site. If the cluster goes out of service due to failures, whether the data is recoverable depends on the failure type.

The recovery process assumes that the only elements that are available for recovery are the stacked volumes that are produced by COPY EXPORT and removed from the site. 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 recovery procedure into your current DR procedures.

Remember: The DR process is a joint exercise that requires your involvement and that of your IBM SSR to make it as comprehensive as possible.

For many clients, the potential data loss or the recovery time that is required with a stand-alone TS7700 is not acceptable because the COPY EXPORT method might take some time to complete. For those clients, the TS7700 grid provides a near-zero data loss and expedited recovery-time solution. With a multi-cluster grid configuration, up to six clusters are installed, typically at two or three sites, and interconnected so that data is replicated among them. The way that the sites are used then differs, depending on your requirements.

In a two-cluster grid, one potential use case is that one of the sites is the local production center and the other site is a backup or DR center, which is separated by a distance that is dictated by your company's requirements for DR. Depending on the physical distance between the sites, it might be possible to have two clusters be both a high availability and DR solution.

In a three-cluster grid, the typical use is that two sites are connected to a host and the workload is spread evenly between them. The third site is strictly for DR and there probably are no connections from the production host to the third site. Another use for a three-cluster grid might consist of three production sites, which are all interconnected and holding the backups of each other.

In a four or more cluster grid, DR and high availability can be achieved. The high availability is achieved with two local clusters that keep RUN or SYNC volume copies, with both clusters attached to the host. The third and fourth (or more) remote clusters can hold deferred volume copies for DR. This design can be configured in a crossed way, which means that you can run two production data centers, with each production data center serving as a backup for the other.

The only connection between the production sites and the DR site is the grid interconnection. There is normally no host connectivity between the production hosts and the DR site's TS7700. When client data is created at the production sites, it is replicated to the DR site as defined through Outboard policy management definitions and storage management subsystem (SMS) settings.

5.3.3 Planning guidelines

As part of planning a TS7700 grid configuration to address this solution, you must consider the following items:

- ▶ Plan for the necessary WAN infrastructure and bandwidth. You need more bandwidth if you are primarily using a Copy Consistency Points of RUN or SYNC because any delays in copy time that are caused by bandwidth limitations result in longer job run times.

If you have limited bandwidth available between sites, use Deferred Copy Consistency Point, or copy only the data that is critical to the recovery of your key operations. The amount of data that is sent through the WAN and the distance it is sent possibly might justify the establishment of a separate, redundant, and dedicated network only for the multi-cluster grid. There are also newer IPEX SAN42B-R switches that are available and IP extension hardware that might help with this issue.
- ▶ A factor to consider in the implementation of Copy Export for DR is that the export does not capture any volumes in the export pool that are not in the TVC of the export cluster. Any data that is migrated to back-end tape is not going to be on the EXPORT COPY volumes.
- ▶ Plan for host connectivity at your DR site with sufficient resources to run your critical workloads. If the cluster that is local to the production host becomes unavailable and there is no access to the DR site's cluster by this host, production cannot run. Optionally, plan for an alternative host to take over production at the DR site.
- ▶ Design and code the Data Facility System Management Subsystem (DFSMS) automatic class selection (ACS) routines to control what MC on the TS7700 is assigned. It is these MCs that control which Copy Consistency Points are used. You might need to consider MC assignment policies for testing your procedures at the DR site that are different from the production policies.
- ▶ Prepare procedures that your operators can run if the local site becomes unusable. The procedures include various tasks, such as bringing up the DR host, varying the virtual drives online, and placing the DR cluster in one of the ownership takeover modes.
- ▶ Perform a periodic capacity planning of your tape setup and host throughput to evaluate whether the disaster setup still can the full production workload in a disaster.

- ▶ If encryption is used in production, ensure that the disaster site supports encryption. The EKs must be available at the DR site or the data cannot be read.
- ▶ Consider how you test your DR procedures. Many scenarios can be set up:
 - Test based on all data from an existing TS7700?
 - Test based on using the Copy Export function and an empty TS7700?
 - Test based on stopping production access to one TS7700 cluster and running production to another cluster?

5.4 High availability and disaster recovery configurations

A few examples of grid configurations are addressed.

5.4.1 Example grid configurations

These examples are a small subset of possible configurations, and are only provided to show how the grid technology can be used. With five-cluster or six-cluster grids, there are many more ways to configure a grid.

Two-cluster grid

With a two-cluster grid, you can configure the grid for DR, high availability, or both. Configuration considerations for two-cluster grids are described. The scenarios that are presented are typical configurations. Other configurations are possible, and might be better suited for your environment.

Disaster recovery configuration

This section provides information that is needed to plan for a TS7700 2-cluster grid configuration to be used specifically for DR purposes.

A natural or human-caused event has made the local site's cluster unavailable. The two clusters are in separate locations, which are separated by a distance that is dictated by your company's requirements for DR. The only connections between the local site and the DR site are the grid interconnections. There is no host connectivity between the local hosts and the DR site cluster.

Figure 5-1 summarizes this configuration.

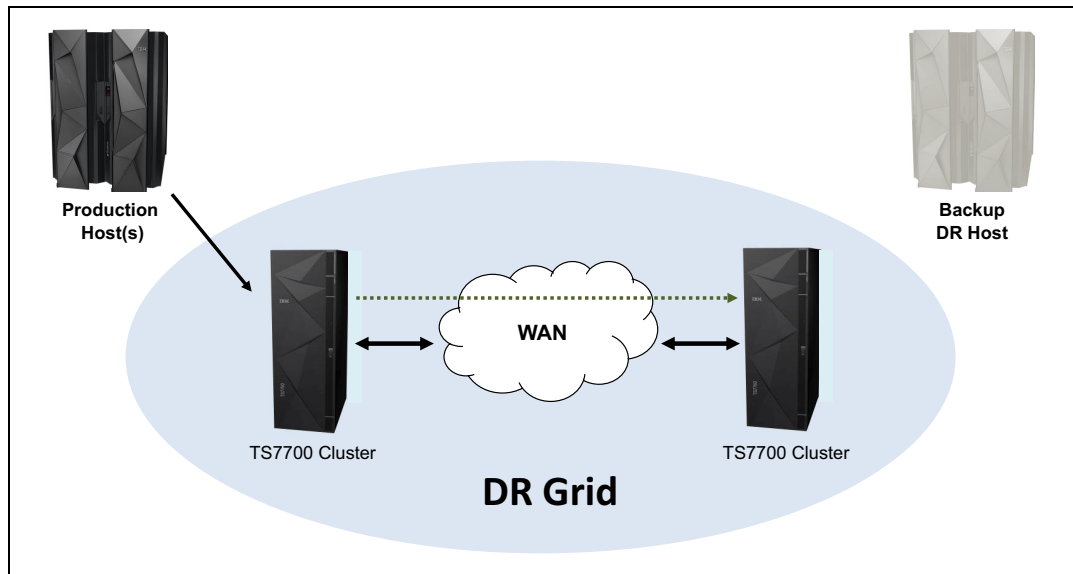


Figure 5-1 Disaster recovery configuration

Consider the following information as part of planning a TS7700 grid configuration to implement this solution:

- ▶ Plan for the necessary WAN infrastructure and bandwidth to meet the copy policy requirements that you need. If you have limited bandwidth available between sites, have critical data copied with a consistency point of RUN, with the rest of the data using the Deferred Copy Consistency Point. RUN or SYNCH are only acceptable copy policies for distances less than 100 kilometers. Distances greater than 100 km must rely on the Deferred Copy Consistency Point.
- ▶ Plan for host connectivity at your DR site with sufficient resources to perform your critical workloads.
- ▶ Design and code the DFSMS ACS routines to control what MC on the TS7700 is assigned, which determines what data gets copied, and by which Copy Consistency Point.
- ▶ Prepare procedures that your operators can run if the local site becomes unusable. The procedures include various tasks, such as bringing up the DR host, varying the virtual drives online, and placing the DR cluster in one of the ownership takeover modes (unless AOTM is configured).

Configuring for high availability

This section provides the information that is needed to plan for a 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 affect the availability of data.

In a high-availability configuration, both clusters are within metro distance of each other. These clusters are connected through a LAN. If one of them becomes unavailable because it has failed, or is undergoing service or being updated, data can be accessed through the other cluster until the unavailable cluster is made available.

As part of planning a grid configuration to implement this solution, consider the following information:

- ▶ Plan for the virtual device addresses in both clusters to be configured to the local hosts. In this way, a total of 512 or 992 virtual tape devices are available for use (256 or 496 from each 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 cluster.
- ▶ Design and code the DFSMS ACS routines and MCs on the TS7700 to set the necessary Copy Consistency Points.
- ▶ Ensure that AOTM is configured for an automated logical volume ownership takeover method in case a cluster becomes unexpectedly unavailable within the grid configuration. Alternatively, prepare written instructions for the operators that describe how to perform the ownership takeover manually, if necessary. See 2.3.34, “Autonomic Ownership Takeover Manager” on page 90 for more details about AOTM.

Figure 5-2 summarizes this configuration.




Figure 5-2 Availability configuration

Configuring for disaster recovery and high availability

You can configure a two-cluster grid configuration to provide both DR and high availability solutions. The assumption is that the two clusters are in separate locations, which are separated by a distance that is dictated by your company’s requirements for DR. In addition to the configuration considerations for DR, you need to plan for the following items:

- ▶ Access to the FICON channels on the cluster at the DR site from your local site’s hosts. This can involve connections that use dense wavelength division multiplexing (DWDM) or channel extender, depending on the distance that separates the two sites. If the local cluster becomes unavailable, you use this remote access to continue your operations by using the remote cluster.

- ▶ Because the virtual devices on the remote cluster are connected to the host through a DWDM or channel extension, there can be a difference in read or write performance when compared to the virtual devices on the local cluster.

If performance differences are a concern, consider using only the virtual device addresses in the remote cluster when the local cluster is unavailable. If that is important, you must provide operator procedures to vary online and offline the virtual devices to the remote cluster.
- ▶ You might want to have separate Copy Consistency Policies for your DR data versus your data that requires high availability.

Figure 5-3 summarizes this configuration.

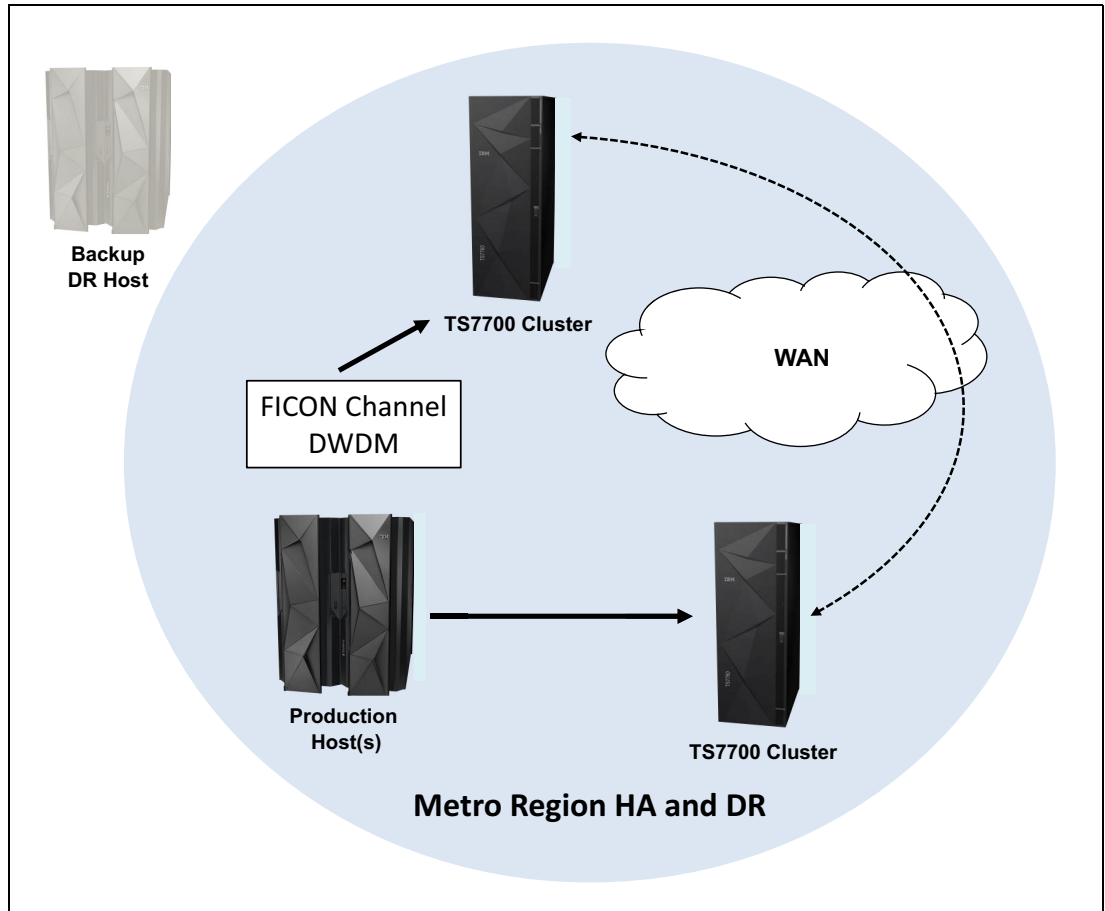


Figure 5-3 Availability and disaster recovery configuration

Three-cluster grid

With a three-cluster grid, you can configure the grid for DR and high availability or use dual production sites that share a common DR site. Configuration considerations for three-cluster grids are described. The scenarios that are 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 5-4 illustrates a combined high availability and DR solution for a three-cluster grid. In this example, Cluster 0 and Cluster 1 are the high-availability clusters and are local to each other (less than 50 kilometers (31 miles) apart). Cluster 2 is at a remote site that is away from the production site or sites. The virtual devices in Cluster 0 and Cluster 1 are online to the host and the virtual devices in Cluster 2 are offline to the host. The host accesses the virtual devices that are provided by Cluster 0 and Cluster 1.

Host data that is written to Cluster 0 is copied to Cluster 1 at RUN time or earlier with Synchronous mode. Host data that is written to Cluster 1 is written to Cluster 0 at RUN time. Host data that is written to Cluster 0 or Cluster 1 is copied to Cluster 2 on a Deferred basis.

The Copy Consistency Points at the DR site (NNR or NNS) are set to create a copy only of host data at Cluster 2. Copies of data are not made to Cluster 0 and Cluster 1. This enables DR testing at Cluster 2 without replicating to the production site clusters.

Figure 5-4 shows an optional host connection that can be established to the remote Cluster 2 by using DWDM or channel extenders. With this configuration, you must define an extra 256 or 496 virtual devices at the host.

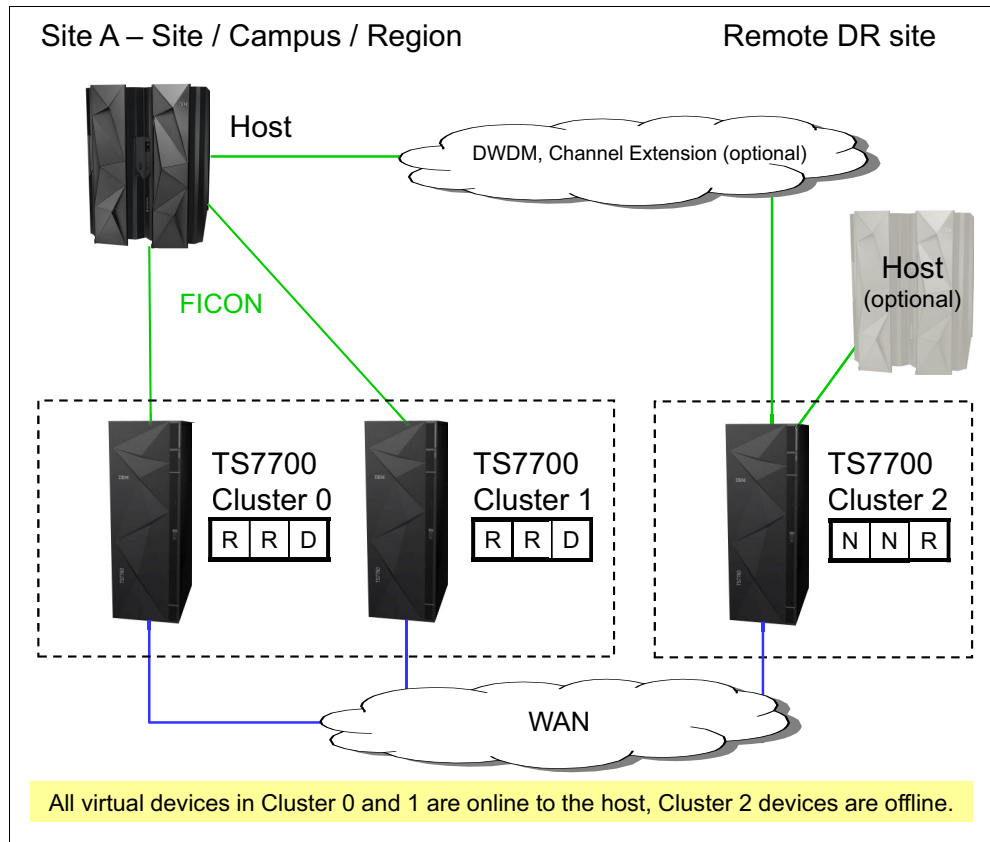


Figure 5-4 High availability and disaster recovery configuration

Dual production site and disaster recovery

Figure 5-5 on page 201 illustrates dual production sites that are sharing a DR site in a three-cluster grid (similar to a hub-and-spoke model). In this example, Cluster 0 and Cluster 1 are separate production systems that can be local to each other or distant from each other. The DR 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 virtual devices that are provided by their respective clusters. Host data that is written to Cluster 0 is not copied to Cluster 1. Host data that is written to Cluster 1 is not written to Cluster 0. Host data that is written to Cluster 0 or Cluster 1 is copied to Cluster 2 on a Deferred basis.

The Copy Consistency Points at the DR site (NNR or NNS) are set to create only a copy of host data at Cluster 2. Copies of data are not made to Cluster 0 and Cluster 1. This enables DR testing at Cluster 2 without replicating to the production site clusters.

Figure 5-5 shows an optional host connection that can be established to remote Cluster 2 using DWDM or channel extenders.

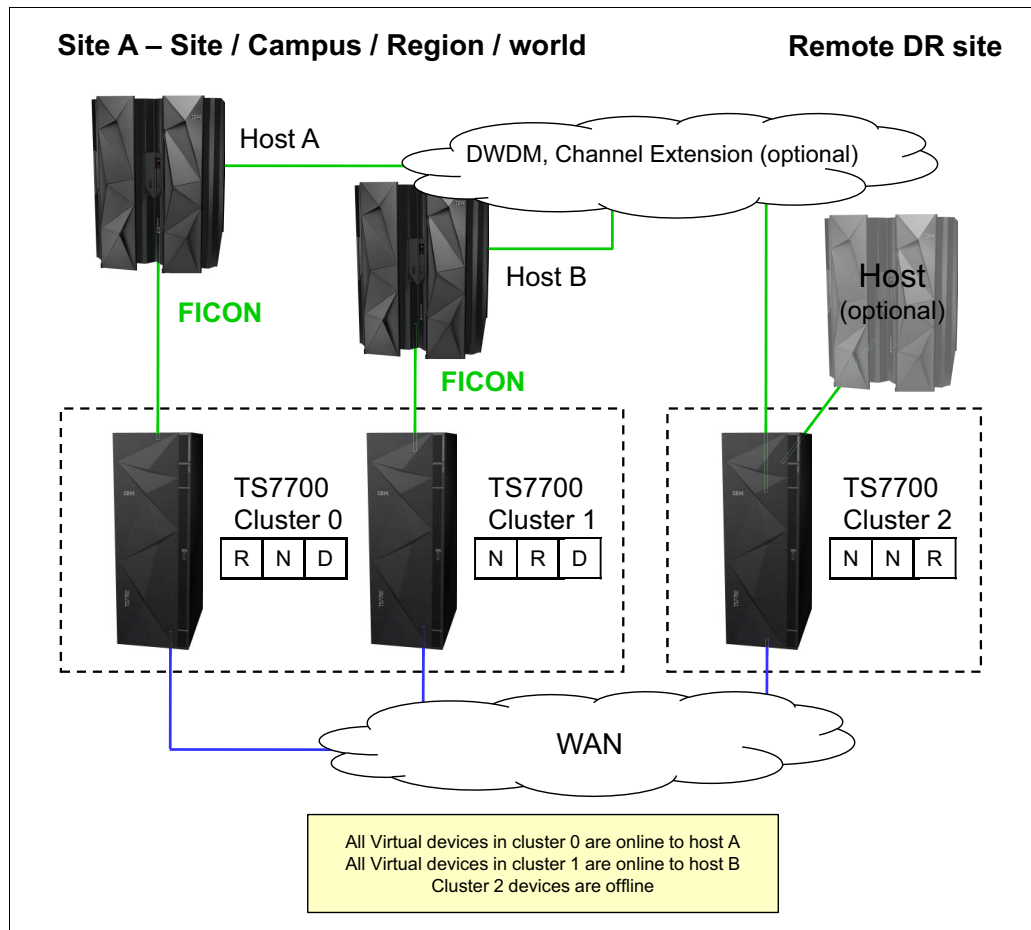


Figure 5-5 Dual production site with disaster recovery

Three-cluster high availability production site and disaster recovery

This model has been adopted by many clients. In this configuration, two clusters are in the production site (same building or separate location within metro area) and the third cluster is remote at the DR site. Host connections are available at the production site (or sites).

In this configuration, each TS7700D replicates to both its local TS7700D peer and to the remote TS7740/TS7700T. Optional copies in both TS7700D clusters provide high availability plus cache access time for the host accesses. At the same time, the remote TS7740/TS7700T provides DR capabilities and the remote copy can be remotely accessed, if needed.

This configuration, which provides high-availability production cache if you choose to run balanced mode with three copies (R-R-D for both Cluster 0 and Cluster 1), is depicted in Figure 5-6.

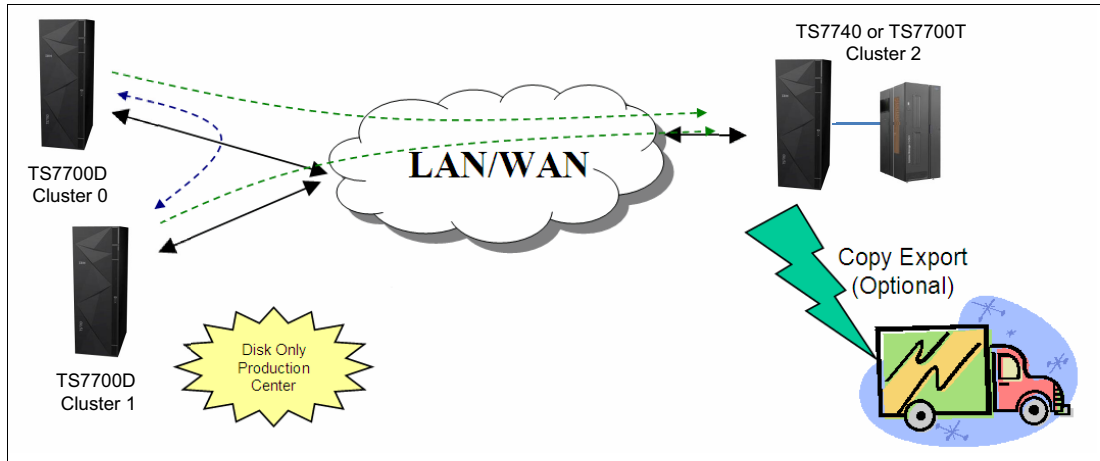


Figure 5-6 Three-cluster high availability and disaster recovery with two TS7700Ds and one TS7740/TS7700T tape library

Another variation of this model uses a TS7700D and a TS7740/TS7700T for the production site, as shown in Figure 5-7, both replicating to a remote TS7740/TS7700T.

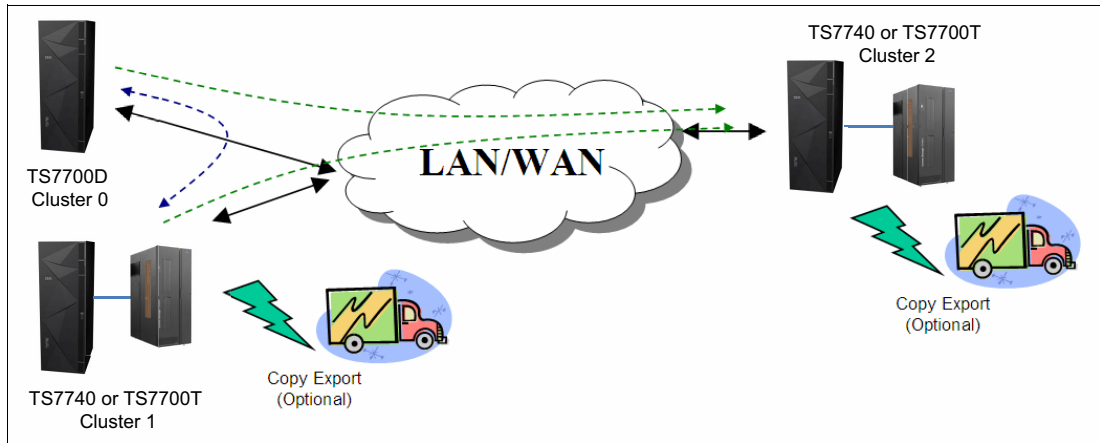


Figure 5-7 Three-cluster high availability and disaster recovery with two TS7740/TS7700T tape libraries and one TS7700D

In both models, if a TS7700D reaches the upper threshold of usage, the PREFER REMOVE data, which has already been replicated to the TS7740/TS7700T, is removed from the TS7700D cache followed by the PREFER KEEP data. PINNED data can never be removed from a TS7700D cache or a TS7700T CP0.

In the example that is shown in Figure 5-7, you can have particular workloads that favor the TS7740/TS7700T, and others that favor the TS7700D, suiting a specific workload to the cluster best equipped to perform it.

Copy Export (shown as optional in both figures) can be used to have an additional copy of the migrated data, if required.

Four-cluster grid

A four-cluster grid that can have both sites for dual purposes is described. Both sites are equal players within the grid, and any site can play the role of production or DR, as required.

Dual production and disaster recovery at Metro Mirror distance

In this model, you have dual production and DR sites. Although a site can be labeled as a high availability pair or DR site, they are equivalent from a technology standpoint and functional design. In this example, you have two production sites within metro distances and two remote DR sites within metro distances between them. This configuration delivers the same capacity as a two-cluster grid configuration, with the high availability of a four-cluster grid. See Figure 5-8.




Figure 5-8 Four-cluster high availability and disaster recovery

You can have host workload balanced across both clusters (Cluster 0 and Cluster 1 in Figure 5-8). The logical volumes that are written to a particular cluster are only replicated to one remote cluster. In Figure 5-8, Cluster 0 replicates to Cluster 2 and Cluster 1 replicates to Cluster 3. This task is accomplished by using copy policies. For the described behavior, copy mode for Cluster 0 is RDRN or SDSN and for Cluster 1 is DRNR or DSNS.

This configuration delivers high availability at both sites, production and DR, without four copies of the same tape logical volume throughout the grid.

If this example was not in Metro Mirror distances, use copy policies on Cluster 0 of RDDN and Cluster 1 of DRND.

Figure 5-9 shows the four-cluster grid reaction to a cluster outage. In this example, Cluster 0 goes down due to an electrical power outage. You lose all logical drives that are emulated by Cluster 0. The host uses the remaining addresses that are emulated by Cluster 1 for the entire production workload.

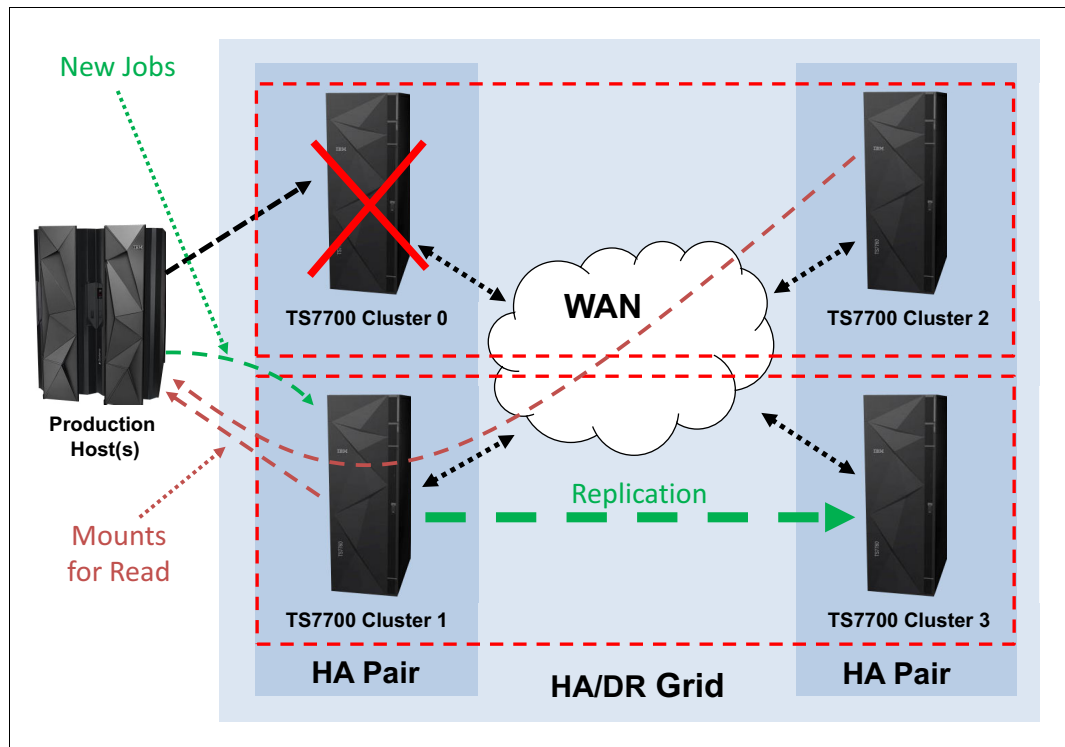


Figure 5-9 Four-cluster grid high availability and disaster recovery - Cluster 0 outage

During the outage of Cluster 0 in the example, new jobs for write use only one half of the configuration (the unaffected partition in the lower part of Figure 5-9). Jobs for read can access content in all available clusters. When power is normalized at the site, Cluster 0 starts and rejoins the grid, reestablishing the original balanced configuration.

In a DR situation, the backup host in the DR site operates from the second high availability pair, which is the pair of Cluster 2 and Cluster 3 in Figure 5-9. In this case, copy policies can be RNRD for Cluster 2 and NRNR for Cluster 3.

If these sites are more than Metro Mirror distance, you can have Cluster 2 copy policies of DNRD and Cluster 3 policies of NDDR.

5.4.2 Restoring the host and library environments

Before you can use the recovered logical volumes, you must restore the host environment. The following steps are the minimum steps that you need to continue the recovery process of your applications:

1. Restore the tape management system (TMS) CDS.
2. Restore the DFSMS data catalogs, including the tape configuration database (TCDB).
3. Define the I/O gen by using the Library IDs of the recovery TS770 tape drives.

4. Update the library definitions in the source control data set (SCDS) with the Library IDs for the recovery TS7700 tape drives in the composite library and distributed library definition windows.
5. Activate the I/O gen and the SMS SCDS.

You might also want to update the library nicknames that are defined through the MI for the grid and cluster to match the library names defined to DFSMS. That way, the names that are shown on the MI windows match those names that are used at the host for the composite library and distributed library.

To set up the composite name that is used by the host to be the grid name, complete the following steps:

1. Select **Configuration** → **Grid Identification Properties**.
2. In the window that opens, enter the composite library name that is used by the host in the grid nickname field.
3. You can optionally provide a description.

Similarly, to set up the distributed name, complete the following steps:

1. Select **Configuration** → **Cluster Identification Properties**.
2. In the window that opens, enter the distributed library name that is used by the host in the Cluster nickname field.
3. You can optionally provide a description.

These names can be updated at any time.

5.5 Disaster recovery testing basics

The TS7700 grid configuration provides a solution for DR needs when data loss and the time for recovery must 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.

Before R3.1, you might decide to run your DR test with Write Protection mode, and choose whether to define write-protect exclusion categories.

5.5.1 Selective write protect for disaster recovery testing

This function enables clients to emulate DR events by running test jobs at a DR location within a TS7700 grid configuration, enabling volumes only within specific categories to be manipulated by the test application. This function prevents any changes to production-written data, which is accomplished by excluding up to 16 categories from the cluster's write-protect enablement.

When a cluster is write-protect-enabled, all volumes that are protected cannot be modified or have their category or storage construct names modified. As in the TS7700 write-protect setting, the option is grid partition scope (a cluster) and configured through the MI. Settings are persistent, except for DR FLASH, and saved in a special repository.

Also, the new function enables any volume that is assigned to one of the categories that are contained within the configured list to be excluded from the general cluster's write-protect state. The volumes that are assigned to the excluded categories can be written to or have their attributes modified.

In addition, those scratch categories that are not excluded can optionally have their Fast Ready characteristics ignored, including Delete Expire and hold processing, enabling the DR test to mount volumes as private that the production environment has since returned to scratch (they are accessed as read-only).

One exception to the write protect is those volumes in the insert category. To enable a volume to be moved from the insert category to a write-protect-excluded category, the source category of insert cannot be write-protected. Therefore, the insert category is always a member of the excluded categories.

Be sure that you have enough scratch space when Expire Hold processing is enabled to prevent the reuse of production scratched volumes when you are planning for a DR test. Suspending the volumes' return-to-scratch processing during the DR test is also advisable.

Because selective write protect is a cluster-wide function, separated DR drills can be conducted simultaneously within one multi-cluster grid, with each cluster having its own independent client-configured settings. Again, DR FLASH is the exception to this statement.

With Release 3.1, a new function, called *FlashCopy for DR Testing*, was introduced. This feature is a major improvement regarding the DR testing possibilities.

Today, three major alternatives exist:

1. DR test without Write Protect Mode
2. Write Protect Mode or Selective Write Protect Mode
3. FlashCopy for Disaster Recovery Testing

For alternatives 1 and 2, you can also decide whether to break the gridlinks. These alternatives are discussed in detail in Chapter 13, "Disaster Recovery Testing" on page 787.

5.6 A real disaster

To clarify what a real disaster means, if you have a hardware issue that, for example, stops the TS7700 for 12 hours, is this a real disaster? It depends.

For a bank, during the batch window, and without any other alternatives to bypass a 12-hour TS7700 outage, this can be a real disaster. However, if the 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.

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.

Several differences exist between a DR test situation and a real disaster situation. In a real disaster situation, you do not have to do anything to be able to use the DR TS7700, which makes your task easier. However, this easy-to-use capability does not mean that you have all the cartridge data copied to the DR TS7700.

If your copy mode is RUN, you need to consider only in-flight tapes that are being created when the disaster happens. You must rerun all these jobs to re-create tapes for the DR site. Alternatively, if your copy mode is Deferred, you have tapes that are not copied yet. To know which tapes are not copied, you can go to the MI in the DR TS7700 and find cartridges that are already in the copy queue. After you have this information, you can, by using your TMS, discover which data sets are missing, and rerun the jobs to re-create these data sets at the DR site.

Figure 5-10 shows an example of a real disaster situation.

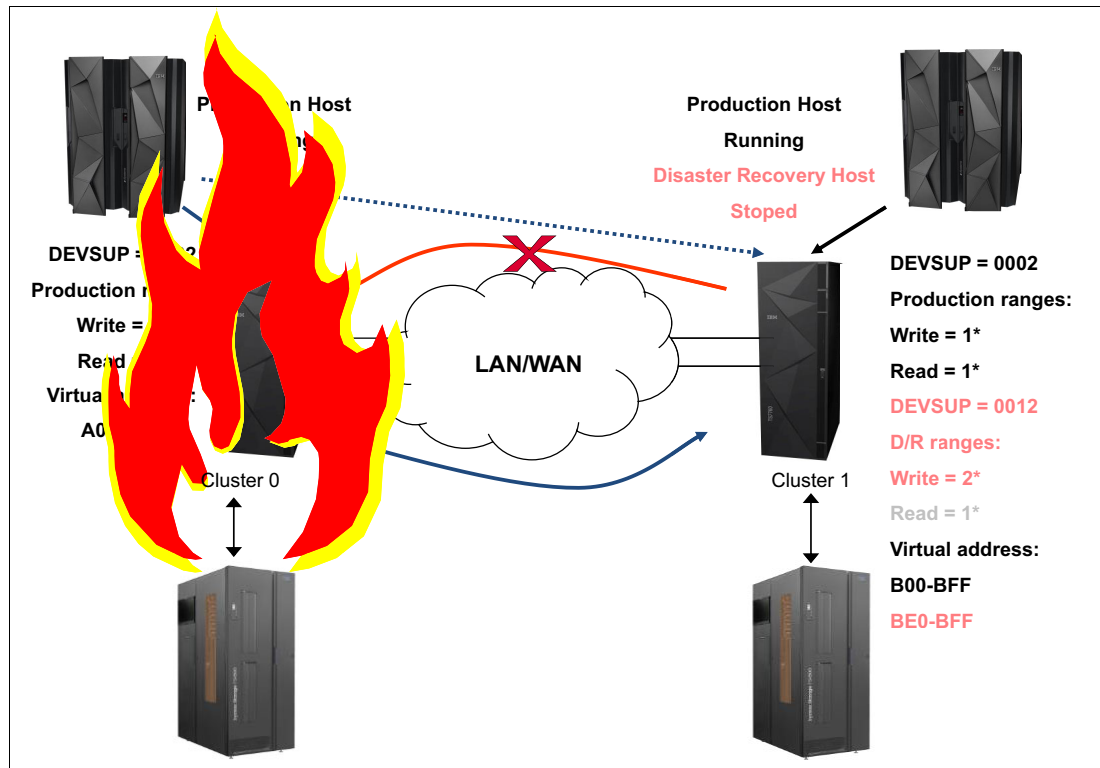


Figure 5-10 Real disaster situation

In a real disaster scenario, the whole primary site is lost. Therefore, you need to start your production systems at the DR site. To do this, you need to have a copy of all your information not only on tape, but all DASD data copied to the DR site.

After you can start the z/OS partitions, from the TS7700 perspective, you must be sure that your hardware configuration definition (HCD) “sees” the DR TS7700. Otherwise, you cannot put the TS7700 online.

You must change ownership takeover, also. To perform that task, go to the MI interface and enable ownership takeover for read and write.

All the other changes that you did in your DR test are not needed now. Production tape ranges, scratch categories, SMS definitions, RMM inventory, and so on, are in a real configuration that is in DASD that is copied from the primary site.

Perform the following changes because of the special situation that a disaster merits:

- ▶ Change your MC to obtain a dual copy of each tape that is created after the disaster.
- ▶ Depending on the situation, 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 that are required to recover your primary site or to create a new site. The old DR site is now the production site, so you must create a DR site, which is beyond the scope of this book.

5.7 Geographically Dispersed Parallel Sysplex for z/OS

The z Systems multisite application availability solution, Geographically Dispersed Parallel Sysplex (GDPS), integrates Parallel Sysplex technology and remote copy technology to enhance application availability and improve DR. The GDPS topology is a Parallel Sysplex cluster that is spread across two sites, with all critical data mirrored between the sites. GDPS manages the remote copy configuration and storage subsystems, automates Parallel Sysplex operational tasks, and automates failure recovery from a single point of control, improving application availability.

5.7.1 Geographically Dispersed Parallel Sysplex considerations in a TS7700 grid configuration

A key principle of GDPS is to have all I/O be local to the system that is running production. Another principle is to provide a simplified method to switch between the primary and secondary sites, if needed. The TS7700 grid configuration provides a set of capabilities that can be tailored to enable it to operate efficiently in a GDPS environment. Those capabilities and how they can be used in a GDPS environment are described in the following sections.

Direct production data I/O to a specific TS7740

The hosts are directly attached to the TS7740 that is local to the host 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 that is remote to a host to cover recovery only when the TS7740 cluster at the GDPS primary site is down. However, during normal operation, the remote virtual devices are set offline in each GDPS host.

The default behavior of the TS7740 in selecting which TVC is used for the I/O is to follow the MC definitions and considerations to provide the best overall job performance. However, it uses 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, complete the following steps:

1. For the MC that is used for production data, ensure that the local cluster has a Copy Consistency Point. If it is important to know that the data is replicated at job close time, specify a Copy Consistency Point of RUN or Synchronous mode copy.

If some amount of data loss after a job closes can be tolerated, 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 MC with separate Copy Consistency Points. In defining the Copy Consistency Points for an MC, 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 selects the TVC local to the TS7740 on which the mount was received if it is available and a Copy Consistency Point other than No Copy is specified for that cluster in the MC 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 selects the TVC local to the TS7740 on which the mount was received if it is available and the cluster has a valid copy of the data, even if the data is only 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, without the next override, it is possible that the remote TVC is selected.

4. Set **Force Volume Copy to Local**. This override has two effects, depending on the type of mount requested. For a private 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 scratch mount, it has the effect of OR-ing the specified MC with a Copy Consistency Point of RUN for the cluster, which forces the local TVC to be used. The override does not change the definition of the MC. It serves only to influence the selection of the I/O TVC or to force a local copy.
5. Ensure that these override settings are duplicated on both TS7740 Virtualization Engines.

Switching site production from one TS7700 to another one

The way that data is accessed by either TS7740 is based on the logical volume serial number. No changes are required in tape catalogs, job control language (JCL), or TMSs. In a failure in a TS7740 grid environment with GDPS, three scenarios can occur:

- ▶ GDPS switches the primary host to the remote location and 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, and the GDPS host and TS7740 cluster are down or inactive:
 - Automatic ownership takeover of volumes, which are then accessed from the remote host, is not possible.
 - Manual intervention is required. Through the TS7740 MI, the administrator must start a manual ownership takeover. To do so, use the TS7740 MI and click **Service** → **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 is then accessed from the remote host, manual intervention is required. Through the TS7740 MI, start a manual ownership takeover. To do so, click **Service** → **Ownership Takeover Mode** in the TS7740 MI.

5.7.2 Geographically Dispersed Parallel Sysplex functions for the TS7700

GDPS provides TS7700 configuration management and displays the status of the managed TS7700 tape drives on GDPS windows. TS7700 tape drives that are managed by GDPS are monitored, and alerts are generated for abnormal conditions. The capability to control TS7700 replication from GDPS scripts, and to window by using **TAPE ENABLE** and **TAPE DISABLE** by library, grid, or site, is provided for managing the TS7700 during planned and unplanned outage scenarios.

The TS7700 provides a capability called Bulk Volume Information Retrieval (BVIR). If there is an unplanned interruption to tape replication, GDPS uses this BVIR capability to automatically collect information about all volumes in all libraries in the grid where the replication problem occurred. In addition to this automatic collection of in-doubt tape information, it is possible to request GDPS to perform BVIR processing for a selected library by using the GDPS window interface at any time.

GDPS supports a physically partitioned TS7700. For more information about the steps that are required to partition a TS7700 physically, see Appendix I, “Case study for logical partitioning of a two-cluster grid” on page 937.

5.7.3 Geographically Dispersed Parallel Sysplex implementation

Before implementing the GDPS support for TS7700, ensure that you review and understand the following topics:

- ▶ 2.2.19, “Copy Consistency Point: Copy policy modes in a stand-alone cluster” on page 46
- ▶ *IBM Virtualization Engine TS7700 Series Best Practices Copy Consistency Points*, which is available at the following website:
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101230>
- ▶ *IBM Virtualization Engine TS7700 Series Best Practices Synchronous Copy Mode*, which is available at the following website:
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102098>

The complete instructions for implementing GDPS with the TS7700 can be found in the GDPS manuals.



Part 2

Implementation and migration

This part provides you with the information that is required to implement IBM TS7720, TS7720T (Tape Attached), and TS7740 R3.3 in your environment, and to migrate from another tape solution to IBM TS7700 R3.3.

This part includes the following chapters:

- ▶ IBM TS7700 implementation
- ▶ Hardware configurations and upgrade considerations
- ▶ Migration



IBM TS7700 implementation

This chapter describes how to implement the IBM TS7700 on IBM z Systems platforms. From a software perspective, differences exist between the TS7760D, TS7760T, TS7740, TS7720D, and the TS7720T. If no specific differences are indicated, the implementation steps apply to all models. Otherwise, the differences are explained in each relevant step.

This chapter includes the following sections:

- ▶ TS7700 implementation
- ▶ TS3500 or TS4500 tape library definitions
- ▶ Setting up the TS7700
- ▶ Hardware configuration definition
- ▶ Setting values for the Missing Interrupt Handler
- ▶ TS7700 software definitions

For more 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 and *IBM TS4500 R3 Tape Library Guide*, SG24-8235.

6.1 TS7700 implementation

The following sections describe the implementation and installation tasks to set up the TS7700. Specific names are used in this chapter if a certain task applies only to one of the five models because there are slight differences between the TS7760D, TS7760T, TS7740, TS7720D, and TS7720T.

The TS7720D and TS7760D do not have a tape library that is attached, so the implementation steps that are related to a physical tape library, for IBM TS4500 or IBM TS3500, do not apply.

You can install the TS7760T, TS7740, or TS7720T together with your existing TS3500 tape library, or install them with a new TS4500 to serve as the physical back-end tape library. When using a TS3500 as the physical back-end tape library to either the TS7760T, TS7740, or TS7720T, the IBM 3953 Library Manager is no longer required because the Library Manager functions are provided by the TS7700 Licensed Internal Code.

6.1.1 Implementation tasks

The TS7700 implementation can be logically separated into three major sections:

- ▶ TS7760T, TS7740, or TS7720T and tape library setup

Use the TS7760T, TS7740, or TS7720T and the TS4500/TS3500 tape library interfaces for these setup steps:

- Defining the logical library definitions of the TS7760T, TS7740, or TS7720T, such as physical tape drives and cartridges by using the TS4500/TS3500 tape library GUI, which is the web browser interface with the TS4500/TS3500 tape library.
- Defining specific settings, such as encryption, and inserting logical volumes into the TS7760T, TS7740, or TS7720T. You can also use the Management Interface (MI) to define logical volumes, management policies, and volume categories.

This chapter provides details about these implementation steps.

- ▶ TS7700 hardware input/output (I/O) configuration definition

This section relates to the system generation. It consists of processes, such as FICON channel attachment to the host, hardware configuration definition (HCD) or input/output configuration program (IOCP) definitions, and missing-interrupt handler (MIH) settings. This activity can be done before the physical hardware installation and it can be part of the preinstallation planning.

- ▶ TS7700 software definition

This is where you define the new virtual tape library to the individual host operating system. In a z Systems environment with Data Facility Storage Management Subsystem (DFSMS)/IBM MVS, this includes updating DFSMS automatic class selection (ACS) routines, object access method (OAM), and your tape management system (TMS) 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.

These three groups of implementation tasks can be done in parallel or sequentially. HCD and host definitions can be completed before or after the actual hardware installation.

6.2 TS4500/TS3500 tape library definitions

Use this section if you are implementing the TS7760T, TS7740 or TS7720T with a TS4500/TS3500 tape library in a z Systems environment. If your TS7700 does not have an associated tape library (TS7760D or TS7720D), see 6.3, “Setting up the TS7700” on page 216.

Your IBM Service Support Representative (IBM SSR) installs the TS7760T, TS7740, or TS7720T hardware, its associated tape library, and the frames. This installation does not require your involvement other than the appropriate planning. For more information, see Chapter 4, “Preinstallation planning and sizing” on page 125.

Clarification: The steps that are described in this section relate to the installation of a new IBM TS4500/TS3500 tape library with all of the required features, such as Advanced Library Management System (ALMS), installed. If you are attaching an existing IBM TS3500 tape library that is already attached to Open Systems hosts to z Systems hosts, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 or *IBM TS4500 R3 Tape Library Guide*, SG24-8235 for extra actions that might be required.

The following tasks are for TS4500/TS3500 library definition. For the detailed procedure, see 9.3.1, “The tape library with the TS7700T cluster” on page 509.

- ▶ Defining a logical library
 - Ensuring that ALMS is enabled
 - Creating a new logical library with ALMS
 - Setting the maximum cartridges for the logical library
- ▶ Adding drives to the logical library
- ▶ Defining control path drives

Each TS7760T, TS7740, or TS7720T requires four control path drives defined.
- ▶ Defining the Encryption Method for the new logical library
- ▶ Defining CAPs
- ▶ Inserting TS7760T, TS7740, or TS7720T physical volumes
- ▶ Assigning cartridges in the TS4500/TS3500 tape library to the logical library partition

This procedure is necessary only if a cartridge was inserted, but a Cartridge Assignment Policy (CAP) was not provided in advance.

6.3 Setting up the TS7700

This section describes the tasks that are required to set up the TS7700.

6.3.1 Definitions for TS7760T TS7740, or TS7720T

If you have a TS7760D or TS7720D, skip to 6.3.2, “TS7700 definitions” on page 216. For the detailed procedures, see 9.3.2, “TS7700T definitions” on page 529.

The tasks that are listed in this section are for TS7760T, TS7740, or TS7720T only:

- ▶ Defining VOLSER ranges for physical volumes
- ▶ Defining physical volume pools:
 - Reclaim threshold setting
 - Inhibit Reclaim schedule

6.3.2 TS7700 definitions

Use the TS7700 MI to continue with the TS7700 Virtualization subsystem setup. For more information, see 9.3.3, “TS7700 definitions” on page 546.

- ▶ Defining scratch categories and logical volume expiration time
- ▶ Defining TS7700 constructs
 - To use the Outboard Policy Management functions, you must define four constructs:
 - Storage Group (SG)
 - Management Class (MC)
 - Storage Class (SC)
 - Data Class (DC)
- ▶ TS7700 licensing
- ▶ Defining Encryption Key Server (EKS) addresses
- ▶ Defining Simple Network Management Protocol (SNMP)
- ▶ Cluster Network Setting
- ▶ Enabling Internet Protocol Security (IPSec)
- ▶ Security Settings
- ▶ Inserting logical virtual volumes

6.4 Hardware configuration definition

This section describes the process of defining the TS7700 through the HCD interface. Usually, HCD definitions are made by IBM z/OS system administrators. A helpful approach is to complete a table with all of the definitions that the administrators will need, and then give the table to the administrators.

Table 6-1 is an example definition table for a stand-alone cluster. In general, all of the blank cells must be completed by system administrators because they know what channels are free, what control unit (CU) numbers are free, and so on.

Table 6-1 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
		10			00-0F		11
		11			00-0F		12
		12			00-0F		13
		13			00-0F		14
		14			00-0F		15
		15			00-0F		16
		16			00-0F		17
		17			00-0F		18
		18			00-0F		19
		19			00-0F		1A
		1A			00-0F		1B
		1B			00-0F		1C
		1C			00-0F		1D
		1D			00-0F		1E
		1E			00-0F		1F

6.4.1 Defining devices through HCD

You can define up to 31 CUs with 16 devices each per cluster in the grid configuration. Use CUADD=0 - CUADD=7 and LIBPORT-IDs of 01 - 08 for the first eight CUs, as shown in Table 6-2.

Table 6-2 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 CUs, use CUADD=8 - CUADD=F and LIBPORT-IDs of 09 - 10, as shown in Table 6-3.

Table 6-3 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 6-1 and Figure 6-2 on page 219 show the two important windows for specifying a tape CU. To define devices by using HCD, complete the following steps:

1. Specify the CU number and the type here (3490), as shown in Figure 6-1. Press Enter.

```

----- 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 6-1 Add the first TS7700 CU through HCD (Part 1 of 2)

- The window that is shown in Figure 6-2 opens. Select the processor to which the CU 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 6-2 Add the first TS7700 CU through HCD (Part 2 of 2)

Tip: When the TS7700 is not attached through Fibre Channel connection (FICON) directors, the link address fields are blank.

- Repeating the previous process, define the 2nd - 16th TS7700 virtual tape CUs, specifying the logical unit address (CUADD)=1 - F, in the Add Control Unit windows. The Add Control Unit summary window is shown in Figure 6-2.
- To define the TS7700 virtual drives, use the Add Device window that is shown in Figure 6-3.

```

----- 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 6-3 Add the first 16 drives through HCD

- After you enter the required information, you can specify to which processors and operating systems the devices are connected to. Figure 6-4 shows the window that is 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 when 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 6-4 HCD Define Device / Processor window

- After you enter the required information and specify to which operating systems the devices are connected, the window in Figure 6-5 is displayed, where you can update the device parameters.

```

CBDPDV13 Define Device Parameters / Features Row 1 of 6
Command ==> _____ Scroll ==> PAGE
Specify or revise the values below.
Configuration ID . . : AB          MVS operating system
Device number . . . : 0440        Number of devices :16
Device type . . . . : 3490

Parameter /
Feature    Value  P Req.  Description
OFFLINE    Yes    Device considered online or offline at IPL
DYNAMIC    Yes    Device supports dynamic configuration
LOCANY     No     UCB can reside in 31 bit storage
LIBRARY    Yes    Device supports auto tape library
AUTOSWITCH No    Device is automatically switchable
LIBRARY-ID CA010 5-digit library serial number
LIBPORT-ID 01    2 digit library string ID (port number)
MTL        No     Device supports manual tape library
SHARABLE   No     Device is Sharable between systems
COMPACT    Yes    Compaction
***** Bottom of data *****
F1=Help    F2=Split    F4=Prompt    F5=Reset    F7=Backward
F8=Forward  F9=Swap     F12=Cancel   F22=Command

```

Figure 6-5 Define Device Parameters HCD window

Tips:

- ▶ If you are defining drives that are installed in a system-managed IBM tape library, such as the TS7700, you must specify `LIBRARY=YES`.
- ▶ If more than one z Systems host will be sharing the virtual drives in the TS7700, specify `SHARABLE=YES`. This forces `OFFLINE` to `YES`. It is up to the installation to ensure the correct serialization from all attached hosts.
- ▶ You must use the composite library ID of the TS7700 in your HCD definitions.
- ▶ The distributed library IDs are not defined in HCD.

To define the remaining TS7700 3490E virtual drives, repeat this process for each CU in your implementation plan.

6.4.2 Activating the I/O configuration

There are differences in the concurrent input/output definition file (IODF) activation process between a new tape library implementation and a configuration change that is made to an existing library. Changes to the virtual devices' address range of an existing library is an example of where concurrent IODF activation is useful.

As an alternative to the procedures described next, you can always perform an initial program load (IPL) or restart of the system.

Installing a new tape library

If you are installing a TS7700 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 completed to get the tape library or TS7700 online without an IPL of your systems:

1. Activate the IODF.
2. Run MVS console command **VARY ONLINE** to vary online the devices in the library. This command creates some of the control blocks. You should see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```
3. Do the final **ACTIVATE**. This action is required to build the eligible device table (EDT) for MVS Allocation.

After activation, you can check the details by using the **DEVSERV QTAPE** command. See 10.1.2, "MVS system commands" on page 605.

Modifications to an existing tape library

When you are modifying an existing tape library so that existing device addresses can be changed, complete 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 console command **VARY ONLINE** to vary online the devices in the library. This creates some of the control blocks. You see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```
4. Do the final **ACTIVATE**.

Alternatively, you can use the **DS QL,nnnn,DELETE** (where *nnnn* is the LIBID) command to delete the library's dynamic control blocks. If you have IODF with LIBID and LIBPORT coded already, and you deleted the library's dynamic control blocks, complete the following steps:

1. Use **QLIB LIST** to see whether the INACTIVE control blocks are 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, after you delete the library's dynamic control blocks, complete the following steps:

1. Run MVS console command **VARY ONLINE** to vary on the devices in the library. This creates some control blocks, and you see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```

2. Activate an IODF that defines all of the devices in the modified library.
3. Use **QLIB LIST** to verify that the ACTIVE control blocks are properly defined.

6.5 Setting values for the Missing Interrupt Handler

The TS7700 emulates 3490E devices and does not automatically communicate the 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 preferable 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 TS1150/TS1140/TS1130/TS1120/3592 drives that the TS7740 manages in the back end. The host knows only about logical 3490E devices.

Table 6-4 summarizes the minimum values, which might need to be increased, depending on specific operational factors.

Table 6-4 Tape device MIH values

Tape device	MIH
TS7700 stand-alone grid with 3490E emulation drives	20 minutes
TS7700 multi-cluster grid with 3490E emulation drives	45 minutes
TS7700 multi-cluster grid with 3490E emulation drives and not using Rewind Unload (RUN) copy policies	20 minutes

Specify the MIH values in PARMLIB member IECIOSxx. Alternatively, you can also set the MIH values using the z Systems 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 this command:

```
SETIOS MIH,DEV=(0A40-0A7F),TIME=45:00
```

- To display the new settings, use this command:

```
D IOS,MIH,DEV=0A40
```

- To check the current MIH time, use this command:

```
D IOS,MIH,TIME=TAPE
```

For more information about MIH settings, see *MVS Initialization and Tuning Reference*, SA23-1380.

During IPL (if the device is defined to be ONLINE) or during the **VARY ONLINE** in process, some devices (such as the IBM 3590/3592 physical tape devices) might present their own MIH timeout values through the *primary/secondary* MIH timing enhancement that is 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 can be used for special operations, such as long-busy conditions or long-running I/O operations.

Any time that a user specifically sets a device or device class to an MIH timeout value that is different from the default for the device class that is set by IBM, that value overrides 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 does 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.

6.6 TS7700 software definitions

This section describes the software definition considerations for implementing the TS7700 in z/OS. The TS7700 must be defined as a new tape library with emulated 3490E Tape Drives from the host system. For more information about defining this configuration, see *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

To use the TS7700, at least one SG must be created to enable the TS7700 tape library virtual drives to be allocated by the storage management subsystem (SMS) 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 SG.

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*, SC23-6867
- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *IBM TS4500 R3 Tape Library Guide*, SG24-8235

If your TMS is DFSMS Removable Media Manager (DFSMSrmm), the following manuals can be useful:

- ▶ *DFSMSrmm Primer*, SG24-5983
- ▶ *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874

If this installation is the first SMS tape library at this host, additional steps are required. The full product documentation for your TMS needs to be consulted, in addition to the OAM PISA listed in the previous bullets.

Complete the following steps to define the TS7700 tape library in an existing z/OS SMStape environment:

1. Use the ISMF Library Management → Tape Library → Define panel to define the tape library as a DFSMS resource. Define the composite library and one or more distributed libraries. Library names cannot start with a “V”. In the following figures, we define one composite library that is named IBMC1 and a single distributed library that is named IBMD1.

Remember: Library ID is the only field that applies for the distributed libraries. All other fields can be blank or left as the default.

Figure 6-6, Figure 6-7, and Figure 6-8 on page 225 illustrate defining the composite tape library.

```

Command ==>                                TAPE LIBRARY DEFINE                                Page 1
SCDS Name . : MIKE.SCDs
Library Name : IBMC1
To Define Library, Specify:
  Description ==> IS7700 GRID COMPOSITE LIBRARY
  Library ID . . . . . CA010      (00001 to FFFFF)
  Console Name . . . . .
  Entry Default Data Class . . . . . DCATLDS
  Entry Default Use Attribute . . . . . S      (P=PRIVATE or S=SCRATCH)
  Eject Default . . . . . K      (P=PURGE or K=KEEP)

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 6-6 Define Composite Tape Library screen 1

```

Command ==>                                TAPE LIBRARY DEFINE                                Page 2
SCDS Name . : MIKE.SCDs
Library Name : IBMC1
Media Type:                                Scratch Threshold
Media1: . . . . . 0      (0 to 999999)
Media2: . . . . . 3000   (0 to 999999)
Media3: . . . . . -      (0 to 999999)
Media4: . . . . . 0      (0 to 999999)
Media5: . . . . . 0      (0 to 999999)
Media6: . . . . . 0      (0 to 999999)
Media7: . . . . . 0      (0 to 999999)
Media8: . . . . . 0      (0 to 999999)
Media9: . . . . . 0      (0 to 999999)
Media10: . . . . . 0     (0 to 999999)
Media11: . . . . . 0     (0 to 999999)
Media12: . . . . . 0     (0 to 999999)
Media13: . . . . . 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 6-7 Define Composite Tape Library screen 2

```

                                TAPE LIBRARY DEFINE                                Page 3
Command ==>
SCDS Name . : MIKE.SCDs
Library Name : IBMC1

Initial Online Status (Yes, No, or Blank):
SYSTEM1 ==> yes

Warning:
When you connect a tape library to a system group rather than a system,
you lose the ability to vary that library online or offline to the
individual systems in the system group. It is strongly recommended that
the tape library be connected to individual systems only.

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 6-8 Define Composite Tape Library screen 3

Figure 6-9 illustrates defining the distributed tape library.

```

                                TAPE LIBRARY DEFINE                                Page 1
Command ==>
SCDS Name . : MIKE.SCDs
Library Name : IBMD1

To Define Library, Specify:
Description ==> TS7700 Distributed Library
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)

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 6-9 Define Distributed Tape Library

2. Using the Interactive Storage Management Facility (ISMF), create or update the DCs, SCs, and MCs for the TS7700. Ensure that these defined construct names are the same as those that you defined at the TS7700 MI.
3. Using ISMF, create the SGs for the TS7700. Ensure that these defined construct names are the same as those that you defined at the TS7700 MI.

The composite library must be defined in the SG. Do *not* define the distributed libraries in the SG.

Tip: At OAM address space initialization, if a distributed library is defined to an SG, the warning message CBR3017I is generated, which indicates that the distributed library is incorrectly defined to the SG.

4. Update the ACS routines to assign the constructs that are needed to use the TS7700 and then convert, test, and validate the ACS routines.
5. Customize your TMS to include the new volume ranges and library name. For DFSMSrmm, this involves EDGRMMxx updates for VLPOOL and LOCDEF statements, in addition to any OPENRULE and PRITITION statements needed. If you leave REJECT ANYUSE(*) in your EDGRMMxx member, you cannot use any tape volume serial numbers (VOLSERs) not previously defined to RMM.

6. Consider whether your current TMS database or control data set (CDS) has sufficient space for the added library, data set, and volume entries. For DFSMSrmm, consult the IBM Redbooks Publication *DFSMSrmm Primer*, SG24-5983 under the topic “Creating the DFSMSrmm CDS.”
7. Modify the SYS1.PARMLIB members DEVSUPxx for new categories, as described in 4.3.4, “Sharing and partitioning considerations” on page 164. The DEVSUPxx default categories should always be changed to prevent disruption of library operations. For more information, see the “Changing the library manager category assignments in an ATLDS” topic in *z/OS V2R2.0 DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.
8. Consider updating COMMANDxx to vary the library online at IPL if you have not already set it to come online during the definition of the library through ISMF. For more information, see 10.1.1, “DFSMS operator commands” on page 602. The OAM address space must be active for the vary to complete successfully.
9. In a Grid environment, it might be desirable to update the ALLOCxx parmlib member SYSTEM TAPELIB_PREF to balance workloads. Although the default algorithm, EQUAL, works well if the libraries under consideration have an equal number of online devices, the recommendation is to use BYDEVICES.
10. Consider the current size of the volume catalogs and the additional space that is required for the new volume and library entries. For a TS7700 with 100,000 volumes, at least 32 cylinders should be allocated. To more precisely estimate the space allocation requirements in the tape volume catalog, use the following steps:
 - a. Estimate the number of tape library entries and tape volume entries to be cataloged in the VOLCAT. Each tape library entry requires 320 bytes and each volume entry requires 275 bytes.
 - b. Divide the total number of bytes by 1024 to determine the number of kilobytes, or by 1,048,576 to determine the number of megabytes.

For more information, see “Defining Names for a Tape Volume Catalog” in *z/OS DFSMS Managing Catalogs*, SC23-6853.
11. Activate the new Source Control Data Set (SCDS) with the **SETSMS SCDS (scdsname)** command.
12. Restart the OAM address space with the **F OAM, RESTART** command. SMS SCDS activation initiates an OAM restart if the **RESTART=YES** parameter is specified in the OAM startup procedure in PROCLIB, and in that case a manual restart is not needed.
13. Define any security profiles required.
14. Vary the composite library and distributed libraries online. See 10.1.1, “DFSMS operator commands” on page 602.
15. Vary the TS7700 virtual drives online, as described in 10.1.1, “DFSMS operator commands” on page 602.



Hardware configurations and upgrade considerations

This chapter includes the following sections:

- ▶ TS7700 hardware components
- ▶ TS7700 component upgrades

This section describes upgrading back-end tape drives in existing TS7740 or TS7700T cluster with data. You might want to upgrade the back-end tape drives to a higher model to have more capacity from the existing media, because the drives are not encryption capable, or for any other reason. TS7740 and TS7700T support the 3592-J1A, TS1120 (3592-E05), TS1130 (3592-E06/EU6), TS1140 (3592-E07), and TS1150 (3592-E08) tape drives.

- ▶ TS7700 upgrade to Release 4.0
- ▶ Adding clusters to a grid
- ▶ Removing clusters from a grid

7.1 TS7700 hardware components

IBM TS7700 Release 4.0 Licensed Internal Code (LIC) runs only on a 3957 model V07/VEB/VEC. Model V07 and VEB are based on an IBM POWER7® processor-based server. Model VEC is based on an IBM POWER8 processor-based server with an I/O expansion drawer that contains Peripheral Component Interface Express (PCIe) adapters. The hardware platform enhances the performance capabilities of the subsystem when compared to the previous implementation. It also makes room for future functions and enhancements.

This section describes the hardware components that are part of the TS7700. These components include the TS7720, TS7760, and TS7740, which are attached to an IBM TS3500 or TS4500 tape library that is configured with IBM 3592 tape drives.

The TS7760D contains the following components:

- ▶ One IBM 3952 F06 Tape Base Frame, which houses the following components:
 - One TS7760 Server
 - One TS7700 Input/Output (I/O) Expansion Drawer (primary and alternate)
 - One TS7760 Encryption Capable 3956-CSA Cache Controller Drawer with up to nine optional TS7760 Encryption Capable 3956-XSA Cache Expansion Drawers
 - Two Ethernet switches
 - One TS3000 Total System Storage Console (TSSC)
- ▶ One or two optional 3952 Model F06 Storage Expansion Frames, housing the following components:
 - One TS7760 Encryption Capable 3956-CSA Cache Controller Drawer
 - Up to 15 optional TS7760 Encryption Capable 3956- XSA Cache Expansion Drawers
 - Up to three cache strings¹ that are housed within all the frames

The TS7760T contains the following components:

- ▶ One IBM 3952 Model F06 Tape Base Frame, which houses the following components:
 - One TS7760T tape-attached Server
 - One TS7700 Input/Output (I/O) Expansion Drawer (primary and alternate)
 - One TS7760 Encryption Capable 3956-CSA Cache Controller Drawer with up to nine optional TS7760 Encryption Capable 3956-XSA Cache Expansion Drawers
 - Two Ethernet switches
 - One TS3000 TSSC
- ▶ One or two optional IBM 3952 Model F06 Storage Expansion Frames, which house the following components:
 - One TS7760 Encryption Capable 3956-CSA Cache Controller Drawer
 - Up to 15 optional TS7760 Encryption Capable 3956- XSA Cache Expansion Drawers
 - Up to three cache strings that are housed within all the frames
- ▶ Connection to a TS3500 or TS4500 tape library with 4 - 16 IBM 3592 tape drives and two Fibre Channel (FC) switches

¹ A TS7760 Cache Controller and up to five attached TS7760 Cache Drawers are referred to as a *string*, with each TS7760 Cache Controller acting as the “head of [the] string.” A single TS7700 can have up to three “strings” attached, the first in the base frame (base string) and the next two in the expansion frames (string 1 and string 2).

The TS7720D contains the following components:

- ▶ One IBM 3952 F05 Tape Base Frame, which houses the following components:
 - One TS7720 Server
 - One TS7700 Input/Output (I/O) Expansion Drawer (primary and alternate)
 - One TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer with up to nine optional TS7720 Encryption Capable 3956-XS9 Cache Expansion Drawers
 - Two Ethernet switches
 - One TS3000 Total System Storage Console (TSSC)
- ▶ One or two optional 3952 Model F05 Storage Expansion Frames, housing the following components:
 - One TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer
 - Up to 15 optional TS7720 Encryption Capable 3956- XS9 Cache Expansion Drawers
 - Up to four cache strings that are housed within all the frames

The TS7720T contains the following components:

- ▶ One IBM 3952 Model F05 Tape Base Frame, which houses the following components:
 - One TS7720T tape-attached Server
 - One TS7700 Input/Output (I/O) Expansion Drawer (primary and alternate)
 - One TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer with up to nine optional TS7720 Encryption Capable 3956-XS9 Cache Expansion Drawers
 - Two Ethernet switches
 - One TS3000 TSSC
- ▶ One or two optional IBM 3952 Model F05 Storage Expansion Frames, which house the following components:
 - One TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer
 - Up to 15 optional TS7720 Encryption Capable 3956- XS9 Cache Expansion Drawers
 - Up to three cache strings that are housed within all the frames
- ▶ Connection to a TS3500 or TS4500 tape library with 4 - 16 IBM 3592 tape drives and two Fibre Channel (FC) switches

The TS7740 contains the following components:

- ▶ One IBM 3952 Model F05 Tape Base Frame, which houses the following components:
 - One TS7740 Server
 - One TS7700 Input/Output (I/O) Expansion Drawer (primary and alternate)
 - One TS7740 Encryption Capable 3956-CC9 Cache Controller Drawer with zero, one, or two TS7740 Encryption Capable 3956-CX9 Cache Expansion Drawers
 - Two Ethernet switches
 - One TS3000 TSSC
- ▶ Connection to a TS3500 or TS4500 tape library with 4 - 16 IBM 3592 tape drives and two Fibre Channel switches

7.1.1 Common components for the TS7700 models

This section lists the components for the TS7700 models.

3952 Tape Frame

The 3952 Tape Frame houses TS7700 controllers and their components. The 3952 Tape Frame that is used with the TS7700 contains the following components:

- ▶ Ethernet switches
- ▶ Optional components
 - TSSC Server
 - Keyboard and monitor
 - Ethernet switch
- ▶ TS7700 Server
 - TS7760 Server (3957-VEC)
 - TS7720 Server (3957-VEB)
 - TS7740 Server (3957-V07)
- ▶ I/O drawers
- ▶ Cache controller
 - TS7760 Cache Controller (3956-CSA)
 - TS7720 Cache Controller (3956-CS9)
 - TS7740 Cache Controller (3956-CC9)
- ▶ Optional cache expansion drawers
 - TS7760 Cache Drawer (3956-XSA)
 - TS7720 Cache Drawer (3956-XS9)
 - TS7740 Cache Drawer (3956-CX9)

The 3952 Tape Frame can be designated as a TS7700 Storage Expansion Frame when ordered with FC 7334, TS7700 Encryption-capable expansion frame.

Any lock on the 3952 Tape Frame prevents access to the TS7700 Emergency Power Off (EPO) switch. If a lock (FRU 12R9307) is installed on the 3952 Tape Frame, an external EPO switch or circuit breaker must be installed near the TS7700 to allow an emergency shutdown. Additionally, the emergency contact label that is included with the Installation Instruction RPQ 8B3585 (Front/Rear and Side Panel Locking Procedure), PN 46X6208, must be completed and affixed to the 3952 Tape Frame door in an immediately visible location. This label must clearly indicate the location of the external EPO switch or circuit breaker.

If a lock is installed on the 3952 Tape Frame and the original key is not available, any 3952 Tape Frame key can be used to open the lock. If no frame key is available and immediate access is required to get inside the frame, you must contact a locksmith to open the lock. If the key is still unavailable after the lock is opened, you can contact your IBM service representative to order a new lock and key set (FRU 12R9307).

For more information, see Table 4-1 on page 127.

The IBM 3952 Model F05 Tape Base Frame provides up to 36U (rack units or Electronics Industry Association (EIA) units) of usable space. The IBM 3952 Model F06 Tape Base Frame provides up to 40U of usable space. The rack units contain the components of the defined tape solution.

The 3952 Tape Base Frame is not a general-purpose frame. 3952 Model F05 is designed to contain the components of specific tape offerings, such as the TS7740, TS7720, and TS7720T. 3952 Model F06 is designed to contain the components of TS7760.

Only components of one solution family can be installed in a 3952 Tape Frame. The 3952 Tape Frame is configured with a Dual AC Power Distribution feature for redundancy.

Note: *Available by RPQ* is the ability for top exit as it relates to cables in the 3952 F06 frame. An example is *RPQ 8B3670*.

Ethernet switches

Previous Ethernet routers are replaced by 1 gigabit Ethernet (GbE) switches in all new TS7700 tape drives. Primary and alternate switches are used in the TS7700 internal network communications.

The communications to the external network use a set of dedicated Ethernet ports on adapters in the 3957 server. Internal network communications (interconnecting TS7700 switches, TSSC, Disk Cache System, and TS3500 or TS4500 when present) use their own set of Ethernet ports on adapters in the I/O Expansion Drawer.

Communications were previously handled by the routers, including Management Interface (MI) addresses and encryption key (EK) management. The virtual Internet Protocol (IP) address that was previously provided by the router's conversion capability is now implemented by virtual IP address (VIPA) technology.

When you replace an existing TS7700 V06/VEA with a new V07/VEB model, the old routers stay in place. However, they are reconfigured and used solely as regular switches. The existing external network connections are reconfigured and connected directly to the new V07/VEB server. Figure 7-1 shows the new 1 Gb switch and the old Ethernet router for reference.




Figure 7-1 New switch (top) and old router (bottom)

TS7700 grid adapters

The connection paths between multiple TS7700 tape drives in a grid configuration are the two grid adapters in slot one of the I/O expansion drawers. The dual-ported 1 gigabit per second (Gbps) Ethernet adapters can be copper RJ45 or optical fiber (shortwave (SW)). These optical adapters have an LC duplex connector.

Depending on your bandwidth and availability needs, TS7700 can be configured with two or four 1-Gb links. Feature Code 1034 (FC1034) is needed to enable the second pair of ports in the grid adapters. These ports can be either fiber SW or copper. Also, there is a choice of two longwave (LW) single-ported Optical Ethernet adapters (FC1035) for two or four 10-Gb links. Your network infrastructure must support 10 Gbps for this selection. The adapter does not scale down to 1 Gbps.

The Ethernet adapters cannot be intermixed within the same cluster, they must be of the same type (same feature code).

Disk encryption

The latest TS7700 cache models, 3956-CC9/CS9, 3956-CS9/XS9, and 3956-CSA/XSA, support full disk encryption (FDE). All cache controllers and cache drawers must be encryption capable in order for FDE to be activated. With FDE, data is secure at the most basic level – the hard drive. FDE covers most data exposures and vulnerabilities all at once.

FDE simplifies security planning and provides unparalleled security assurance with government-grade encryption. FDE uses the Advanced Encryption Standard (AES) 128 encryption to protect the data. This algorithm is approved by the US government for protecting secret-level classified data. Data is protected through the hardware lifecycle and enables return of defective drives for servicing. It also allows for quick decommission or repurposing of drives with instant cryptographic erase.

FDE preserves performance because the encryption is hardware-based in the disk drive. FDE doesn't slow the system down because the encryption engine matches the drive's maximum port speed and scales as more drives are added. In order for drive-level encryption to provide this value, the key that enables encryption must be protected and managed.

There are two types of keys that are used with the FDE drives. The data encryption key is generated by the drive and never leaves the drive, so it always stays secure. It is stored in an encrypted form within the drive and performs symmetric encryption and decryption of data at full disk speed with no effect on disk performance. Each FDE drive uses its own unique encryption key, which is generated when the disk is manufactured and regenerated when required by the SSR.

The lock key or security key is a 32-byte random number that authenticates the drive with the CC9/CS9/CSA Cache Controller by using asymmetric encryption for authentication. When the FDE drive is *secure enabled*, it must authenticate with the CC9/CS9/CSA Cache Controller or it does not return any data and remains locked.

After the drive is authenticated, access to the drive operates like an unencrypted drive. One security key is created for all FDE drives attached to the CC9/CS9/CSA cache controller and CX9/XS9/CSA Cache Expansion drawers. The authentication key is generated, encrypted, and hidden in the subsystem (NVS RAM) of the CC9/CS9/CSA Cache Controller in each of the two CECs. The TS7700 stores a third copy in the Vxx persistent storage disks. A method is provided to securely export a copy to DVD.

The authentication typically occurs only after the FDE starts, where it will be in a “locked” state. If encryption was never enabled (the lock key is not initially established between the CC9/CS9/CSA Cache Controller and the disk), the disk is considered unlocked with access unlimited, as in a non-FDE drive.

The lock key or security key is set up by the SSR using the SMIT panels. There are two feature codes that are required to enable FDE. Feature Code 7404 is required on all 3956-CC9, 3956-CX9, 3956-CS9, 3956-XS9, 3956-CSA, and 3956-XSA cache drawers.

In addition, the following feature codes are required:

- ▶ FC7730 is required on the 3952-F05 base frame for a TS7740.
- ▶ FC7331 is required on the 3952-F05 base frame for a TS7720.
- ▶ FC7332 is required on the 3952-F05 Expansion frame for a TS7720.
- ▶ FC7333 is required on the 3952-F06 base frame for a TS7760.
- ▶ FC7334 is required on the 3952-F06 Expansion frame for a TS7760.

Through the SMIT menus, the SSR can “erase” the cache subsystem disks by requesting the FDE drives to erase their data encryption key and generate a new one.

Starting with R3.3 code, the TS7700 can use external key management with 3956-CC9 and 3956-CS9 cache types. Similarly, R4.0 adds the same functionality with 3956-CSA. Feature codes 5276 and 5277 are required to enable the external management. Systems that are already set up with Local encryption key management can be converted to external key management.

For more information about IBM Security Key Lifecycle Manager, go to the following website:

<http://www.ibm.com/software/products/key-lifecycle-manager>

Back-end drive fiber switch placement using Cisco switch (16 Gbps)

Release 4.0 contains support for TS7740/TS7720T/TS7760T to attach to a TS4500 library. R4.0 also contains support for a new 16 Gbps Cisco fiber switch that can be used to communicate with the back-end tape drives. The TS7700 16 Gbps Cisco fiber switch can be housed in a 3584-LXX or 3584-DXX frame. TS7740/TS7720T/TS7760T can use the new 16 Gbps fiber switch attached to either a TS3500 or TS4500.

The switches can be in a frame that contains some of the associated back-end drives, or can reside in a frame that does not contain any of the associated drives. The switches are placed at the bottom of the tape library frame. The fiber patch panel must be removed from the frame if it has one. Any fiber cables for drives in the frame with the fiber switches not associated with the TS7740/TS7720T/TS7760T must be run directly to the drives.

A frame that contains the back-end switches can still house up to 12 or 16 drives (Based on TS3500 or TS4500). Feature code 4879 supplies the mounting hardware for the back-end switches and a pair of dressed eight fiber cable trunks to connect the back-end switches to the associated back end drives in the frame.

Only eight pairs are supplied in the trunks because the preferred practice for TS7740/TS7720T/TS7760T drive placement states that the drives should be split evenly between two frames. Drives that do not attach to the back-end switches must be cabled directly to the drives, because the patch panel has been removed.

Note: The TS7760T does not support 4 Gbps and 8 Gbps fiber switches for connection to the back-end drives. Currently, a TS7760T must use the Cisco 16 Gbps fiber switch to connect to the back-end drives.

Table 7-1 shows the feasible combination of TS7700, TSx500, and the necessary switches.

Table 7-1 Combination of TS7700, TSx500, and the necessary switches

Machine with code R4.0	TS3500 with 4 or 8 Gb drive switch	TS3500 with 16 Gb drive switch^a	TS4500 with 16 Gb drive switch
TS7760T	N/A	Homogeneous with all generation drives	Heterogeneous with TS1150 and TS1140
TS7720T/TS7740	All configurations as supported in R3.3	All configurations as supported in R3.3	Heterogeneous with TS1150 and TS1140

a. TS3500 top rack required due to size of switches

7.1.2 TS7760 components

This section provides topics for the TS7760 components.

Figure 7-2 displays the frame layout for a manufacturing installed TS7760.




Figure 7-2 Single frame layout of a TS7760 with a manufacturing installed 3957-VEC, 3956-CSA, and 3956-XSA

Figure 7-3 displays the frame layout for a TS7760 Storage Expansion Frame.

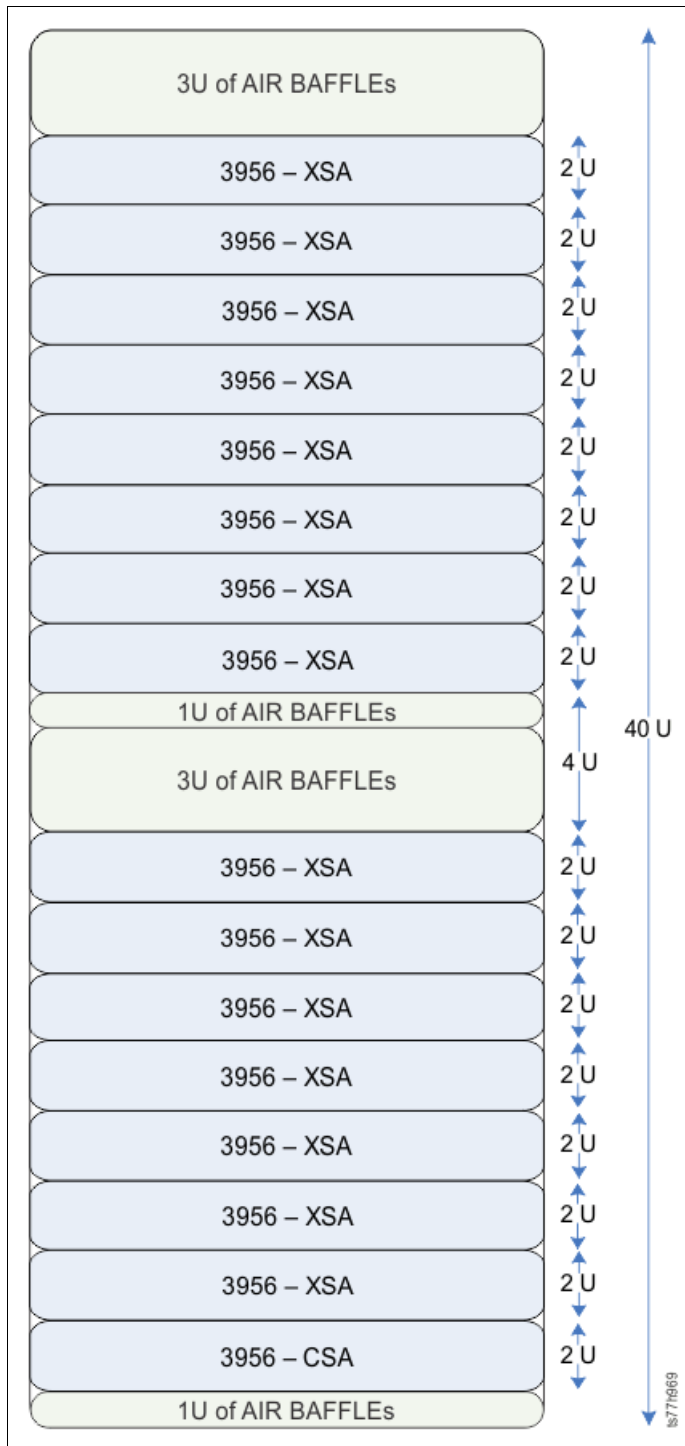


Figure 7-3 Layout of a TS7760 Storage Expansion Frame with 3956-CSA and 3956-XSA

The TS7760 Storage Expansion Frame is a 3952 Tape Frame that is designated as a cache expansion frame for use with a fully configured TS7700 Base Frame.

The TS7760 Storage Expansion Frame enables expansion of the TS7760 Cache by attaching up to two more storage frames. Each frame contains one additional TS7760 Cache Controller, which can attach to a maximum of 15 additional TS7760 Cache Drawers, resulting in a maximum addition of 16 TS7760 Cache units within each TS7760 Storage Expansion Frame. A maximally configured TS7700 Base Frame with two attached and maximally configured TS7760 Storage Expansion Frames contains 42 total cache units.

The distance between a TS7760 Storage Expansion Frame and the TS7700 Base Frame cannot exceed 10 meters. This distance enables connection of the frames by using a 30-meter cable.

The TS7760 Storage Expansion Frame consists of the following components:

- ▶ One 3952 Tape Frame
- ▶ One TS7760 Cache Controller (3956-CSA), containing 12 DDMs, each of which have a storage capacity of 4 TB
- ▶ Optional attachment to up to 15 TS7760 Cache Drawers (3956-XSA), each containing 12 DDMs with a storage capacity of 4 TB

TS7760 Server model (3957-VEC)

The TS7700 Server comprises a server and two drawers for I/O adapters. The TS7700 Server controls virtualization processes such as host connectivity and device virtualization, and hierarchical storage management (HSM) functions such as storage, replication, and organization of data across physical media and libraries.

The TS7700 Server (3957-VEC) offers the following features:

- ▶ Two 10-core 3.42 GHz POWER8 processor cards
- ▶ Processor card and memory configuration (using only 2 x 16 GB DDR3 DIMMs):
 - 32 GB total memory with 1 processor card and 20 cores
- ▶ An additional SAS controller with support for RAID 0, 1, 5, 6, and 10
- ▶ 8 SFF 300 GB SAS internal drives using RAID 0
- ▶ 1 Gb or 10 Gb Ethernet
- ▶ Four USB ports:
 - Two USB 3.0 ports for general use
 - Two USB 2.0 ports for the FSP service processor
- ▶ One system (serial) port with RJ45 connector
- ▶ Two Hardware Management Console (HMC) ports
- ▶ Extended Error Handling (EEH) and hot plug support on PCI expansion slots
- ▶ Two 6-drive SAS bays
- ▶ One slim line DVD RAM drive

Each Expansion Unit I/O adapter drawer offers the following features:

- ▶ Six additional hot-pluggable PCIe cartridge style slots (used to house FICON adapters for host attachment, Fibre Channel adapters for cache drawer attachment, Fibre Channel adapters for tape communication, and Ethernet adapters for grid communication).
- ▶ Redundant AC power
- ▶ Redundant cooling
- ▶ Concurrent maintenance of:
 - PCIe or PCI-X adapters
 - Two power supplies
 - Two fans

Figure 7-4 and Figure 7-5 show the TS7700 Server System Unit.



Figure 7-4 TS7700 Server (3957-VEC) System Unit (front view)




Figure 7-5 TS7700 Server (3957-VEC) System Unit (rear view)

Figure 7-6 shows the TS7700 Server Expansion Unit I/O drawer.



Figure 7-6 TS7700 Server Expansion Unit I/O drawer (rear view)

TS7760 Cache Controller

The TS7760 Cache Controller is a self-contained 2U enclosure that mounts in the 3952 Tape Frame.

The TS7760 Cache Controller provides dynamic disk pools-protected virtual volume disk storage for fast retrieval of data from cache. The TS7760 Cache Controller offers the following features:

- ▶ Two Fibre Channel processor cards
- ▶ CPU microprocessor

- ▶ Two battery backup units (one for each processor card)
- ▶ Two AC power supplies with imbedded enclosure cooling units
- ▶ 12 DDMs, each with a storage capacity of 4 TB (3.63 TiB)
- ▶ Supports Advanced Encryption Standard (AES) 128-bit encryption
- ▶ Optional attachment to a maximum of nine TS7760 Cache Drawers in a TS7760 Base Frame
- ▶ Optional expansion of cache capabilities (up to 15 more TS7760 Cache Drawers) when a fully populated TS7760 Base Frame attaches to a fully populated TS7760 Storage Expansion Frame
- ▶ 12 Gb SAS port
- ▶ Dual active 16 Gb Fibre Channel connectivity to server 3957-VEC

Figure 7-7 shows the TS7760 Cache Controller from the front and Figure 7-8 shows the TS7760 Cache Controller from the rear.



Figure 7-7 TS7760 Cache Controller 3956-CSA (front view)



Figure 7-8 TS7760 Cache Controller 3956-CSA (rear view)

TS7760 Cache Drawer

The TS7760 Cache Drawer is a self-contained 2U enclosure that mounts in the 3952 Tape Frame.

The TS7760 Cache Drawer expands the capacity of the TS7760 Cache Controller by providing additional dynamic disk pools-protected disk storage. Each TS7760 Cache Drawer offers the following features:

- ▶ Two Environmental Services Module (ESM) cards
- ▶ Two AC power supplies with embedded enclosure cooling units
- ▶ 12 DDMs, each with a storage capacity of 4 TB (3.63 TiB)
- ▶ Supports Advanced Encryption Standard (AES) 128-bit encryption
- ▶ Attachment to the TS7760 Cache Controller

Figure 7-9 shows the TS7760 Cache drawer from the front and Figure 7-10 shows the TS7760 Cache drawer from the rear.



Figure 7-9 TS7760 Cache Drawer 3956-XSA (front view)




Figure 7-10 TS7760 Cache Drawer 3956-XSA (rear view)

7.1.3 TS7720 components

The TS7720 provides most of the benefits of the TS7740 without physical tape attachment.

The TS7720 consists of a 3952 Model F05 Encryption Capable Base Frame and one or two optional 3952 Model F05 Encryption Capable Storage Expansion Frames. FC5272 enables FDE on the VEB. FC7404 is needed to enable FDE on each cache drawer. After it is enabled, FDE cannot be disabled.

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

- ▶ One TS7720 Server, 3957 Model VEB.
- ▶ One TS7700 I/O Expansion Drawer (primary and alternative).
- ▶ One TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer. The controller drawer has 0 - 9 TS7720 Encryption Capable 3956- XS9 Cache Expansion Drawers. The base frame must be fully configured before you can add a first storage expansion frame.
- ▶ Two Ethernet switches.

The 3952 Model F05 Storage Expansion Frame houses one TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer. Each controller drawer can have 0 - 15 TS7720 Encryption Capable 3956-XS9 Cache Expansion Drawers. The first expansion frame must be fully configured before you can add a second storage expansion frame.

The base frame, first expansion frame, and second expansion frame are not required to be of the same model and type. Only when the base frame is of the CS9 type is it required to be fully populated when you add an expansion frame. When you add a second expansion frame, the first expansion frame must be fully populated if it contains CS9 technology.

Using 3 TB HDDs, the maximum configurable capacity of the TS7720 at Release 3.2 or later with the 3952 Model F05 Storage Expansion Frame is 1007.86 TB of data before compression.

Figure 7-11 shows the TS7720 Base Frame components.




Figure 7-11 TS7720 Base Frame components

Figure 7-12 shows the TS7720 Expansion Frame components.




Figure 7-12 TS7720 Expansion Frame components

TS7720 Server model (3957-VEB)

The server consists of an IBM System POWER7 processor-based server and an expansion I/O drawer (primary and alternative) containing PCIe adapters. This replaces the original IBM POWER® 5 ++ and the I/O drawer from the V06/VEA version. The TS7700 Server controls virtualization processes, such as host connectivity and device virtualization. It also controls the internal hierarchical storage management (HSM) functions for logical volumes and replication.

Figure 7-13 shows the front view of the TS7700 Server models 3957-VEB.



Figure 7-13 TS7700 Server models 3957-VEB (front view)

The TS7700 Server VEB offers the following features:

- ▶ Rack-mount (4U) configuration.
- ▶ One 3.0-gigahertz (GHz) 8-core processor card.
- ▶ 16 GB of 1066 MHz error-checking and correcting (ECC) memory (32 GB when 8-Gb Fibre Channel connection (FICON) is present).
- ▶ The following integrated features:
 - Service processor
 - Quad-port 10/100/1000 megabits (Mb) Ethernet
 - IBM EnergyScale™ technology
 - Hot-swap capability and redundant cooling
 - Two system (serial) ports
 - Two Hardware Management Console (HMC) ports
 - Two system power control network (SPCN) ports
 - One slim bay for a DVD-RAM
- ▶ Five hot-swap slots:
 - Two PCIe x8 slots, short card length (slots 1 and 2)
 - One PCIe x8 slot, full card length (slot 3)
 - Two PCI-X DDR slots, full card length (slots 4 and 5)

The hot-swap capability is only for replacing an existing adapter with another of the same type. It is not available if you change adapter types in a machine upgrade or change.
- ▶ SAS HDD: The TS7700 uses eight HDDs. Four disks mirror one SAS adapter, and the other four backup drives are assigned to a separate SAS adapter.
- ▶ A SAS card is used for the mirroring and SAS controller redundancy. It has an external cable for accessing the mirrored disks.

Each I/O expansion Drawer offers six extra PCI Express adapter slots:

- ▶ One or two 4 Gb FICON adapters per I/O Expansion Drawer, for a total of two or four FICON adapters per cluster. Adapters can work at 1, 2, or 4 Gbps. FICON card must be of the same type within one cluster.
- ▶ One or two 8 Gb FICON adapters per I/O Expansion Drawer, for a total of two or four FICON adapters per cluster. Adapters can work at 2, 4, or 8 Gbps.

The FICON card must be of the same type within one cluster. 8 Gb FICON card requires FC 3462, 16 GB memory upgrade. All servers with 8 Gb FICON adapters require 32 GB of memory, an increase of 16 GB of memory to the default server configuration.

- ▶ Grid Ethernet card (PCI Express). Grid Ethernet can be copper or fiber (1 or 10 Gbps).
- ▶ 8 Gbps Fibre Channel to disk cache (PCI Express).
- ▶ 8 Gbps Fibre Channel PCI Express connection to tape in a TS7740 and TS7720T or 8 Gbps Fibre Channel PCI Express for connection to TS7720 Expansion frames.

TS7720 Cache Controller (3956-CS9)

The TS7720 Encryption Capable 3956-CS9 Cache Controller Drawer is a self-contained 2U enclosure. It mounts in the 3952 Tape Base Frame and the optional 3952 Storage Expansion Frame. Figure 7-14 shows the TS7720 Cache Controller Drawer from the front (left side) and rear (right side). The rear view details the two separated controllers that are used for access redundancy and performance (Controller A on left and Controller B on the right side).



Figure 7-14 TS7720 Encryption Capable Cache Controller, 3956-CS9 (front and rear views)

The TS7720 Cache Controller provides RAID 6 protection for virtual volume disk storage, enabling fast retrieval of data from cache.

The TS7720 Cache Controller Drawer offers the following features:

- ▶ Two 8 Gbps Fibre Channel processor cards
- ▶ Two battery backup units (one for each processor card)
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ Twelve disk drive modules (DDMs), each with a storage capacity of 3 TB, for a usable storage capacity of 23.86 TB
- ▶ Configurations with only CS9 controllers support one, two, or three TS7720 Cache Controllers:
 - All configurations provide one TS7720 Cache Controller in the 3952 Tape Base Frame. The 3952 Tape Base Frame can have 0 - 9 TS7720 Encryption Capable SAS Cache Drawers, 3956 Model XS9.
 - All configurations with the optional 3952 Storage Expansion Frame provide one extra TS7720 Encryption Capable Cache Controller, 3956 Model CS9. When the second is added, an extra set of 8 Gb FC adapters is also added. The 3952 Storage Expansion Frame can have 0 - 15 TS7720 Encryption Capable SAS Cache Drawers, 3956 Model XS9.

TS7720 Cache Drawer (3956-XS9)

The TS7720 Encryption Capable Cache Drawer is a self-contained 2U enclosure. It mounts in the 3952 Tape Base Frame and in the optional 3952 Storage Expansion Frame. Figure 7-15 shows the TS7720 Cache Drawer from the front (left side) and rear (right side). It offers attachment to the TS7720 Encryption Capable Cache Controller.



Figure 7-15 TS7720 Encryption Capable Cache Drawer (front and rear views)

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

- ▶ Two 8 Gb Fibre Channel processor cards
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ Eleven DDMs, each with a storage capacity of 3 TB, for a usable capacity of 24 TB per drawer

The TS7720 disk-only can be used to write tape data that does not need to be copied to physical tape, which enables access to the data from the Tape Volume Cache (TVC) until the data expires.

The TS7720T enables a TS7720 to act like a TS7740 to form a virtual tape subsystem to write to physical tape. Full disk and tape encryption are supported. It contains the same components as the TS7720 disk-only. In a TS7720 disk-only configuration, the Fibre Channel ports are used to communicate with the attached cache, while in a TS7720T configuration two of the Fibre Channel ports are used to communicate with the attached tape drives.

7.1.4 TS7740 components

The TS7740 combines the TS7700 with a tape library to form a virtual tape subsystem to write to physical tape. Since Release 3.2, TS7740 plant-built configurations include these components:

- ▶ One TS7740 Server, 3957 Model V07.
- ▶ One TS7740 Encryption Capable 3956-CC9 Cache Controller Drawer.
- ▶ The controller drawer has a maximum of three attached Encryption Capable 3956-CX9 Cache Expansion Drawers.

The total usable capacity of a TS7740 with one 3956-CC9 and two 3956-CX9s is approximately 28 TB before compression.

The Model CX9s can be installed at the plant or in an existing TS7740.

Figure 7-16 shows a summary of the TS7740 components.




Figure 7-16 TS7740 components

TS7740 Server model (3957-V07)

The detailed hardware of the TS7740 server 3957-V07 is that same as the TS7720 server 3957-VEB. For more information, see “TS7720 Server model (3957-VEB)” on page 242.

TS7740 Cache Controller (3956-CC9)

The TS7740 Encryption Capable Cache Controller is a self-contained 2U enclosure that mounts in the 3952 Tape Frame.

Figure 7-17 shows the front and rear views of the TS7740 Encryption Capable 3956-CC9 Cache Controller Drawer.



Figure 7-17 TS7740 Encryption Capable Cache Controller Drawer (front and rear views)

Figure 7-18 shows a diagram of the rear view, detailing the two separated controllers that are used for access redundancy and performance (Controller A and Controller B).




Figure 7-18 TS7740 Encryption Capable Cache Controller (rear view)

The TS7740 Encryption Capable Cache Controller Drawer provides RAID 6-protected virtual volume disk storage. This storage temporarily holds data from the host before writing it to physical tape. When the data is in the cache, it is available for fast retrieval from the disk.

The TS7740 Cache Controller Drawer offers the following features:

- ▶ Two 8 Gbps Fibre Channel processor cards
- ▶ Two battery backup units (one for each processor card)
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ Twenty-two DDMs, each possessing 600 GB of storage capacity, for a usable capacity of 9.45 TB
- ▶ Optional attachment to one or two TS7740 Encryption Capable 3956-CX9 Cache Expansion Drawers

TS7740 Cache Expansion Drawers (3956-CX9)

The TS7740 Encryption Capable Cache Expansion Drawer is a self-contained 2U enclosure that mounts in the 3952 Tape Frame.

Figure 7-19 shows the front view and the rear view of the TS7740 Encryption Capable 3956-CX9 Cache Expansion Drawer.




Figure 7-19 TS7740 Encryption Capable Cache Drawer (front and rear views)

The TS7740 Encryption Capable Cache Expansion Drawer expands the capacity of the TS7740 Cache Controller Drawer by providing extra RAID 6 disk storage. Each TS7740 Cache Expansion Drawer offers the following features:

- ▶ Two Environmental Service Modules (ESMs)
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ 22 DDMs, each with 600 GB of storage capacity, for a total usable capacity of 9.58 TB per drawer
- ▶ Attachment to the TS7740 Encryption Capable 3956-CC9 Cache Controller Drawer

7.1.5 TS7700 tape library attachments, drives, and media

In TS7700 configurations, the TS7740, TS7720T, and TS7760T are used with an attached tape library. The TS7740, TS7720T, and TS7760 must have their own logical partitions (LPARs) within the TS3500 or TS4500 tape library, with dedicated tape drives and tape cartridges.

Tape libraries

A TS7740, TS7720T, or a TS7760T attached to a TS3500 or TS4500 tape library interfaces directly with tape drives in the library.

When attached to a TS3500 or TS4500 tape library, the TS7700 can attach only to 3592 Tape Drives. Up to 16 3592 Tape Drives can be attached.

Communication, control, and data signals travel along Fibre Channel connections between the TS7700 and tape drives contained in the TS3500 or TS4500 tape library. A pair of Fibre Channel switches routes the data to and from the correct tape drive.

Note: TS1140 EH7 tape drives and TS1150 EH8 tape drives are used with the TS4500 tape library. All other tape drives are used with the TS3500 tape library.

Tape drives

The 3592 Tape Drives supported for use with the TS7740, TS7720T, and TS7760T include:

- ▶ TS1150 Tape Drive
- ▶ TS1140 Tape Drive
- ▶ TS1130 Tape Drives
- ▶ TS1120 Tape Drives (in native mode and emulating 3592 J1A Tape Drives)
- ▶ 3592 J1A Tape Drives

Tape media

The TS7740 and TS7720T support the following types of media:

- ▶ 3592 Tape Cartridge (JA)
- ▶ 3592 Expanded Capacity Cartridge (JB)
- ▶ 3592 Advanced Type C Data (JC)
- ▶ 3592 Advanced Type D Data (JD)
- ▶ 3592 Economy Tape Cartridge (JJ)
- ▶ 3592 Advanced Type K Economy (JK)
- ▶ 3592 Advanced Type L Economy (JL)

For more information, see “Tape drives and media support (TS7740,TS7720T, and TS7760T)” on page 131.

7.1.6 TS3000 Total System Storage Console

The TS3000 TSSC connects to multiple Enterprise Tape Subsystems, including TS3500/TS4500 tape libraries, 3592 Controllers, and the TS7700. The TS3000 TSSC is a required component for the TS7700. It can be a new console or an existing TSSC. A new TSSC can be installed in the TS7700 3952 Tape Base Frame or another existing rack.

All of these devices are connected to a dedicated, private local area network (LAN) that is owned by TSSC. Remote data monitoring of each one of these subsystems is provided for early detection of unusual conditions. The TSSC sends this summary information to IBM if something unusual is detected and the Call Home function has been enabled.

Note: For Call Home and remote support since TS7700 R3.2, an internet connection is necessary.

For IBM TS7700 R4.0, the following features are available for installation:

- ▶ FC 2704, TS3000 System Console expansion 26-port Ethernet switch/rackmount
- ▶ FC 2715, Console attachment
- ▶ FC 2725, Rackmount TS3000 System Console
- ▶ FC 2748, Optical drive

For more information, see Appendix A, “Feature codes and RPQ” on page 823.

7.1.7 Cables

This section describes the cable feature codes for attachment to the TS7700, extra cables, fabric components, and cabling solutions.

Required cable feature codes

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

A TS7700 Server with the FICON Attachment features (FC 3441, FC 3442, FC 3443, FC 3438, or FC 3439) can attach to FICON channels of IBM z Systems components by using FICON cable features ordered on the TS7700 Server. A maximum of eight FICON cables, each 31 meters, can be ordered.

One cable must be ordered for each host system attachment by using the following cable features:

- ▶ FC 3442 and FC 3443, 4-Gb FICON Long-Wavelength Attachment feature: The FICON long-wavelength adapter that is included with FC 3442 (4-Gb FICON Long-Wavelength Attachment) or FC 3443 (4 Gb FICON 10-kilometer (km) Long-Wavelength Attachment) has an LC Duplex connector. It can connect to FICON long-wavelength channels of z Systems components by using a 9-micron single-mode fiber cable.

The maximum fiber cable length is 4 KM (2.48 miles) for FC 3442 and 10 KM (6.2 miles) for FC 3443. If standard host attachment cables (31 m) are required, they can be specified with FC 0201 (9-micron LC/LC 31-meter fiber cable) or FC 0203 (50 micron LC/LC 31-meter fiber cable).

- ▶ FC3441 (4 Gb FICON Short-Wavelength Attachment) feature: The FICON shortwave-length adapter that is included with FC 3441 has an LC Duplex connector. It can connect to FICON short-wavelength channels of z Systems components by using a 50-micron or 62.5-micron multimode fiber cable. At 4 Gbps, the maximum fiber cable length that is allowed by 50-micron cable is 150 m, or 55 m if you use a 62.5-micron cable.

If standard host attachment cables are required, they can be specified with FC 0203 - 50 Micron LC/LC 31-meter fiber cable and FC 3438, 8 Gb FICON Short-Wavelength Attachment.

Requirement: 8-Gb FICON adapters require FC 3462 (16-GB memory upgrade) and TS7700 Licensed Internal Code R3.1 or later.

- ▶ Wavelength Attachment provides one short-wavelength FICON adapter with an LC Duplex connector for attachment to a FICON host system shortwave (SW) channel by using a 50 micron or 62.5-micron multimode fiber cable. Each FICON attachment can support up to 512 logical channels. At 8 Gbps speed, the total cable length cannot exceed the following lengths:

- 150 meters using 50-micron OM3 (2000 MHz*km) Aqua blue-colored fiber
- 50 meters using 50-micron OM2 (500 MHz*km) Orange-colored fiber
- 21 meters using 62.5-micron OM1 (200 MHz*km) Orange-colored fiber

If standard host attachment cables are required, they can be specified with FC 0201 (9-micron LC/LC 31-meter fiber cable) or FC 0203 (50-micron LC/LC 31-meter fiber cable).

- ▶ FC 3439 (8 Gb FICON Long Wavelength Attachment) provides one long-wavelength FICON adapter, with an LC Duplex connector, for the attachment to a FICON host system long wave channel that uses a 9-micron single-mode fiber cable. The total cable length cannot exceed 10 km. Each FICON attachment can support up to 512 logical channels.

If standard host attachment cables are required, they can be specified with FC 0201 (9-micron LC/LC 31-meter fiber cable) or FC 0203 (50 micron LC/LC 31-meter fiber cable).

Requirement: FC 3401 (Enable 8 Gb FICON dual port) enables the second port on each installed 8-Gb FICON adapter. With FC 3401, two instances of FC 0201 or FC 0203 are required for each FC 3438 or FC 3439.

Extra cables, fabric components, and cabling solutions

Conversion cables from SC Duplex to LC Duplex are available as features on z Systems platforms if you are using cables with SC Duplex connectors that now require attachment to fiber components with LC Duplex connections. Extra cable options, along with product support services, such as installation, are offered by IBM Global Technology Services®.

See the *IBM Virtualization Engine TS7700 Introduction and Planning Guide*, GA32-0568, for Fibre Channel cable planning information.

If Grid Enablement (FC4015) is ordered, Ethernet cables are required for the copper/optical 1 Gbps and optical LW adapters to attach to the communication grid.

7.2 TS7700 component upgrades

Several field-installable upgrades give an existing TS7700 more functions or capacities. This section reviews the TS7700 component FC upgrades.

7.2.1 TS7700 concurrent system component upgrades

Concurrent system upgrades can be installed while the TS7700 is online and operating. The following component upgrades can be made concurrently to an existing, onsite TS7700:

- ▶ Enable 1 TB pending tape capacity (FC 5274).

Each instance of FC 5274 (Enable 1 TB Pending Tape Capacity) enables up to 1 TB of data to be pending migration to physical tape. Up to 10 instances can be installed.

- ▶ Incremental disk cache capacity enablement (TS7740 only).

You can add a 1 TB (0.91 tebibytes (TiB)) increment of disk cache to store virtual volumes, up to 28 TB (25.46 TiB). Use FC 5267 (1 TB cache enablement) to achieve this upgrade.

- ▶ Enable 8 Gb FICON second port. Use FC 3401.

- ▶ Incremental data throughput.

You can add a 100 MBps increment of peak data throughput, up to your system's hardware capacity. When the maximum number of performance increments are installed, the system no longer restricts performance. Use FC 5268 (100 MBps increment) to achieve this upgrade.

Peak data throughput increments of 100 MBps are available as transferred from a host to a vNode before compression. If additional peak data throughput capacity is needed, up to 10 more increments can be ordered for the 3957-V07/VEB when FC 3441/3442/3443 (4 Gb FICON) is installed, or up to 24 can be ordered for the 3957-V07/VEB/VEC when FC 3438/3439 (8 Gb FICON) is installed (to complement FC 9268, Plant installation of 100 MBps throughput installed at the plant).

If the maximum number of increments is installed, the TS7700 places no limits on data transfers through the cluster. This installation is performed by you through the TS7700 Management Interface by entering the license key that is obtained with the purchase of FC 5268, 100 MBps increment.

Note: All host data transfers through the TS7740 Cluster are considered for the data transfer limit regardless of which TS7740 Cluster initiated or received the data transfer.

- ▶ Selective Device Access Control.

You can grant exclusive access to one or more logical volume ranges by only certain logical control units (LCUs) or subsystem IDs within a composite library for host-initiated mounts, ejects, and changes to attributes or categories. Use FC 5271, Selective Device Access Control (SDAC) to add this upgrade.

Each instance of this feature enables the definition of eight selective device access groups. The default group provides a single access group, resulting in nine total possible access groups. This feature is available only with a Licensed Internal Code level of 8.20.0.xx or later.

Consideration: The feature must be installed on all clusters in the grid before the function becomes enabled.

► Increased logical volumes.

The default number of logical volumes that is supported is 1,000,000. You can add support for extra logical volumes in 200,000 volume increments by using FC 5270. Up to a total of 4,000,000 logical volumes are supported by the maximum quantity of 15 FC 5270 components.

Remember: The number of logical volumes that are supported in a grid is set by the cluster with the smallest number of FC 5270 increments installed.

When joining a cluster to an existing grid, the joining cluster must meet or exceed the currently supported number of logical volumes of the existing grid.

When merging one or more clusters into an existing grid, all clusters in the ending grid configuration must contain enough FC 5270 increments to accommodate the sum of all post-merged volumes.

► Dual-port grid connection.

You can concurrently enable the second port of each dual port, 1 Gb grid connection adapter in the following TS7700 Server configurations:

- On a 3957-V07, 3957-VEB, or 3957-VEC when FC 1036 (1 Gb grid dual port copper connection) or FC 1037 (1 Gb dual port optical SW connection) are present.
- On a 3957-VEC when FC 1038 (10 Gb dual port optical LW connection) is present.

Use FC 1034 (Enable dual port grid connection) to achieve these upgrades.

Note: A non-concurrent install is required to change grid adapters to a different configuration.

► Disk encryption.

You can encrypt the DDMs within a TS7700 disk storage system.

► TS7700 Storage Expansion frame.

You can add up to two cache expansion frames to a fully configured TS7760 using FC 9323 (Expansion frame attachment) and applying FC 7334 (TS7700 Encryption-capable expansion frame) to a 3952 F06 Tape Frame.

You can add up to two cache expansion frames to a fully configured TS7720 using FC 9323 (Expansion frame attachment) and applying FC 7323 (TS7720 Storage expansion frame) to a 3952 F05 Tape Frame. For cache upgrade requirements and configurations, see 7.2.4, “TS7720 Cache upgrade options” on page 259.

Note: The adapter installation (FC 5241, Dual port FC HBA) is non-concurrent.

7.2.2 TS7700 non-concurrent system component upgrades

A multi-cluster GRID configuration can enable practically all changes or upgrades to be concurrent from a client's standpoint, putting one individual member in service at a time. In a stand-alone cluster configuration, non-concurrent upgrades require the TS7700 to be brought offline before installation. In certain instances, the targeted component must be reconfigured before the upgrade takes effect. The component upgrades listed in the following sections must be made non-concurrently to an existing TS7700:

► TS7720 /TS7760 enable tape attach

For TS7720 /TS7760 servers (VEB/VEC hardware), use the following features to upgrade and enable the server to attach to a tape library:

- FC 5273 (TS7720/TS7760 Tape Attach enablement)
- FC 5274 (Enable 1 TB Pending Tape Capacity)

Before installing the features, ensure that all hardware and microcodes prerequisites are met. In addition, you should plan the future usage of the TS7700T and prepare the necessary information ahead. After the upgrade is installed, you need to update the information on the MI.

- Number of tape partitions: Depending on the usage decision, you might want to introduce multiple tape partitions, either for different customers, different LPARs (like prod, test, and development), or for different workload types.
- Size of the tape partitions: CP1 is the only tape partition that will be assigned during the upgrade. The origin tape partition size is 3 TB. Consider that currently all data is still in CP0, and if you do not redirect the workload to the CP1 the origin value is acceptable. However, you should adjust the size before you redirect the workload to CP1.
- Delay premigration limit: In each tape partition, you can define how much data with a “delay premigration” attribute can be stored.
- Storage Class updates: To direct workload to the Tape partitions, consider whether you want to introduce either new storage classes, or redirect the workload of existing storage classes to the CP1 partition.
- Storage Group: You might want to use different physical pools in the TS7700T. If so, you need dedicated storage groups to direct the data to the physical pool.
- Setup of the physical pool: You might want to specify the number of premigration drives, or the reclaim value in a physical pool. In addition, you might also want to introduce a Copy export pool. If so, you also need to update the appropriate management class.
- Set inhibit reclaim schedule: Depending on your workload, you might want to inhibit the reclaim at different schedules in the day or week.
- Adjust Settings: Adjust at least PMPRIOR and PMTHLVL (use the amount of FC 5274 for the PMTHLVL). We also strongly advise you to update the ALERT settings to ensure that alerts will be reported on several threshold levels.
- Update your automation to pick up the new Alert messages.
- After the upgrade, all data is still in CP0. If you want to move this data to a CPx partition, IBM provides a **LI REQ PARTRFSH** command to do so. However, plan the movement carefully. All data moved from a CP0 counts immediately to the premigration queue. Moving too much data concurrently fills up the premigration queue and can lead to throttling.

For detailed information about the **LI REQ PARTRFSH** command, see the *IBM TS7700 Series z/OS Host Command Line Request User's Guide*:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

► 8 Gb FICON adapters

You can install up to two 8 Gb FICON adapters or exchange adapters for another type (SW-to-LW or LW-to-SW) to connect a TS7700 Server (3957-V07, 3957-VEB, or 3957-VEC) to a host system. FICON adapter replacement is non-concurrent when used with a 3957-V07 or 3957-VEB. Use FC 3438 (8 Gb FICON Short Wavelength Attachment) or FC 3439 (8 Gb FICON Long Wavelength Attachment) for this installation.

You can also use FC 3401 (Enable 8 Gb FICON dual port) to enable a second 8 Gb FICON adapter port for double the number of host connections. The enablement of the second port is concurrent.

When replacing the 4 Gbps FICON cards with 8 Gbps FICON cards, make sure that you also examine the number of performance increments. If you are exceeding 1,000 MBps on your 4 Gbps machine, and you install 8 Gbps FICON without more throughput increments, you could actually reduce your peak performance. It is very important to consider purchasing additional throughput increments when you upgrade to 8 Gbps FICON.

► 4 Gb FICON adapters

You can install Fibre Channel (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 for SW (4 km (2.48 miles)) or LW (10 km (6.2 miles)). Replacement of an existing FICON adapter requires the removal of the original feature and addition of the new feature.

Use FC 3441, FICON short-wavelength attachment, FC 3442, FICON long-wavelength attachment, and FC 3443, FICON 10-km long-wavelength attachment for these upgrades.

► Ethernet adapters for grid communication:

– SW fiber Ethernet

You can add a 1 Gb shortwave fibre Ethernet adapter for grid communication between TS7700s.

For a TS7760 Server (3957-VEC) use FC 1037, 1 Gb dual port optical SW connection to achieve this upgrade.

For a 3957-V07 or 3957-VEB use FC 1037, 1 Gb dual port optical SW connection to achieve this upgrade.

– LW fiber Ethernet

You can add a longwave fibre Ethernet adapter for grid communication between TS7700s.

For a 3957-VEC use FC 1038, 10 Gb dual port optical LW connection to achieve this upgrade.

For a 3957-V07 or 3957-VEB use FC 1035, 10 Gb grid optical LW connection to achieve this upgrade.

Consideration: These 10 Gb adapters cannot negotiate down to run at 1 Gb. They must be connected to a 10 Gb capable network connection.

- Copper Ethernet

You can add a 1 Gbps copper Ethernet adapter for grid communication between TS7700 tape drives. On a 3957-V07, 3957-VEB, or 3957-VEC use FC 1036, 1 Gbps grid dual port copper connection to achieve this upgrade.

Clarification: On a TS7700, you can have two 1 Gbps copper Ethernet adapters *or* two 1 Gbps SW fiber Ethernet adapters *or* two 10 Gbps LW fiber Ethernet adapters (3957-V07, VEB, and VEC only) installed. Intermixing different types of Ethernet adapters within one cluster is not supported.

- TS7700 Server dual copper/optical Ethernet Adapter Card Conversion

You can convert a dual port grid Ethernet adapter in a TS7700 Server for a dual port adapter of the opposite type, by ordering FC 1036 (dual port copper) in exchange for a dual port optical Ethernet adapter FC 1037, or vice versa. In a similar way, you can order the 10 Gb grid LW adapter (FC 1035) in exchange for the 1 Gbps adapters (FC 1036 and FC 1037) and vice versa.

When you upgrade four 1 Gbps grid links to two 10 Gbps grid links, please consult with your IBM SSR. Each grid link can communicate only to the corresponding grid link in other clusters in the grid. When one cluster upgrades four 1 Gbps grid links to two 10 Gbps grid links, only these two can communicate to the cluster with 4 times 1 Gbps.

The other two grid links will not be used any more, and will report link degradation hardware messages. Consider disabling these unused grid links until all the clusters upgrade to two 10 Gbps grid links.

- ▶ TS7720 Server Fibre Channel host bus adapter installation

You can install two Fibre Channel interface cards in the TS7760 Server (3957-VEC) to connect the TS7760 Server to the disk arrays in the TS7760 Storage Expansion Frame. Use FC 5242, Dual Port 16 Gb Fibre Channel HBA to achieve this installation.

You can install two Fibre Channel interface cards in the TS7720 Server (3957-VEB) to connect the TS7720 Server to the disk arrays in the TS7720 Storage Expansion Frame. Use FC 5241, Dual port FC HBA to achieve this installation.

- ▶ TS7700 Server physical memory upgrade

You can add 16 GB of physical memory to a 3957-V07 or 3957-VEB that contains 16 GB, for a resulting total of 32 GB of physical memory. Use FC 3462, 16 GB memory upgrade to achieve this upgrade.

- ▶ TS7740 to TS7760T frame replacement

This replaces a 3952 F05 frame containing a TS7740 3957-V06 (F05 with FC 5628) or a 3952 F05 frame containing a TS7740 3957-V07 (F05 with FC 5629) with 3952 F06 frame containing a TS7760 tape attach 3957-VEC.

16 Gb switches are required for TS7760T connectivity and R2.1 or above microcode is mandatory. Existing J1A/TS1120/TS1130/TS1140/TS1150 and the accompanying tape media can be retained.

- ▶ Additional Virtual Devices (FC5275)

Each instance of FC 5275 enables up to 16 additional virtual devices (virtual tape drives). The maximum number of features that is supported is 15 for a total of 496 virtual devices. If FC 5275 is installed on a cluster in a mixed GRID microcode level, the additional devices will not be available to the host until all cluster members are at 8.32.x.x or later.

When all cluster members are at 8.32.x.x or later, the cluster member with FC 5275 needs to be taken offline and back online. If this was done previously, while the cluster member had FC 5275 installed, there is no need to take the cluster back offline again.

To gain access to the new devices, the IODF needs to be modified. Additional control units with new libport-IDs need to be defined. This can either be done before or after the installation of the feature codes. For more information, see Appendix H, “Extra IODF examples” on page 925.

After the installation and activation of the FCs and the activation of the new IODF, you need to restart the OAM address space on all attached LPARs where the additional drives will be used to refresh the unit control blocks in OAM. Then you can vary online the drives to the LPARs.

If FC 5271 (Selective Device Access Enablement) is installed on the cluster the customer must need to include the new Library Port IDs into the Access port groups using the TS7700 Management Interface. If you do not include these in the appropriate access groups, the z/OS can select the device for a mount, but the TS7700 does not enable you to mount the virtual volume. The following message is displayed:

```
CBR4175I Volume volser library library-name access group denies mount
```

7.2.3 TS7760 Cache upgrade options

This section describes the TVC upgrade options that are available for the TS7760. If you want to implement encryption, see the feature codes in Appendix A, “Feature codes and RPQ” on page 823.

For the data storage values in TB versus TiB, see 1.6, “Data storage values” on page 12.

TS7760 Base frame minimum cache configuration is one CSA Cache Controller with 31.42 TB. Up to nine XSA Cache Drawers can be added to the base frame. Single drawer increment is supported.

Expansion frame minimum cache configuration is one CSA Cache Controller. Up to fifteen XSA Cache Drawers can be added to the expansion frame. Single drawer increment is supported.

Up to two Expansion frames can be installed.

Figure 7-20 shows the minimum and maximum cache configuration of TS7760.

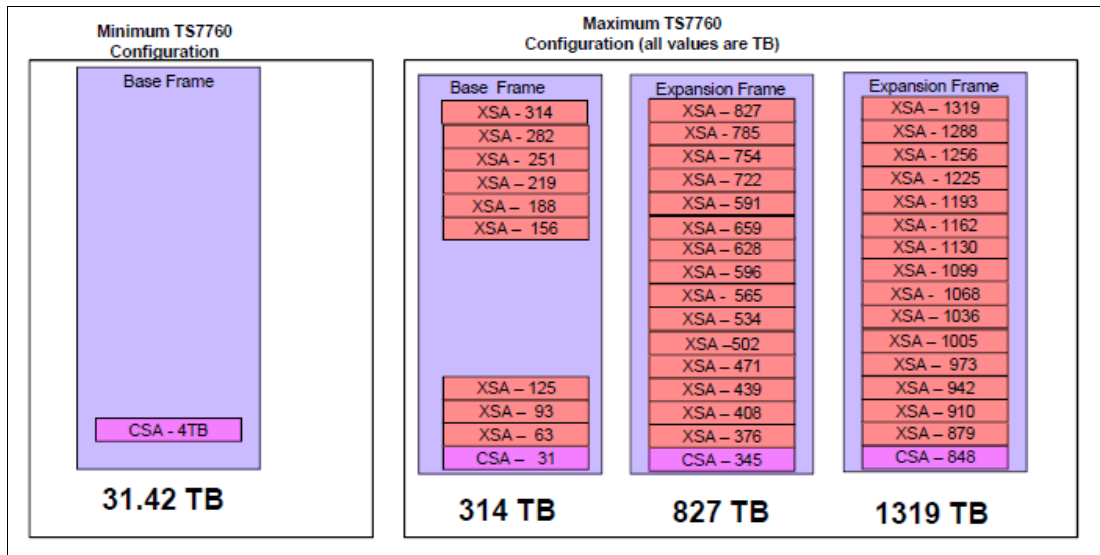


Figure 7-20 TS7760 Cache capacity

Table 7-2 shows the resulting usable capacity that is associated with each upgrade configuration.

Table 7-2 TS7760 Storage Expansion Frame configurations

Cache configuration in a new TS7760	Cache units ^a in each TS7760 Storage Expansion Frame cache controller (3956-CSA) plus optional cache drawers (3956-XSA)	First TS7760 Storage Expansion Frame		Second TS7760 Storage Expansion Frame	
		Total cache units (including TS7760 Base Frame)	Available capacity	Total cache units (including TS7760 Base Frame)	Available capacity
1 TS7760 Cache Controller (3956-CSA) 9 TS7760 Cache Drawers (3956-XSA)	1 (controller only)	11	345.07 TB (313.84 TiB)	27	847.60 TB (770.88 TiB)
	2	12	376.49 TB (342.42 TiB)	28	879.03 TB (799.47 TiB)
	3	13	407.92 TB (371 TiB)	29	910.45 TB (828.05 TiB)
	4	14	439.35 TB (399.59 TiB)	30	941.88 TB (856.64 TiB)
	5	15	470.78 TB (428.17 TiB)	31	973.31 TB (885.22 TiB)
	6	16	502.21 TB (456.76 TiB)	32	1004.74 TB (913.80 TiB)
	7	17	533.64 TB (485.34 TiB)	33	1036.17 TB (942.39 TiB)
	8	18	565.06 TB (513.92 TiB)	34	1067.60 TB (970.97 TiB)
	9	19	596.49 TB (542.51 TiB)	35	1099.02 TB (999.56 TiB)
	10	20	627.92 TB (571.09 TiB)	36	1130.45 TB (1028.14 TiB)
	11	21	659.35 TB (599.67 TiB)	37	1161.88 TB (1056.72 TiB)
	12	22	690.78 TB (628.26 TiB)	38	1193.31 TB (1085.31 TiB)
	13	23	722.21 TB (656.84 TiB)	39	1224.74 TB (1113.89 TiB)
	14	24	753.63 TB (685.43 TiB)	40	1256.17 TB (1142.48 TiB)
	15	25	785.06 TB (714.01 TiB)	41	1287.59 TB (1171.06 TiB)
	16	26	816.49 TB (742.59 TiB)	42	1319.02 TB (1199.64 TiB)

a. The term *Total cache units* refers to the combination of cache controllers and cache drawers.

7.2.4 TS7720 Cache upgrade options

This section describes the TVC upgrade options that are available for the TS7720. If you want to implement encryption, see the feature codes in Appendix A, “Feature codes and RPQ” on page 823.

For the data storage values in TB versus TiB, see 1.6, “Data storage values” on page 12.

TS7720 existing frame operating with a 3956-CS7/CS8 controller drawer can be expanded by populating any CS9 expansion frames that already exist, or by adding another expansion frame that contains CSA/XSA cache drawers. Empty CS7/CS8 slots remain empty.

A TS7720 existing frame operating with a 3956-CS9 controller drawer can be expanded by adding XS9 expansion drawers, or by adding CSA-based expansion frames after the existing CS9 frames are completely populated.

Note: CS9/XS9 MES will be withdrawn in the future, so consult with your IBM service support representative for the upgrade options.

7.2.5 TS7740 Tape Volume Cache upgrade options

This section describes the TVC upgrade options that are available for the TS7740. If you want to introduce encryption, see the feature codes in Appendix A, “Feature codes and RPQ” on page 823. Incremental features help tailor storage costs and solutions to your specific data requirements.

Subsets of total cache and peak data throughput capacity are available through incremental features FC 5267, 1 TB cache enablement, and FC 5268, 100 MBps increment. These features enable a wide range of factory-installed configurations, and enable you to enhance and update an existing system.

They can help you meet specific data storage requirements by increasing cache and peak data throughput capability to the limits of your installed hardware. Increments of cache and peak data throughput can be ordered and installed concurrently on an existing system through the TS7740 MI.

Incremental disk cache capacity enablement

Incremental disk cache capacity enablement is available in 1 TB (0.91 TiB) increments in a TS7740 cluster. Disk cache is used for these types of data:

- ▶ Data that is originated by a host through the vnodes of a local or remote cluster
- ▶ Data recalled from a physical tape drive associated with the cluster
- ▶ Data that is copied from another cluster

The capacity of the system is limited to the number of installed 1 TB increments, but the data that is stored is evenly distributed among all physically installed disk cache. Therefore, larger drawer configurations provide improved cache performance even when usable capacity is limited by the 1 TB installed increments. Extra cache can be installed up to the maximum capacity of the installed hardware.

Table 7-3 on page 260 displays the maximum physical capacity of the TS7740 Cache configurations and the instances of FC 5267, 1 TB cache enablement, required to achieve each maximum capacity. Install the cache increments by using the TS7740 MI.

Considerations:

- ▶ A minimum of one instance of FC5267, 1 TB cache enablement, can be ordered on the TS7740 Cache Controller, and the required amount of disk cache capacity is 1 TB.
- ▶ Enough physical cache must be installed before you add extra 1-TB cache increments.
- ▶ Cache increments become active within 30 minutes.
- ▶ FC5267, 1 TB Cache Enablement, is not removable after activation.

Table 7-3 shows the maximum physical capacity of the TS7740 Cache configurations by using the 3956-CC9 cache controller.

Table 7-3 Supported TS7740 Cache configurations that use the 3956-CC9 cache controller

Configuration	Physical capacity	Maximum usable capacity	Maximum quantity of FC5267 ^a
1 TS7740 Cache Controller (3956-CC9)	9.6 TB	9.45 TB (8.59 TiB)	10
1 TS7740 Cache Controller (3956-CC9) 1 TS7740 Cache Drawer (3956-CX9)	19.2 TB	19.03 TB (17.30 TiB)	19
1 TS7740 Cache Controller (3956-CC9) 2 TS7740 Cache Drawer (3956-CX9)	28.8 TB	28.60 TB (26.02 TiB)	28

a. Number of instances that are required to use maximum physical capacity

7.2.6 Upgrading drive models in an existing TS7740 or TS7700T

This section describes upgrading back-end tape drives in an existing TS7740 or TS7700T cluster with data. You might want to upgrade the back-end tape drives to a higher model to have more capacity from the existing media because the drives are not encryption capable, or for any other reason. TS7740 or TS7700T supports the 3592-J1A, TS1120 (3592-E05), TS1130 (3592-E06/EU6), TS1140 (3592-E07), and TS1150 (3592-E08) tape drives.

Consideration: Drive model changes can be made only in an upward direction (from an older to a newer model). Fallback to the older models is not supported.

Hardware configuration and limitations

For more information about tape drives and supported media, see Chapter 4.1.2, “TS7700 specific limitations” on page 134.

Note: Throughout this section, the term *TS7700* refers to either the TS7740, the TS7720 Tape Attach, or the TS7760 Tape Attach.

Restrictions for use with TS1140 Tape Drives

TS1140 Tape Drives are supported in new TS7700 orders from manufacturing, and with existing TS7700s attached to a library. The following additional media restrictions apply when a library attached to a TS7740 or TS7700T contains TS1140 Tape Drives:

- ▶ JA and JJ media are supported for read-only operations. If JA or JJ media exist or are installed in a library that contains TS1140 Tape Drives, the following actions occur:
 - Online processing succeeds, but all JA and JJ media is marked read-only for reclamation.

Note: One main purpose of reclamation is to increase the number of available physical scratch volumes in the pool. When TS1140 Tape Drives are installed, JJ and JA media reclamation reduces, instead of increases, the number of available scratch volumes. Reclamation of a JA or JJ cartridge does not occur if the TS7700 is in a low scratch state (fewer than 15 available scratch volumes) for the pool.

For example, if borrowing is enabled and there are JA physical volumes to be reclaimed in pool 1, the sum of available scratch tapes in pool 1 and the common scratch pool 0 must be greater than 15 for reclamation of the JA physical volumes to occur. If the system contains TS1140 or TS1150 tape drives, the system requires at least 15 scratch physical volumes to run reclamation for sunset media.

- JA and JJ media can be ejected by using the TS7700 Management Interface after their active data is reclaimed onto newer media.

Note: JA and JJ media should not be inserted if the volumes do not exist in the TS7700 database.

- ▶ If JB media contains data that is written in E05 format, it is marked full and is supported as READ-ONLY data. After the data is reclaimed or written in E06 or E07 format, it is supported for read/write operations. The IBM Encryption Key Manager is not supported for use with TS1140 Tape Drives. If encryption is used, either the IBM Security Key Lifecycle Manager (ISKLM) or Security Key Lifecycle Manager (SKLM) must be used.
- ▶ 3592 EU6 Tape Drives cannot be converted to TS1140 Tape Drives.

Restrictions for use with TS1150 Tape Drives (Homogeneous Configuration)

TS1150 Tape Drives are supported in new TS7700 orders from manufacturing, and with existing TS7700s attached to a library. The following additional media restrictions apply when a library attached to a TS7740 or TS7700T contains TS1150 Tape Drives:

- ▶ JA, JJ, and JB media are not supported.
- ▶ The IBM Encryption Key Manager is not supported for use with a 3592 E08 Tape Drive. If encryption is used, either the IBM Security Key Lifecycle Manager (ISKLM) or Security Key Lifecycle Manager (SKLM) must be used.
- ▶ TS1140 Tape Drives cannot be converted to TS1150 Tape Drives.

Restrictions for use with TS1150 Tape Drives (Heterogeneous Configuration)

TS1150 Tape Drives are supported in new TS7700 orders from manufacturing, and with existing TS7700s attached to a library. The following additional media restrictions apply when a library attached to a TS7740, TS7720 Tape Attach, or TS7760 Tape Attach contains TS1150 Tape Drives:

- ▶ TS1150 Tape Drives can be intermixed with one other TS11xx drive type in a library that is attached to a TS7740 or TS7700T for migration purposes.
- ▶ JA, JJ, and JB media are supported for read-only operations. If JA, JJ, or JB media exist or are installed in a library that contains TS1150 Tape Drives, the following actions occur:
 - Online processing succeeds, but all JA, JJ, and JB media is marked read-only for reclamation.

Note: One main purpose of reclamation is to increase the number of available physical scratch volumes in the pool. When TS1150 Tape Drives are installed, JJ, JA, and JB media reclamation reduces, instead of increases, the number of available scratch volumes. Reclamation of a JA, JJ, or JB cartridge does not occur if the TS7700 is in a low scratch state (fewer than 15 available scratch volumes) for the pool.

For example, if borrowing is enabled and there are JA physical volumes to be reclaimed in pool 1, the sum of available scratch tapes in pool 1 and the common scratch pool 0 must be greater than 15 for reclamation of the JA physical volumes to occur. If the system contains TS1140 or TS1150 tape drives, the system requires at least 15 scratch physical volumes to run reclamation for sunset media.

- JA, JJ, and JB media are read by the TS11xx drive.
- JA, JJ, and JB media can be ejected by using the TS7700 Management Interface after their active data is reclaimed onto newer media.

Note: JA, JJ, and JB media should not be inserted if the volumes do not exist in the TS7700 database.

- ▶ The IBM Encryption Key Manager is not supported for use with a TS1150 Tape Drive. If encryption is used, either the IBM Security Key Lifecycle Manager (ISKLM) or the Security Key Lifecycle Manager (SKLM) must be used.
- ▶ TS1140 Tape Drives cannot be converted to TS1150 Tape Drives.

Considerations for upgrading tape drives

This section describes considerations that you need to consider when you upgrade your back-end tape drives.

Upgrading to homogeneous TS1150 tape drive configuration

Here are the limitations for homogeneous tape drive configuration:

- ▶ A maximum of 16 tape drives is supported.
- ▶ A minimum of 4 tape drives is required.
- ▶ JJ, JA, and JB cannot be read or written by TS1150 drives. These media are called *sunset media*.
- ▶ TS7700 does not go online if there is a logical volume in sunset media and no read-compatible drives are available.

For existing TS7740 or TS7700T customers, it is possible to replace all existing tape drives with TS1150 tape drives if all active data on tape is on JK or JC cartridges, which implies that the tape drives that are currently installed are TS1140s. If the existing TS7740 or TS7700T has TS1140 drives installed but uses JB media, data must first be migrated to JK or JC cartridges.

Another possible use case is to use the existing TS7700T to complete the following steps:

1. First, change the storage class definitions to point to the resident-only partition.
2. Then, run a **PARTRFSH** command to move the logical volumes from the tape-attached partition to the resident-only partition.
3. Then, recall all of the migrated data from the legacy tapes into the cache resident partition and allow the MES to have TS1150 tape drives installed.
4. After the MES, you can again change the storage class, run a **PARTRFSH**, and push the logical volumes back to backend tapes.

After the first TS1150 is installed and configured, the TS7700 detects the new drive generation during the online process and acknowledges the cartridge types. Make sure that all of the logical volumes in sunset media, such as JJ, JA, or JB, are migrated to JC or JK media before TS1150 installation. If TS7700 detects logical volumes in sunset media, the online process fails. TS7700 comes online if no logical volumes exist in sunset media. They can be ejected from the TS7700 MI after TS7700 becomes online.

When a TS7700 has all TS1150 tape drives, the following 3592 media types can be used as scratch media for read/write:

- ▶ JK - Advanced Type K Economy (ATKE) (900 GB)
- ▶ JC - Advanced Type C Data (ATCD) (7000 GB)
- ▶ JL - Advanced Type L Economy (ATLE) (2000 GB)
- ▶ JD - Advanced Type D Data (ATDD) (10000 GB)

Empty sunset media, such as JJ, JA, or JB, are marked as sunset read-only after TS7700 comes online. This media cannot be used as scratch tapes.

For a storage pool that is not a copy export pool, the E08 recording format is used from the beginning when writing tapes. If the storage pool is a copy export pool, the recording format must be selected through the TS7700 MI.

Upgrading to limited heterogeneous drive configuration

Here are the limitations for heterogeneous tape drive configuration:

- ▶ A maximum of 16 tape drives is supported.
- ▶ Up to 14 TS1150 tape drives are allowed.
- ▶ A minimum of four TS1150 tape drives is required.
- ▶ A minimum of two previous generation 3592 tape drives to be used for data migration over time.
- ▶ Previous generation 3592 tape drives are defined as read-only drives and used for reading logical volumes from JA, JB, and JJ media.
- ▶ JA, JB, and JJ cannot be read or written by TS1150. They cannot be used as scratch media.
- ▶ At least 15 empty scratch media (JC, JK, JD, or JL) are necessary to support data migration from sunset media to newer media

Support for limited heterogeneous tape drives seamlessly helps customers move from older media and drives to JK, JC, JL, and JD media, and TS1150 tape drives. This option enables you to add TS1150 tape drives to the existing TS7700 so that all new workloads can be directed to the TS1150 tape drives, and to leave at least one of the legacy tape drive generations (3592-J1A, TS1120, TS1130, or TS1140) to handle legacy media.

Note: Only one generation of legacy tape drives can be together with the newly installed TS1150.

During the TS7700 online processing, media types JJ, JA, and JB are marked as sunset read-only. Those volumes are mounted only for recall or idle reclamation according to the reclaim percentage pool settings. Make sure at least 15 scratch JC, JK, JD, or JL media are inserted to run reclamation of sunset media. After the reclaim process moves the data to TS1150 supported media, the operator can eject sunset media by using the TS7700 MI.

In heterogeneous drive configuration, legacy tape drives are defined as read-only drives. They are used for reading logical volumes from sunset media, and are not used for writing new data. However, there are two exceptions when read-only drives are used for writing data:

- ▶ One is Secure Data Erase (SDE) for a sunset media. If SDE is enabled, previous generation 3592 tape drives write a repeating pattern to the legacy media to erase data.
- ▶ The other is a Copy Export operation. If the legacy media exist in a Copy Export pool, previous generation 3592 tape drives write a DB backup to the legacy media.

Drive change scenarios

The drive change is performed by your IBM SSR. Work with your IBM SSR when you plan for cluster downtime.

Clarification: This section provides high-level information about the subject. Do not use this information as a step-by-step guide, but work with your IBM SSR to prepare for update.

Complete the following steps:

1. Before you change the tape drives, stop all host activity to this cluster. If this TS7740 or TS7700T is part of a grid, vary on the logical drives in the other clusters to provide tape resources to the host. If this is a stand-alone cluster, plan for a 3 - 4 hour outage.
2. Vary off all logical drives in this cluster. The IBM SSR places this cluster in service mode (if part of a grid) and offline. All physical drives must be unloaded and emptied of cartridges. Next, from the TS3500/TS4500 Management Interface (MI), the drives must be unassigned from the TS7740 or TS7700T Logical Library.
3. Now, drives can be removed and the new drives can be installed in their places. The IBM SSR installs new drives, and checks the configuration and firmware code level by using the TS3500/TS4500 MI or another tool. If necessary, the IBM SSR updates the firmware level of the new drives. Also, 4 Gb fiber switches can be replaced with the new 8 Gb switches, if ordered. If a TS7760T is used, 16 Gb switches are necessary.
4. New drives are assigned back to the TS7740 or TS7700T Logical Library by the TS3500/TS4500 MI, and the control paths are correctly assigned. Drive fiber cables are reconnected, and connections to the switches are verified. If everything appears correct, the IBM SSR runs a drive reconfiguration at the TS7740 or TS7700T cluster.
5. After the reconfiguration, all new drives and paths must be available and healthy. If not, the IBM SSR acts upon the errors to correct them. This completes the drive upgrade change. Now, the TS7740 or TS7700T can be taken out-of-service mode and varied online by the IBM SSR.

Remember: In the previous scenario, all cartridges in the filling state are closed if the new drives do not support writing in the original tape format. Otherwise, the cartridges continue to be written in the same format to the end. Scratch tapes that are in use after the change are automatically reinitialized to the new tape format.

You can apply the same procedure when changing the tape emulation mode in the TS7740 or TS7700T from 3592-J1A emulation to TS1120-E05 native mode. All steps apply except the steps that relate to changing drives physically and changing drive emulation mode within the TS3500. Drive emulation is changed in the TS3500 web interface (see Figure 9-154 on page 519 for a reference) by using a specific command in the TS7740 or TS7700T by the IBM SSR. The media format is handled as described in the previous scenario.

Migrating TS7740 or TS7700T data from sunset media type after upgrading heterogeneous drive configuration

Consideration: Restrictions apply to some configurations. For the valid combinations of media type and drive models within the TS7740 or TS7720T, see , “Tape drives and media support (TS7740,TS7720T, and TS7760T)” on page 131.

This procedure can be helpful when upgrading your tape drives to the TS1150 3592-E08 tape drives, or when replacing the existing media cartridges with a newer type to increase the storage capacity of your library. The E08 drives cannot read JA, JB, or JJ tapes, so you must have JC, JK, JL, or JD media for the TS7740 or TS7700T to support new logical volumes that are written to TS1150 drives. In addition, you must have at least two drives of a sunset generation that are available to read logical volumes from JA, JB, or JJ tapes.

In this scenario, coming from a previous 3592 tape drive model to the E08, all JA, JB, or JJ media are *sunset*, which means that after reclaiming the active logical volumes still contained in it, the JA, JB, or JJ media can be ejected from the library. In this case, you must have a working pool of stacked volumes of the supported media, such as JC, JK, JD, or JL. Your data, formerly in a JA, JB, or JJ media, is forcibly migrated into the supported media.

There are two alternatives to introduce the new media. You can either use one physical volume pool or two physical volume pools. In the second scenario, complete the following steps:

1. Create a range of physical volumes in the TS7740 or TS7700T for the new media, as shown in Figure 9-79 on page 410.
2. Create a Cartridge Assignment Policy (CAP) for the new range and assign it to the correct TS7740 or TS7700T logical partition (LPAR), as described in “Defining Cartridge Assignment Policies” on page 520.
3. Insert the new cartridges in the TS3500/TS4500 tape library, as described in “Inserting TS7700T physical volumes” on page 521.
4. Assign an existing pool or pools of physical volumes in the TS7740 or TS7700T to the new media type, as described in “Defining physical volume pools in the TS7700T” on page 531.
5. Modify the Storage Group (SG) in the TS7740 or TS7700T constructs to point to the new cartridge pools, as described in “Defining TS7700 constructs” on page 551.
6. Modify the reclaim settings in the existing media pool by using the new pools as the target pools, as described in “Defining physical volume pools in the TS7700T” on page 531.

These settings cause the TS7740 or TS7700T to start using the new media for stacking newly created logical volumes. Existing physical volumes are reclaimed into the new JD media, becoming empty.

You might prefer to keep your pool definitions unchanged throughout the media change process. In this case, you can just run the previous steps 1-3. No further change is necessary if you are migrating from JJ/JA/JB cartridges to JD cartridges. If you are migrating from JC/JK to JD/JL cartridges, you should set the new media as “First media” in the Pool Properties table. This way, cartridges of the previous media type are not available for selection in the common scratch pool.

You can keep the previous media type as the secondary media type as a precaution to not run out of scratch media. For more information, see “Defining physical volume pools in the TS7700T” on page 531. After the old-type cartridges are emptied, they can be ejected from the tape library.

Clarification: You might use the new Host Console Request Resident on Recall for Sunsetting RRCLSUN (ReCaLI SUNset) to expedite the replacement of the sunset media with newer media. In this case, ensure that the common scratch pool has the new media type available, and that the storage pools are set to borrow from the common scratch pool. Otherwise, the storage pools run out of scratch.

This function invalidates the logical volume on the sunset media just after recall, regardless of whether the logical volume is updated. As a result, any recalled volume is premigrated to newer media. The library request command is shown:

```
LI REQ, lib_name,RRCLSUN ENABLE/DISABLE/STATUS
```

Where:

ENABLE Activates the force residency on recall function.

DISABLE Deactivates the force residency on recall function.

STATUS Displays the current setting.

If you are changing existing drives to new drive models that use the same media type, use the Library Request (**LI REQ**) command to accelerate the media type conversion. Use this process to reclaim capacity from your existing media. In this scenario, you are not changing the existing cartridges already in use. There are no changes that are needed regarding the existing physical volume pools.

To accelerate the media type conversion, modify the pool property to set a high value to Sunset Media Reclaim Threshold Percentage by using TS7700 MI. Whenever the active data percentage of sunset media goes below the threshold, the sunset media is reclaimed and active data is migrated to the newer media.

Figure 7-21 shows how the Sunset Media Reclaim Threshold Percentage is set in the Physical Volume Pool properties.




Figure 7-21 Sunset Media Reclaim Threshold Percentage

If you want to limit the service resource of data migration, there is a new Host Console Request command in Release 3.3 that is available to influence the performance of the replacement of the sunset media.

Clarification: Use Host Console Request, Reclaim Maximum Tasks Limit For Sunset Media (RCLMSMAX) so that the TS7700 has fewer reclaims. You change the format of the sunset media to the newest one, and use the service resource for other activities in the cluster.

This function provides a method to limit the maximum number of concurrent reclamation tasks that run against sunset media. The maximum is the number of installed sunset drives - 1. You can set the maximum by specifying "0" (this is the default value). Here is the library request command:

```
LI REQ, <lib_name>, SETTING, RECLAIM, RCLMSMAX, <number_of_drives>
```

Low Sunset Drive warning PDRVSLOW/PDRVSCRT

The new keywords **PDRVSLOW** and **PDRVSCRT** generate messages that indicate that the available sunset physical drives have fallen below the low warning limit and the critical warning limit.

The values that can be set for these new keywords are the same as the existing keywords **PDRVLOW** and **PDRVCRIT**, except that the existing keywords are used for non-sunset drives and the new keywords are for sunset drives.

Clarification: Here are the parameters for the low sunset drive alert:

▶ **SETTING2,ALERT,PDRVSLOW**

Gives an “orange” alert if only *x* drives from sunset drives are available. The parameter can be set to 0 (no alert) or 3 - *the number of installed sunset drives*.

▶ **SETTING2,ALERT,PDRVSCRT**

Gives a “red” alert if only *x* drives from sunset drives are available. It can be set to 0 (no alert) or 3 - *the number of installed sunset drives*.

Here is the library request command:

```
LI REQ, <lib_name>, SETTING2, ALERT, [PDRVSLOW|PDRVSCRT], <number_of_drives>
```

7.2.7 Frame replacement of old hardware with new hardware

This scenario describes frame replacement of the following cases:

- ▶ Frame replacement of TS7740 (V06 or V07) with TS7760T (VEC)
- ▶ Frame replacement of TS7720T (VEB) with TS7760T (VEC)

You might want to replace old hardware (V06, V07, or VEB) with the new hardware (VEC) to have significant performance improvement and increased disk cache capacity over the old ones. R4.0 PGA1 supports frame replacement of cluster at R2.1 and higher code level.

To replace the old hardware, all the logical volumes need to be pre-migrated to physical volumes, transferred to the new frame, and recalled from the physical volumes. There are several considerations before perform frame replacement:

- ▶ If the frame you want to replace is TS7720T (VEB), all the private logical volumes in CP0 need to be moved to CPx before frame replacement.
- ▶ If the frame you want to replace is TS7720T (VEB), all the resident logical volumes with delayed premigration setting need to be premigrated before frame replacement.

All the physical volumes that you want to migrate to the new frame must be able to be read by the new tape drives. See 4.1.2, “TS7700 specific limitations” on page 134.

7.3 TS7700 upgrade to Release 4.0

Release 4.0 can be installed on a previous V07 or VEB-based TS7700 system. The following prerequisites exist:

- ▶ Install FC 3462 (32 GB Memory Upgrade).
- ▶ The TS7700 must be at a minimum microcode level of R3.2 or later.

Existing V06 and VEA systems do *not* support Release 3.1 or higher levels of Licensed Internal Code.

7.3.1 Planning for the upgrade

A multi-cluster GRID configuration can enable practically all changes or upgrades to be concurrent from a client's standpoint, putting one individual cluster into service at a time. The Release 4.0 Licensed Internal Code upgrade is a disruptive activity in a stand-alone cluster. A Licensed Internal Code update is done by an IBM service support representative (IBM SSR). Preinstallation planning and a scheduled outage are necessary.

When you are updating code on a cluster in a grid configuration, plan an upgrade to minimize the time that a grid operates clusters at different code levels. Also, the time in service mode is important.

Before you start a code upgrade, all devices in this cluster must be varied offline. A cluster in a grid environment must be put into service mode and then varied offline for the code update. You might consider making more devices within other clusters in the grid available because you are losing devices for the code upgrade.

Consideration: Within the grid, some new functions or features are *not* usable until all clusters within the grid are updated to the same Licensed Internal Code (LIC) level and feature codes.

The MI in the cluster that is 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 before you perform the Licensed Internal Code upgrade.

Important: Ensure that you check the D/T3957 Preventive Service Planning (PSP) bucket for any recommended maintenance before performing the LIC upgrade.

PSP buckets are at the following address. Search for D/T3957:

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

7.4 Adding clusters to a grid

The TS7700 cluster can be installed in stand-alone or multi-cluster grid configurations. This section describes the available options and the required steps to add a cluster to a grid, merge a cluster into a grid, or merge a grid with another grid. Adding clusters to a grid is a concurrent process from the client's standpoint. None of the existing clusters must go into service when joining a new cluster to the grid.

7.4.1 TS7700 grid upgrade concept

A TS7700 grid refers to two, three, four, five, or six TS7700 clusters that can be physically separated and are interconnected by using an Internet Protocol network.

Migrations to a TS7700 multi-cluster grid configuration require the use of the Internet Protocol network. In a two-cluster grid, the grid link connections can be direct-connected (in a point-to-point mode) to clusters that are located within the supported distance for the adapters present in the configuration.

See “TS7700 grid interconnect LAN/WAN requirements” on page 136 for distances that are supported by different grid adapters and cabling options. For separated sites or three or more cluster grids, be sure that you have the network prepared at the time that the migration starts. The TS7700 provides two or four independent 1 Gbps copper (RJ-45) or SW fiber Ethernet links (single-ported or dual-ported) for grid network connectivity.

Alternatively, on a 3957-V07 /VEB server, two 10 Gbps LW fiber Ethernet links or 3957-VEC with four 10 Gbps LW fiber Ethernet links can be provided. Be sure to connect each one through an independent WAN interconnection to be protected from a single point of failure that disrupts service to both WAN paths from a node. See 4.1.3, “TCP/IP configuration considerations” on page 136 for more information.

Grid upgrade terminology

The following terminology is used throughout the Grid configuration sections:

► Join

Join is the process that is performed when an empty cluster is joined to another cluster or clusters to create a grid or a larger grid. The empty cluster is referred to as the *joining cluster*. The cluster or clusters to which it is joined must have a chosen cluster to act as the existing cluster. The existing cluster can be a new empty cluster, an existing stand-alone cluster, or a cluster that is a member of an existing grid. There are many combinations of code levels and configurations that are supported when joining an empty cluster.

► Merge

Merge is the process that is performed in the following situations:

- Merging a cluster with data to another stand-alone cluster with data (to create a grid)
- Merging a cluster with data to an existing grid
- Merging a grid with data to another existing grid

The *merging cluster* can be a stand-alone cluster or it can be a cluster in an existing grid. Similarly, the existing cluster can be a stand-alone cluster or it can be a cluster in an existing grid.

Note: An RPQ is required before you can implement a five-cluster or six-cluster configuration. If you need a configuration with more than four clusters, contact your IBM sales representative to submit the RPQ.

7.4.2 Considerations when adding a cluster to the existing configuration

Figure 7-22 shows an example of joining or merging a new cluster to an existing stand-alone configuration.

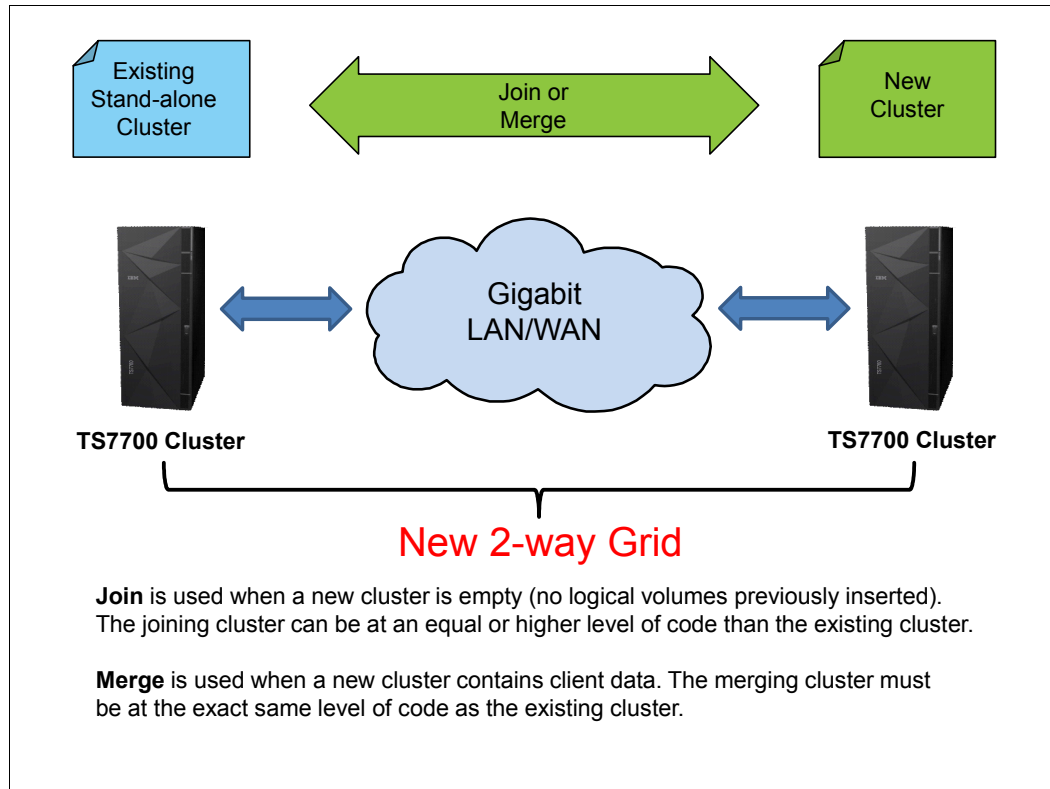


Figure 7-22 Example of a join or merge of a new cluster

Figure 7-23 shows an example of merging or joining a new cluster to an existing grid configuration. The example shows a join or merge of a new cluster to an existing 5-cluster grid.

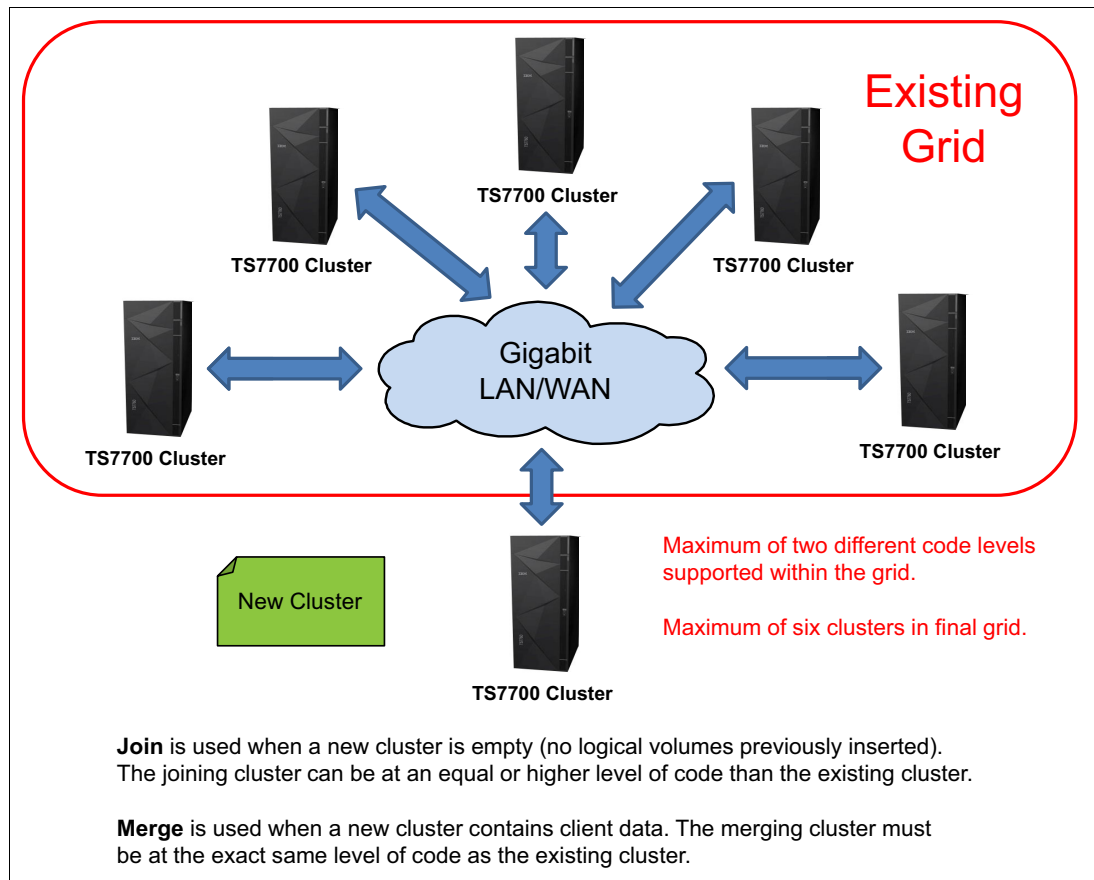


Figure 7-23 Join or merge a new cluster to a multi-cluster grid

Preparation

When performing a join, the actual data does not get copied from one cluster to another. This process instead creates only placeholders for all of the logical volume data in the final grid. When joining to an existing grid, the process is initiated to a single cluster in the grid and the information is populated to all members of the grid.

TS7700 constructs, such as Management Class (MC), Data Class (DC), Storage Class (SC), and Storage Group (SG), are copied over from the existing cluster or grid to the joining cluster.

Host configuration changes

It is important to consider the host configuration changes that are needed before you attempt to use the newly joined cluster. For more information, see 4.3.1, "Host configuration definition" on page 161 and 6.4, "Hardware configuration definition" on page 216.

- ▶ All HCDs, subsystem IDs, and Port IDs must be updated, and the cabling must be done correctly.
- ▶ Define the new distributed library ID to the storage management subsystem (SMS). Check with the IBM SSR for the appropriate library sequence number (LIBRARY-ID). Management and data policy planning.

Plan to define the following management and data policies after the TS7740 Cluster join is complete:

- ▶ Define stacked volume ranges
- ▶ Define reclaim threshold percentage

Logical volume considerations

Ensure that the joining cluster has at least the same number of FC5270 components installed as in the existing cluster or grid.

Licensed Internal Code supported levels and feature code for join

Release 4.0 supports the ability to have V07/VEB/VEC clusters (new from manufacturing or empty through a manufacturing clean-up process) join an existing grid with a restricted mixture of Release 2.1, Release 3.x clusters, and Release 4.0. There can be three total code level differences across both targets and the joining system during the MES where R2.1 can be the lowest of the three levels.

The joining cluster must be at an equal or later code level than the existing clusters. One or more Release 4.0, Release 3.x, or Release 2.1 clusters can exist in the grid if the total of all levels, including the joining cluster, does not exceed three unique levels.

When you join one cluster to a cluster in an existing grid, all clusters in the existing grid are automatically joined. Before you add an empty cluster to an existing cluster or grid, ensure that you have addressed the following restrictions for the join process:

- ▶ The joining cluster must be empty (contain no data, no logical volumes, and no constructs).
- ▶ If the existing cluster to be joined to is a member of a grid, it must be the current code level of any member in the grid.
- ▶ The joining cluster must be at an equal or later code level than the existing clusters.
- ▶ The joining cluster and existing cluster must have FC 4015 installed.
- ▶ The joining cluster must support at least the number of logical volumes that are supported by the grid by using FC 5270.
- ▶ The joining cluster must contain FC 5271 if the existing cluster to be joined has this feature code installed.
- ▶ If the joining cluster has FC 1035 installed, the client's infrastructure must support 10 Gb.

Join steps

Complete the following steps to join the cluster:

1. Arrange for these join cluster tasks to be performed by the IBM SSR:
 - a. Verify the feature code.
 - b. Establish the cluster index number on the joining cluster.
 - c. Configure the grid IP address on both clusters and test.
 - d. Configure and test Autonomic Ownership Takeover Manager (AOTM) when needed. See Chapter 2, "Architecture, components, and functional characteristics" on page 15 for more information.
2. Change HCD channel definitions.

Define the new channels and the device units' addresses in HCD.

3. Change SMS and tape configuration database (TCDB).

With the new grid, you need one composite library and up to six distributed libraries. All distributed libraries and cluster IDs must be unique. You must now define the new added distributed library in SMS. Ensure that you enter the correct Library-ID that was delivered by the IBM SSR.

4. Activate the input/output definition file (IODF) and the SMS definitions and issue an object access method (OAM) restart (if it was not done after the SMS activation).

Consideration: If the new source control data set (SCDS) is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be generated:

```
CBR3006I Library <library-name> with Library ID <library-ID> unknown in I/O configuration.
```

5. Vary devices online to all connected hosts. After a new cluster is joined to a cluster in an existing grid, all clusters in the existing grid are automatically joined. Now, you are ready to validate the grid.

6. Modify Copy Policies and Retain Copy mode in the MC definitions according to your needs. Check all constructs on the MI of both clusters and ensure that they are set properly for the grid configuration. See 2.3.25, “Copy Consistency Point: Copy policy modes in a multi-cluster grid” on page 80 for more information.

7. Review your family definitions, and decide if the cluster needs to be included in one of the families. In specific situations, you might want to introduce a new family, for example if a new site is populated.

8. Review your SDAC definition, and include the new LIBPORT ID statements if necessary.

9. Review the cluster settings with the **LI REQ SETTING** for the new distributed library. Enable the alerts according to the TS7700 model, and review COPYFSC and RECLPG0 settings (especially if the new cluster is used for DR purposes).

10. For a disk-only model, check the **REMOVE** and **RMVTHR**.

11. For a tape-attached model, check the same settings, but also review the **PMTLVL** setting and set it to the amount of installed FC 5274, for example, 6 x FC5274 = 6000. For more information about these settings, see Chapter 11, “Performance and monitoring” on page 635.

12. If the joined cluster is a tape attach, also define the tape partitions or resize CP1 and revisit the storage classes. Also review the Storage groups and the inhibit reclaim schedule. If you use multiple physical pools, you might also want to influence the number of maximum used premigration drives or the reclaim value. For more information, see Chapter 11, “Performance and monitoring” on page 635.

13. Run test jobs to read and write to volumes from all of the clusters.

14. Test the write and read capabilities with all of the clusters, and validate the copy policies to match the previously defined Copy Consistency Points and other constructs.

15. Consider creating a new MC for BVIR purposes to be able to run specific cluster reports.

Population of a new cluster (COPYRFSH)

If you want part or all of the existing logical volumes to be replicated to the new cluster, this can be done in different ways. In general, the new rules of the management policies need to be retrieved for every logical volume that will be replicated to the new cluster. Ensure that the management classes do not have **Retain mode copy** selected before you start the population process.

The following methods can be considered:

- ▶ Mount or demount to an LVOL
- ▶ **COPYRFSH**, a **LI REQ** command, based on a single logical volume

To produce a new copy, the data needs to be in cache. If your source cluster is a TS7740 or a TS7700T you should consider sorting the logical volumes in a copy order that maps to the physical volume layout. This will improve the performance of the copy action. The **COPYRFSH** processing enables you to specify a source cluster.

Also, prestaging the data to the cache helps to improve the performance. To simplify these actions, IBM provides some support in the “TAPE TOOL” suite. For more information about the performance, see Chapter 11, “Performance and monitoring” on page 635.

The tools are available at the following URL:

<ftp://ftp.software.ibm.com/storage/tapetool/>

7.4.3 Considerations for merging an existing cluster or grid into a grid

Figure 7-24 shows a grid merge scenario that involves a two-cluster grid and a three-cluster grid being merged into a five-cluster grid.

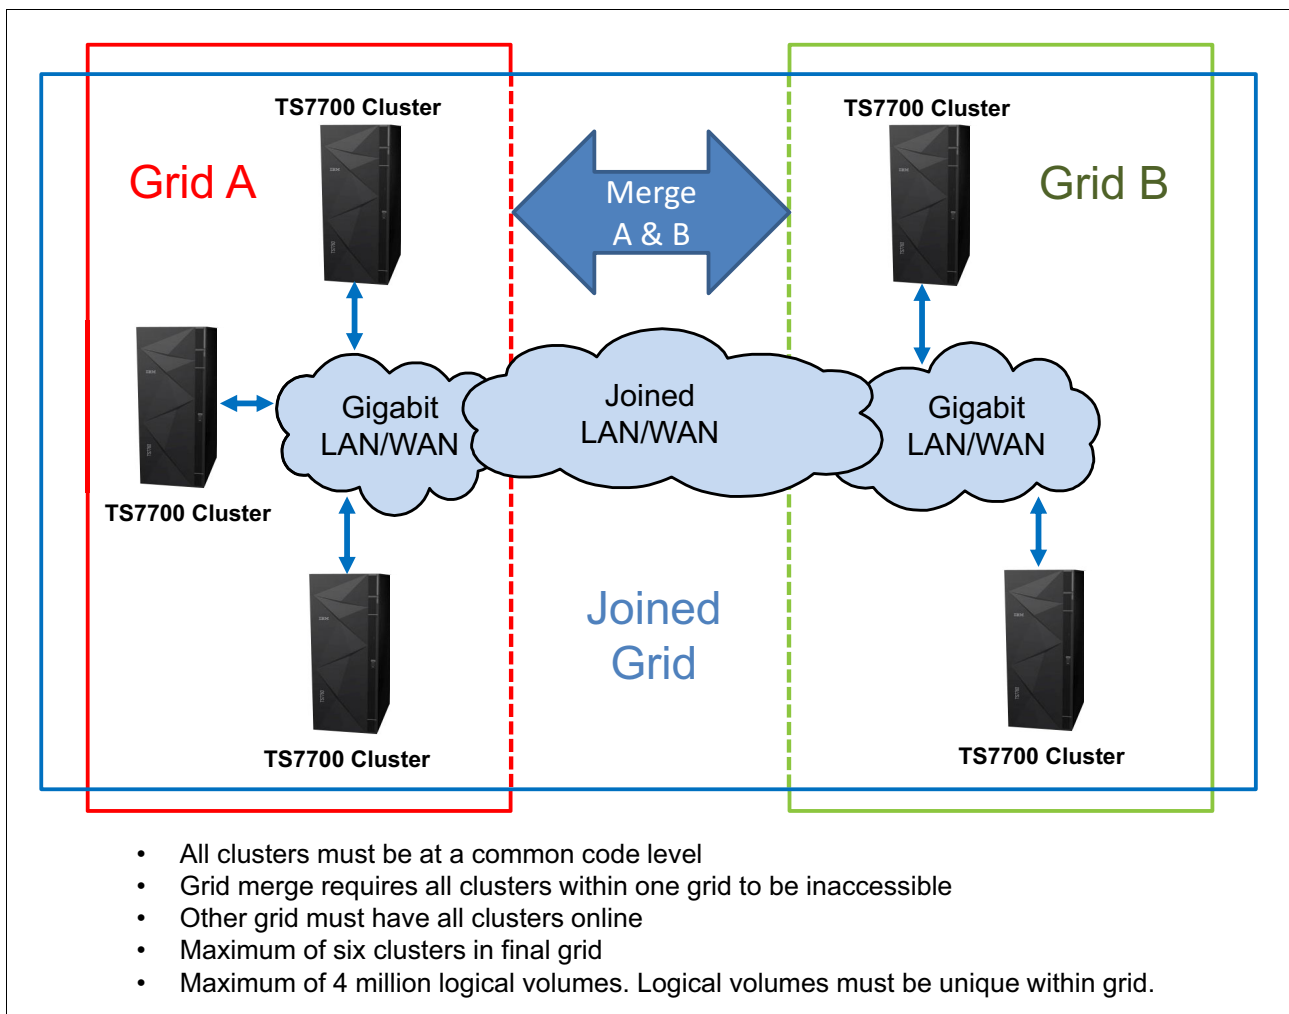


Figure 7-24 Grid merge example

Preparation

You can add an existing TS7700 Cluster to another existing TS7700 Cluster to form a grid for the first time or to create a larger grid. You can also merge a stand-alone cluster to an existing grid, or merge two grids together.

You can merge two existing TS7700 grids to create a larger grid. This solution enables you to keep redundant copies of data in both grids during the entire merge process versus needing to remove one or more clusters first and exposing them to a single copy loss condition.

When performing a merge, it is important to note that the actual data does not get copied from one cluster to another. This process creates place holders for all of the logical volumes in the final grid. When merging grids, the process is initiated to a single cluster in the grid and the information is populated to all members of the grid.

Schedule this process during a low activity time on the existing cluster or grid. The grid or cluster that is chosen to be inaccessible during the merge process has its indexes changed to not conflict with the other grid or cluster. Check with your IBM SSR for planning information.

Ensure that no overlapping logical volume ranges or physical volume ranges exist. The merge process detects that situation. You need to check for duplicate logical volumes and, on TS7740, TS7720T, or TS7760T clusters, for duplicate physical volumes. Logical volume ranges in a TS7700 must be unique. If duplicate volumes are identified during the merge process, the process stops before the actual merge process begins.

Host configuration changes

If you merge clusters or grids together, you must plan which LPAR will have access to which clusters and which device ranges in the grid in advance. These changes need to be prepared in each LPAR (HCD, SMS, and TCDB):

- ▶ Define the new distributed Library ID to SMS. Check with the IBM SSR for the appropriate ID number.
- ▶ The Tape Management System (TMS) and volume category (volcat) definitions must be updated within their respective SGs. These updates are necessary to maintain continued access to the original volumes that were created when the systems were configured as stand-alone clusters.
- ▶ Review your DEVSUPxx members in all connected LPARs to ensure that no duplicate scratch or private categories are defined.
- ▶ All HCDs, subsystem IDs, and Port IDs must be updated, and the cabling must be correct.

There are conditions during a **MERGE** of two existing grids where the newly merged clusters are not correctly recognized in the new grid and do not come online to the host. This is typically only seen if the newly merged grid reuses one of the existing composite library names and a dynamic activate is issued for the changes to the IODF rather than an IPL. The following messages can be issued during **VARY ONLINE** for the library:

```
CBR3715I REQUEST FOR LIBRARY libname FAILED. NO PATHS AVAILABLE FOR I/O.  
CBR3002E LIBRARY libname NO LONGER USEABLE.
```

If the **DS QLIB, libname** console command is issued against merged distributed library names that are now part of the new grid, they might show up erroneously as part of the old grid and continue to be associated with a composite library name that is no longer in use.

Note: The following settings are defined in the ACTIVE configuration:

```
LIBID PORTID DEVICES  
COMPOSITE LIBID <libname>
```


If the libname is in the old composite library name that is no longer being used to identify the grid, then the situation can be resolved by issuing the console command **DS QLIB, libname, DELETE**. The libname that is used in this command is the old composite library name that was previously displayed in response to the **DS QL, libname** command.

This command flushes the old composite name out of the device services control blocks and then the newly merged distributed libraries should come online to the host. Another **DS QL, libname** command can be issued to verify that the correct composite is now being displayed. An alternative solution is to perform an IPL of all the host systems that are reporting this condition.

Management and data policy planning

Check the following information:

- ▶ Constructs in a cluster, which are already in the existing cluster, will be updated with the content of the existing cluster.
- ▶ Constructs in a cluster, which exist in the existing cluster but not in the merging cluster, will be copied.
- ▶ Constructs in a cluster, which exist in the merging cluster but not in the existing cluster, will be kept, but not copied to the existing cluster or grid.

Figure 7-25 shows the MC definition of the merging cluster and the two-cluster grid before the merge.




Figure 7-25 Management Class definition before merge

Figure 7-26 shows the MC definition of the merging cluster and the two-cluster grid after the merge.

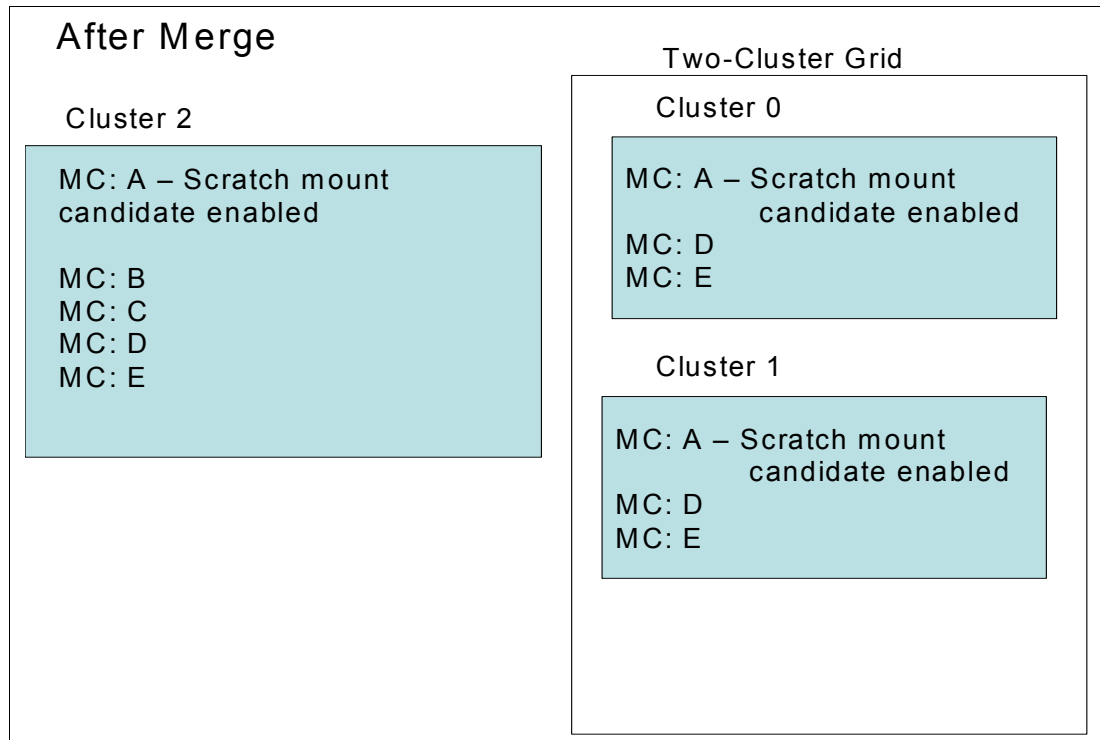


Figure 7-26 MC definition after merge

If categories and constructs are already defined on the merging cluster, verify that the total number of each category and construct that will exist in the grid does not exceed 256. If necessary, delete existing categories or constructs from the joining or merging clusters before the grid upgrade occurs. Each TS7700 grid supports a maximum of 256 of each of the following categories and constructs:

- ▶ Scratch Categories
- ▶ Management Classes
- ▶ Data Classes
- ▶ Storage Classes
- ▶ Storage Groups

Logical volume considerations

The TS7700 default number of supported logical volumes is 1,000,000. With Release 3.0, you can add support for more logical volumes in 200,000 volume increments by using FC 5270, up to a total of 4,000,000 logical volumes (2,000,000 million maximum on V06/VEA). The number of logical volumes that are supported in a grid is set by the cluster with the smallest number of FC 5270 increments installed.

If the current combined number of logical volumes in the clusters to be joined exceeds the maximum number of supported logical volumes, some logical volumes must be moved to another library or deleted to reach the allowed grid capacity. To maximize the full number of logical volumes that is supported on the grid, all clusters must have the same quantity of FC 5270 components that are installed. If feature counts do not match and the final merged volume count exceeds a particular cluster's feature count, further inserts are not allowed until the feature counts on those clusters are increased.

Minimum Licensed Internal Code level and feature code for merge

Before the merge of a cluster or a grid into another grid, the following restrictions apply to the merge process:

- ▶ When you merge from one cluster or a grid to another grid, all clusters in the existing grids are automatically merged. The merging cluster must be offline, and the cluster to be merged is only online if it has 8.21.x.x code or higher.
- ▶ A grid-to-grid merge is only supported at 8.21.x.x code or higher, and both grids must operate at the same Licensed Internal Code level.
- ▶ Merges are only supported when all clusters in the resulting grid are at the exact same code level.
- ▶ FC 4015, Grid enablement, must be installed on all TS7700 clusters that operate in a grid configuration.
- ▶ Both existing clusters and merging clusters must contain enough features to accommodate the total resulting volume count post merge, or the merge will fail.
- ▶ The merging cluster must contain FC 5271 if the cluster to be merged has it installed.
- ▶ If the merging cluster has FC 1035 installed, the client's infrastructure must support 10 Gb.

Merge steps

Complete these steps to merge all of the clusters or grids into a grid:

1. Arrange for these merge cluster tasks to be performed by the IBM SSR:
 - a. Verify the feature code.
 - b. Configure the grid IP address on all clusters and test.
 - c. Configure and test AOTM, when needed. For more information, see Chapter 2, "Architecture, components, and functional characteristics" on page 15.
2. Change HCD channel definitions.

Define the new channels and the device units' addresses in HCD. For more information about HCD, see 4.3.1, "Host configuration definition" on page 161 and 6.4, "Hardware configuration definition" on page 216.
3. Change SMS and TCDB.

With the new grid, you need one composite library and up to six distributed libraries. All distributed libraries and cluster IDs must be unique. You must now define the new added distributed library in SMS. Make sure to enter the correct Library-ID that was delivered by the IBM SSR.
4. Activate the IODF and the SMS definitions and issue an OAM restart (if it was not done after the SMS activation). If you are merging a cluster that was previously part of an existing grid, you might need to delete the services control blocks of that cluster's devices using the **DS QL,nnnn,DELETE** command, where *nnnn* is the LIBID of the cluster.
5. Vary devices online to all connected hosts. After a cluster is merged to a cluster in an existing grid, all clusters in the existing grid are automatically merged. Now, you are ready to validate the grid.
6. Run test jobs to read and write to volumes from all of the clusters. Remember, you must verify all LPARs in the sysplex.
7. Modify copy policies and Retain Copy mode in the MC definitions according to your needs. Check all constructs on the MI of both clusters and ensure that they are set correctly for the new configuration. For more information, see 2.3.25, "Copy Consistency Point: Copy policy modes in a multi-cluster grid" on page 80.

8. Test the write and read capabilities with all of the clusters and validate the copy policies to match the previously defined Copy Consistency Points.
9. If you want part or all of the existing logical volumes to be replicated to the new cluster, the same methods can be used as after a join processing. See “Population of a new cluster (COPYRFSH)” on page 274.

7.5 Removing clusters from a grid

FC 4016, Remove Cluster from Grid, delivers instructions for a one-time process to remove/unjoin a cluster (either TS7720, TS7740, or TS7760) from a grid configuration. It can be used for removing one cluster from a two-cluster to six-cluster grid. Subsequent invocations can be run to remove multiple clusters from the grid configuration.

After the removal, FC 4017 Cluster Cleanup can be run. FC 4017 is required if the removed cluster is going to be reused. A Cluster Cleanup removes the previous data from cache and returns the cluster to a usable state, similar to a new TS7700 from manufacturing, keeping the existing feature codes in place. Both feature codes are one-time use features.

You can delay the cluster cleanup for a short period while the TS7700 grid continues operation to ensure that all volumes are present after the removal of the TS7700 cluster.

The client is responsible for determining how to handle the volumes that have only a Copy Consistency Point at the cluster that is being removed (eject them, move them to the scratch category, or activate an MC change on a mount/demount to get a copy on another cluster). This process needs to be done before you start the removal process. A new Bulk Volume Information Retrieval (BVIR) option Copy Audit or COPYRFSH is provided for generating a list of inconsistent volumes to help you.

The removal of the cluster from the grid is concurrent with client operations on the remaining clusters, but some operations are restricted during the removal process. During this time, inserts, ejects, and exports are inhibited. Generally, run the removal of a cluster from the grid during off-peak hours.

No data, on cache or tapes, on the removed cluster is available after the cluster is removed with the completion of FC 4016. The cluster cannot normally be rejoined with the existing data. However, there is a special service offering to rejoin a cluster with existing data, if this particular operation is wanted. Contact your IBM sales representative for details.

No secure erase or low-level format is done on the tapes or the cache as part of FC 4016 or FC 4017. If the client requires data secure erase of the TVC contents, it is a contracted service for a fee. Consider delaying the cluster cleanup for a short time while the TS7700 grid continues operation to ensure that all volumes are present after the removal of the TS7700 cluster.

7.5.1 Reasons to remove a cluster

This section describes several reasons for removing a cluster.

Data center consolidation

A client is consolidating data centers by collecting the data from remote data centers and using the TS7700 grid to move the data to their centralized data center. In this scenario, the client potentially has two clusters at the primary data center for high availability.

The third cluster is at a remote data center. To consolidate the data center, it is necessary to copy the data from the third cluster to the existing grid in the primary data center. The third cluster is joined with the two existing clusters and the data is copied with grid replication.

After all of their data is copied to the primary data center TS7700 tape drives, the client can remove the third cluster from the remote data center and clean up the data from it. This TS7700 can now be relocated and the process can be repeated.

TS7700 reuse

A client has a multi-site grid configuration, and the client no longer requires a TS7700 at one site. The client can remove this cluster (after all required data is copied, removed, or expired) and use this resource in another role. Before the cluster can be used, it must be removed from the grid domain and cleaned up by using FC4017.

7.5.2 High-level description of the process

The following high-level preparation activities occur before the removal of a cluster from an existing domain:

- ▶ You must determine whether there are any volumes that are only available on the cluster to be removed (for example, MCs defined to have only a copy on one cluster, or auto removal from TS7720 or TS7760). Before the removal, you must create consistent copies on other clusters in the domain. See the BVIR Copy Audit function that is described in the *IBM TS7700 Series Bulk Volume Information Retrieval Function User's Guide* at the following website:
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101094>
- ▶ If volumes that only have a valid copy on the cluster are to be removed, you must determine how to handle those volumes by performing one or more of the following tasks:
 - Eject the logical volumes (see 10.1.7, “Ejecting logical volumes” on page 626).
 - Move the volumes to a scratch category.
 - Activate an MC change on the volume with a mount or unmount to get a copy made on another cluster.
- ▶ Ensure that there are no volumes in the damaged category. You can use the Repair Logical Volumes menu under the MI window to repair them.
- ▶ Remove the cluster from the family definition.
- ▶ Modify MCs so that the removed cluster is no longer the target for copies.
- ▶ If you have Licensed Internal Code level 8.6.0.x or higher installed and the cluster that is being removed is part of a cluster family, the cluster must be removed from the family before the removal by using the TS7700 MI.
- ▶ After the removal, ensure that there is no management class on any other cluster where NO COPY is the only remaining copy consistency policy.
- ▶ Adjust VEHSTATS runs to reflect the new configuration.

Note: If at least one remaining cluster in the grid is at code level 8.7.0.134 or later, you can use one of those clusters to perform the removal of another cluster by using FC4016 even if the grid contains two different code versions. Consult with your IBM SSR for the code version prerequisites for FC4016.

A copy consistency check is run at the beginning of the process. Do not skip consistency checks unless it is a disaster recovery (DR) unjoin or you can account for why a volume is inconsistent. Failure to do this can result in data loss when the only valid copy was present on the removed cluster.

After a cluster is removed, you might want to modify the host configuration to remove the LIBPORT IDs associated with the removed cluster.



Migration

This chapter explains aspects of migrating to a TS7700 environment from an IBM Virtual Tape Server (VTS) or from other tape drive technologies. It presents various options that can be tailored to your current environment.

Guidance is provided to help you achieve the migration scenario that best fits your needs. For this reason, methods, tools, and software products that can help make the migration easier are highlighted.

Updates to new models that are introduced with TS7700 Release 4.0 are also described in this section.

This chapter includes the following sections:

- ▶ Migration to a TS7700
- ▶ Moving data in and out of the TS7700
- ▶ Migration of DFSMSHsm-managed data
- ▶ DFSMSrmm and other tape management systems
- ▶ IBM Spectrum Protect
- ▶ DFSMSdss
- ▶ Object access method
- ▶ Database backups

8.1 Migration to a TS7700

This section covers various aspects of migrating from an existing tape technology to the TS7700. Depending on the source configuration, which can be an IBM VTS, IBM, or original equipment manufacturer (OEM) native tape drives, or some other possibility, one of the descriptions applies partially or completely to your migration scenario.

Migrations to a TS7700D by client can be done only by using the host. The TS7700D does not have any attached back-end tape drives. Therefore, data must be copied into the TS7700D by using host programs.

Migration of VTS Model B10 or B20 hardware to a TS7740 or TS7700T, 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 or TS7700T if the VTS system contains only 3592-formatted data. The outboard migration is offered as IBM Data Migration Services for Tape Systems. Outboard migration provides the following functions:

- ▶ Planning for migration, including consideration related to hardware and z/OS
- ▶ Project management for this portion of the project
- ▶ Assistance with the integration into a complete change plan, if required
- ▶ The actual migration of data from a 3494 VTS B10 or B20 to the TS7740 or TS7700T

Work with your IBM service support representative (IBM SSR) for more details about IBM Migration Services for Tape Systems. These services are available from an IBM migration team and can assist you in the preparation phase of the migration. The migration team performs the migration on the hardware.

When migrating data from a VTS to a new TS7740 or TS7700T installed in the same tape library, the process is called *data migrate without tape move*. If a source VTS is attached to one tape library and the new target TS7740 or TS7700T is attached to another tape library, the process is called *data migration with tape move*. When a source VTS is migrated to an existing TS7740 or TS7700T, or two VTSs are migrated to the same target TS7740 or TS7720T, the process is called *merge*.

If data will be moved inside the same grid (after a join of a cluster, or a merge), COPYRFSH is the preferred method. For more information, see “Population of a new cluster (COPYRFSH)” on page 274.

Migration from VTS with 3590 Tape Drives, or native tape drives to the TS7740 or TS7700T, always requires host involvement to copy the data into the TS7700. For more information about the methods you can use, see 8.2, “Moving data in and out of the TS7700” on page 288.

The hardware migration scenarios have the following aspects:

- ▶ Software changes in storage management subsystem (SMS), hardware configuration definition (HCD), tape configuration database (TCDB), and tape management system (TMS)

Tip: The TMS can be Removable Media Management (DFSMSrmm) or other products from other vendors.

- ▶ Connectivity for the new configuration
- ▶ Migration of the Library Manager database from an existing tape library to the TS7700

- ▶ Physical swap of the B10/B20 VTS to the TS7700 hardware
- ▶ Migration of the database from B10/B20 VTS to TS7700

Information is provided about the TS7700 family replacement procedures that are available with the new hardware platform and the TS7700 R4.0 Licensed Internal Code (LIC) level. With the availability of the new generation hardware, an upgrade path is provided for existing TS7700 users to migrate to this new hardware.

Upgrading tape drive models in an existing TS7740 or TS7700T to get more capacity from your existing media, or to provide encryption support, is further addressed in this section. It details the hardware upgrade procedure and the cartridge migration aspects.

8.1.1 Grid to Grid Migration

The GGM tool is a service offering from IBM. You can use it to copy logical volumes from one grid to another grid while both grids have a separated grid network. After the GGM is set up by an IBM SSR, the data from the logical volumes is transferred from one grid to the other grid through the existing IP addresses for the gridlinks (see Figure 8-1). Much like Join and Copy Refresh processing, there is no host I/O with the FICON adapters.

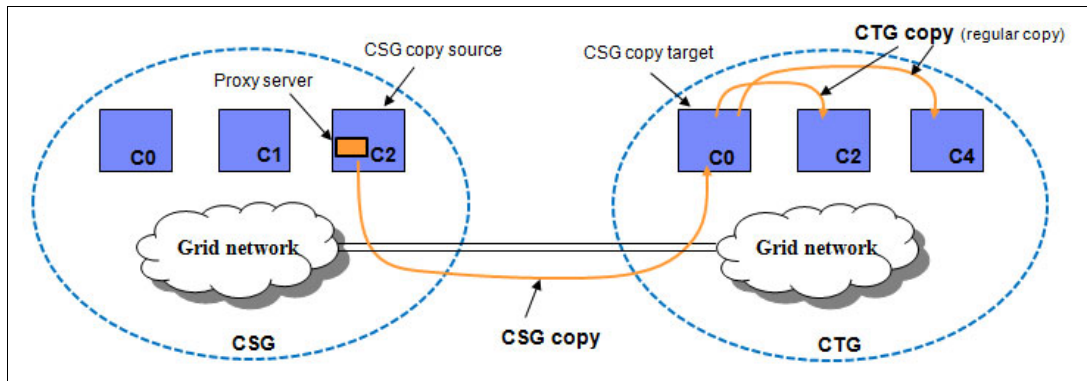


Figure 8-1 TS7700: Grid to Grid Migration overview

The following list describes the main components in GGM overview in Figure 8-1:

- CSG** Copy Source Grid.
- CTG** Copy Target Grid.
- Proxy server** A binary program that is installed in the CSG through `vtd_exec`, which enables the CTG to communicate with the CSG.

The GGM tool should be considered if the following situations are true:

- ▶ There are already six clusters installed in the grid.
- ▶ The Join and Copy Refresh processing cannot be used (there are floor space requirements, microcode restrictions, or other considerations).
- ▶ Source and Target grid belongs are maintained by different providers.

General information for GGM

GGM includes the following information:

- ▶ GGM is based on **LI REQ** commands. All commands are run on the *Cluster Target Grid* (CTG). The CTG must have at least R3.3 installed.
- ▶ It is single logical volume-based. Commands need to be run for each logical volume that needs to be copied.
- ▶ A volume to be copied by GGM must reside in the *Cluster Source Grid* (CSG). If no copy resides in this cluster, the logical volume cannot be copied, even if other clusters in the source grid have a copy. The minimum microcode level for a CSG is R2.1.
- ▶ The logical volume serial number is kept during the copy process, so all volume ranges that are copied must be unique in the grids.
- ▶ The CTG pulls a copy of the logical volume from the CSG. When this copy is pulled, depending on the MC content in the CTG, further copies might be produced in the CTG.
- ▶ After a logical volume is successfully copied to all clusters in the CTG, a broadcast message can be issued to all attached LPARs. Depending on the logical scenario, the customer needs to run actions manually then.
- ▶ After a logical volume is accessed for update (either by an append or a housekeeping process), the logical volume cannot be copied again.
- ▶ GGM has no automatic interaction to the TMS. Depending on the scenario, the workload profile, the amount of data, and the used TS7700 models, IBM recommends using different approaches to select the sequence of the copies.
- ▶ GGM is based on a *deferred copies* mechanism. Therefore, for the CTG, you need to review your CPYCNT parameter, because the Cluster Source Grid (CSG) might throttle the GGM copies using the values defined for DCOPYT and DCTAVGTD. The throttling might have an effect on the performance of the GGM copies.
- ▶ The data placement in the CSG has an effect on the performance of GGM. Recalls from physical tape can slow down the process of GGM. Whenever possible, we suggest using a TS7700D as the CSG.
- ▶ The GGM tool also provides several different options, such as how the new data (new device categories) and the old data (keep or delete the data in the source grid) is treated.

In general, we distinguish between two major logical scenarios for GGM.

Data movement with GGM

The actual goal is to move the data from one grid to another, while the same LPARs with the origin TCDB and TMS are used. After a single volume is copied, the information in the overlying TCDB and TMC needs to be changed after a successful migration of the single logical volume. These changes are the responsibility of the customer, and must be processed manually. In this scenario, you need to consider that a multivolume file can be accessed only if all logical volumes belonging to the multivolume file are located in the same grid.

We suggest switching the production workload to the new grid before you start the GGM process. This has several benefits:

- ▶ The CSG is not subject to Deferred copy throttling, because no host I/O is processed.
- ▶ If the CSG is a TS7740 or a TS7700T CPx, prestaging of the data speeds up the GGM copy processing, and the cache (especially in a TS7740) is no longer needed for host I/O.
- ▶ Data with a short lifecycle can expire in the source grid without being copied at all to reduce the amount of data that needs to be copied.

Data duplication with GGM

This scenario is mainly applicable if dedicated LPARs or workload of an LPAR needs to be separated from an existing environment. This is usually necessary if a customer decides to change to a different service provider, or if a company sells a subsidiary. In this case, the data is copied over to a new grid; however, production is still running to the existing grid until a so-called *cutover time*.

In this case, normally the TCDB and TMS are not changed during the copy process, but need to be adjusted at cutover time, or before if cutover tests are made.

Also, it is necessary to copy the same logical volume multiple times, because the lifecycle processing of “create - expire - delete - create” is running in the origin system.

In this case, you should consider using the lifecycle information (expiration, retention) as input to avoid that data with a very short lifetime is copied.

To ensure that during tests of the new environment the original copied data to the new grid is only read but not modified, you should put all clusters in the target grid in write protect mode for the origin categories and use different logical volume ranges and categories for the testing. While the cluster is in write protect, no GGM copies can be performed.

Cluster Source Grid considerations

Whenever possible, you should use a TS7700 disk-only model or the CP0 in a TS7700T as the CSG. This simplifies the migration project, because you all data is in cache, and you do not need to take care about your back end resources.

Using a CSG where all data resides in cache enables you to concentrate to copy the data by lifecycle information, or if needed by application purposes, especially for multivolume data sets.

If only a TS7740 or a TS7700T can be chosen, you also need to consider the backend resources and the available cache for the GGM copy processing. In this case, we strongly advise you to do a prestaging of the data, and copy the data based on physical volume, to avoid too many back end movements. Consider that this approach might not match with the TMS selection for retention or application purposes.

If the data will be recalled by the GGM process itself, the recall process might take much longer, and affect your overall GGM performance. Prestaging the data helps you to improve this performance. Consider changing your RECLPG0 value to allow recalled data to reside with the original storage class settings in the cluster. Otherwise, recalled data might already be already migrated before GGM could start the copy processing. For more information, see the following GGM white paper. In addition to a detailed description of the process and the necessary definitions, it also contains all necessary **LI REQ** commands.

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5328>

In addition, several supporting tools to create the necessary input control statements, and the necessary TCDB entry changes and TMC entry changes, are provided at the IBM Tape Tool website:

<ftp://public.dhe.ibm.com/storage/tapetool>

For more information, ask your local IBM SSR.

8.2 Moving data in and out of the TS7700

When moving data into the TS7700, it is not possible to simply move the cartridges out of a 3494-Bxx VTS and insert them into a stand-alone cluster TS7700 and copy the control data sets (CDSs). This migration approach is supported only for the scenarios that are described in the hardware migration scenarios for the TS7740 or TS7700T, which are described in previous TS7700 IBM Redbooks publications (R1.7, R2.0, R3.0, R3.1, R3.2, and R3.3), which are still available at:

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

In all other scenarios, migrating data into the TS7700 requires that the TS7700 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 a TS7700, the process requires host involvement to copy the data into the TS7700.

This section describes techniques for moving data in and out of the TS7700. You can start using the TS7700 by moving data into it. The best method depends on the application you want to manage with the TS7700.

There are two methods:

- | | |
|----------------------|---|
| Phased method | This method consists of using the TS7700 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. It is considered a quick method because it swiftly puts all the data that you want to move under TS7700 control. |

Hints about how to move data out of the TS7700 are provided in 8.2.5, “Moving data out of the TS7700” on page 293. However, the TS7700 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 that must be moved back out.

8.2.1 Phased method of moving data

The data movement techniques that are outlined here depend more on changes in parameters, routines, or procedures than on overt data movement.

Selecting 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 can select data based on data set name, by application, or by any other variable that you can use in the automatic class selection (ACS) routines. You can also select data based on type, such as System Management Facilities (SMF) data or DASD DUMP data.

Updating the applicable parameters

If you select DFSMSHsm-owned data, review the ARCCMDxx member according to the guidelines in 8.3, “Migration of DFSMSHsm-managed data” on page 296 and update the following definitions:

- ▶ Data Class (DC) ACS routines (if used)
- ▶ Management Class (MC) ACS routines (if used)
- ▶ Storage Class (SC) ACS routines (required)
- ▶ Storage Group (SG) ACS routines (required)
- ▶ For Basic Tape Library Support (BTLS), the unit parameter in the JCL

For DFSMSDss, update the following definitions:

- ▶ DC ACS routines (if used)
- ▶ MC ACS routines (if used)
- ▶ SC ACS routines (required)
- ▶ SG 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 by updating the following definitions:

- ▶ DC ACS routines (if used)
- ▶ MC ACS routines (if used)
- ▶ SC ACS routines (required)
- ▶ SG ACS routines (required)

For data other than DFSMSHsm and DFSMSDss, if you are using SMS tape, update the ACS routines to include the data that you want to move. You decide which data to filter 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.

Updating the tape management system

Although you are not overtly copying data in this option, ensure that you update the TMS catalog or CDS to reflect the changes that you expect. Check the retention rules and limits, and update as needed. If you change data set names when moving to TS7700, you must validate changes against retention rules in your TMS. See 8.4, “DFSMSrmm and other tape management systems” on page 304 for more information.

Watching the data move to the TS7700

Data movement that uses this option does not involve overt actions, such as COPY, RECYCLE, or DUMP. When you activate the ACS routines that contain the code for the TS7700, all new data allocations for the data that you selected are written to the TS7700. 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, such as SMF records or DASD DUMPS, and you can also select data types that create many small data sets. By observing how the TS7700 handles each type of data, you become familiar with the TS7700, its functions, and capabilities.

8.2.2 Quick method of moving data

The steps that are outlined in this section involve overt actions on your part to move data into the TS7700. As with the techniques that are outlined in 8.2.1, “Phased method of moving data” on page 288, you choose the data that you want to move to the TS7700.

Selecting the data to copy

The data that 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 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 can select data that you keep for seven years. Probably the best method for selecting data to copy in the TS7700 is based on the data set name, application by application.

Certain applications have knowledge of the VOLSER where the data is stored. There are special considerations for these applications. If you change the VOLSER on which the data is stored, the application has no way of knowing where the data is. For more information about this topic, see "Implementing Outboard Policy Management for non-z/OS hosts" on page 843.

An easy method is to obtain information from the TMS database. Reports can give you details about the data you have in the tape shop, which helps you select the input volumes.

If you are using DFSMSrmm, you can easily acquire data from a Removable Media Management (RMM) EXTRACT file, which is normally created as part of the regular maintenance. Then, using a **REXX EXEC** or **ICETOOL** JCL program, you extract the needed information, such as data set name, VOLSER, and file sequence of the input volumes.

Moving data to a new TS7700

Although you must use the Tape Copy Utility Tool to move data to TS7700, not all z/OS data can be moved by using this tool. In general, this tool is able to move any data on tape, except those products that manage their own data, for example, DFSMSHsm, DFSMdfp object access method (OAM), or IBM Tivoli Storage Manager.

When using SMS tape, the first step is to update the ACS routines to direct all new data to the TS7700. With this change, new data on tapes gets created in the TS7700 so that moving it again later is not necessary.

If you move DFSMSHsm-owned data, you can use the OAM recycle process, OAM Storage Management Component (OSMC), or the OAM **MOVEVOL** utility to move the data to the TS7700. Use a **COPYDUMP** job to move DFSMSdss data to the TS7700. The utility to use depends on the data selected. In most cases, it is sequential data that can be copied by using the **IEBGENER** utility, DITTO/ESA. If you have DFSORT, **ICEGENER** and **ICETOOL** perform better.

You must use a specific utility when the input data is in a special format, for example, DFSMSdss dump data. DFSMSdss uses blocks up to 256 KB blocksize and only the proper DSS utility, such as **COPYDUMP**, can copy with that blocksize. Be careful when copying multifile and multivolume chains.

In general, all other data (except data that is owned by an application, such as DFSMSHsm) belongs to batch and backup workloads. Use **EXPDT** and **RETPD** from the DFSMSrmm EXTRACT file to discover which tapes have a distant expiration date, begin moving these tapes, and leave the short-term retention tapes to the last phase of data movement. They likely will be moved by the everyday process.

Updating the tape management system with the correct retention information

When the manual copy operation has been successful, it might be necessary to update the TMS 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. For detailed information about this command, see *z/OS DFSMSrmm Managing and Using Removable Media, SC23-6873*.

To avoid this time-consuming process, use a tape copy tool because they can make all the necessary changes in a TMS.

Updating 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. This can be done with **IDCAMS DELETE NOSCRATCH** or within TSO with the 'U' command followed by **DEFINE NONVSAM** or a 'C' command in TSO. For more information, see *z/OS DFSMS Access Method Services Commands, SC23-6846*.

Tape copy tools recatalog tapes during movement without the need for manual intervention.

Releasing the input volume for SCRATCH processing

This final step must be done after you are sure that 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 generation data groups (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. Then, call IDCAMS to update the integrated catalog facility (ICF) catalog. For more information, see *z/OS DFSMS Access Method Services Commands, SC23-6846*.

When the operation completes and all errors are corrected, use the RMM **DELETEVOLUME** command to release the input volumes. For more information about RMM commands and REXX variables, see *z/OS DFSMSrmm Managing and Using Removable Media, SC23-6873*. If you are using a TMS other than RMM, see the appropriate product functions to obtain the same results.

Migrating data inside the TS7700 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, see the following sections:

- ▶ 8.3, "Migration of DFSMSHsm-managed data" on page 296
- ▶ 8.5, "IBM Spectrum Protect" on page 306

With DFSMSHsm, you can change the ARCCMDxx tape device definitions to an esoteric name with TS7700 virtual drives (in a BTLS environment) or change SMS ACS routines to direct DFSMSHsm data in the TS7700. 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 and then starting the space reclamation process.

If you are moving DB2 data into the TS7700, ensure that, when copying the data, the DB2 catalog is also updated with the new volume information. You can use the DB2 **MERGECOPY** utility to speed up processing, using TS7700 virtual volumes as output.

In general, DB2 image copies and Archlog are not retained for a long time. After all new write activity goes to the TS7740 or TS7720T, you can expect that this data is moved by the everyday process.

8.2.3 Products to simplify the task

You might want to consider using a product that is designed to copy data from one medium to another. The first choice is the IBM offering that interacts with DFSMSrmm called *Tape Copy Tool* (see Table 8-1 on page 293). 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. This tool supports any tape media that are supported by DFSMSrmm. The input tape media can be different from the output tape media.

Do not use the tool to copy tape data sets owned by Hierarchical Storage Manager (DFSMSHsm), IBM Tivoli Storage Manager, or similar products, where information of old VOLSERs is kept within the product and not reflected after a copy is made. This challenge typically applies to products where tapes are not cataloged in an ICF catalog, but kept in the product's own database.

The DFSMSrmm Tape Copy Tool cannot be used when you have a TMS other than DFSMSrmm. You must choose another Tape Copy Tool from Table 8-1 on page 293.

Consider the following factors when you evaluate a tape copy product:

- ▶ Interaction with your TMS
- ▶ 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 wildcards or other criteria, such as expiration or creation date

Table 8-1 on page 293 lists several common tape copy products. You can choose one of these products or perhaps use your own utility for tape copy. 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.

Table 8-1 Selection of tape copy tools

Product name	Vendor name	For more information
Tape Copy Tool/ DFSMSrmm	IBM	Contact your IBM SSR for more information about this service offering. Do not confuse this with the Tape Analysis Tools that are mentioned in 11.16.2, "Tools download and installation" on page 721, which can be download from IBM for no extra fee.
Tape Optimizer	IBM	http://www.ibm.com/software/tivoli/products/tape-optimizer-zos/
Beta55	Beta Systems Software AG	http://www.betasystems.com
CA-1/TLMS Copycat	Computer Associates International, Inc.	http://www.cai.com
Tape/Copy	Rocket Software	http://www.rocketsoftware.com/products/rocket-tapecopy-tape-migration
TelTape	Cartagena Software Ltd.	http://www.cartagena.com
Zela	Software Engineering of America	http://www.seasoft.com/home/products/solutions-for-system-z/tape-management/zela?
FATScopy	Innovation	http://www.fdr.com/newsviaemail/pdf/FATSCOPY_NVE4.pdf

In addition to using one of these products, consider using IBM Global Technology Services (GTS) to assist you in planning and moving the data into the TS7700.

8.2.4 Combining methods to move data into the TS7700

You will most likely want to use a combination of the phased and quick methods for moving data into the TS7700. One approach is to classify your data as *static* or *dynamic*.

Static data is information that will be around for a long time. These data can be moved into the TS7700 only with the quick method. You must decide how much of this data will be moved into the TS7700. One way to decide this is to examine expiration dates. You can then set a future time when all volumes, or a subset, are copied into the TS7700. There might be no reason to copy volumes that are going to expire in two months. By enabling these volumes to 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.

8.2.5 Moving data out of the TS7700

There are many reasons why you might want to move data out of the TS7700. The most common reason is for disaster recovery or data interchange. You can move data out of the TS7700 in three ways:

- ▶ Host-based copy tools
- ▶ Copy Export and Copy Export Merge
- ▶ DFSMSHsm aggregate backup and recovery support

Host-based copy tools

You can use a host-based tool to copy the data from the TS7700 to the target.

With this method, the data is reprocessed by the host and copied to another medium. This method is described in 8.2.1, “Phased method of moving data” on page 288. The only difference is that you need to address the TS7700 as input and the non-TS7700 drives as output.

Copy Export and Copy Export Merge

You can use the Copy Export function to copy data from one TS7700 to another empty TS7740 or TS7700T. With this function, a copy of the selected logical volumes that are in the TS7700 can be removed and taken offsite.

In addition to the traditional role in a disaster recovery scenario, Copy Export has been enhanced to enable you to merge Copy Export content from one grid to another. This can help you to move workloads without using your host. This option can save you some tape subsystem outages. Client-initiated Copy Export can be done only to an empty TS7700, but Copy Export Merge can move data into an active TS7700 cluster. It is only available as a Service offering.

This service provides an IBM storage specialist to help you plan for and implement the merge of logical tape volumes that are copy-exported from an IBM TS7740 or TS7700T cluster to another TS7740 or TS7700T cluster that is part of a different IBM TS7700 Grid. The source and the target TS7740/TS7700T can be a stand-alone TS7740 or TS7700T or part of a TS7700 Grid. This helps enhance disaster recovery capability by enabling data recovery to an existing TS7740 or TS7700T without effecting existing data. Contact your IBM SSR for this service offering.

You can find more details about Copy Export and Copy Export Merge in this book in Chapter 12, “Copy Export” on page 757.

DFSMSHsm aggregate backup and recovery support

The third way is to copy the data with the DFSMSHsm aggregate backup and recovery support (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. Then, start addressing other tape drives outside the TS7700, or use the Copy Export function. In this way, you obtain an exportable copy of the data that can be put in an offsite location.

You can use this process to perform these tasks:

1. Creating a selection data set.
2. Defining an aggregate group.
3. Running the ABACKUP VERIFY command.
4. Running the ABACKUP run command.

Creating 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 that 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 placeholders. 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, obtain a list of files on logical volumes that the ABARS function copies to non-TS7700 drives.

You can also use the Copy Export function to move the ABARS tapes to a data recovery site outside of the library.

Defining an aggregate group

Define an aggregate group and related MC to specify exactly which data sets are to be backed up.

Define the aggregate group and MC used for aggregate backup to DFSMS through ISMF panels.

The *aggregate group* lists the selection data set names, instruction data set name, and extra control information that is used by the aggregate backup to determine which data sets to back up.

Running the ABACKUP VERIFY command

You can use the **ABACKUP** command to verify the contents of the aggregate backup without backing up any data sets. This is the same as performing a test run of aggregate backup. The following example shows the **ABACKUP** command:

```
HSEND ABACKUP agname VERIFY UNIT(non_TS7700_unit) PROCESSIONLY(USERTAPE)
```

With the PROCESSIONLY(USERTAPE) keyword, only tape data sets are processed. In this way, you can be sure that only the input data from TS7700 logical volumes is used.

Running the ABACKUP run command

When you are ready, start the actual backup by using the following command:

```
HSEND ABACKUP agname run UNIT(non_TS7700_unit) PROCESSIONLY(USERTAPE)
```

When you enter the **ABACKUP** command with the run 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 that is 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 data, which can be used for disaster recovery and stored offsite using other physical tapes. Consider using the Copy Export function, which enables you to move a copy of the original logical volumes to an offsite 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 being required on the host. However, Copy Export requires an empty TS7700 at your disaster site.

For more information, see the following resources:

- ▶ For Copy Export, see 12.2, “Implementing and running Copy Export” on page 770.
- ▶ For Copy Export Recovery, see 12.3, “Using Copy Export Recovery” on page 780.
- ▶ For using the DFSMSHsm ABARS function, see the *z/OS DFSMSHsm Storage Administration*, SC23-6871.

8.3 Migration of DFSMSHsm-managed data

DFSMSHsm is an application that can use the full cartridge capacity, but for various reasons, you might want to consider using the TS7700 rather than native physical drives for DFSMSHsm data. For example, when writing Migration Level 2 (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 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 Migration Level 2 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 several partial tapes beyond DFSMSHsm needs if recall takeaway activity occurs frequently.

Excess partial tapes created by recall takeaway activity result in poor use 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. However, piggy-back recalls have a positive effect by reducing the number of mounts that are required to run several recalls. You must also consider that multiple recalls from the same tape must be performed serially by the same recall task.

If those same data sets are on separate tapes, the recalls can potentially be performed in parallel, given enough recall tasks. In addition, the persistence of the virtual tape in the Tape Volume Cache (TVC) after it has been unmounted enables DFSMSHsm to run ML2 recalls from the disk cache without requiring that a physical tape be mounted.

Other reasons also exist for directing DFSMSHsm data into a TS7700. The number of native drives limits the number of DFSMSHsm tasks that can run concurrently. With the large number of up to 496 virtual drives in a stand-alone cluster configuration or 992 virtual drives in a two-cluster grid configuration, you can dedicate a larger number of virtual drives to each DFSMSHsm function and enable 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, 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 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 the TS7700 tape drive's high throughput and from its large TVC size, which enables long periods of peak throughput.

DFSMSHsm data is well-suited for the TS7700 because of the appropriate tailoring of those parameters that can affect DFSMSHsm performance. The subsequent sections describe this tailoring in more detail.

For more information, see the *z/OS DFSMSHsm Storage Administration Guide*, SC23-6871.

8.3.1 Volume and data set sizes

The size of user data sets is important when you choose between a TS7700 and 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 that a single data set can span:

- ▶ DASD data sets are limited to 59 volumes, except for partitioned data sets (PDS) and partitioned data set extended (PDSE) data sets, which are limited to one volume.
- ▶ A data set on a virtual I/O (VIO)-simulated device is limited to 65,535 tracks and to one volume.
- ▶ Tape data sets are limited to 255 volumes, but the limit for data sets that are backed up and migrated with DFSMSHsm is 254.

Table 8-2 lists the maximum data set sizes that are supported by DFSMSHsm in z/OS environments.

Table 8-2 Maximum supported data set sizes

Storage medium	Maximum volume size	Maximum number of volumes	Maximum Data set size
DASD: IBM System Storage DS8000®	Standard volumes of 54 GB, EAV sizes are user determined	59	3.18 TB
Tape: TS1120 z/OS V1.13	700 GB x 2.5 compression	40	70 TB
Tape: TS1120 z/OS V2.1 and higher	700 GB x 2.5 compression	254	444.5 TB
Tape: TS1140 / JC z/OS V2.1	4 TB x 2.5 compression	254	2540 TB
Tape: TS7700 z/OS V1.13	25 GB x 2.5 compression	40	3.81 TB
Tape: TS7700 z/OS V2.1 and higher	25 GB x 2.5 compression	254	15.875 TB

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 (with z/OS V2R1, this limit is now 254). 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 supports logical volume sizes of 400, 800, 1000, 2000, 4000, 6000, and 25000 MiB. In z/OS V1.13, with a maximum of 40 volumes that are supported and assuming a compression ratio of 2.5:1, the maximum user data set size for 800 MiB volumes is 80 GiB:

$$800 \text{ MiB} \times 2.5 \times 40 = 80000 \text{ MiB}$$

In z/OS V2.1 and higher, the limit is 254 volumes for HSM user data sets, so the maximum user data set becomes 508,000 MiB:

$$800 \text{ MiB} \times 2.5 \times 254 = 508000 \text{ MiB}$$

Assume that you have a very large data set of 300 GiB. This data set does not fit on 40 volumes of 800 MiB each, but it can fit on 6000 MiB large virtual volumes, as shown in the following example:

$$6000 \text{ MiB} \times 2.5 \times 40 = 600000 \text{ MiB}$$

However, in z/OS V2.1 and higher, this data set can fit on 800 MiB volumes. Any single user data set larger than 3.81 TiB at z/OS 1.13 or 15.875 TiB in z/OS 2.1 and higher, is a candidate for native 3592 tape drives. Assuming a compression rate of 2.5:1, they might not fit onto the supported number of volumes. In this case, consider using native 3592-E06 (TS1130) or 3592-E07 (TS1140) tape drives rather than TS7700.

IDCAMS DCOLLECT BACKUPDATA can be used to determine the maximum size of backed-up data sets in DFSMSHsm inventory. **MIGRATEDATA** can be used to determine the maximum size of migrated data sets in DFSMSHsm inventory.

Important: DFSMSHsm can consist of more than one address space on a single LPAR (Multi-Address Space HSM or MASH), or you can have multiple HSMs sharing a single set of CDSs, called an HSMplex. In either case, you can define commands in the ARCCMDxx member of your DFSMSHsm parmlib to apply only to specific DFSMSHsm hosts by using the ONLYIF statement, or you can have commands apply to all HOSTs in an HSMplex.

Each instance of DFSMSHsm can have a unique MIGUNIT specified. For instance, one host can specify MIGUNIT(3590-1) and another MIGUNIT(TS7700). The same is true for BUUNIT.

The DFSMSHsm host that has 3590-1 specified as a migration or backup unit should process only space management or automatic backup for the SGs where your large data sets, such as z/FS, are. The other DFSMSHsm hosts can then migrate and back up SGs containing the smaller data sets to the TS7700.

To direct a command to a specific instance of DFSMSHsm, you can use an **MVS MODIFY** command with the started task name of the instance of DFSMSHsm that you want to process the command. For example, “F DFSMS2, BACKDS...” or “F DFSMS2, BACKVOL SG(SG)...”.

The following commands affect which output device is used by a specific function:

- ▶ SETSYS TAPEMIGRATION(ML2TAPE(TAPE(unittype)))
- ▶ SETSYS RECYCLEOUTPUT(MIGRATION(unittype))
- ▶ SETSYS BACKUP(TAPE(unittype))
- ▶ SETSYS RECYCLEOUTPUT(BACKUP(unittype))

Migration to a different logical volume size

To ensure that DFSMSHsm starts using larger data sets, you must mark as full any empty or partially filled tapes that are written by using the previous logical volume size. To identify these tapes, enter the following DFSMSHsm command:

```
LIST TTOC SELECT(NOTFULL)
```

Each tape that is 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 301. If a data set spans across DFSMSHsm volumes, it becomes a *connected set* in DFSMSHsm terms.

However, a key point is that if the data set spans, DFSMSHsm uses Force end-of-volume (FEOV) processing to get the next volume mounted. Therefore, the system thinks that the volume is part of a multivolume set regardless of whether DFSMSHsm identifies it as a connected set. Because of the end-of-volume (EOV) processing, the newly mounted DFSMSHsm volume uses the same DC 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 multivolume set and enabling the use of a new DC and other SMS constructs. If the volume is not marked full, the existing multivolume set continues to grow and to use the old constructs.

Use the SETSYS PARTIALTAPE MARKFULL option because it reduces the number of occasions in which DFSMSShsm appends to a partial tape. This results not only in the need to mount a physical tape, but also in the invalidation of the existing virtual tape, which eventually must be reclaimed by the TS7700.

This 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 DC and other SMS constructs. From then on, each new migration or backup volume uses the new size.

8.3.2 TS7700 implementation considerations

This section summarizes DFSMSShsm implementation considerations regarding the TS7700.

Mount wait time

You can direct DFSMSShsm data into a TS7700 (TS7760, TS7720, or TS7740). For a TS7740 or TS7700T, modify your DFSMSShsm mount wait timer to be 12 minutes. This modification enables possibly needed extra time on specific mounts for the TS7740 or TS7700T to stage the data back into cache. Member IECIOSxx is in PARMLIB. Consider defining a special missing-interrupt handler (MIH) value that is named MIH MOUNTMSG=YES,MNTS=10:00 to ensure mount-pending messages if delays in specific mounts occur. The value of 10 can be adjusted to your specific value.

Logical volume size

Consider using large logical volumes, such as 6000 MiB or 25000 MiB, for backup and smaller logical volumes for migration, especially if you are using the TS7740 or TS7700T. 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 very large data sets are migrated that might result in reaching the 40 or 254-volume limits, should use the maximum capacity of the virtual volume. Write your ACS routines to select a different SMS DATACLAS for backup and migration activities that is based on the optimum volume size.

See Table 8-3 on page 302 when you customize the ARCCMDxx SETSYS parameters. HSM is aware of the large virtual volume capacity; it is not necessary to use high PERCENTFULL values to tune capacity of tapes from a DFSMSShsm point of view. The maximum PERCENTFULL value that can be defined is 110% but it is no longer necessary to go above 100%.

Other applications might have a similar existing TAPECAPACITY-type specification or a PERCENTFULL-type specification to enable applications to write beyond the default volume sizes for MEDIA1 (cartridge system tape) and MEDIA2 (enhanced capacity cartridge system tape).

In OAM's Object Tape Support, the TAPECAPACITY parameter in the SETOAM statement of the CBROAMxx PARMLIB member is used to specify the larger logical volume sizes. Because OAM also obtains the size of the logical volume from the TS7700, defining TAPECAPACITY in the CBROAMxx PARMLIB member is not necessary. For more information about Outboard policy management, see the *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

Multisystem considerations

If multiple TS7700 tape drives are eligible for a request, also consider that the same logical volume size is used for the request across all libraries. When you view the volumes through your TMS, the TMS might continue to display the volume capacity based on the default volume size for the media type. However, the volume usage (or a similar parameter) shows how much data has been written to the volume, which reflects its larger capacity.

Scratch volumes

The default volume size is overridden at the library through the DC policy specification, and is assigned or reassigned when the volume is mounted for a scratch mount or rewritten from load point as a specific mount. Using a global scratch pool, you benefit from a fast mount time by establishing your scratch categories, as explained in “Defining scratch categories” on page 548. Consider using the following definitions to benefit from the fast scratch mount times:

- ▶ `SETSYS SELECTVOLUME(SCRATCH)`: Requests DFSMSHsm to use volumes from the common scratch pool.
- ▶ `SETSYS TAPEDELETION(SCRATCHTAPE)`: Defines that DFSMSHsm returns tapes to the common scratch pool.
- ▶ `SETSYS PARTIALTAPE(MARKFULL)`: Defines that an DFSMSHsm task will mark the last tape that it used in a cycle to be full, avoiding a specific mount during the next cycle.

When you use the **MARKFULL** parameter, the stacked volume contains only the written data of each logical volume that is 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 might leave unused while it tries to eliminate the 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 otherwise enables. The amount of unused media can vary 0 - *nnnn* physical megabytes, but roughly averages 50% of the median data set size. For example, if you specify 4000 MiB and your median data set size is 2 MiB, on average, only 1 MiB of media is unused per cartridge.

Installations that currently experience an excessive number of spanning data sets need to consider specifying a larger value in the **SETSYS TAPESPANSIZE** command. Using a high value reduces tape spanning. In a TS7700, this value reduces the number of virtual volumes that need to be recalled to satisfy DFSMSHsm recall or 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 MiB that does not fit on the remaining space of a virtual volume is started on a fresh new virtual volume.

8.3.3 DFSMSHsm task-related considerations

To better understand the use of DFSMSHsm with TS7700, this section summarizes the DFSMSHsm functions that use tapes and analyzes the benefit of tape virtualization for these functions.

Backups of DFSMSHsm control data sets

DFSMSHsm CDSs can be backed up easily in a TS7700 by using virtual volumes rather than physical volumes, which might otherwise be underused.

Volume dumps

When using TS7700 as output for the DFSMSHsm AUTODUMP function, *do not* specify the following parameters:

```
DEFINE DUMPCLASS(dclass STACK(nn))  
BACKVOL SG(sgname) | VOLUMES(volses) DUMP(dclass STACK(10))
```

These parameters were introduced to force DFSMSHsm to use the capacity of native physical cartridges. If used with TS7700, 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 TS7740 or TS7700T 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 TS7740 or TS7700T.

For example, if your installation often has more than 10 tape recall tasks at one time, you probably need 12 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 for DFSMSHsm backup and recovery functions benefits you without affecting DFSMSHsm performance. However, review your DFSMSHsm performance requirements before moving DFSMSHsm BACKUP to the TS7700.

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 as the target for the TAPECOPY operation from an original volume that is not a TS7700 volume might cause problems in DFSMSHsm because TS7700 virtual volumes have different volume sizes.

Use the information in Table 8-3 to tailor your TAPECOPY environment.

Table 8-3 TAPECOPY usage

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full to be defined (assuming 2:1 compression)
TS7700 (CST): 400 MB	3490E (CST)	100%
TS7700 (ECCST): 800 MB	3490E (ECCST)	100%
3490E (CST): 400 MB	TS7700 CST: 400 MB	45%

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full to be defined (assuming 2:1 compression)
3490E (ECCST): 800 MB	TS7700 (ECCST): 800 MB	45%
TS7700 (CST): 400 MB	TS7700 (CST): 400 MB	100%
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(CST): 6 GB	TS7700 (CST): 6 GB	100%
TS7700 (ECCST): 800 MB	TS7700 (ECCST): 800 MB	100%
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%
TS7700 (ECCST): 6 GB	TS7700 (ECCST): 6 GB	100%

For example, if you are planning to put DFSMSHsm alternative copies into a TS7700, a tape capacity of 45% might not be enough for the input non-TS7700 ECCST cartridges. TAPECOPY fails if the (virtual) output cartridge encounters EOV before the input volume has been copied completely.

However, using TS7700 logical volumes as the original and 3490E native as the TAPECOPY target might cause EOV at the alternative volume because of the higher LZ data compression algorithm, IBMLZ1, compression seen on the virtual drive compared to the improved data-recording capability (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..)
```

DUPLEX TAPE

For duplexed migration, both output tapes must be of the exact same size and unit type. A preferred practice is to use a multi-cluster grid and the new Synchronous mode copy support, and enable the hardware to run the duplex rather than the DFSMSHsm software function. This method also enables you to more easily manage the disaster side. You can use Geographically Dispersed Parallel Sysplex (GDPS) and switch to the remote DASD side and the tape VOLSER itself does not need to be changed. No **TAPEREPL** or **SETSYS DISASTERMODE** commands are needed.

When HSM writes ML2 data to tape, it deletes the source data as it goes along, but before the RUN is sent to the TS7700. Therefore, until the copy is made, only one copy of the ML2 data might exist. The reason is because the TS7700 grid, even with a Copy Consistency Point of [R,R], makes a second copy at RUN time.

By using the appropriate MC settings in SMS, you can ensure 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 hierarchical storage management (HSM) duplexing or the new Synchronous copy mode option support in the TS7700. Both ways create 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 cluster to avoid loss of data if one of the clusters experiences a disaster. You can use the Copy Consistency Points to ensure that each copy of the duplexed data is sent to a separate cluster.

RECYCLE

The DFSMSHsm RECYCLE function reduces the number of logical volumes inside the TS7700, but when started, it can cause bottlenecks in the TS7740 or TS7700T recall process. If you have a TS7740 or TS7700T with four physical drives, use a maximum of two concurrent DFSMSHsm RECYCLE tasks. If you have a TS7740 or TS7700T with six physical drives, use no more than five concurrent DFSMSHsm RECYCLE tasks.

Select the RECYCLEPERCENT and consider the following information:

- ▶ You will free logical volumes on a stacked volume with hundreds of other logical volumes.
- ▶ The space that is occupied by the logical volume is freed up only if and when the logical volume is used (overwritten) again, unless you are using Expired Volume Management.
- ▶ To RECYCLE, the TS7700 must load the input volumes into the TVC.

Use a RECYCLEPERCENT value that depends on the logical volume size, for example:

- ▶ 5 for 1000 MiB, 2000 MiB, 4000 MiB, or 6000 MiB volumes
- ▶ 10 for 400 MiB or 800 MiB volumes

You can use the following commands to limit which volumes can be selected for DFSMSHsm RECYCLE processing. For instance, you might want to limit RECYCLE to only your old technology, and exclude the newer tape technology from RECYCLE until the conversion is complete. You can use the following commands to limit which tape volume ranges are selected for RECYCLE:

- ▶ RECYCLE SELECT (INCLUDE (RANGE (nnnnn:mmmmm)))
- ▶ RECYCLE SELECT (EXCLUDE (RANGE (nnnnn:mmmmm)))

You can also use the SETSYS RECYCLEOUTPUT to determine which tape unit to use for the RECYCLE output tapes. You can use your ACS routines to route the RECYCLEOUTPUT unit to the wanted library by using the &UNIT variable.

See *IBM z/OS DFSMSHsm Primer*, SG24-5272 for more information about implementing DFSMSHsm.

8.4 DFSMSrmm and other tape management systems

No changes are required to any TMS to support basic TS7700. You review only the retention and movement criteria for the data in the TS7700. You must check your daily tape management process to delete any step that relates to EJECT activities.

DFSMSrmm accepts logical volume capacity from an open close end-of volume (OCE) module. DFSMSrmm can now always list the actual reported capacity from TS7700.

To start the low-on-scratch procedure, DFSMSrmm uses these messages:

- ▶ CBR3660A
- ▶ CBR3792E
- ▶ CBR3794A

Note: Prior to APAR OA49373, the CBR3660A message was deleted when the scratch count was 2X+1 above the threshold. With this APAR (z/OS V2R1 and above), when the CBR3660A is deleted it can be customized using the CBROAMxx PARMLIB member and the SETTLIB command.

When you direct allocations inside the TS7700, the vital record specifications (VRSs), or vault rules, indicate to the TMS 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 must be exported. For Copy Export, the required location is only used for stacked volumes that have been Copy Exported.

Other TMSs 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 enable 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 CBRSPCLS. For more information about the enhanced return-to-scratch process, see *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874.

Stacked volumes cannot be used by the host; they are managed exclusively by the TS7740 or TS7720T. Do not enable 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 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 TMSs.

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 CDS. 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 or more virtual volumes, the size of the RMM CDS and the amount of secondary space available must be checked. RMM uses 1 MB for every 1,000 volumes defined in its CDS. An extra 1,000,000 volumes need 1,000 MB of space. However, do not add all the volumes initially. See "Inserting virtual volumes" on page 546 for more information.

To increase the size of the RMM CDS, you must quiesce RMM activities, back up the CDS, and 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 Implementation and Customization Guide*, SC23-6874. You should consider using VSAM extended format in your CDS. Extended format and Multivolume support almost any growth rate in the Configuration data set.

Other TMSs, 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. Therefore, they must stop to define more cartridges.

Additionally, some TMSs do not enable the specification of tape volumes with alphanumeric characters or require user modifications to do so. See the correct product documentation for this operation.

In both RMM and the other TMSs, the virtual volumes do not have to be initialized. The first time that a VOLSER is used, TS7700 marks the virtual volume with VOL1, HDR1, and a tape mark, as though it had been done by EDGINERS or IEHINITT.

8.5 IBM Spectrum Protect

IBM Spectrum™ Protect (Tivoli Storage Manager family) provides backup, snapshot, archive, recovery, space management, bare machine recovery, and disaster recovery capabilities. Throughout this publication, the Tivoli Storage Manager name will be referenced. 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.

Tivoli Storage Manager 6.1, released in 2009, had no Tivoli Storage Manager Server support for z/OS. IBM Tivoli Storage Manager for z/OS Media V6.3 and IBM Tivoli Storage Manager for z/OS Media Extended Edition V6.3 are replacement products for Tivoli Storage Manager V5.5 and Tivoli Storage Manager Extended Edition for z/OS V5.5, with new functions available in Tivoli Storage Manager V6, while maintaining the ability to access Fibre Channel connection (FICON)-attached storage on a z/OS system.

IBM Tivoli Storage Manager for z/OS Media and IBM Tivoli Storage Manager for z/OS Media Extended Edition, introduced with Version 6.3, are designed to enable IBM Tivoli Storage Manager V6.3 servers that are running on IBM AIX® and Linux on z Systems to access various FICON-attached tape libraries on z/OS, including the TS7700 family.

Tip: Beginning with Version 7.1.3, IBM Tivoli Storage Manager is now IBM Spectrum Protect™. Some applications, such as the software fulfillment systems and IBM License Metric Tool, use the new product name. However, the software and its product documentation continue to use the Tivoli Storage Manager product name. To learn more about the rebranding transition, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21963634>

For the current Tivoli Storage Manager supported levels for Linux on z Systems, see IBM Knowledge Center:

<http://www.ibm.com/support/docview.wss?uid=swg21243309>

Figure 8-2 shows a high-level drawing of the data flow in a Tivoli Storage Manager for z/OS Media environment.

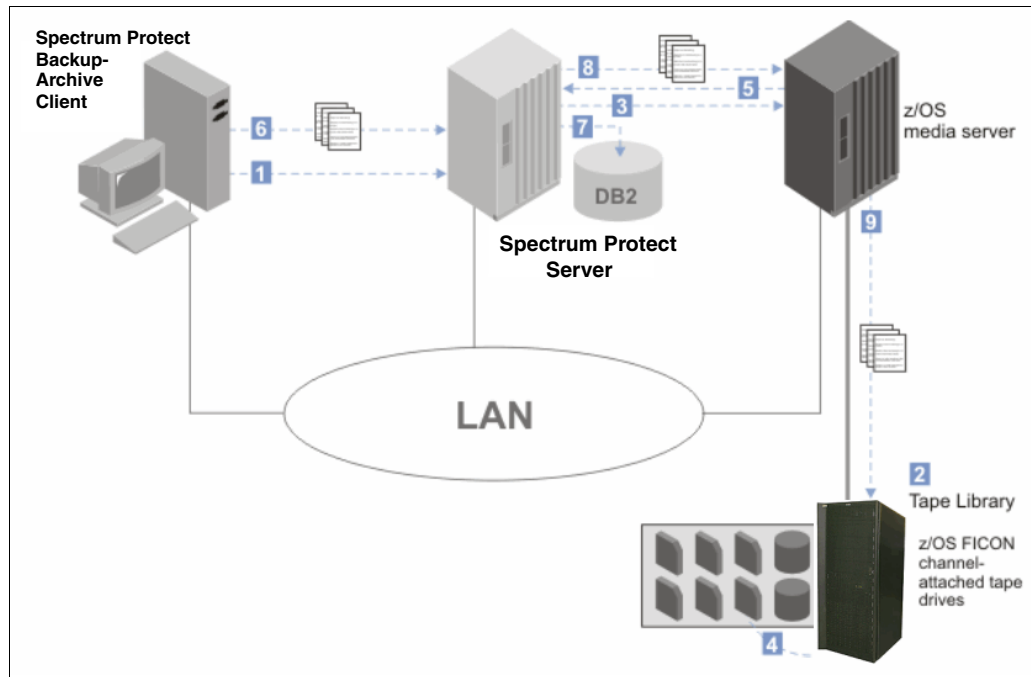


Figure 8-2 Data flow from the backup-archive client to z/OS media server storage

The following numbers correspond to the numbers in Figure 8-2:

1. The Tivoli Storage Manager backup-archive client contacts the Tivoli Storage Manager server.
2. The Tivoli Storage Manager server selects a library resource and volume for the backup operation.
3. The Tivoli Storage Manager server contacts the z/OS media server to request a volume mount.
4. The z/OS media server mounts the tape volume.
5. The z/OS media server responds to the Tivoli Storage Manager server that the mount operation is complete.
6. The backup-archive client begins sending data to the Tivoli Storage Manager server.
7. The Tivoli Storage Manager server stores metadata in the database and manages the data transaction.
8. The Tivoli Storage Manager server sends the backup-archive client data to the z/OS media server.
9. The z/OS media server writes the data to z/OS storage.

If you plan to store IBM Tivoli Storage Manager data in the TS7700, consider the following suggestions for placing data on your TS7700:

- ▶ Use TS7740 or TS7700T for IBM Tivoli Storage Manager Archiving for archiving and backing up large files or databases for which you do not have a high-performance requirement during backup and restore. TS7740 or TS7700T is ideal for IBM Tivoli Storage Manager archive or long-term storage because archive data is infrequently retrieved. Archives and restorations for large files can see less effect from the staging.

Small files, such as individual files on file servers, can see performance effects from the TS7740 or TS7700T staging. If a volume is not in cache, the entire volume must be staged before any restore can occur.

- ▶ 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 TS7740 or TS7700T can be slower because all volumes must be staged to the cache.

Periodically, set IBM Tivoli Storage Manager reclamation on by setting the threshold to a lower value to regain the use of TS7700 volumes with a small amount of valid data that will not expire for a longer period. IBM Tivoli Storage Manager reclamation must be scheduled for off-peak hours.

- ▶ Use collocation to reduce the number of TS7700 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 separate physical tapes. Collocation with TS7740 or TS7700T does not minimize the physical tapes that are used, but minimizes the number of logical volumes that is used.

Collocation with TS7740 or TS7700T can improve the restore time for large amounts of data. TS7740 or TS7700T does not ensure physical tape separation when collocation is used because separate logical volumes can be on the same physical tape.

- ▶ Use TS7740 or TS7700T for IBM Tivoli Storage Manager database backups that are to be used for recovery from local media, and use TS7740 or TS7700T at a recovery site or native drives for backups that are to be used for recovery from offsite media. IBM Tivoli Storage Manager requires a separate tape for every backup of the IBM Tivoli Storage Manager database, so many logical volumes with less data is created.

When using the TS7700, you do not have to worry about the unused capacity of logical volumes.

- ▶ Use TS7740 or TS7700T for backups of primary pools, noting that similar considerations apply to copy storage pools. If only one copy pool is used for local backups, that storage pool must not be in the TS7740 or TS7700T, because there is no guarantee that data in the copy storage pools is on separate physical volumes.

If storage pools for local and offsite backups are used, the copy storage pools for local backups can be in the TS7740 or TS7700T. The copy storage pools for offsite backups must use native drives or a TS7740 or TS7700T at the recovery site.

- ▶ Use TS7700 in server-to-server configurations for multiple IBM Tivoli Storage Manager server implementations. If you are using an IBM Tivoli Storage Manager server-to-server configuration, the data from your remote IBM Tivoli Storage Manager servers is stored as virtual volumes in the TS7700, which appear as sequential media volumes on the source server and are stored as archive files on a target server. These are ideal candidates for a TS7700.

Native or virtual drives

When only one stand-alone TS7740 or TS7700T is available, you might choose native drives for data that will be used for frequent individual file restores or require high performance for backup and restore without any delays because of staging activity. IBM Tivoli Storage Manager uses the 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 multiple TS7740 or TS7700T tape drives and can IMPORT/EXPORT between the TS7740 or TS7700T tape drives.

IBM Tivoli Storage Manager parameter settings

The settings for the following parameters can affect the performance of IBM Tivoli Storage Manager with TS7700:

- ▶ **MAXSCRATCH** (storage pool definition). As for DFSMSHsm, IBM Tivoli Storage Manager must 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 scratch mounts.
- ▶ **MOUNTLimit** (device class definition). With one TS7700, you have up to 496 virtual drives available. The number of drives available for IBM Tivoli Storage Manager use can probably be increased, considering TS7700 performance. Set the 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, 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 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 that is 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 that are 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 scratch mounts.

For more information about setting up Tivoli Storage Manager, see the following web page:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.common.doc/t_installing_srv.html

8.6 DFSMSdss

This section describes the uses of DFSMSdss with the TS7700.

8.6.1 Full volume dumps

DFSMSdss full volume dumps can use the TS7700. Ensure that 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 throughput capabilities, the TS7700 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 TS7740 or TS7700T, you fill the stacked cartridge without changing JCL by using multiple virtual volumes. The TS7740 or TS7700T then moves the virtual volumes that are created onto a stacked volume.

The only problem that you might experience when using TS7700 for the data set Services (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 use an SMS DATACLAS specification, Volume Count *nn*, to enable more than five volumes. A better method is to choose a 25000 MiB logical volume through your SMS DATACLAS. This method prevents unneeded multivolume files.

Using the **COMPRESS** keyword of the **DUMP** command, you obtain a software compression of the data at the host level. Because data is compressed at the TS7700 before being written into the TVC, host compression is not required unless channel use is high already.

8.6.2 Stand-Alone Services

DFSMSdss Stand-Alone Services provide a stand-alone restore function that enables you to restore vital system packs without needing to rely on a z Systems environment.

Stand-Alone Services support the 3494 and 3584 (TS3500) tape library and the VTS. You can use it to restore from native and virtual tape volumes in a TS7700. With Stand-Alone Services, you specify the input volumes on the **RESTORE** command and send the necessary mount requests to the 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 are provided as a replacement to the previous DFDSS V2.5 and DFSMS 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 **BUILD SA** command of DFSMSdss. Create a new virtual tape as non-labeled and then put the stand-alone program on it.

For more information about how to use the TS7700 MI to set a device in stand-alone mode, see “Modify Virtual Volumes window” on page 399.

Complete these steps to use an IPL of the Stand-Alone Services program from a virtual device and restore a dump data set from virtual volumes:

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 the virtual device from which you will load the Stand-Alone Services program in stand-alone mode by selecting **Virtual Drives** on the TS7700 MI of the cluster where you want to mount the logical volume.
3. Follow the sequence that is described for stand-alone mount in “Virtual tape drives” on page 381.
4. Load the Stand-Alone Services program from the device you 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.

5. When the IPL is complete, enter the Stand-Alone Services **RESTORE** command from the specified input device. Example 8-1 shows a group of statements for using this command.

Example 8-1 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 that is 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 request 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 to unmount it and to mount the next volume.

6. When the restore is complete, unload and unmount the IPL volume from the virtual device by using the TS7700 MI's Setup Stand-alone Device window.
7. In the Virtual Drives window in Figure 9-60 on page 381, click **Actions** → **Unmount logical volume** to unload the virtual drive and finish the Stand-alone Mount operation.

Stand-Alone Services send the necessary mount and unmount orders to the library. If you are using another stand-alone restore program that does not support the mounting of library resident volumes, you must set the source device in stand-alone mode and manually instruct the TS7700 to mount the volumes by using the Setup Stand-alone Device window.

For more information about how to use Stand-Alone Services, see the *z/OS DFSMSdss Storage Administration*, SC23-6868.

8.7 Object access method

Tape cartridges provide a low-cost storage medium for storing primary and backup copies of OAM objects.

Enabling objects to be stored on tape volumes with DASD and optical media provides flexibility and more efficiency within the storage management facility.

OAM stores objects on a TS7700 as they are stored in a normal TS3500 tape library, with up to 496 virtual drives and many virtual volumes available.

When using the TS7740 or TS7700T, 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.

The OAM **TAPECAPACITY** parameter is no longer needed when you use the TS7700 with OAM because OAM now obtains the capacity of the logical volume from the library.

Virtual volumes in a TS7700 can be used to store your object data (OBJECT or OBJECT BACKUP SG data). With the data in cache, the TS7700D (or TS7700T CP0) can be ideal for your primary OBJECT SG needs. Your OBJECT BACKUP SGs can then be in a TS7700 or TS7700T CPx, depending on your recovery needs.

As with DFSMSHsm, the Synchronous mode copy option can be used to replicate your data to other clusters in the grid. This replicated copy (and others in the grid) is then managed by the TS7700. The replication capabilities in the grid can be used in addition to any OAM-managed backup copies.

A virtual volume can contain multiple OAM objects. To optimize the use of TS7700 storing OAM object data, consider the following suggestions:

- ▶ Review the **MOUNTWAITTIME** parameter when using TS7740 or TS7700T CPx to store OAM object tape data. The default (5 minutes) probably needs to be increased. Twelve minutes is a better number in case you must recall a logical volume to read object data and there are other recall requests queued at the time. The TS7740 or TS7700T CPx might need to stage the data back into cache, which accounts for the extra mount time.
- ▶ Review the **MAXTAPERETRIEVETASKS** and **MAXTAPESTORETASKS** parameters when using TS7700 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 to store OAM data.

8.8 Database backups

Using a TS7700 as output confers several advantages to database backups. This section provides a detailed description of these benefits for database products, such as DB2.

8.8.1 DB2 data

DB2 uses tapes for storing archive logs and for storing image copies. Either one can be created in multiple copies to be stored both onsite for local recovery purposes and offsite for disaster recovery purposes. To use DB2 tape data with the TS7700, use the approaches that are described in this section.

Archive logs

DB2 tracks 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, make backups of archive logs. You can use general backup facilities or the DB2 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, by using Tape Mount Management (TMM). The other copy can be written directly to tape and then moved to an offsite location.

With TS7700, you can write the local archive log directly inside the TS7700. Avoiding the use of TMM saves DASD space, saves DFSMSshm CPU cycles, and simplifies the process. The disaster recovery copy can be created by using Copy Export capabilities in the TS7740 and TS7700T, or by using native tape drives, so that it can be moved offsite.

The size of an archive log data set varies from 150 MB to 1 GB. The size of a virtual volume on a TS7700 can be up to 25000 MiB, so be sure that your archive log is directed to a virtual volume that can hold the entire log. Use a single volume when unloading an archive log to tape. The size of a virtual volume on a TS7700 can be up to 75000 MiB, assuming a 3:1 compression ratio.

Tailoring the size and number of active log DASD data sets enables 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 is not a problem with the TS7700 because it requires no manual operation. Even with the TS7740 (or TS7700T), the archive logs stay in the TVC for some time and are 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 that is 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 that are processed on real IBM 3480 (if improved data-recording capability (IDRC) is enabled) and IBM 3490 tape drives.

On a TS7700, 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 that TS7700 can provide for DB2 operations is the availability of up to 496 (stand-alone cluster) or 2976 virtual drives (six-cluster grid configuration) because DB2 often needs many drives concurrently to run 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 changed since the last full image copy was taken. Incremental image copies are typically taken daily. 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 offsite storage.

The size of the table spaces to be copied varies from a few megabytes to several gigabytes. The TS7700 solution is best for small-sized 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 and 496 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 enables you to run multiple copies concurrently and use the MERGECOPY DB2 utility without effect. An advisable solution is to run a full image copy of the DB2 databases once a week outside the TS7700, and run the incremental image copies daily by using TS7700. The smaller incremental copy fits better with the TS7700 volume sizes.

8.8.2 CICS and IMS

As with DB2, both IBM CICS® and IMS™ use tapes to store logs and image copies of databases.

CICS is only a data communication product. 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 tracks database changes in its journal data sets. IMS tracks 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, you can create the local copy directly on TS7700 virtual volumes, and then copy those volumes to non-TS7700 tape drives, or to a remote TS7700.

Having a local copy of the logs that is written inside the TS7700 enables you faster recovery because the data stays in the TVC for some time.

When recovering a database, you can complete back out operations in less time with the TS7700 because when reading logs from tape, IMS uses the slow read backward operation (100 KBps) on real tape drives. With the TS7700, 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. Both perform much better than on physical drives. The reason is not just that the data is in the TVC, but the TS7700 code also fully buffers the records in the reverse order that they are on the volume when in read backwards mode.

Another benefit TS7700 provides to recovery operations is the availability of up to 496 virtual drives per cluster. This configuration enables you to mount several logs concurrently and 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, 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 to copy one or more databases to tape.

With the TS7700, 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 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, you must use an image copy that is taken before running the batch job to recover the database. Otherwise, you can use logs and checkpoints, which enable you to restart from a consistent database image that was taken during the batch execution processing. Using TS7700, you can access these image copies and logs at a higher speed.

The TS7700 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.

8.8.3 Batch data

The following applications write to tape and benefit from using the TS7700:

- ▶ 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 of this data benefits from using the TS7700 for output.

With TS7700, the application can write one file per volume, by using only part of the volume capacity. The TS7740 or TS7700T takes care of completely filling the stacked cartridge for you, without JCL changes.

The only step that you must remember is that if you need to move the data offsite, you must address a device outside the local TS7700, or use other techniques to copy TS7700 data onto other movable tapes, as described in 8.2.5, “Moving data out of the TS7700” on page 293.



Part 3

Operation

This part describes daily operations and the monitoring tasks related to the IBM TS7700 with R3.3. It also provides you with planning considerations and scenarios for disaster recovery, and for disaster recovery testing.

This part includes the following chapters:

- ▶ Operation
- ▶ Host Console operations
- ▶ Performance and monitoring
- ▶ Copy Export
- ▶ Disaster Recovery Testing



Operation

This chapter provides information about how to operate and configure the IBM TS7700 by using the Management Interface (MI).

This chapter includes the following sections:

- ▶ User interfaces
- ▶ TS7700 Management Interface
- ▶ Common procedures
- ▶ Basic operations
- ▶ Tape cartridge management
- ▶ Cluster intervention scenarios
- ▶ TS7700 Management Interface considerations

This chapter also includes information about these topics:

- ▶ TS3500 tape library GUI
- ▶ TS4500 tape library management GUI
- ▶ Call Home and Electronic Customer Care

For general guidance regarding TS3500 or TS4500 tape libraries, see the following IBM Redbooks publications:

- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *IBM TS4500 R3 Tape Library Guide*, SG24-8235

9.1 User interfaces

To successfully operate the TS7700, you must understand its concepts and components. This chapter combines the components and functions of the TS7700 into two groups:

- ▶ The logical view
- ▶ The physical view

Each component and each function belong to only one view.

The logical view is named the *host view*. From the host allocation point of view, there is only one library, called the *composite library*. The logical view includes virtual volumes and virtual tape drives.

Currently, a composite library can have up to 2976 virtual addresses for tape mounts, considering a six-cluster grid with support for 496 virtual devices in each cluster (available with FC5275 and z/OS APAR). Read more about this in Chapter 2, “Architecture, components, and functional characteristics” on page 15. The host is only aware of the existence of the underlying physical libraries because they are defined through Interactive Storage Management Facility (ISMF) in a z/OS environment. The term *distributed library* is used to denote the physical libraries and TS7700 components that are part of one cluster of the multi-cluster grid configuration.

The *physical view* shows the hardware components of a stand-alone cluster or a multi-cluster grid configuration. In a TS7700 tape-attached model, it includes the TS3500 or TS4500 tape library and the configured drives. Supported drives are 3592 J1A, TS1120, TS1130, TS1140, or TS1150 tape drives with the TS3500 tape library.

Release 4.0 introduces support for the TS4500 tape library when attached to models TS7740-V07, TS7720T-VEB, and TS7760. TS3500 tape library can still be attached to all TS7700 tape attach models.

TS4500 tape library supports TS1140 (3592 EH7) and TS1150 (3592 EH8) tape drive models on TS7700 tape attach configurations.

Release 3.3 introduced support for TS1150 along with heterogeneous support for two different tape drive models at the same time, as described in 7.1.5, “TS7700 tape library attachments, drives, and media” on page 248.

Release 3.2 introduced support for 496 devices per cluster (available with feature code 5275), making up for 2976 virtual tape devices in a grid of six fully configured clusters. z/OS also needs an authorized program analysis report (APAR) to grow from the previous limit of 2048 to 4096 per grid or composite library.

Before R3.2 of Licensed Level of Code (LIC), a composite library can have up to 1536 virtual addresses for tape mounts, considering a six-cluster grid (256 devices or 16 logical control units (LCUs) per cluster).

The following operator interfaces for providing information about the TS7700 are available:

- ▶ Object access method (OAM) commands are available at the host operator console. These commands provide information about the TS7700 in stand-alone and grid environments. This information represents the host view of the components within the TS7700. Other z/OS commands can be used against the virtual addresses. This interface is described in Chapter 10, “Host Console operations” on page 601.

- ▶ Web-based management functions are available through web-based user interfaces (UIs). The following browsers can be used to access the web interfaces:
 - Firefox 24.x, 31.x, and 38.x
 - Microsoft Internet Explorer Version 9.x, 10.x, and 11
 - Chrome 39.x and 42.x
 - Microsoft Edge 25.xEnable cookies and disable the browser's function of blocking windows for the MI usage. Unsupported web browser versions might cause some MI windows to not display correctly.
 - ▶ Considering the overall TS7700 implementation, two different web-based functions are available:
 - The tape library GUI, which enables management, configuration, and monitoring of the configured tape library. With R4.0, the TS3500 and TS4500 tape libraries are used in TS7700 implementations for tape-attached models.
- Note:** TS4500 tape library is supported on TS7760 and TS7700 V07/VEB tape attach configurations with R4.0.
- The TS7700 MI is used to run all TS7700 configuration, setup, and monitoring actions.
 - ▶ Call Home Interface: This interface is activated on the TS3000 System Console (TSSC) and provides helpful information to IBM Service, Support Center, and Development personnel. It also provides a method to connect IBM storage systems with IBM remote support, also known as Electronic Customer Care (ECC). No user data or content is included in the call home information.

9.1.1 The tape library management GUI

The tape library management GUI web interface enables the user to monitor and configure most of the library functions from the web. The tape library GUI can be started from the tape library expanded page on TS7700 MI by clicking the tape library image there. Starting with R4.0, the tape attach TS7700 can be configured with the TS3500 and TS4500 tape libraries.

Figure 9-1 shows the TS3500 tape library GUI initial window with the System Summary.



Figure 9-1 TS3500 tape library GUI initial window

Figure 9-2 shows the TS4500 Management GUI initial Summary Screen. Notice that GUI general appearance and warning messages presentation are similar to the TS7700 and other IBM storage products MIs.



Figure 9-2 The TS4500 tape library management GUI

The tape library management GUI windows are used during the hardware installation phase of the TS7700 tape attach models. The installation activities are described in 9.3.1, “The tape library with the TS7700T cluster” on page 509.

9.1.2 Call Home and Electronic Customer Care

The tape subsystem components include several external interfaces that are not directly associated with data paths. Rather, these interfaces are associated with system control, service, and status information. They support customer interaction and feedback, and attachment to IBM remote support infrastructure for product service and support.

These interfaces and facilities are part of the IBM System Storage Data Protection and Retention (DP&R) storage system. The main objective of this mechanism is to provide a safe and efficient way for the System Call Home (Outbound) and Remote Support (Inbound) connectivity capabilities.

For a complete description of the connectivity mechanism and related security aspects, see *IBM Data Protection and Retention (DP&R) - System Connectivity and Security*, found at:

<https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102531>

The Call Home function generates a service alert automatically when a problem occurs with one of the following components of the TS7700 subsystem:

- ▶ TS7720
- ▶ TS7740
- ▶ TS7760
- ▶ TS3500 tape library
- ▶ TS4500 tape library

Error information is transmitted to the TSSC for service, and then to the IBM Support Center for problem evaluation. The IBM Support Center can dispatch an IBM SSR to the client installation. Call Home can send the service alert to a window service to notify multiple people, including the operator. The IBM SSR can deactivate the function through service menus, if required.

See Figure 9-3 for a high-level view of call home and remote support capabilities.

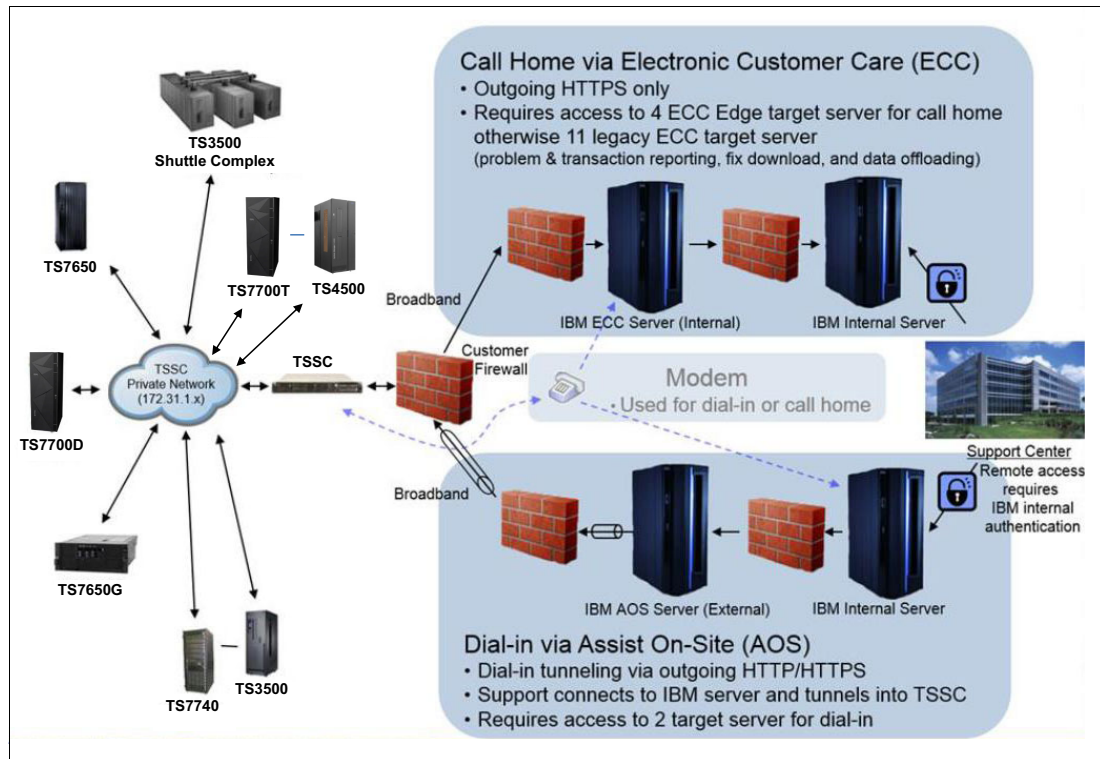


Figure 9-3 Call home and remote support functions

The TSSC can be ordered as a rack mount feature for a range of products. Feature Code 2725 provides the enhanced TS3000 TSSC. Physically, the TS3000 TSSC is a standard rack 1U mountable server that is installed within the 3592 F05 or F06 frame.

Feature code 2748 provides an optical drive, needed for the Licensed Internal Code changes and log retrieval. With the new TS3000 TSSC provided by FC2725, remote data link or call home by using an analog telephone line and modem is no longer supported. Dial-in function through Assist On-site (AOS) and Call Home with ECC functions are both available using HTTP/HTTPS broadband connection.

Electronic Customer Care

ECC provides a method to connect IBM storage systems with IBM remote support. The package provides supports for dial-out communication for broadband Call Home and modem connections. All information that is sent back to IBM is Secure Sockets Layer (SSL) or Transport Layer Security (TLS) encrypted. Modem connectivity protocols follow similar standards as for direct connected modems, and broadband connectivity uses the HTTPS protocol.

Note: Modem option is no longer offered with the latest TSSC FC 2725. All modem-based call home and remote support is planned to be discontinued by the end of 2017.

ECC is a family of services featuring problem reporting by opening a problem management record (PMR), sending data files, and downloading fixes. The ECC client provides a coordinated end-to-end electronic service between IBM business operations, its IBM Business Partners, and its clients.

The ECC client runs electronic serviceability activities, such as problem reporting, inventory reporting, and fix automation. This becomes increasingly important because customers are running heterogeneous, disparate environments, and are seeking a means to simplify the complexities of those environments.

The TSSC enables the use of a proxy server or direct connection. Direct connection implies that there is not an HTTP proxy between the configured TS3000 and the outside network to IBM. Selecting this method requires no further setup. ECC supports customer-provided HTTP proxy. Additionally, a customer might require all traffic to go through a proxy server. In this case, the TSSC connects directly to the proxy server, which initiates all communications to the Internet.

Note: All inbound connections are subject to the security policies and standards that are defined by the client. When a Storage Authentication Service, Direct Lightweight Directory Access Protocol (LDAP), or RACF policy is enabled for a cluster, service personnel (local or remote) are *required* to use the LDAP-defined service login.

Important: Be sure that local and remote authentication is allowed, or that an account is created to be used by service personnel, before enabling storage authentication, LDAP, or RACF policies.

The outbound communication that is associated with ECC call home can be through an Ethernet connection, a modem, or both, in the form of a failover setup. Modem is not supported in the new TS3000 TSSC. 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 TS7700 Autonomic Ownership Takeover Manager (AOTM)
- ▶ The Internal Ethernet Connection, which is used for the local attached subsystem's subnet

Note: The AOTM and ECC interfaces should be in different TCP/IP subnets. This avoids both communications from using the same network connection.

All of these connections are set up using the Console Configuration Utility User Interface that is on the TSSC. TS7700 events that start a Call Home are displayed in the Events pane under the Monitor icon.

Assist On-site

Enhanced support capabilities include the introduction of Assist On-site (AOS) to expand maintenance capabilities. Assist On-site allows IBM support personnel to remotely access local TSSC and the Tape Subsystems under it to identify and resolve technical issues in real time. Assist On-site facilitates problem determination and solution by providing a powerful suite of tools that enables IBM support team to quickly identify and fix issues with the system.

AOS uses the same network as broadband call home, and works on either HTTP or HTTPS. Although the same physical Ethernet adapter is used for these functions, different ports must be opened in the firewall for the different functions. For more information, see 4.1.3, "TCP/IP configuration considerations" on page 136. 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 associated with the target TS7700, which requires physical access by the IBM SSR to the TSSC or the client through the customer interface.
- ▶ Unattended mode, also called *Lights Out mode*, enables a remote support session to be established without manual intervention at the TSSC associated with the target TS7700.

All AOS connections are outbound, so no connection is initiated from the outside to the TSSC. It is always the TSSC that initiates the connection. In unattended mode, the TSSC checks whether there is a request for a session when it connects to the regional AOS relay servers, periodically. When a session request exists, the AOS authenticates and establishes the connection, allowing remote access to the TSSC.

Assist On-site uses current security technology to ensure that the data that is exchanged between IBM Support engineers and the TSSC is secure. Identities are verified and protected with industry-standard authentication technology, and Assist On-site sessions are kept secure and private by using randomly generated keys for session, plus advanced encryption.

Note: All authentications are subject to the authentication policy that is in effect, as described in 9.2.9, “The Access icon” on page 446.

9.2 TS7700 Management Interface

The TS7700 MI is the primary interface to monitor and manage the TS7700. The TS7700 GUI is accessed via TCP/IP, by entering the TS7700 IP address in your web browser. The following web browsers are currently supported:

- ▶ Mozilla Firefox ESR 24.x, 31.x, and 38.x
- ▶ Microsoft Internet Explorer 9, 10, and 11
- ▶ Google Chrome 39.x and 42.x
- ▶ Microsoft Edge 25.x

The current TS7700 graphical user interface (GUI) implementation has an appearance and feel similar to other MI adopted in other IBM Storage products.

9.2.1 Connecting to the Management Interface

To connect to the TS7700 MI, complete the following steps:

1. The TS7700 must first be installed, configured, and online.
2. In the address field of a supported web browser, enter `http://x.x.x.x` (where `x.x.x.x` is the virtual IP address that was assigned during installation). Press Enter or click **Go** in the web browser.
3. The virtual IP is one of three IP addresses that are provided during installation. To access a specific cluster, enter the cluster IP address as shown in Example 9-1, where Cluster 0 is accessed directly.

Example 9-1 IP address to connect to Cluster 0 in a grid

`http://x.x.x.x/0/Console`

4. If a local name server is used, where names are associated with the virtual IP address, then the cluster name rather than the hardcoded address can be used for reaching the MI.
5. The login window for the MI displays as shown in Figure 9-4. Enter the default login name as `admin` and the default password as `admin`.

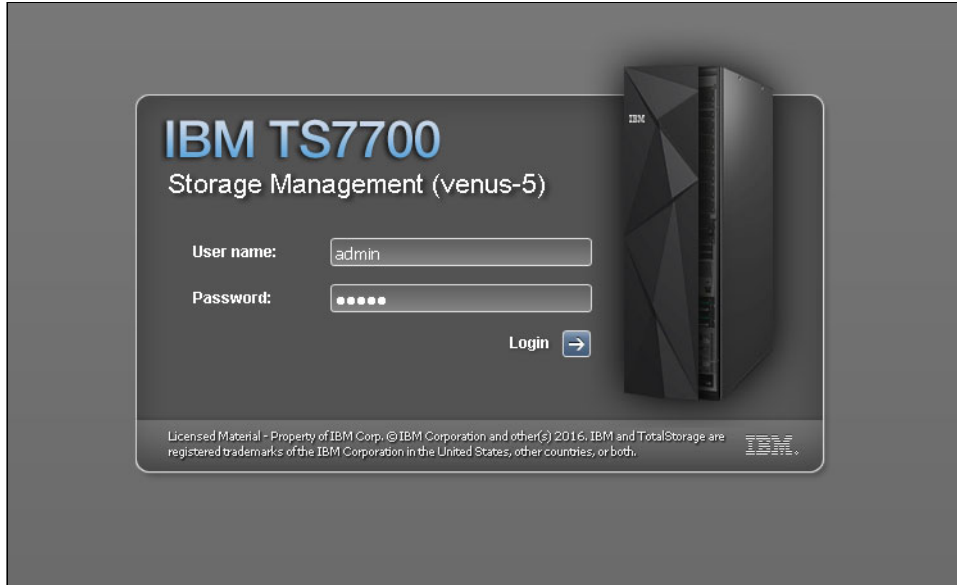


Figure 9-4 TS7700 MI login

After logging in, the user is presented to the Grid Summary page, as shown in Figure 9-5 on page 328.

After security policies are implemented locally at the TS7700 cluster or by using centralized role-base access control (RBAC), a unique user identifier and password can be assigned by the administrator. The user profile can be modified to provide only functions applicable to the role of the user. All users might not have access to the same functions or views through the MI.

For more information, see 9.2.9, “The Access icon” on page 446.

Figure 9-5 shows a visual summary of the TS7700 Grid. It shows a three-cluster grid, the components of it and health status of components. The composite library is depicted as a data center, with all members of the grid on the raised floor. Notice that the TS7760 (the cluster on the left) has a distinct visual appearance when compared to the TS7740 and TS7720 at center and right of the picture.

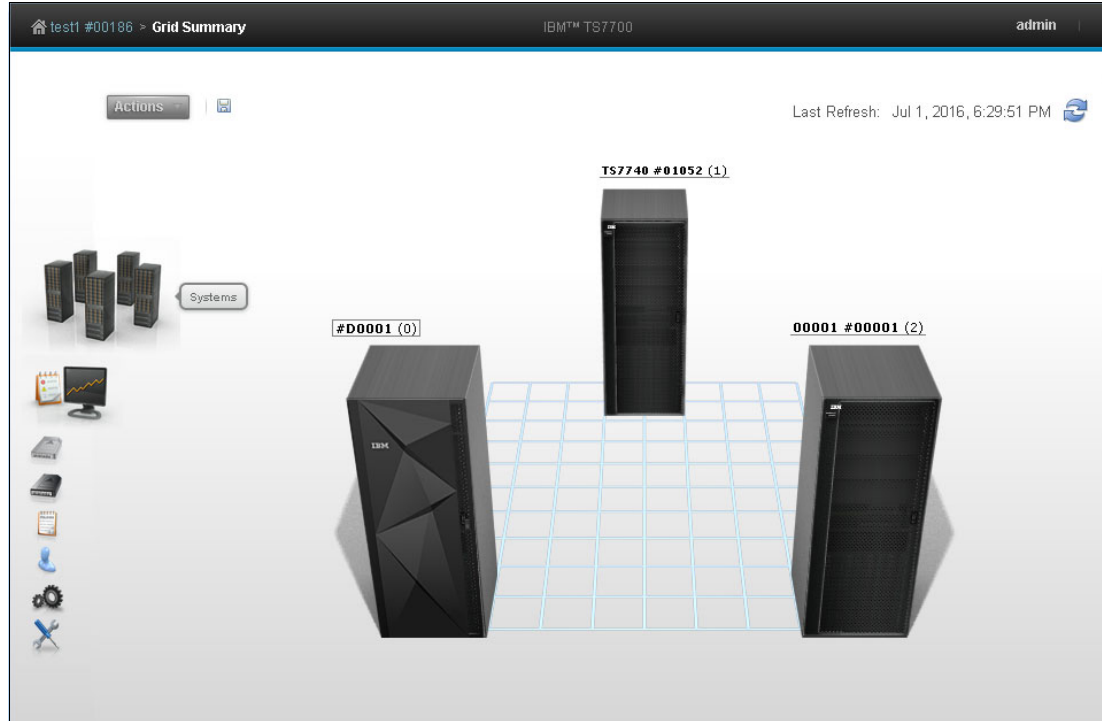


Figure 9-5 MI Grid summary

Each cluster is represented by an image of the TS7700 frame, displaying the cluster's nickname and ID, and the composite library name and Library ID.

The health of the system is checked and updated automatically at times that are determined by the TS7700. Data that is displayed in the Grid Summary window is not updated in real time. The Last Refresh field, in the upper-right corner, reports the date and time that the displayed data was retrieved from the TS7700. To populate the summary with an updated health status, click the Refresh icon near the Last Refresh field in the upper-right corner of Figure 9-5.

The health status of each cluster is indicated by a status sign next 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.

Library control with the TS7700 Management Interface

The TS7700 MI can also link to the TS3500 or TS4500 tape library GUI, which interacts with the physical tape library. In environments where the tape library is separated from the LAN-attached hosts or web clients by a firewall, the ports that are shown in Table 9-1 should be open for correct functioning.

Table 9-1 Network interface firewall

Function	Port	Direction (from library)	Protocol
Tape Library Management Interface	80	Inbound	TCP/IP
Simple Network Management Protocol (SNMP) traps	161/162	Bidirectional	User Datagram Protocol (UDP)/IP
IBM Encryption Key Manager	1443	Outbound	Secure Sockets Layer (SSL)
IBM Encryption Key Manager	3801	Outbound	TCP/IP
LDAP	389/636	Bidirectional	TCP/IP and UDP
HTTPS	443	Bidirectional	TCP/IP

For more information, see the topic *Infrastructure Requirements* under the Planning section in the TS7700 R4.0 IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/hydra_c_ichome.html

9.2.2 Using the TS7700 Management Interface

This section describes how to use the TS7700 management interface (MI).

Login window

Each cluster in a grid uses its own login window, which is the first window that opens when the cluster URL is entered in the browser address field. The login window shows the name and number of the cluster to be accessed. After logging in to a cluster, other clusters in the same grid can be accessed from the same web browser window.

Navigating between windows

Navigation between MI window can be done by clicking active links on a window or on the banner, or by selecting a menu option or icon.

Banner

The banner is common to all windows of the MI. The banner elements can be used to navigate to other clusters in the grid, run some user tasks, and locate additional information about the MI.

See Figure 9-6 for an example of the banner elements and available tasks.



Figure 9-6 Management Interface Banner

Status and event indicators

Status and alert indicators occur at the bottom of each MI window. These indicators provide a quick status check for important cluster and grid properties. Grid indicators provide information for the entire grid. These indicators are displayed on the left and right corners of the window footer, and include tasks and events.

Figure 9-7 shows some examples of status and events that can be displayed from the Grid Summary window. Also, the procedure to dock or undock the LI REQ command panel is highlighted in the picture.

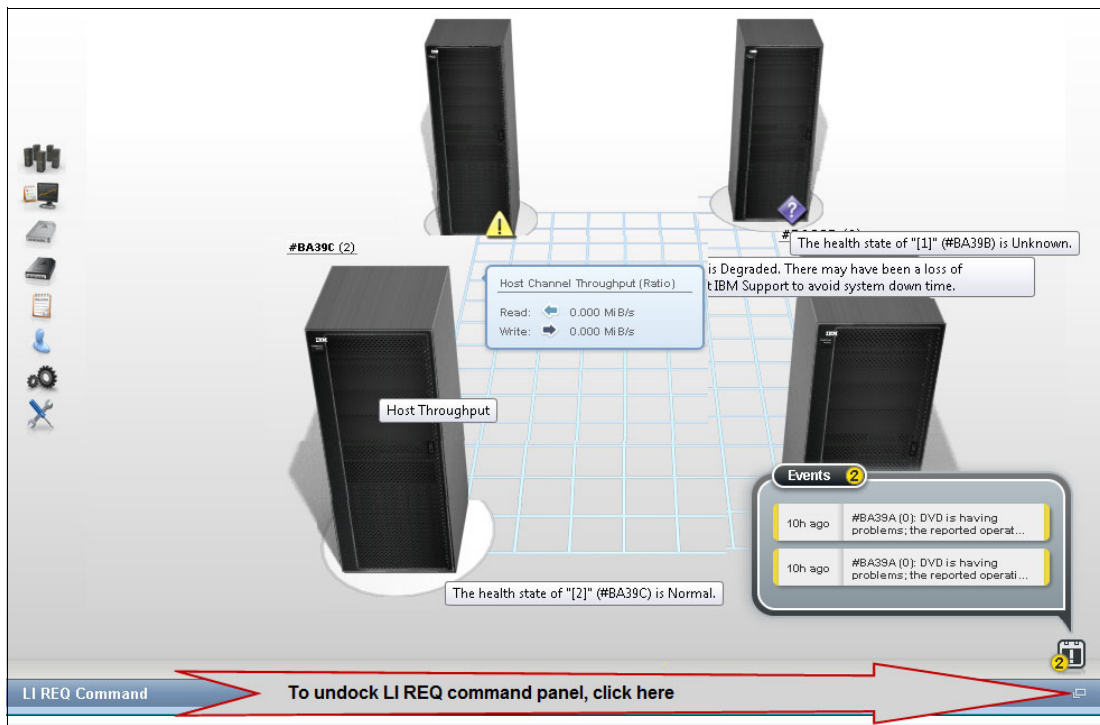


Figure 9-7 Status and Events indicators in the Grid Summary pane

All cluster indicators provide information for the accessing cluster only, and are displayed only on MI windows that have a cluster scope. These three indicators occur in the middle of the window footer and include the following information:

- ▶ Physical Cache
- ▶ Copy Queues
- ▶ Health Status

Figure 9-8 shows a Cluster Summary panel, and some examples of status, events, and messages that can be seen such as:

- ▶ Cache licensed capacity and total physical used capacity by cache partition
- ▶ Cluster health, name, model, code level, disk encryption status, and family information.
- ▶ Server model and health, cache model and health, Ethernet switches health
- ▶ Tape library model, capacity, and health

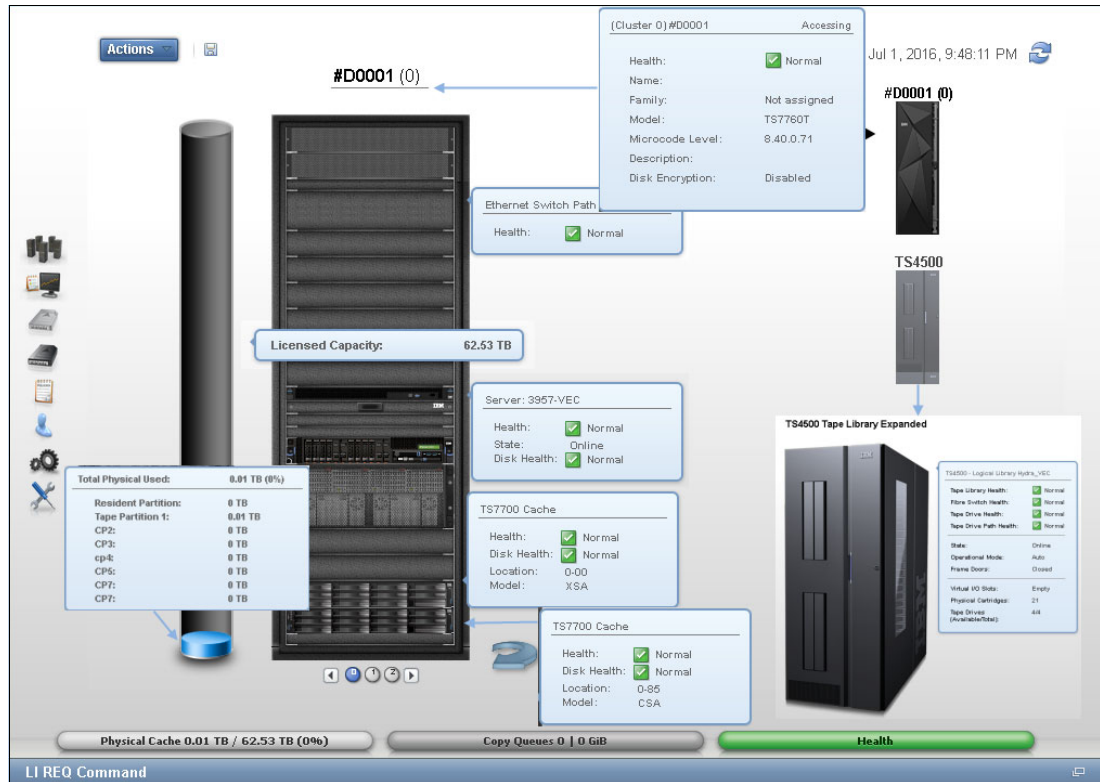


Figure 9-8 Cluster Summary panel and some available information

MI also provides ways to filter, sort, and change the presentation of different tables in the MI. For example, the user can hide or display a specific column, modify its size, sort the table results, or download the table row data in a comma-separated value (CSV) file to a local directory.

For a complete description of tasks, the behavior of health and status icons, and a description of how to optimize the table presentations, see the *Using the Management Interface* topic in the TS7700 R4.0 IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_ua_using_the_management_interface.html

Library Request Command window

The LI REQ command pane in the MI expands the interaction of the system administrator with the TS7700 subsystem. By using the LI REQ panel, a standard **LI REQ** command can be run by the Storage Administrator directly from the MI to a grid (also known as *Composite Library*), or to a specific Cluster (also known as *Distributed Library*), with no need to be logged in to the z/OS host system.

The LI REQ panel is minimized and docked at the bottom of the MI window. The user must only click it (at the lower right end) to open the LI REQ command pane. Figure 9-9 shows the new LI REQ command panel and operation.

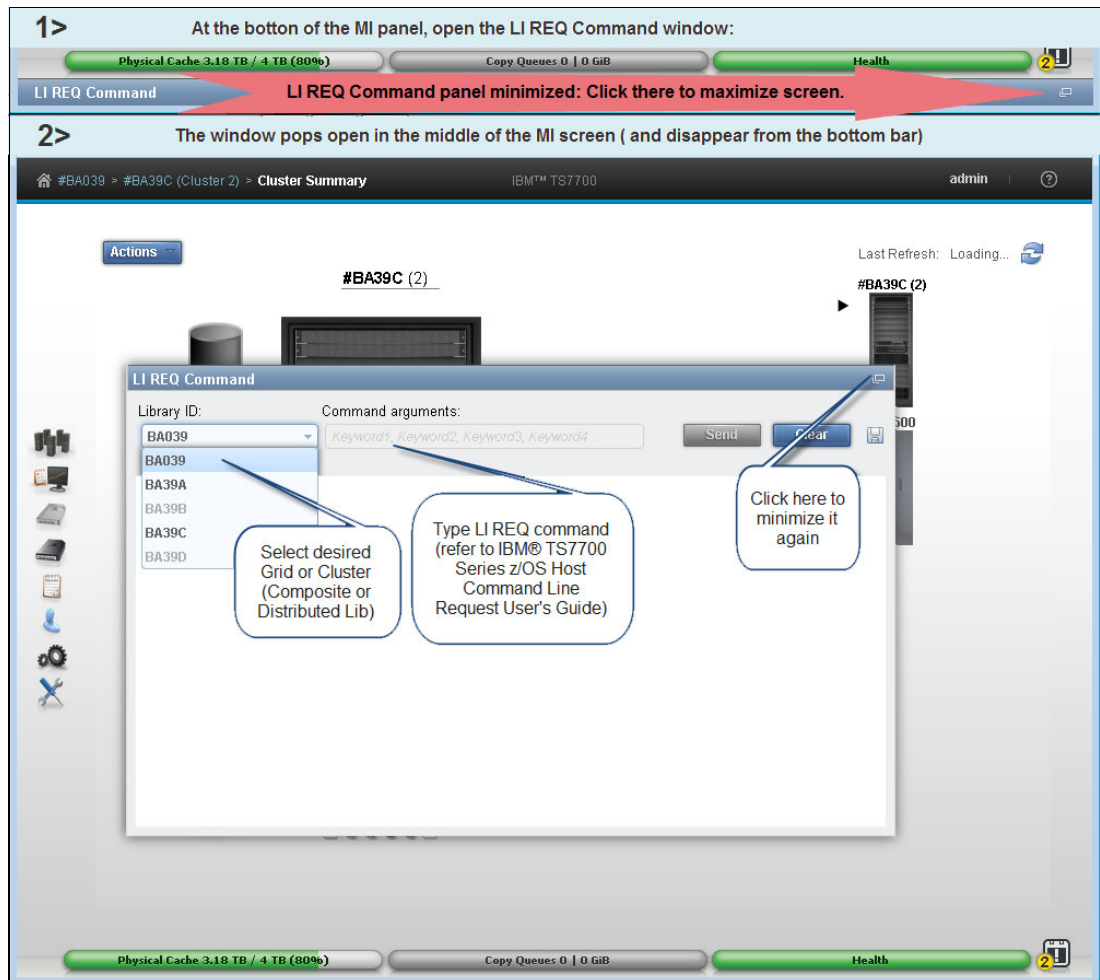


Figure 9-9 LI REQ Command window and usage

By default, the only user role that is allowed to run LI REQ commands is the Administrator. LI REQ commands are logged in to *tasks*.

Remember: The LI REQ option shows only in the bottom of the MI windows for users with the Administrator role, and is not displayed on the host console.

Figure 9-10 shows an example of a library request command reported in the Tasks list, and shows how to get more information about the command by selecting **Properties** clicking **See details** in the MI window.




Figure 9-10 LI REQ command log and information

Important: LI REQ commands that are issued from this window are not presented in the host console logs.

For a complete list of available LI REQ commands, their usage, and respective responses, see the current *IBM TS7700 Series z/OS Host Command Line Request User's Guide* (WP101091), found at:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101091>

Standard navigation elements

This section of the TS7700 MI provides functions to manage and monitor the health the TS7700. Listed next are the expandable interface windows that are shown on the left side of the MI Summary window. The exception is the systems window, which is displayed only when the cluster is part of a grid.

More items might also show, depending on the actual cluster configuration:

- Systems icon** This window shows the cluster members of the grid and grid-related functions.
- Monitor icon** This window gathers the events, tasks, and performance information about one cluster.
- Light cartridge icon** Information that is related to virtual volumes is available here.
- Dark cartridge icon** Information that is related to physical cartridges and the associated tape library are under this window.
- Notepad icon** This window contains the constructs settings.
- Blue man icon** Under the Access icon, all security-related settings are grouped.

Gear icon Cluster general settings, feature licenses, overrides, SNMP, write protect mode, and backup and restore settings are under the Gear icon.

Tool icon Ownership takeover mode, network diagnostics, data collection, and other repair/recovery-related activities are under this icon.

MI Navigation

Use this window (Figure 9-11) for a visual summary of the TS7700 MI Navigation.

TS7700 MI Navigation

- 1.0 Systems**
 - 1.1 Grid Summary (All Systems)
 - Action: Grid Identification Properties
 - Action: Lower removal Threshold
 - Action: Cluster Families
 - Display: Copy queue
 - Display: Host throughput
 - Display: Throttling
 - 1.2 Cluster Summary
 - Action: Cluster Identification Properties
 - Action: Change Cluster Status (Service / prep shutdown / Force Service)
 - Display: Physical Library
 - Display: Cache expansion frame (real layout)
 - Display: Cluster (real layout)
 - Display: Port view (cluster Nodes)
 - Display: FICON, Port Acceleration
 - Display: Tape Volume Cache
 - Display: All hardware and cluster specific info
- 2.0 Monitor**
 - 2.1 System (Cluster Summary)
 - 2.2 Events
 - 2.3 Performance
 - Page: Historical Summary
 - Page: Virtual Mounts
 - Page: Physical Mounts
 - Page: Host Throughput
 - Page: Cache throttling
 - Page: Cache Utilization
 - Page: Grid Network Throughput
 - Page: Pending Updates
 - 2.4 Tasks
- 3.0 Virtual**
 - 3.1 Incoming Copy Queue
 - 3.2 Recall Queue
 - 3.3 Virtual Tape Drives
 - 3.4 Virtual Volumes (all pages)
 - 3.5 Categories
- 4.0 Physical**
 - 4.1 Physical Volume Pools
 - 4.2 Physical Volumes (all pages)
 - 4.3 Physical Tape Drives
 - 4.4 Physical Media Inventory
- 5.0 Constructs**
 - 5.1 Storage Groups
 - 5.2 Management Classes
 - 5.3 Storage Classes
 - 5.4 Data Classes
- 6.0 Access**
 - 6.1 Security Settings
 - Direct LDAP
 - 6.2 Roles & Permissions
 - 6.3 SSL Certificates
 - 6.4 InfoCenter Settings
- 7.0 Settings**
 - 7.1 Cluster Network Settings
 - 7.2 Feature Licenses
 - 7.3 SNMP
 - 7.4 Library Port Access Groups
- 7.5 Cluster Settings**
 - Page: Copy Policy Override
 - Page: Encryption Key Servers
 - Page: Inhibit reclaim Schedule
 - Write Protect Mode
 - Backup Settings
 - Restore Settings
- 7.6 Copy Export Settings**
- 8.0 Service**
 - 8.1 Ownership Takeover Mode
 - 8.2 Repair Virtual Volumes
 - 8.3 Network Diagnostics
 - 8.4 Data Collection
 - 8.5 Copy Export Recovery pages (Standalone)
- 9.0 Banner**
 - 9.1 Help Items
 - Help
 - Learning and Tutorial
 - Info Center
 - About
 - 9.2 User Items
 - Log off
 - Change Password
 - 9.3 Breadcrumbs
- 10.0 Login**
 - Cluster name/ID
- 11.0 Status Pods**
 - Grid Scope Tasks and Events
 - Cluster Scope Cache capacity, Copy Queue and Health

Figure 9-11 TS7700 MI Navigation

9.2.3 The Systems icon

The TS7700 MI windows that are gathered under the **Systems** icon can help to identify quickly cluster or grid properties, and assess the cluster or grid “health” at a glance.

Tip: The **Systems** icon is only visible when the accessed TS7700 Cluster is part of a grid.

Grid Summary window

The Grid Summary window is the first window that opens in the web interface when the TS7700 is online, and the cluster that is currently being accessed by MI is part of a grid. This window can be used to quickly assess the health of all clusters in the grid, and as a starting point to investigate cluster or network issues.

Note: If the accessing cluster is a stand-alone cluster, the Cluster Summary window is shown upon login instead.

This window shows a summary view of the health of all clusters in the grid, including family associations, host throughput, and any incoming copy queue. Figure 9-12 shows an example of a Grid Summary window, including the pop-up windows.

Grid Summary window includes information on:

- ▶ Cluster throttling
- ▶ Host throughput rate (sampled before compression by host adapters within cluster)
- ▶ Copy queue size and type
- ▶ Running tasks and events

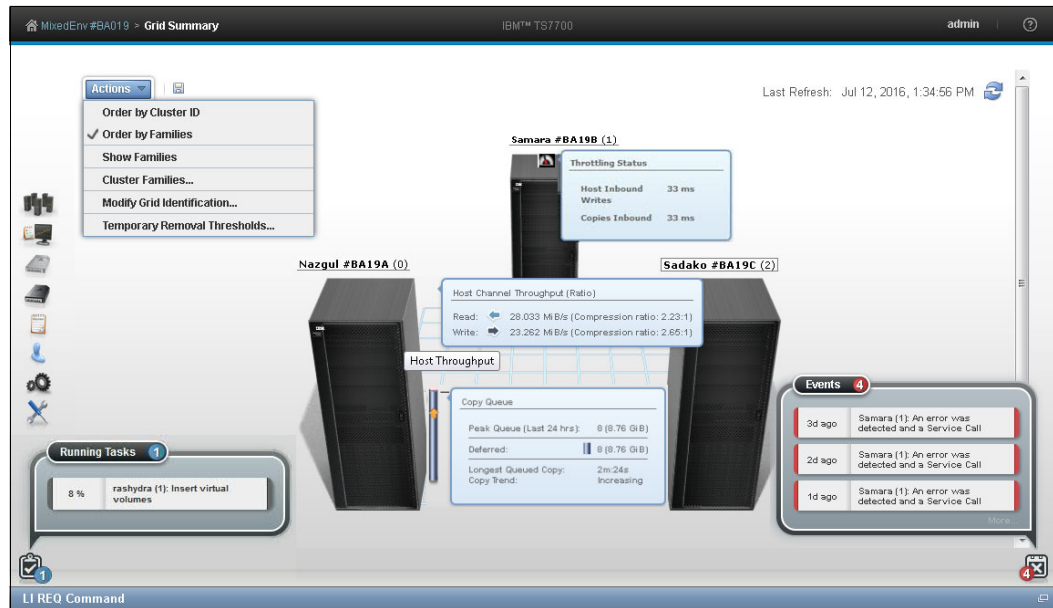


Figure 9-12 Grid Summary and pop-up windows

With R4.0, there is a diskette icon on the right of the Actions button. Clicking the icon saves a CSV-formatted file with a summary of the grid components information. See Figure 9-13 for an example of the grid summary spreadsheet.

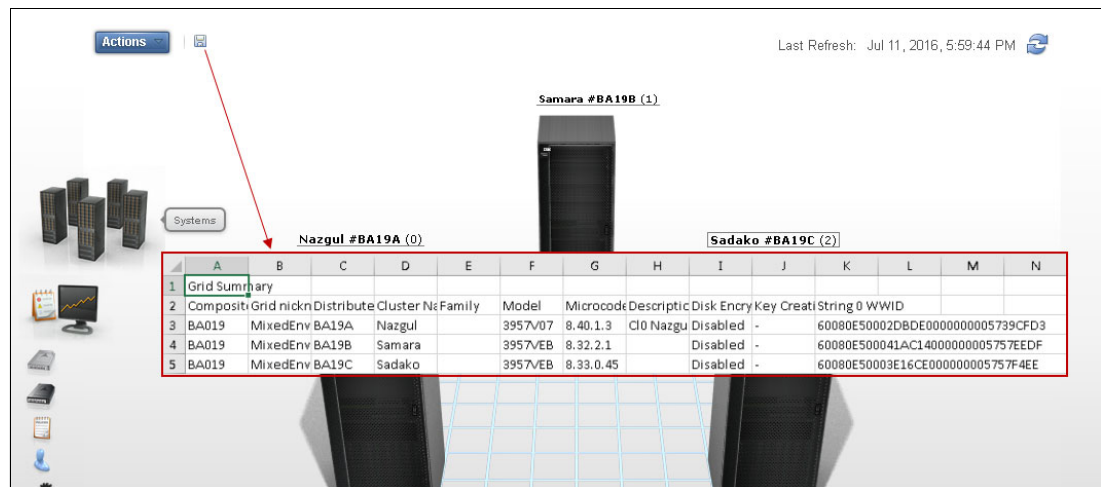


Figure 9-13 Grid information spreadsheet

Actions menu

Use this menu to change the appearance of clusters on the Grid Summary window or grid identification details. When the grid includes a disk-only cluster, this menu can also be used to change removal threshold settings for it or resident partitions (CPO) of tape-attached clusters. See Figure 9-14 for the Actions menu window.

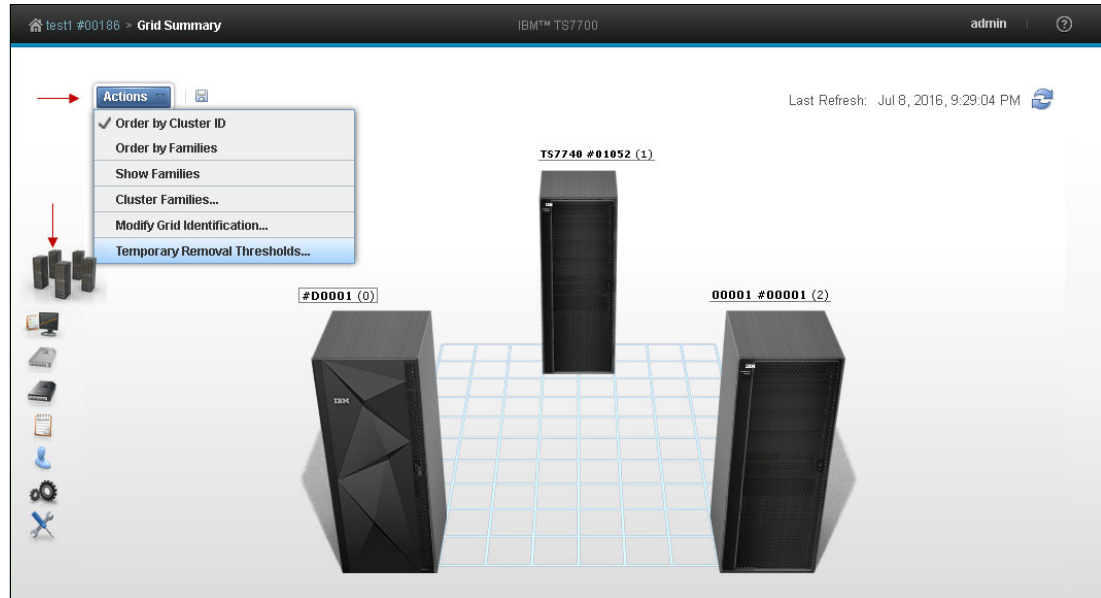


Figure 9-14 Grid Summary window and Actions list

The following tasks are on this menu:

- ▶ Order by Cluster ID

Select this option to group clusters according to their cluster ID number. Ordered clusters are shown first from left to right, then front to back. Only one ordering option can be selected at a time.

Note: The number that is shown in parentheses in breadcrumb navigation and cluster labels is always the cluster ID.

- ▶ Order by Families

Select this option to group clusters according to their family association.

- ▶ Show Families

Select this option to show the defined families on the grid summary window. Cluster families are used to group clusters in the grid according to a common purpose.

- ▶ Cluster Families

Select this option to add, modify, or delete cluster families used in the grid.

Cluster Families window

To view information and run actions that are related to TS7700 cluster families, use the window that is shown in Figure 9-15.

Data transfer speeds between TS7700 clusters sometimes vary. The cluster family configuration groups clusters so that Licensed Internal Code can optimize grid connection performance between the grouped clusters. You can read more about cluster family functions in 2.3.6, “Cluster family concept” on page 65.

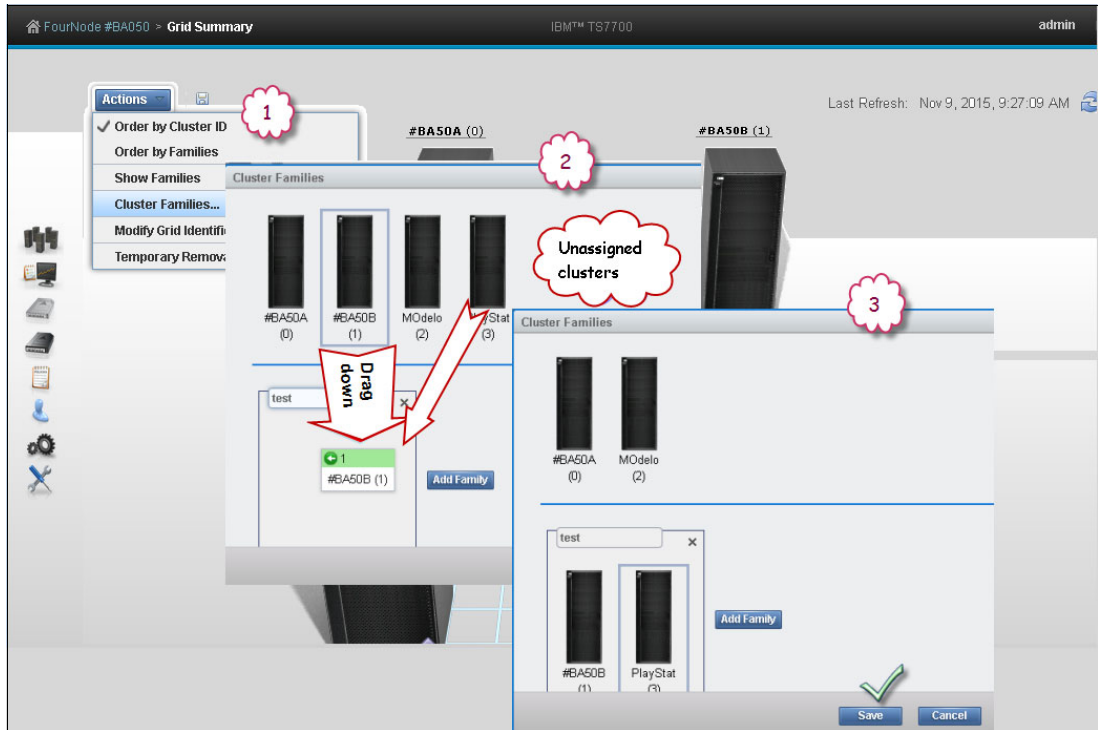


Figure 9-15 MI Add Cluster Families: Assigning a cluster to a family

To view or modify cluster family settings, first verify that these permissions are granted to the assigned user role. If the current user role includes cluster family permissions, select **Modify** to run the following actions:

- ▶ 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 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 by using the method that is described in the *Move a cluster* function, described next.
- ▶ Move a cluster: One or more clusters can be moved by dragging, 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 these methods:
 - Clicking the cluster with the mouse.
 - Using the Spacebar key on the keyboard.

- Pressing and holding the Shift key while selecting clusters to select multiple clusters at one time.
 - Pressing the Tab key on the keyboard to switch between clusters before selecting one.
- Move the selected cluster or clusters:
- Click and hold the mouse on the cluster, and drag the selected cluster to the destination cluster family or the Unassigned Clusters area.
 - Using the arrow keys on the keyboard to move the selected cluster or clusters right or left.

Consideration: An existing cluster family cannot be moved within the Cluster Families window.

- ▶ Delete a family: To delete an existing cluster family, click the **X** in the upper-right corner of the cluster family to delete it. If the cluster family to be deleted 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 that are made to the Cluster Families window and return it to read-only mode.

Remember: Each cluster family must contain at least one cluster. An attempt to save a cluster family that does not contain any clusters results in an error message. No changes are made, and the Cluster Families window remains in edit mode.

Grid Identification properties window

To view and alter identification properties for the TS7700 grid, use the window that is shown in Figure 9-16. In a multigrid environment, use this window to identify clearly a particular composite library, making it easier to distinguish, operate, and manage this TS7700 grid (avoiding operational mistakes due to ambiguous identification).

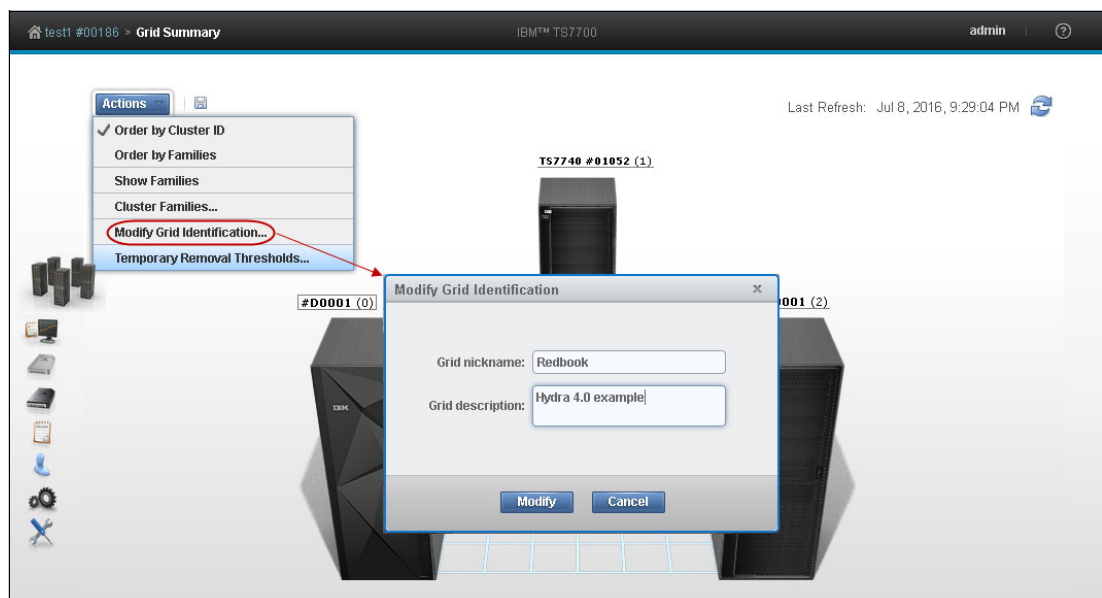


Figure 9-16 MI Grid Identification: Giving a name and description to the grid

The following information, related to grid identification, is displayed in Figure 9-16 on page 338. To change the grid identification properties, edit the available fields and click **Modify**. The following fields are available:

- ▶ Grid nickname: The grid nickname must be 1 - 8 characters in length and composed of alphanumeric characters with no spaces. The characters at (@), period (.), dash (-), and plus sign (+) are also allowed.
- ▶ Grid description: A short description of the grid. Up to 63 characters can be used.

Lower removal threshold

Select **Temporary Removal Threshold** from the **Actions** menu in the Grid summary view to lower the removal threshold for any disk-only cluster or cache resident partition of a tape attach cluster in the grid. For more information about removal policies, see 4.2.6, “TS7720 and TS7760 cache thresholds and removal policies” on page 154.

Figure 9-17 shows the Temporary Removal Threshold window.

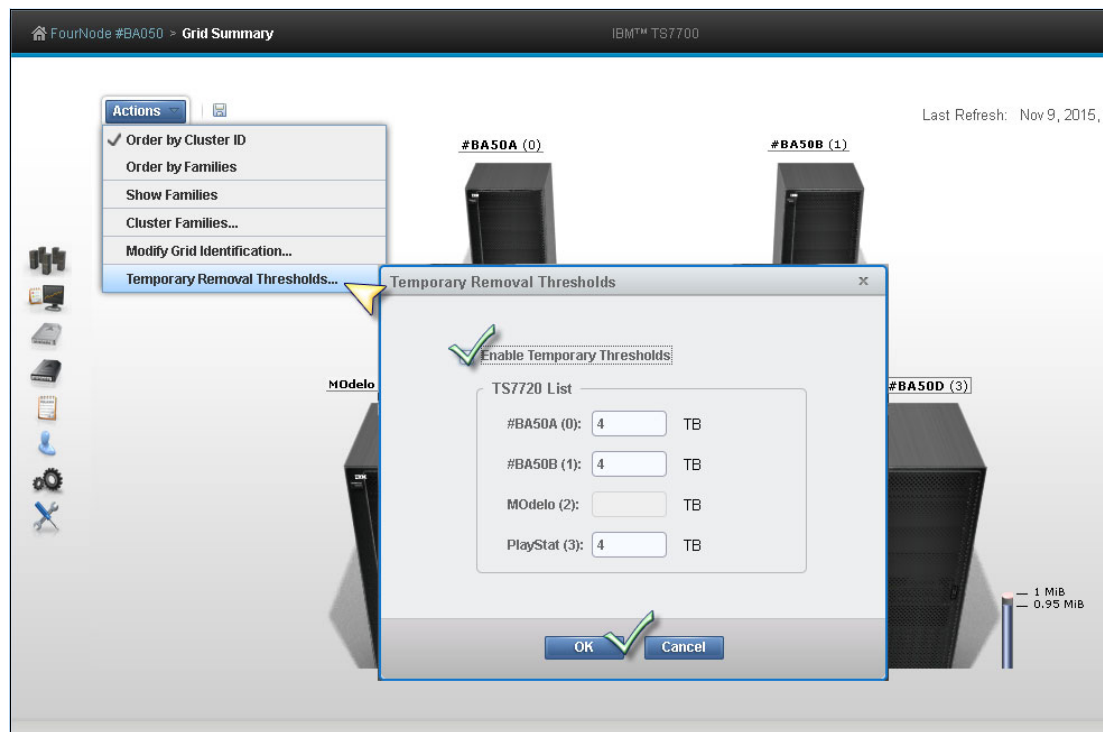


Figure 9-17 Setting the Temporary Removal Threshold

Grid health and details

In the Grid Summary view, cluster is in a normal state (healthy) when there is no warning or degradation icon that is displayed at the lower left side at the cluster’s representation in the MI. Hovering the mouse pointer over the lower right corner of the cluster’s picture in the Grid Summary window shows a message stating The health state of the [cluster number] [cluster name] is Normal, confirming that this cluster is in a normal state.

Exceptions in the cluster state are represented in the Grid Summary window by a little icon at the lower right side of the cluster’s picture. Use the main view of the Grid Summary window to compare the details and health status of all clusters in the grid. Additional information about the status can be viewed by hovering your cursor over the icon.

Figure 9-18 shows the appearance of the degraded icon, and the possible reasons for degradation to happen. The complete list of the icons and their meaning can be found at the TS7700 IBM Knowledge Center, which can be accessed from the MI window by hovering your cursor over the question mark symbol at the right side of the banner, and clicking IBM Knowledge Center. Alternatively, see the TS7700 4.0 IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_managing.html


Icon	Indicated health status
	<p>Warning or Degraded</p> <p>The connection between the grid and cluster is working but one of the cluster's redundant parts has stopped functioning. The following operational states can cause a degraded health status for a cluster:</p> <p>Out of Empty Stacked Volumes</p> <p>The cluster does not contain any empty scratch virtual volumes. Refer to HYDME1099W in Management interface messages for additional information.</p> <p>Copies Disabled by the System</p> <p>Copies are disabled on the cluster because the cluster is out of physical scratch volumes or out of cache space. Refer to HYDME1100W in Management interface messages for additional information.</p> <p>Copies Disabled by the Host</p> <p>Copies have been disabled using a Host Console Request. Refer to HYDME1101W in Management interface messages for additional information.</p> <p>Immediate Deferred Copies</p> <p>Copy operations on the cluster are deferred until after a rewind-unload operation occurs at the host. Refer to HYDME1102W in Management interface messages for additional information.</p> <p>All Storage Cells Full in Physical Library</p> <p>No empty I/O slots exist in a physical library attached to the cluster. Refer to HYDME1131W in Management interface messages for additional information.</p> <p>Limited Free Space in Cache</p> <p>Cache resources are limited and in danger of reaching maximum capacity. Refer to HYDME1132W in Management interface messages for additional information.</p> <p>Out of Cache Resources</p> <p>The cluster is out of cache space. Refer to HYDME1133W in Management interface messages for additional information.</p>

Figure 9-18 Warning or Degraded Icon meanings

The following list includes the other possible statuses for a cluster that can be found in the MI:

- ▶ Failed
- ▶ Service or Service Prep
- ▶ Unknown
- ▶ Offline
- ▶ Write Protect Mode

Figure 9-19 shows the icons.



Figure 9-19 Other cluster status icons

For more information, hover over the status icon with the mouse pointer. See the IBM TS7700 R4.0 IBM Knowledge Center, available locally in the MI window by clicking the question mark icon at the right of the banner.

For a complete list of icons and meanings, see the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7700_ua_grid_summary_main.html

In the Grid Summary pane, there is an indicator that indicates that throttling activity is occurring for a determined cluster within the grid. See Figure 9-20 for the visual reference of the throttling indicator. See Chapter 11, “Performance and monitoring” on page 635 for practical considerations on this topic, what it means, and what can be done to avoid it.

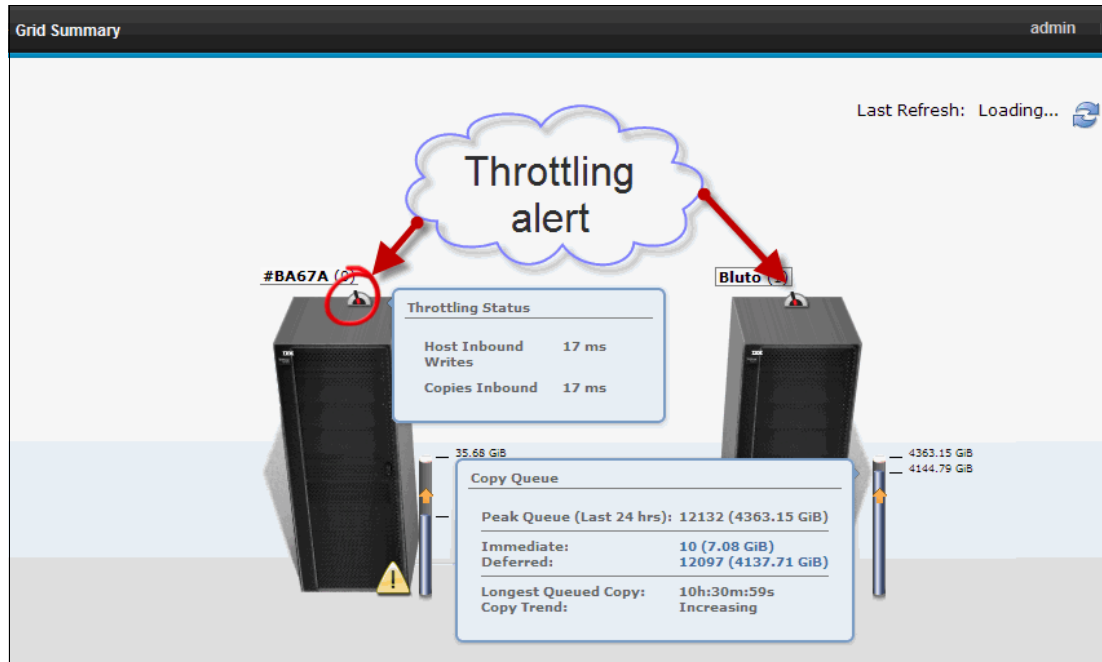


Figure 9-20 Clusters throttling in a two-cluster grid

Also, for an in-depth explanation of the throttling mechanism, where it is applied in the TS7700 and how it affects the subsystem performance, see the *IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance* white paper:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101465>

Cluster Summary window

By clicking the icon of an individual cluster in the grid, or by selecting a specific cluster in the cluster navigation element in the banner, the Cluster Summary window can be accessed. In a stand-alone configuration, this is the first icon available in the MI.

Figure 9-21 shows an example of the Cluster Summary window.

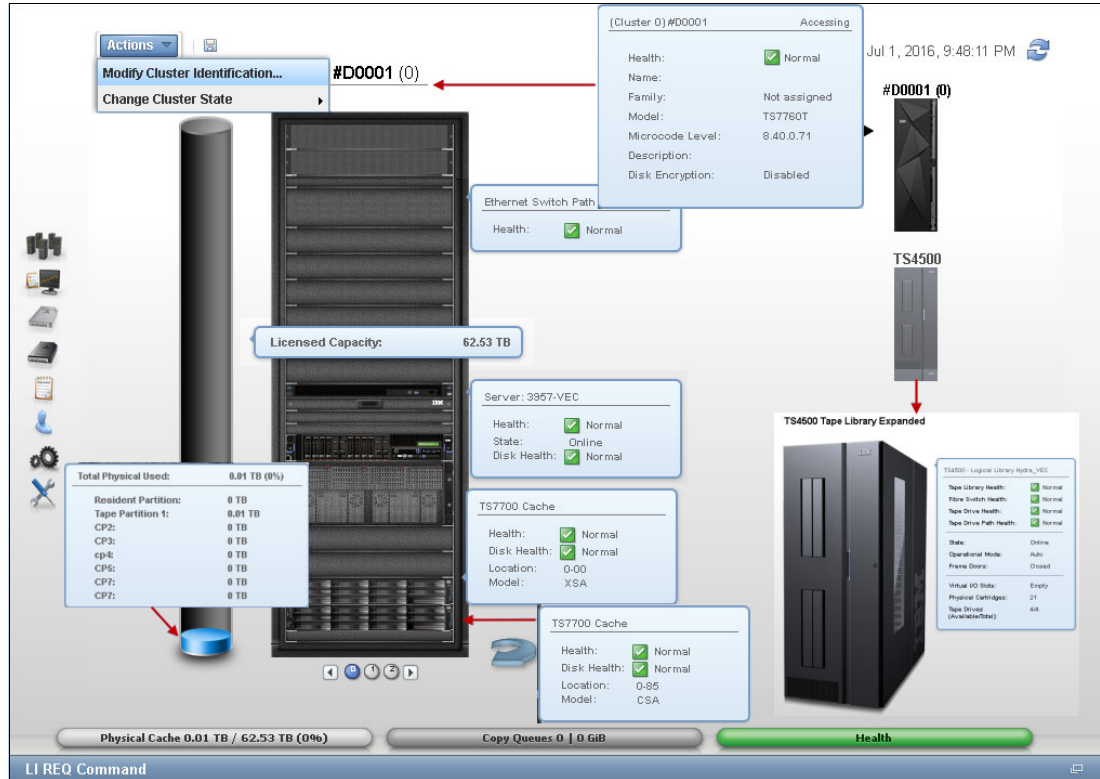


Figure 9-21 Cluster Summary window with TS7760T and TS4500 tape library

In the Cluster Summary window, the following options, under the **Actions** menu, are available:

- ▶ **Modify Cluster Information**
- ▶ **Change Cluster State** → **Force Shut Down**
- ▶ **Change Cluster State** → **Service Prep**

The Cluster Information can be displayed by hovering the cursor over the components, as shown in Figure 9-21. In the resulting box, the following information is available:

- ▶ Cluster components health status
- ▶ Cluster Name
- ▶ Family to which this cluster is assigned
- ▶ Cluster model
- ▶ Licensed Internal Code (LIC) level for this cluster
- ▶ Description for this cluster
- ▶ Disk encryption status
- ▶ Cache size and occupancy (Cache Tube)

With R4.0 there is a diskette icon to the right of the **Actions** button. Clicking the icon downloads a CSV-formatted file with the meaningful information about that cluster.

Figure 9-22 shows an example of the cluster summary spreadsheet.

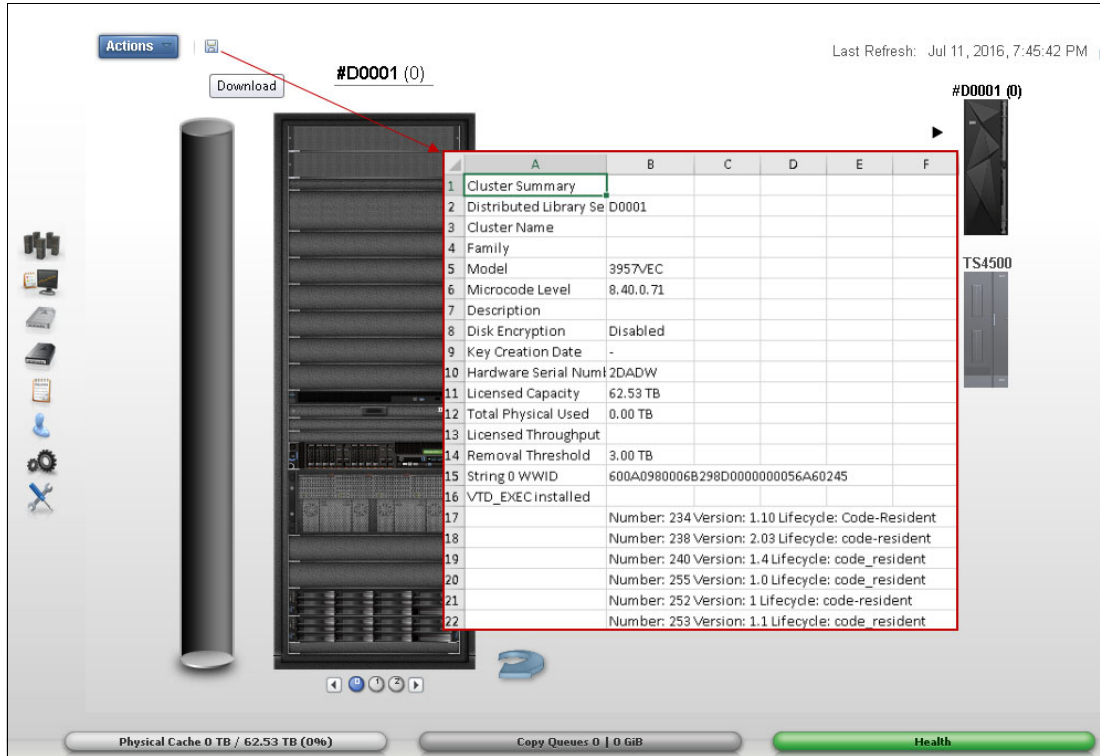


Figure 9-22 Downloading cluster information spreadsheet

Cluster Actions menu

By using the options under this menu, the user can change the state or settings of a cluster. Also, when the selected cluster is a tape attach TS7700 (a tape library is present), this menu can be used to change the Copy Export settings.

From the **Action** menu, the Cluster State can be changed to a different one to perform a specific task, such as preparing for a maintenance window, performing a disaster recovery drill, or moving machines to a different IT center. Depending on the current cluster state, different options display.

Table 9-2 describes options available to change the state of a cluster.

Table 9-2 Options to change the cluster state

If the current state is	Selection	Restrictions and notes
Online	Service Prep	All following conditions must be met first: <ul style="list-style-type: none"> ▶ The cluster is online. ▶ No other clusters in the grid are in service prep mode. ▶ At least one other cluster must remain online. Caution: If only one other cluster remains online, a single point of failure exists when this cluster state becomes service prep mode. Select Service Prep to confirm this change.
	Force Shutdown	Select Force Shutdown to confirm this change. Important: After a shutdown operation is initiated, it cannot be canceled.

If the current state is	Selection	Restrictions and notes
Service Pending	Force Service	Use this option when an operation stalls and is preventing the cluster from entering Service Prep. Select Force Service to confirm this change. All but one cluster in a grid can be placed into service mode, but it is advised that only one cluster be in service mode at a time. If more than one cluster is in service mode, and service mode is canceled on one of them, that cluster does not return to normal operation until service mode is canceled on <i>all</i> clusters in the grid.
	Return to Normal	Select this option to cancel a previous service prep change and return the cluster to the normal online state. Select Return to Normal to confirm this change.
	Force Shutdown	Select Force Shutdown to confirm this change. Important: After a shutdown operation is initiated, it cannot be canceled.
Shutdown (offline)	User interface not available	After an offline cluster is powered on, it attempts to return to normal. If no other clusters in the grid are available, skip hot token reconciliation can be tried.
Online-Pending or Shutdown Pending	Menu disabled	No options to change state are available when a cluster is in a pending state.

Going offline and coming online considerations

Whenever a member cluster of a grid goes offline or comes back online, it needs to exchange information with other peer members regarding the status of the logical volumes that are controlled by the grid. Each logical volume is represented by a so-called *token*, which contains all of the pertinent information regarding that volume, such as creation date, whose cluster it belongs to, which cluster is supposed to have a copy of it, and what kind of copy it should be.

Each cluster in the grid keeps its own copy of the collection of tokens, representing all of the logical volumes that exist in grid, and those copies are kept updated at the same level by the grid mechanism. When coming back online, a cluster needs to reconcile its own collection of tokens with the peer members of the grid, making sure that it represents the status of the grid inventory. This reconcile operation is also referred to as *token merge*.

Here are some items to consider when going offline and coming online:

► Pending token merge

A cluster in a grid configuration attempts to merge its token information with all of the other clusters in the grid as it goes online. When no other clusters are available for this merge operation, the cluster attempting to go online remains in the going online, or blocked, state indefinitely as it waits for the other clusters to become available for the merge operation. If a pending merge operation is preventing the cluster from coming online, there is an option to skip the merge step.

Click **Skip Step** to skip the merge operation. This button is only available if the cluster is in a blocked state while waiting to share pending updates with one or more unavailable clusters. When you click **Skip Step**, pending updates against the local cluster might remain undetected until the unavailable clusters become available.

► Ownership takeover

If ownership takeover was set at any of the peers, the possibility exists that old data can surface to the host if the cluster is forced online. Therefore, before attempting to force this cluster online, it is important to know whether any peer clusters have ever enabled ownership takeover mode against this cluster while it was unavailable. In addition, if this cluster is in service, automatic ownership takeover from unavailable peers is also likely and must be considered before attempting to force this cluster online.

If multiple clusters were offline and must be forced back online, force them back online in the reverse order that they went down in (for example, the last cluster down is the first cluster up). This process ensures that the most current cluster is available first to educate the rest of the clusters forced online.

► Autonomic Ownership Takeover Manager (AOTM)

If it is installed and configured, it attempts to determine whether all unavailable peer clusters are 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 in a failed state, the forced online cluster can be taking ownership of volumes that it need not take ownership of. If AOTM discovers that all unavailable peers failed and network issues are not to blame, this cluster is then forced into an online state.

After it is online, AOTM can further enable ownership takeover against the unavailable clusters if the AOTM option is enabled. Additionally, manual ownership takeover can be enabled, if necessary.

► Shutdown restrictions

To shut down a cluster, it is necessary to be logged in to this system. To shut down another cluster, log out of the current cluster and log in to the cluster to shut down. For more information, see “Cluster Shutdown window” on page 348.

Note: After a **shutdown** or **force shutdown** action, the targeted cluster (and associated cache) are powered off. A manual intervention is required on site where the cluster is physically located to power it up again.

A cluster shutdown operation that is started from the TS7700 MI also shuts down the cache. The cache must be restarted before any attempt is made to restart the TS7700 cluster.

Service mode window

Use the window that is shown in Figure 9-23 to put a TS7700 cluster into service mode, whenever required by a service action or any disruptive activity on a cluster that is a member of a grid. See Chapter 2, “Architecture, components, and functional characteristics” on page 15 for more information.

Remember: Service mode is only possible for clusters that are members of a grid.

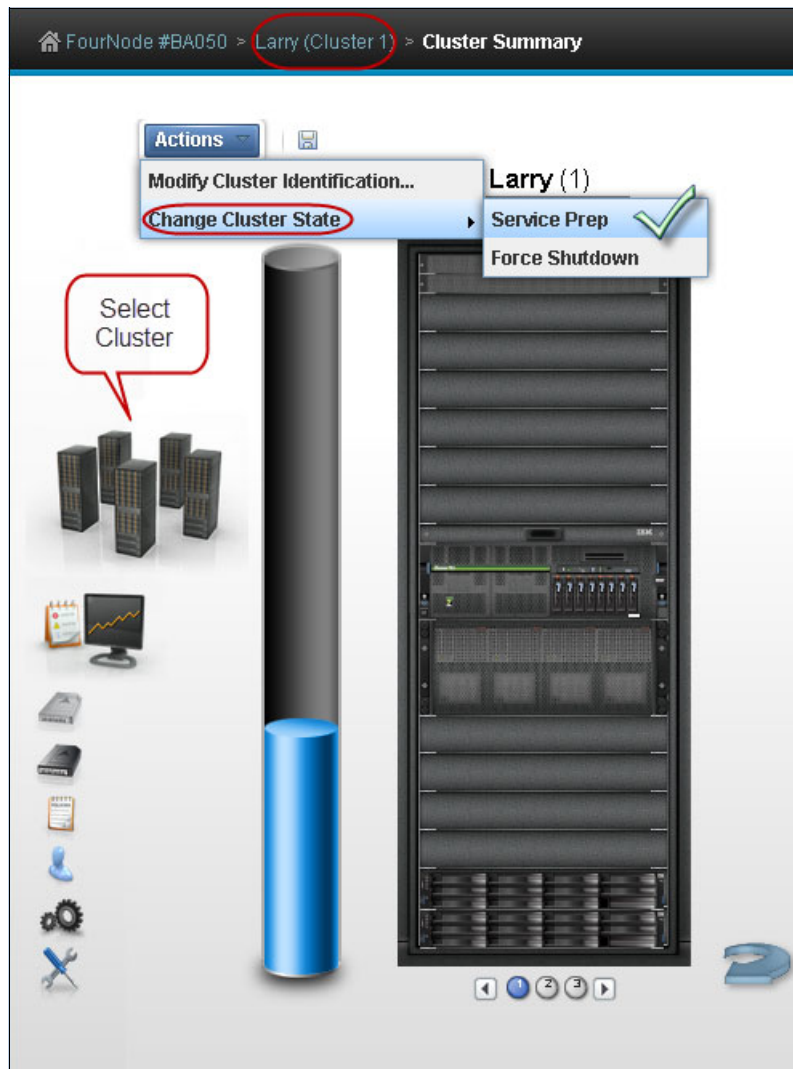


Figure 9-23 Cluster Summary: Preparing for service

Service mode enables the subject cluster to leave the grid gracefully, surrendering the ownership of its logical volumes as required by the peer clusters in the grid to attend to the tasks being performed by client. The user continues operating smoothly from the other members of the grid automatically, if consistent copies of volumes that reside in this cluster also exist elsewhere in the grid, and the host also has access to those clusters.

Before changing a cluster state to Service, the user needs to vary offline all logical drives that are associated to this cluster, and vary on logical drives on the remaining cluster. These actions ensure that the host has mounting points to continue operation. No host access is available in a cluster that is in service mode.

Important: Forcing Service Mode causes jobs that are currently mounted or that use resources that are provided by targeted cluster to fail.

Whenever a cluster state is changed to Service, it enters first in service preparation mode, and then, when the preparation stage finishes, it goes automatically into service mode.

During the service preparation stage, the cluster monitors the status of current host mounts, sync copy mounts targeting local Tape Volume Cache (TVC), monitors and finishes up the copies that are currently in execution, and makes sure that there are no remote mounts targeting local TVC. When all running tasks have ended, and no more pending activities are detected, the cluster finishes the service preparation stage and enters Service mode.

In a TS7700 grid, service preparation can occur on only one cluster at any one time. If service prep is attempted on a second cluster before the first cluster has entered in Service mode, the attempt will fail. After service prep has completed for one cluster, and that cluster has entered in service mode, 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. Although all clusters except one can be in Service mode at the same time within a grid, the preferred approach is having only one cluster in service mode at a time.

Be aware that when multiple clusters are in service mode simultaneously, *they need to be brought back to Normal mode at the same time*. Otherwise, the TS7700 will not get to ONLINE state, waiting until the remaining clusters also leave service mode. Only then, those clusters merge their tokens and rejoin the grid as ONLINE members.

Remember: If more than one cluster is in Service mode, and service is canceled on one of them, that cluster does not return to an online state until Service mode is canceled on all other clusters in this grid.

For a disk-only TS7700 cluster or CP0 partition in a grid, click **Lower Threshold** to lower the required threshold at which logical volumes are removed from cache in advance. See “Temporary removal threshold” on page 157 for more information about the Temporary Removal Threshold. The following items are available when viewing the current operational mode of a cluster.

Cluster State can be any of the following states:

- ▶ **Normal:** The cluster is in a normal operation state. Service prep can 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 the mode that the cluster is in, a different action is presented by the button under the Cluster State display. This button can be used to place the TS7700 into service mode or back into normal mode:

- ▶ **Prepare for Service Mode:** This option puts the cluster into service prep mode and enables the cluster to finish all current operations. If allowed to finish service prep, the cluster enters Service mode. This option is only available when the cluster is in normal mode. To cancel service prep mode, click **Return to Normal Mode**.

- ▶ **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 **Return to Normal Mode** is selected.

A window opens to confirm the decision to change the Cluster State. Click **Service Prep** or **Normal Mode** to change to new Cluster State, or **Cancel** to abandon the change operation.

Cluster Shutdown window

Use the window that is shown in Figure 9-24 to shut down remotely a TS7700 cluster for a planned power outage or in an emergency.

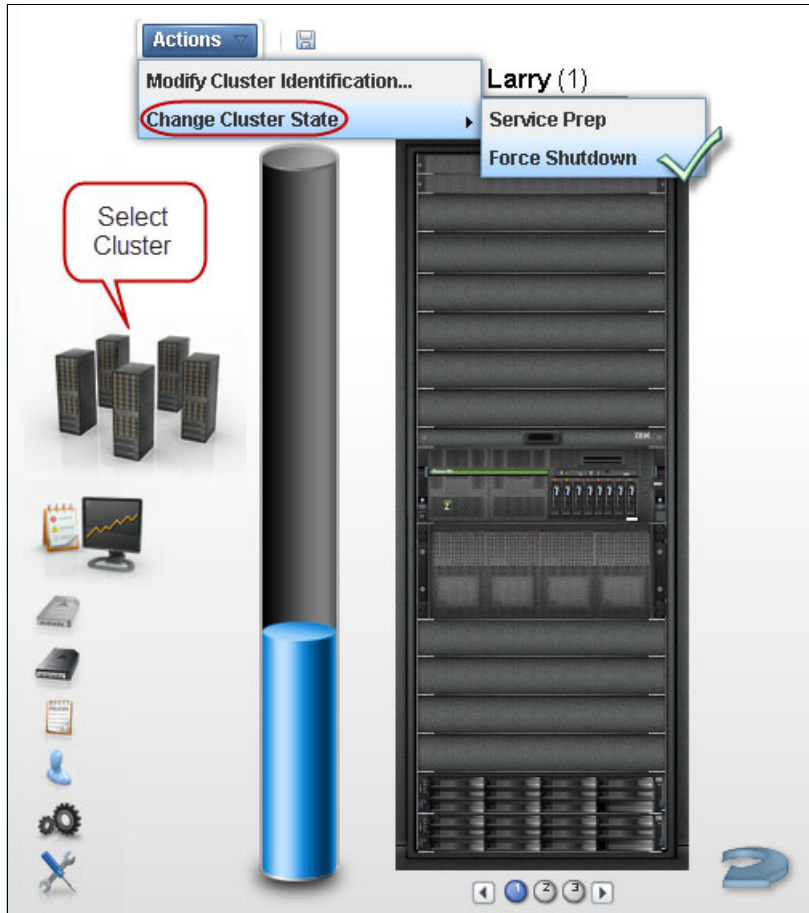


Figure 9-24 MI Cluster: Forcing a cluster to shutdown

This window is visible from the TS7700 MI whether the TS7700 is online or in service. If the cluster is offline, MI is not available, and the error HYDME0504E The cluster you selected is unavailable is presented.

Note: After a **shutdown** or **force shutdown** action, the targeted cluster (and associated cache) are powered off. A manual intervention is required on the site where the cluster is physically located to power it up again.

Only the cluster where a connection is established can be shut down by the user. To shut down another cluster, drop the current cluster connection and log in to the cluster that must be shut down.

Before the TS7700 can be shut down, decide whether the circumstances provide adequate time to perform a clean shutdown. A clean shutdown is not mandatory, but it is suggested for members of a TS7700 grid configuration. A clean shutdown requires putting the cluster in Service mode first. Make sure that no jobs or copies are targeting or being sourced from this cluster during shutdown.

Jobs that use this specific cluster are affected, but also copies are canceled. Eligible data that has not yet been copied to remaining clusters cannot be processed during service and downtime. If the cluster cannot be placed in Service mode, use the **force shutdown** option.

Attention: A forced shutdown can result in lost access to data and job failure.

A cluster shutdown operation that is started from the TS7700 MI also shuts down the cache. The cache must be restarted before any attempt is made to restart the TS7700.

If **Shutdown** is selected from the action menu for a cluster that is still online, as shown at the top of Figure 9-24 on page 348, a message alerts the user to put the cluster in service mode first before shutting down, as shown in Figure 9-25.

Note: For normal situations, set the cluster into service mode before shutdown is always recommendable.

It is still possible to force a shutdown without going into service by entering the password and clicking the **Force Shutdown** button if needed (for example, during a DR test to simulate a cluster failure. In this case, placing a cluster in service does not apply).

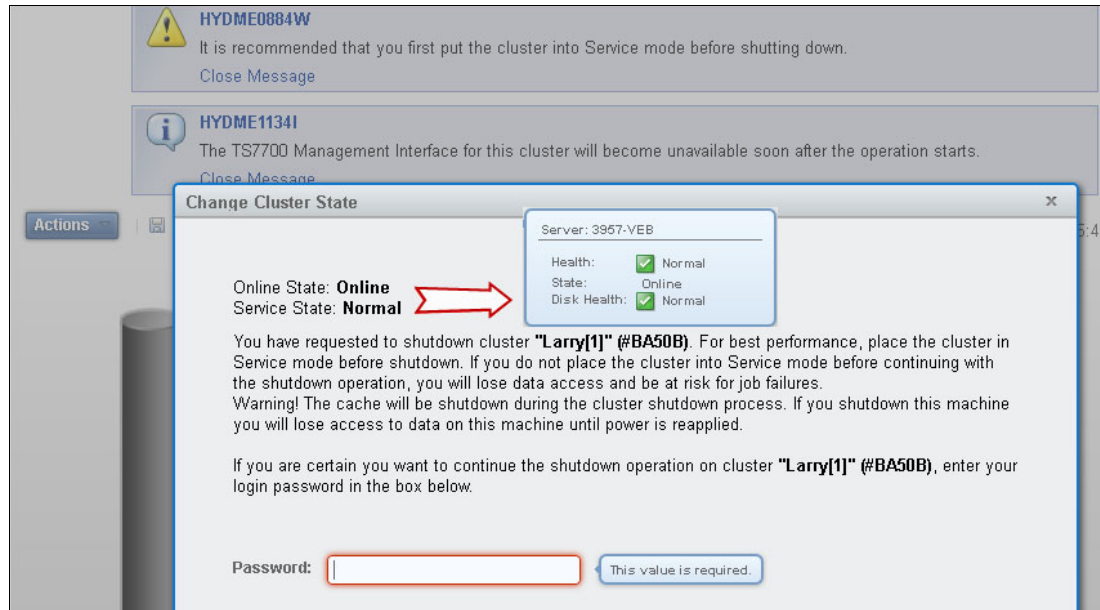


Figure 9-25 Warning message and Cluster Status during forced shutdown

In Figure 9-25, the Online State and Service State fields in the message show the operational status of the TS7700 and appear over the button that is used to force its shutdown. The lower-right corner of the picture shows the cluster status that is reported by the message.

The following options are available:

- ▶ **Cluster State.** The following values are possible:
 - **Normal.** The cluster is in an online, operational state and is part of a TS7700 grid.
 - **Service.** The cluster is in service mode or is a stand-alone system.
 - **Offline.** The cluster is offline. It might be shutting down in preparation for service mode.
- ▶ **Shutdown.** This button initiates a shutdown operation:
 - Clicking **Shutdown** in Normal mode. If **Shutdown** is selected while in normal mode, a warning message suggesting that you set the cluster to Service mode before proceeding opens, as shown in Figure 9-25 on page 349. To place the cluster in service mode, select **Modify Service Mode**. To continue with the force shutdown operation, provide the password and click **Force Shutdown**. To abandon the shutdown operation, click **Cancel**.
 - Clicking **Shutdown** in Service mode. When selecting **Shutdown** while in Service mode, you are prompted for a confirmation. Click **Shutdown** to continue, or click **Cancel** to abandon the shutdown operation.

Important: After a shutdown operation is initiated, it cannot be canceled.

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 following list shows the shutdown sequence:

1. Going offline
2. Shutting down
3. Powering off
4. Shutdown completes

Verify that power to the TS7700 and to the cache is shut down before attempting to restart the system.

A cluster shutdown operation that is started from the TS7700 MI also shuts down the cache. The cache must be restarted first and allowed to achieve an operational state before any attempt is made to restart the TS7700.

Cluster Identification Properties window

To view and alter cluster identification properties for the TS7700, use the window that is shown in Figure 9-26 on page 351, which can be used to distinguish this distributed library.

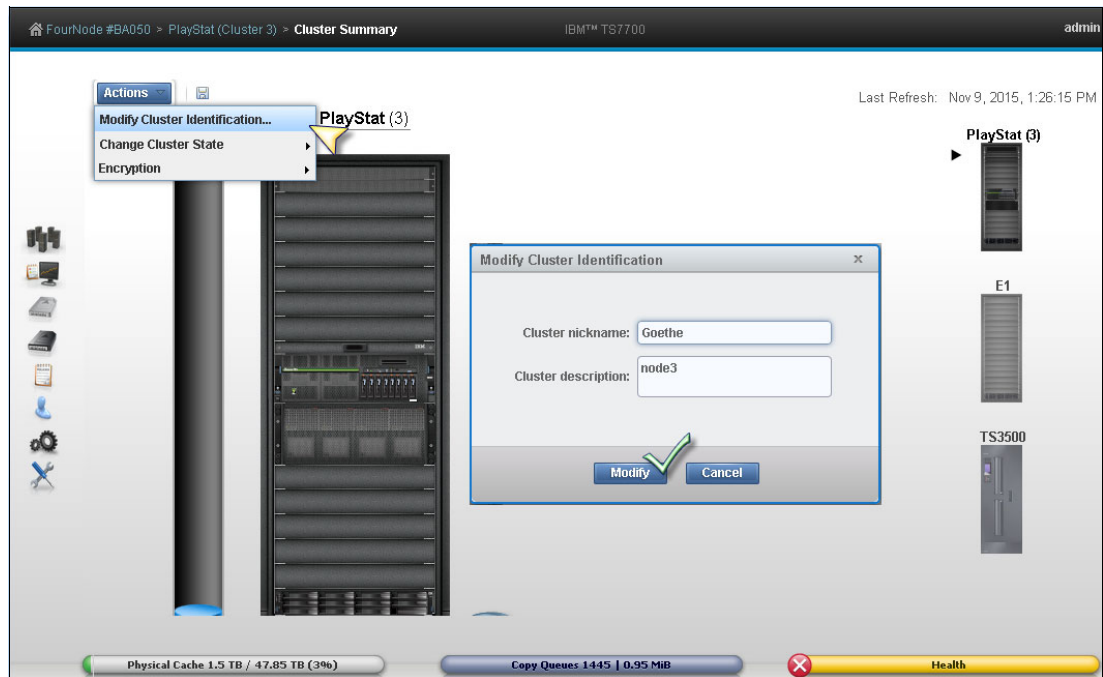


Figure 9-26 MI Cluster: Modifying cluster identification information

The following information that is related to cluster identification is displayed. To change the cluster identification properties, edit the available fields and click **Modify**. The following fields are available:

- ▶ **Cluster nickname:** The cluster nickname must be 1 - 8 characters in length and composed of alphanumeric characters. Blank spaces and the characters at (@), period (.), dash (-), and plus sign (+) are also allowed. Blank spaces cannot be used in the first or last character position.
- ▶ **Cluster description:** A short description of the cluster. Up to 63 characters can be used.

Note: Copy and paste might bring in invalid characters. Manual input is preferred.

Cluster health and detail

The health of the system is checked and updated automatically from time to time by the TS7700. The information status that is reflected on this window is not in real time; it shows the status of the last check-out. To repopulate the summary window with the updated health status, click the **Refresh** icon. This operation takes some minutes to complete. If this cluster is operating in Write Protect Mode, a lock icon is shown in the middle right part of the cluster image.

Figure 9-27, Cluster Summary page, depicts a TS7740-V07 with a TS3500 tape library attached. Within cluster front view page, cluster badge (top of the picture) brings a general description about the cluster, such as model, name, family, Licensed Internal Code level, cluster description, and cache encryption status. Refer to Figure 9-21 for the same MI page with the TS7760 with TS4500 tape library attached.

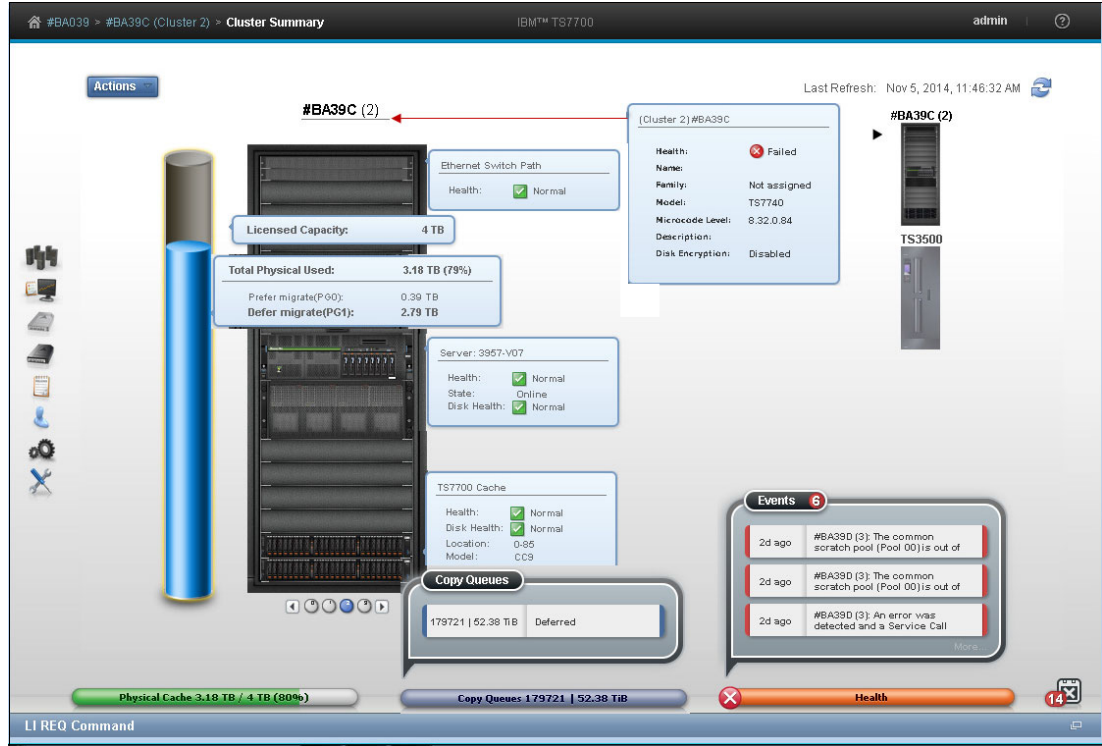


Figure 9-27 Cluster Summary: Front view of a TS3500-attached TS7740 and health details

Hovering the cursor over the locations within the picture of the frame shows the health status of different components, such as the network gear (at the top), TVC controller and expansion enclosures (bottom and halfway up), and the engine server along with the internal 3957-Vxx disks (the middle section). The summary of cluster health shows at the lower-right status bar, and also at the badge health status (over the frame).

Figure 9-28 shows an example of the cache tube display in a multi-partitioned TS7720T (Tape Attach).




Figure 9-28 Display of the cache tube in a multi-partitioned TS7720T

Figure 9-29 on page 354 shows the back view of the cluster summary window and health details. The components that are depicted in the back view are the Ethernet ports and host Fibre Channel connection (FICON) adapters for this cluster. Under the Ethernet tab, the user can see the ports that are dedicated to the internal network (the TSSC network) and those that are dedicated to the external (client) network. The assigned IP addresses are displayed. Details about the ports are shown (IPv4, IPv6, and the health). In the grid Ethernet ports, information about links to the other clusters, data rates, and cyclic redundancy check (CRC) errors are displayed for each port in addition to the assigned IP address and Media Access Control (MAC) address.

The host FICON adapter information is displayed under the Fibre tab for a selected cluster, as shown in Figure 9-29. The available information includes the adapter position and general health for each port.

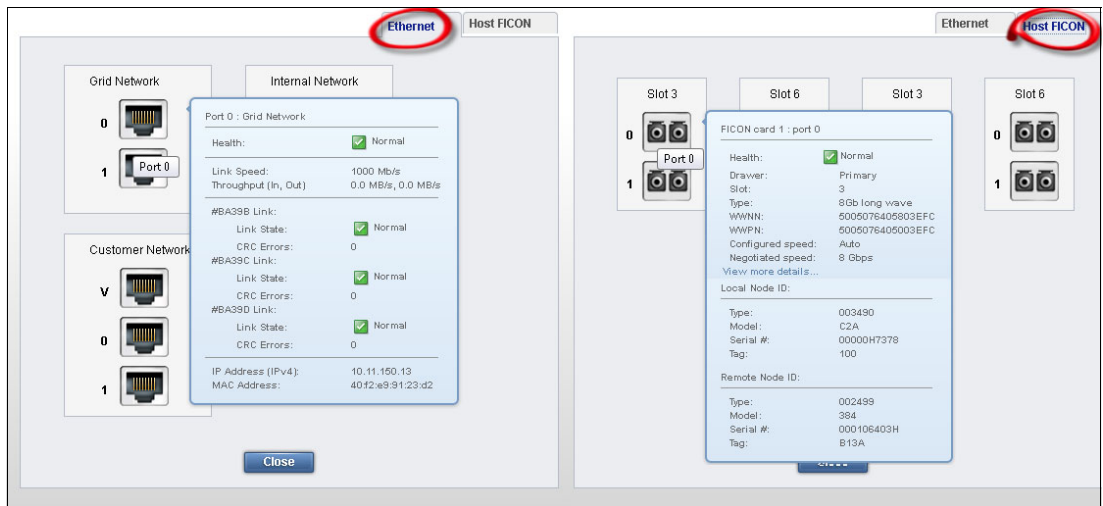


Figure 9-29 Back view of the cluster summary with health details

To display the different area health details, hover the cursor over the component in the picture.

Cache expansion frame

The expansion frame view displays details and health for a cache expansion frame that is attached to the TS7720 cluster. To open the expansion frame view, click the small image corresponding to a specific expansion frame below the **Actions** button.

Tip: The expansion frame icon is only displayed if the accessed cluster has an expansion frame.

Figure 9-30 shows the Cache Expansion frame details and health view through the MI.

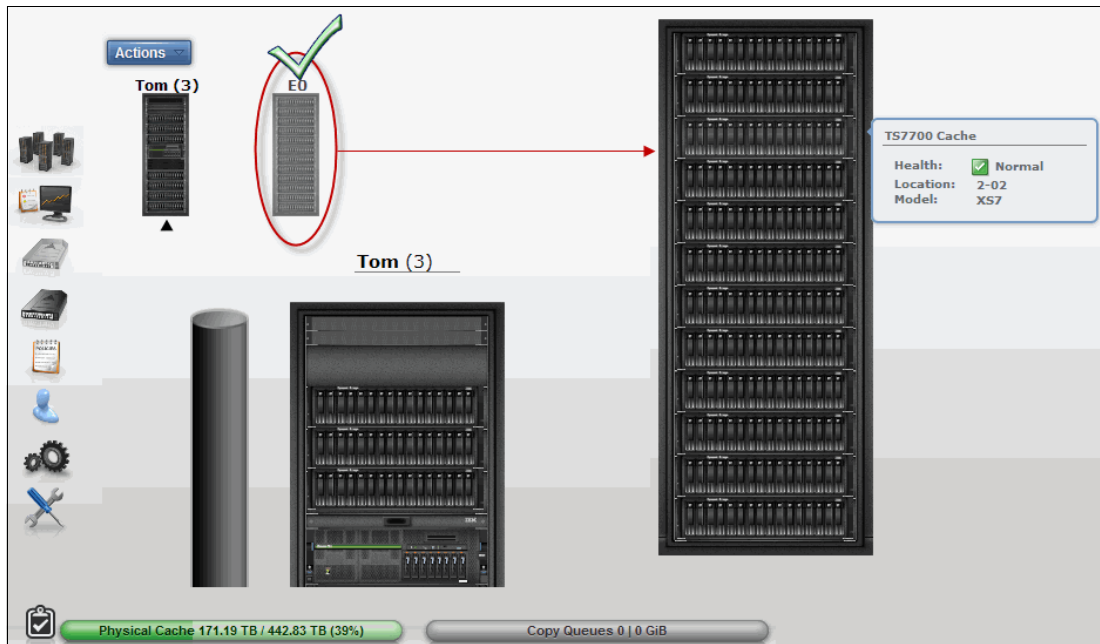


Figure 9-30 Cache expansion frame details and health

Physical library and tape drive health

Click the physical tape library icon, which is shown on a TS7700 tape-attached Cluster Summary window, to check the health of the tape library and tape drives. Figure 9-31 shows a TS4500 tape library that is attached to a TS7760 cluster.

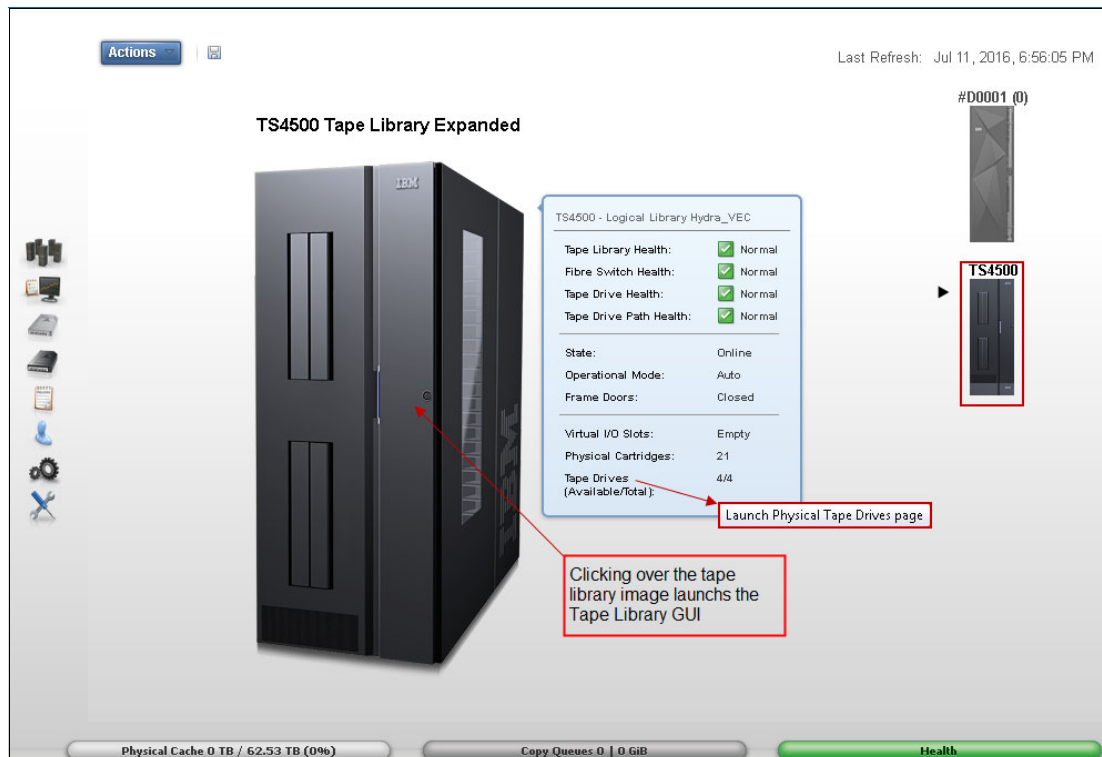


Figure 9-31 TS4500 tape library expanded page and links

Consideration: If the cluster is not a tape-attached model, the tape library icon does not display on the TS7700 MI.

The library details and health are displayed as explained in Table 9-3.

Table 9-3 Library health details

Detail	Definition
Physical library type - virtual library name	The type of physical library (type is always TS3500) accompanied by the name of the virtual library established on the physical library.
Tape Library Health Fibre Switch Health Tape Drive Health	The health states of the library and its main components. The following values are possible: <ul style="list-style-type: none"> ▶ Normal ▶ Degraded ▶ Failed ▶ Unknown
State	Whether the library is online or offline to the TS7700.
Operational Mode	The library operational mode. The following values are possible: <ul style="list-style-type: none"> ▶ Auto ▶ Paused
Frame Door	Whether a frame door is open or closed.
Virtual I/O Slots	Status of the I/O station that is used to move cartridges into and out of the library. The following values are possible: <ul style="list-style-type: none"> ▶ Occupied ▶ Full ▶ Empty
Physical Cartridges	The number of physical cartridges assigned to the identified virtual library.
Tape Drives	The number of physical tape drives available, as a fraction of the total. Click this detail to open the Physical Tape Drives window.

The Physical Tape Drives window can be accessed from the tape library expanded window, as shown in Figure 9-31 on page 355. Click the **Tape Drives** item in the health report, to see the Physical Tape Drives page.

The Physical Tape Drives window looks similar to the example in Figure 9-32.

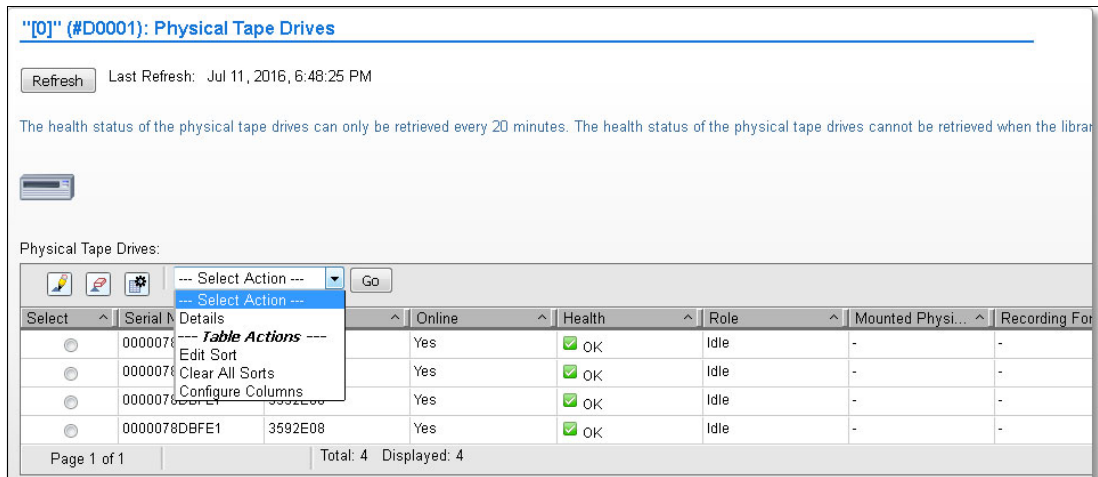


Figure 9-32 Physical Tape Drives window

The Physical Tape Drives window shows all the specific details about a physical tape drive, such as its serial number, drive type, whether the drive has a cartridge mount on it, and for what is it mounted. To see the same information, such as drive encryption and tape library location, about the other tape drives, select a specific drive and choose **Details** in the **Select Action** menu. The detailed drive information window is shown in Figure 9-33.

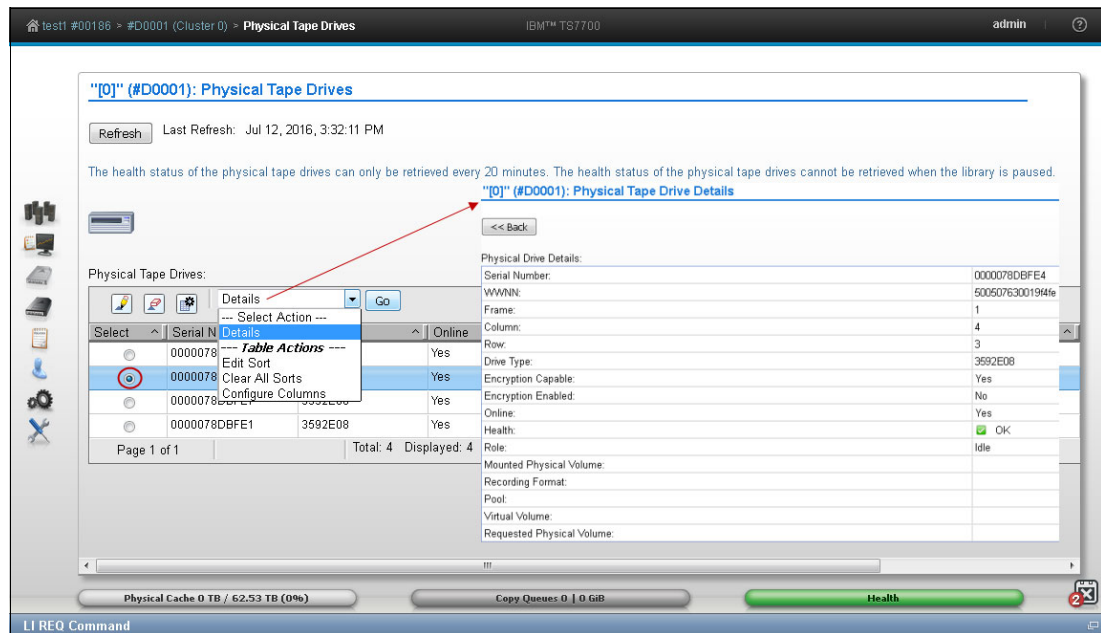


Figure 9-33 Physical Tape Drive Details and navigation

9.2.4 The Monitor icon

The collection of items under the **Monitor** icon in the MI provides means to monitor tasks, events, and performance statistics within the TS7700. Figure 9-34 shows the **Monitor** icon in the TS7700 MI.



Figure 9-34 The monitor icon

Events

Use this window that is shown in Figure 9-35 on page 359 to view all meaningful events that occurred within the grid or a stand-alone TS7700 cluster. Events encompass every significant occurrence within the TS7700 grid or cluster, such as a malfunctioning alert, an operator intervention, a parameter change, a warning message, or some user-initiated action.

To send future events to the host operational system, host notification must be enabled. Although events are grid-wide, enabling or disabling host notification affects only the currently accessed cluster when in a grid configuration. Also, task events are not sent to the host.

Information is displayed on the Events table for 30 days after the operation stops or the event becomes inactive.

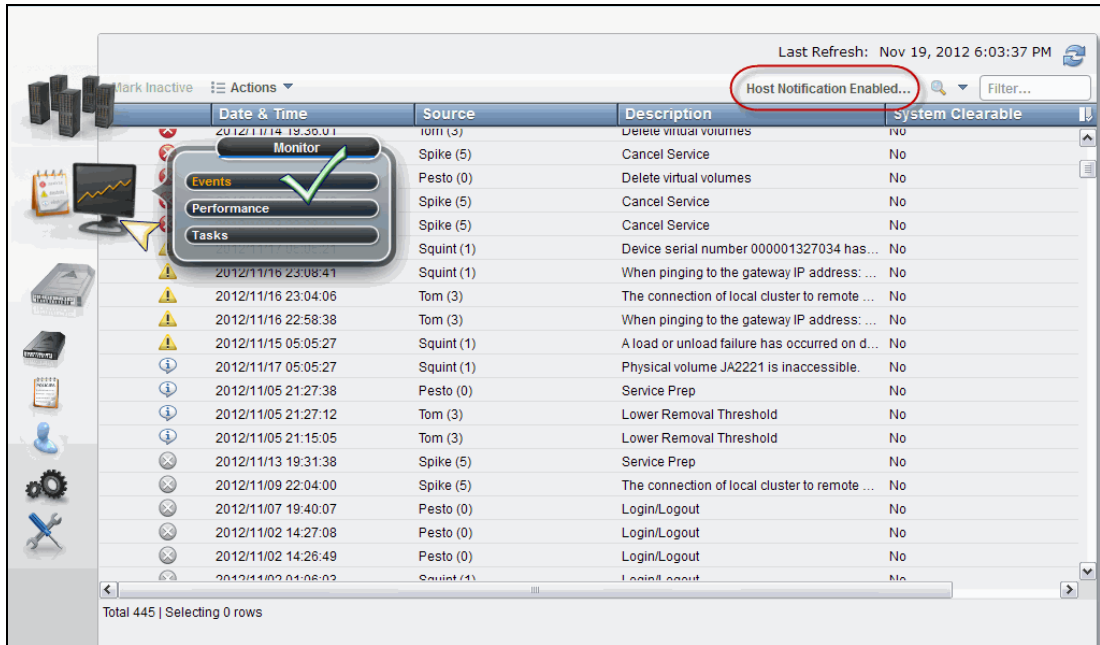


Figure 9-35 TS7700 MI Events window

Note: The Date & Time column refers the time of the events to the local time on the computer where the MI was initiated. If the DATA/TIME is modified in the TS7700 from Coordinated Universal Time during installation, the event times are offset by the same difference in the Events display on the MI. Coordinated Universal Time in all TS7700 clusters should be used whenever possible.

The Event window can be customized to meet your needs. Simply select which columns should show in the Events window. Figure 9-36 shows an example.

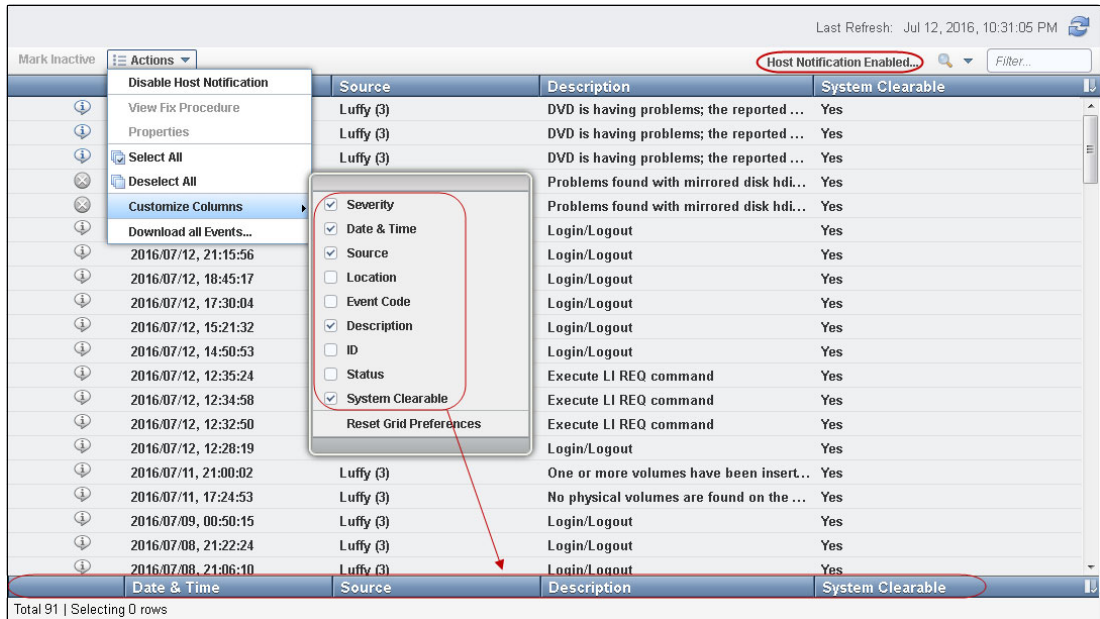


Figure 9-36 Customizing the Events window

Figure 9-37 shows the alerts, tasks, and event values and associated severity icons in the Events window in the MI.



Column name	Description
<p>Severity icon Icons in this column display in color if the event is active.</p>  <p>If an event is inactive, its icon is displayed in greyscale.</p> 	<p>Severity of the event. Possible values include:</p> <p>Error (Priority 0)</p> <p>The highest priority event. In the table it is represented by a red circle containing a white X. This level of intervention indicates:</p> <ul style="list-style-type: none"> • Failure of one or more jobs • Inaccessibility of virtual volumes • Suspension of mount operations • Degraded performance <p>It is important to examine the message immediately and act as needed.</p> <p>Warning (Priority 1)</p> <p>The second highest priority event. In the table, it is represented by a yellow triangle containing a black exclamation point (!). This level of intervention does not pose an immediate threat to operations, but indicates:</p> <ul style="list-style-type: none"> • An installation or configuration problem or change exists • Possible loss of redundant resources • Possible affect on drive cleaning operations <p>You should examine the message as time allows and take action as needed.</p> <p>Information (Priority 2)</p> <p>The lowest priority event. In the table, it is represented by a blue baloon containing a white, lowercase letter i. This level of intervention is an informational message about IBM Virtualization Engine TS7700 or library operations.</p> <p>You can log this event for historical review.</p> <p>Active events are sorted to occur at the top of the table and inactive events occur below them. Active and inactive events are further sorted according priority to this order:</p> <ol style="list-style-type: none"> 1. Error 2. Warning 3. Information

Figure 9-37 Alerts, tasks, and event values and associated severity icons

Table 9-4 describes the column names and descriptions of the fields, as shown in the Event window (see Figure 9-35 on page 359).

Table 9-4 Field name and description for the Events window

Column name	Description
Date & Time	Date and time the event occurred.
Source	Cluster where the event occurred.
Location	Specific location on the cluster where the event occurred.
Description	Description of the event.
ID	The unique number that identifies the instance of the event. This number consists of the following values: <ul style="list-style-type: none"> ▶ A locally generated ID, for example: 923 ▶ The type of event: E (event) or T (task) An event ID based on these examples appears as 923E.
Status	The status of an alert or task. If the event is an alert, this value is a fix procedure to be performed or the status of a call home operation. If the event is a task, this value is its progress or one of these final status categories: <ul style="list-style-type: none"> ▶ Canceled ▶ Canceling ▶ Completed ▶ Completed, with information ▶ Completed, with warning ▶ Failed
System Clearable	Whether the event can be cleared automatically by the system. The following values are possible: <ul style="list-style-type: none"> ▶ Yes. The event is cleared automatically by the system when the condition that is causing the event is resolved. ▶ No. The event requires user intervention to clear. The event needs to be cleared or deactivated manually after resolving the condition that is causing the event.

Table 9-5 lists actions that can be run on the Events table.

Table 9-5 Actions that can be run on the Events table

To run this task	Action
Deactivate or clear one or more alerts	<ol style="list-style-type: none"> 1. Select at least one but no more than 10 events. 2. Click Mark Inactive. If a selected event is normally cleared by the system, confirm the selection. Other selected events are cleared immediately. A running task can be cleared but if the task later fails, it is displayed again as an active event.
Enable or disable host notification for alerts	Select Actions → [Enable/Disable] Host Notification . This change affects only the accessing cluster. Tasks are not sent to the host.
View a fix procedure for an alert	Select Actions → View Fix Procedure . A fix procedure can be shown for only one alert at a time. No fix procedures are shown for tasks.

To run this task	Action
Download a comma-separated value (CSV) file of the events list	Select Actions → Download all Events .
View more details for a selected event	<ol style="list-style-type: none"> 1. Select an event. 2. Select Actions → Properties.
Hide or show columns on the table	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Click the check box next to a column heading to hide or show that column in the table. Column headings that are checked display on the table.
Filter the table data	<p>Follow these steps to filter by using a string of text:</p> <ol style="list-style-type: none"> 1. Click in the Filter field. 2. Enter a search string. 3. Press Enter. <p>To filter by column heading:</p> <ol style="list-style-type: none"> 1. Click the down arrow next to the Filter field. 2. Select the column heading to filter by. 3. Refine the selection.
Reset the table to its default view	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Click Reset Table Preferences.

9.2.5 Performance

This section introduces the performance and statistic windows that are available in the TS7700 MI.

All graphical views, except the Historical Summary, are from the last 15 minutes. The Historical Summary presents a customized graphical view of the different aspects of the cluster operation, in a 24-hour time frame. This 24-hour window can be slid back up to 90 days, which covers three months of operations.

Historical Summary

Figure 9-38 shows the Throughput View for the Historical Summary in **Monitor** → **Performance** MI operation in a tape-attached cluster.

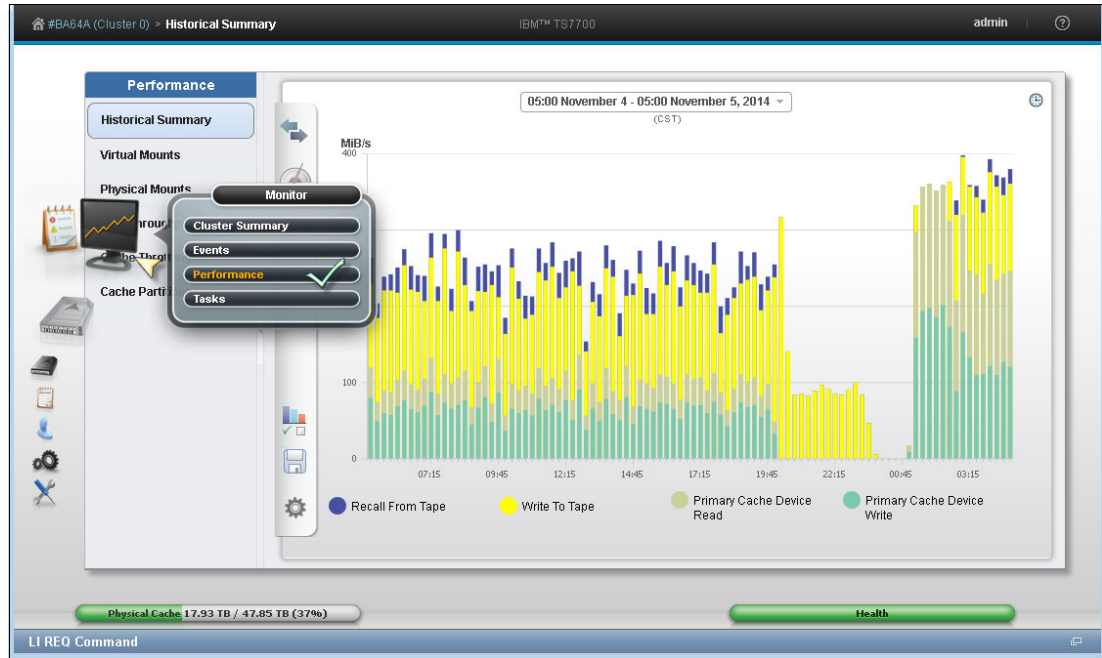


Figure 9-38 Performance window operation, throughput view

The MI is enhanced to accommodate the functions that are introduced by the code. Figure 9-39 shows the Performance Historical Summary and related chart selections that are available for this item.

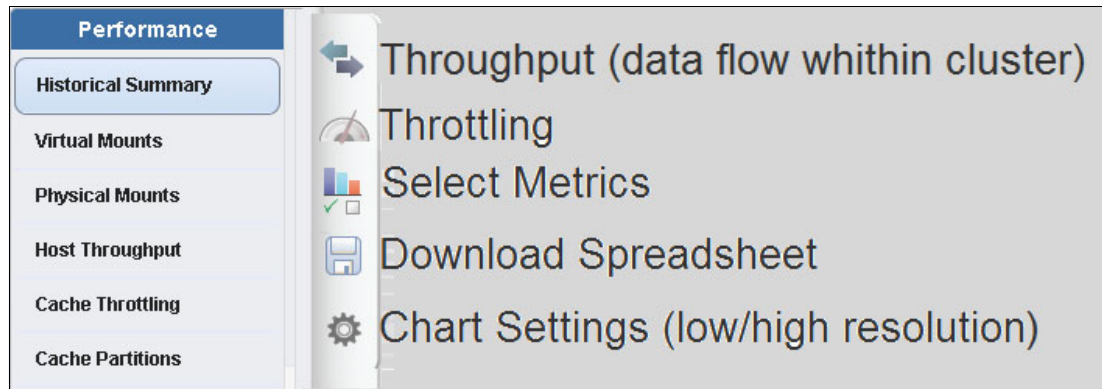


Figure 9-39 Performance options and chart selections

Figure 9-40 on page 364 shows another Historical Summary sample from the same cluster, selecting the Throttling view. The performance data came from a TS7700 tape attach cluster.

Notice that the chart shows in orange the information about the host throttling that is applied to the resident partition (CP0), where the brown line represents the host write throttling values that are applied to the tape-attached partitions (CP1 - CP7). Because all tape attached partition feeds into the same premigration queue, sharing physical tape drive resources, they are all equally affected by the same host throttling value.

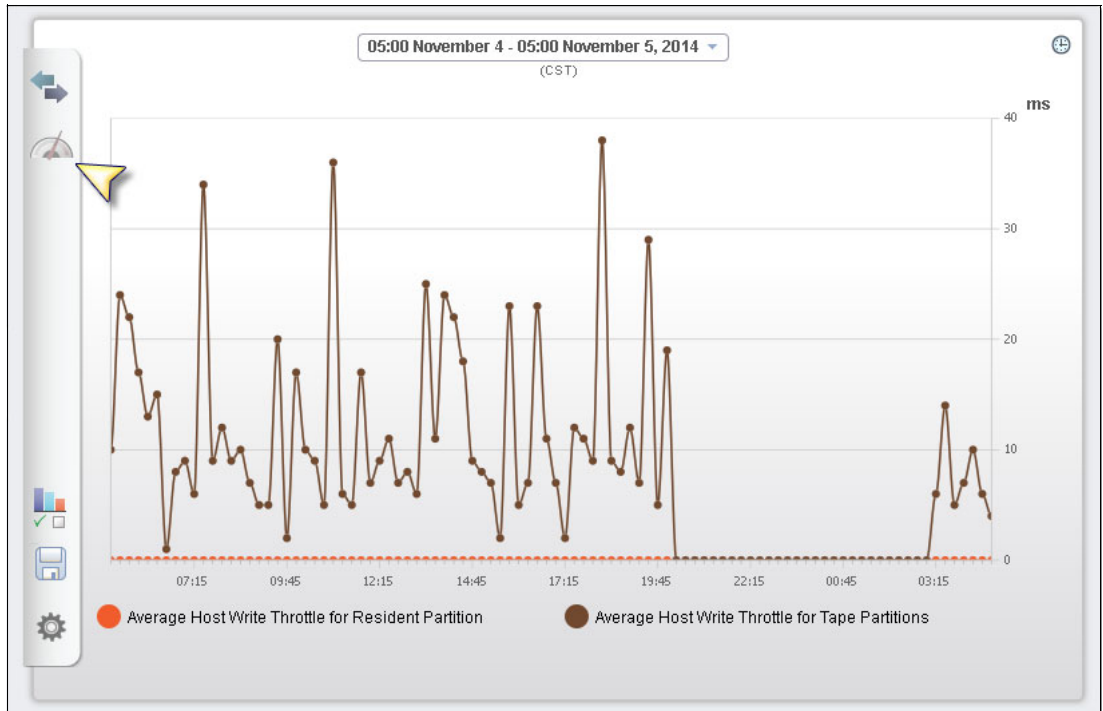


Figure 9-40 Historical Summary showing the Throttling view for a stand-alone tape-attached cluster

Clicking the **Download Spreadsheet** icon that is shown on the left of Figure 9-40 saves the raw data for the graph in a comma-separated value (CSV) file. Downloadable data is also limited to a 24-hour period from the start date and time that are defined in the window.

By using the **Select Metrics** icon, the user is able to select up to 10 data sets of different statistics to populate the chart, depending on what aspect of the cluster's performance is under scrutiny. The Select Metrics window is shown (not all options in the picture) in Figure 9-41.

Note: Up to 10 different statistic data sets can be selected for the same graph view.

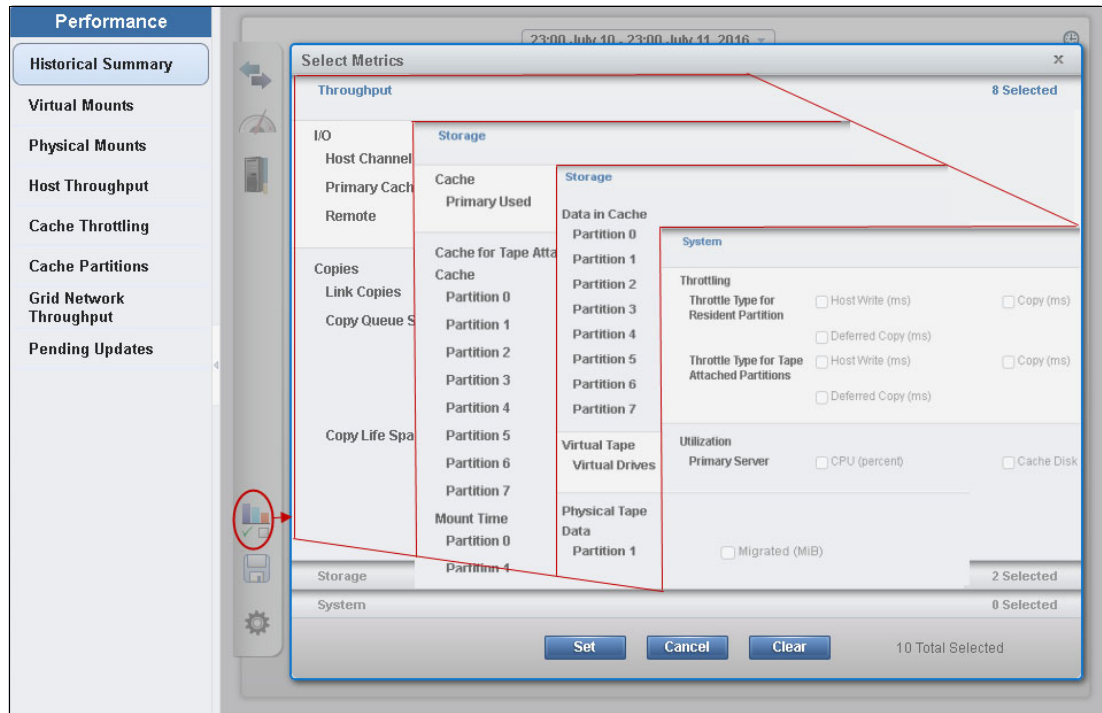


Figure 9-41 Select Metrics window and options

For an explanation of the values and what to expect in the resulting graphs, see Chapter 11, “Performance and monitoring” on page 635. Also, for a complete description of the window and available settings, see the TS7700 R4.0 IBM Knowledge Center. The TS7700 R4.0 IBM Knowledge Center is available both locally on the TS7700 MI (by clicking the question mark icon at the upper right corner of the window) and on the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_ua_performance.html

Also, see *IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance*, WP101465, which is available on the IBM Techdocs Library website:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101465>

IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance is an in-depth study of the inner workings of the TS7700, and the factors that can affect the overall performance of a stand-alone cluster or a TS7700 grid. Also, it explains throttling mechanisms and available tuning options for the subsystem to achieve peak performance.

Virtual Mounts

The Virtual Mounts statistics for the last 15 minutes of activity are displayed in a *bar graphs and table* format per cluster. Figure 9-42 shows an example of a virtual mounts graph for a TS7720T cluster.

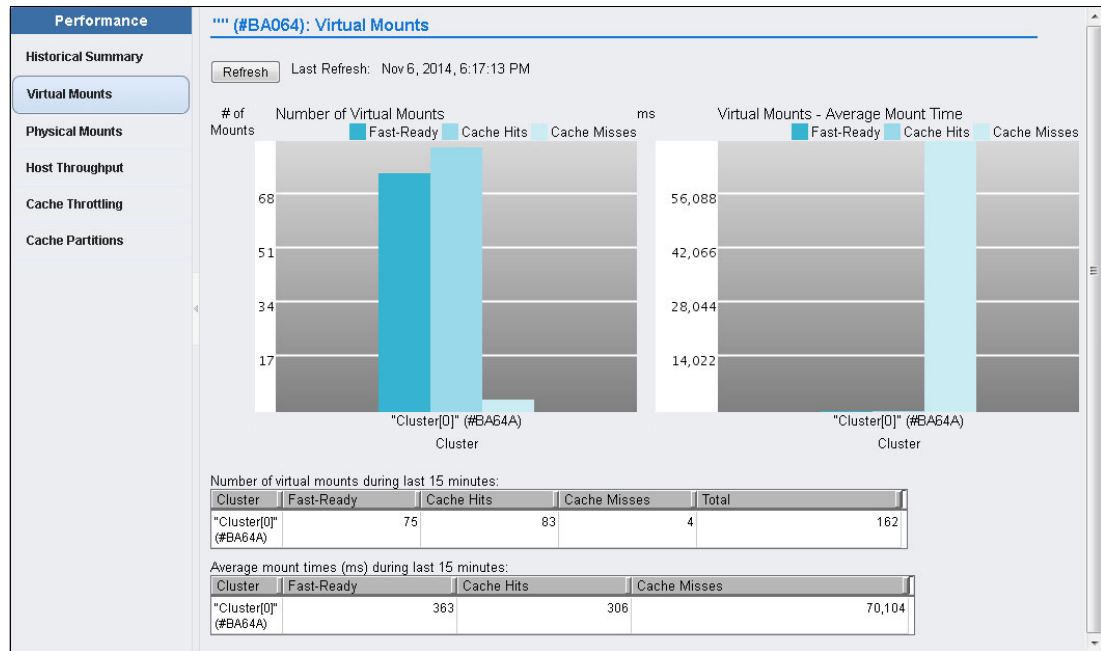


Figure 9-42 Virtual Mounts performance window

In a grid configuration, the Virtual Mounts chart displays the activity for all members in the grid. See the TS7700 Customer Information for more details.

Physical Mounts

The Physical Mounts statistics for the last 15 minutes of activity are displayed in a bar graphs and table format per cluster. This window is available and active when the selected TS7700 is attached to a physical tape library. When a grid possesses a physical library but the selected cluster does not, MI displays the following message:

The cluster is not attached to a physical tape library.

This window is not visible on the TS7700 MI if the grid does not possess a physical library (no tape attached member).

Figure 9-43 shows an example of a physical mounts window for a four-cluster grid.

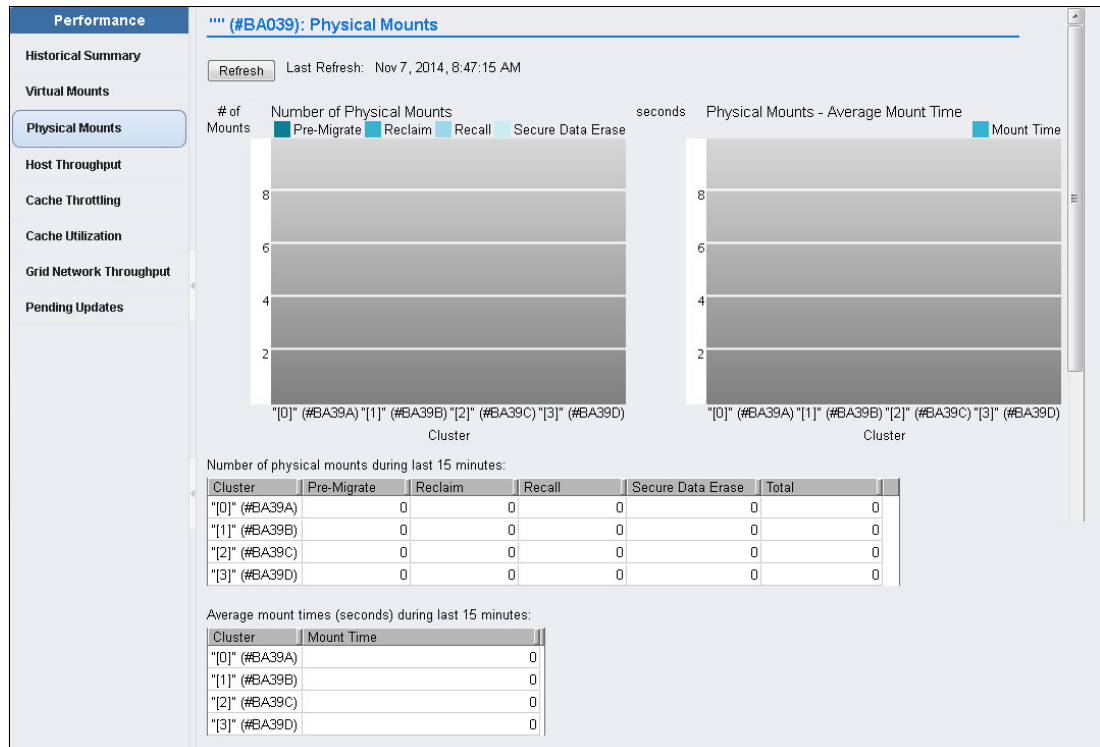


Figure 9-43 Physical mounts statistic display

Host Throughput

The host throughput for data transfer activity is shown for the last 15 minutes in a bar graph and tables. The throughput is shown for all clusters in a grid. Figure 9-44 shows the Host Throughput window in MI.

Notice the hyperlink, which enables the user to single out the throughput numbers of a specific host adapter in a specific cluster for a deeper look.

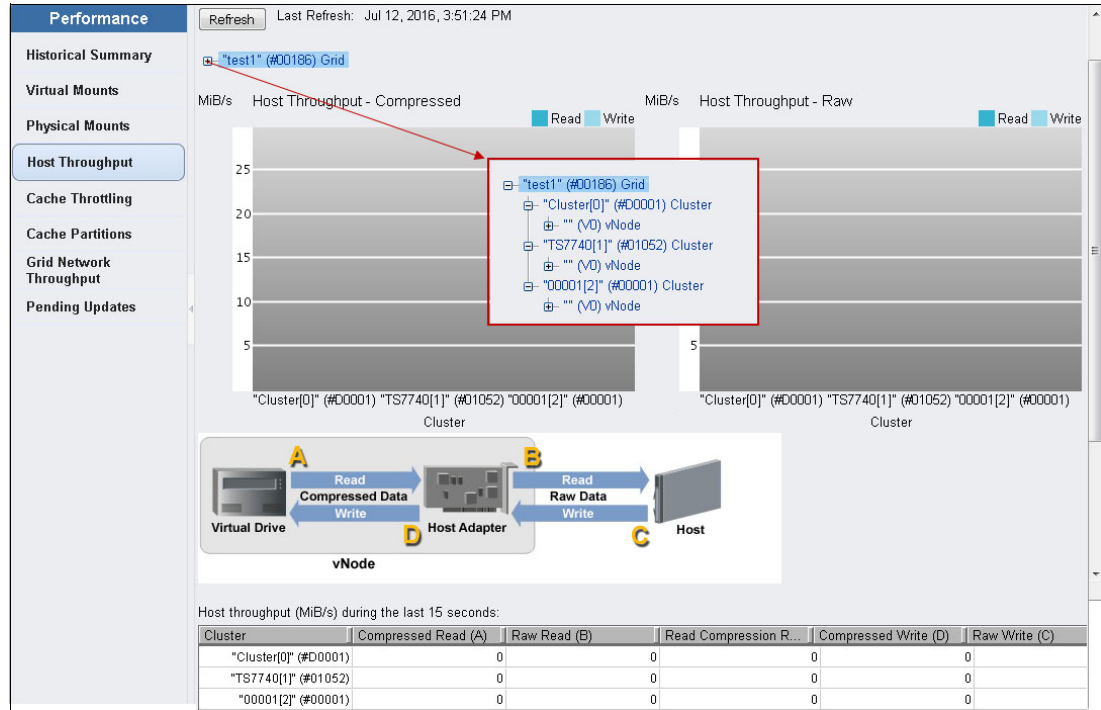


Figure 9-44 The Host Throughput window

For a complete description of the window, see the TS7700 R4.0 IBM Knowledge Center, which is available both locally at TS7700 MI (by clicking the question mark icon at the upper right corner of the window) and on the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_ua_per_host_throughput.html

Cache Throttling

This window shows the statistics of the throttling values that are applied on the host write operations and RUN copy operations throughout the grid.

Figure 9-45 is an example of the Cache Throttling window.

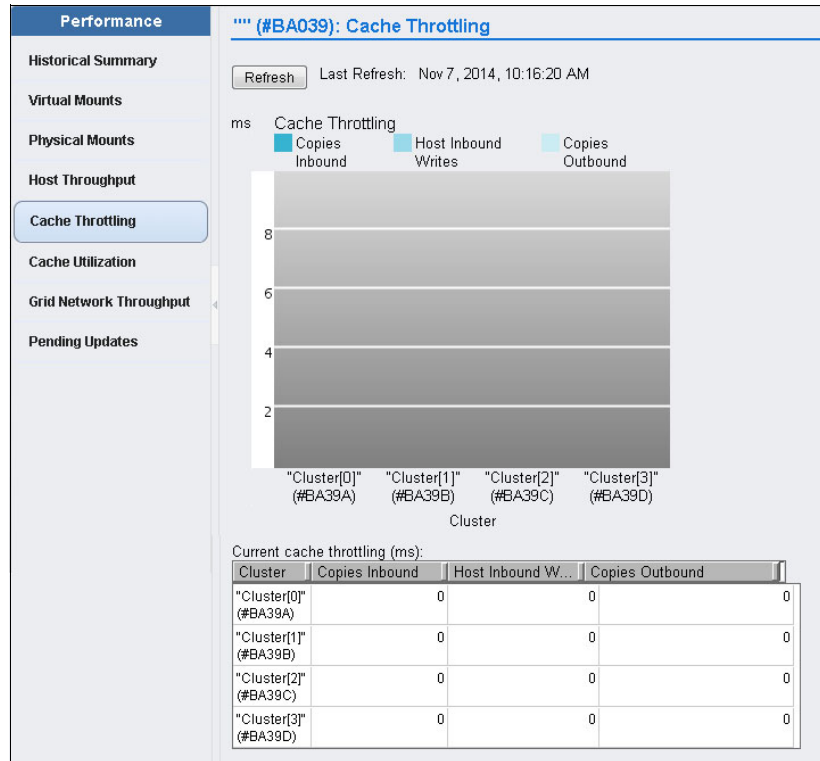


Figure 9-45 Cache Throttling window

For a complete description of the window, see the TS7700 R4.0 IBM Knowledge Center, either locally at TS7700 MI (by clicking the question mark icon at the upper right corner of the window) or on the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_ua_per_cache_throttling.html

Cache Utilization

Cache utilization statistics are presented for clusters that have one resident-only or tape-only partition, and for clusters with partitioned cache. Models TS7720 disk-only and TS7740 have only one resident or tape partition, which accounts for the entire cache. For the TS7720T (tape attach) cluster model, up to eight cache partitions (one CP0 cache resident and up to seven tape partitions) can be defined, and represented in the cache utilization window.

Figure 9-46 shows an example of cache utilization (single partition), as displayed in a TS7720 disk-only or TS7740 cluster.

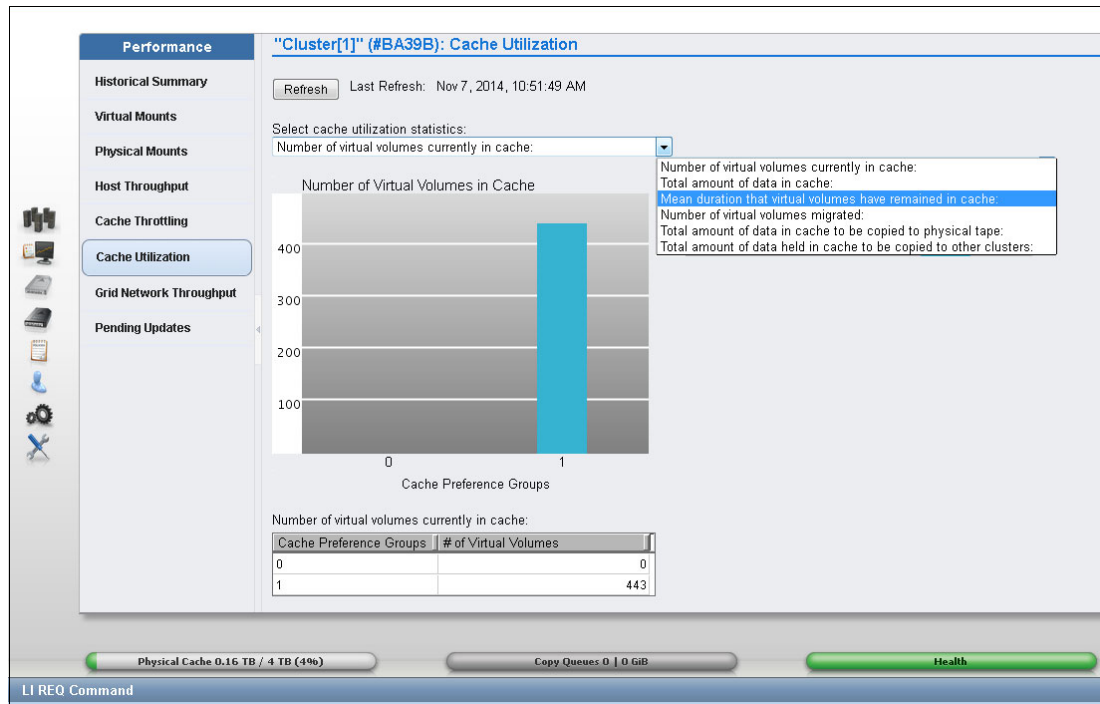


Figure 9-46 TS7700 Cache Utilization window

Cache Partition

The Cache Partition window presents the cache use statistics for the TS7720 or TS7760 tape-attached models, in which the cache is made up of multiple partitions. Figure 9-47 shows a sample of the Cache Partition (multiple partitions) window. This window can be reached by using the **Monitor** icon (as described here) or by using the **Virtual** icon. Both ways direct to the same window. In this window, the user can display the already existent cache partitions, but also can create a new partition, reconfigure an existing one, or delete a partition as needed.

Tip: Consider limiting the MI user roles who are allowed to change the partition configurations through this window.




Figure 9-47 Cache Partitions window

For a complete description of the window, see the TS7700 IBM Knowledge Center, either locally on the TS7700 MI (by clicking the question mark icon) or on the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_ua_cache_utilization.html

Grid Network Throughput

This window is only available if the TS7700 cluster is a member of a Grid. The Grid Network Throughput window shows the last 15 minutes of cross-cluster data transfer rate statistics, which are shown in megabytes per second (MBps). Each cluster of the grid is represented both in the bar graph chart and in the tables. Figure 9-48 shows an example of the Grid Network Throughput window.

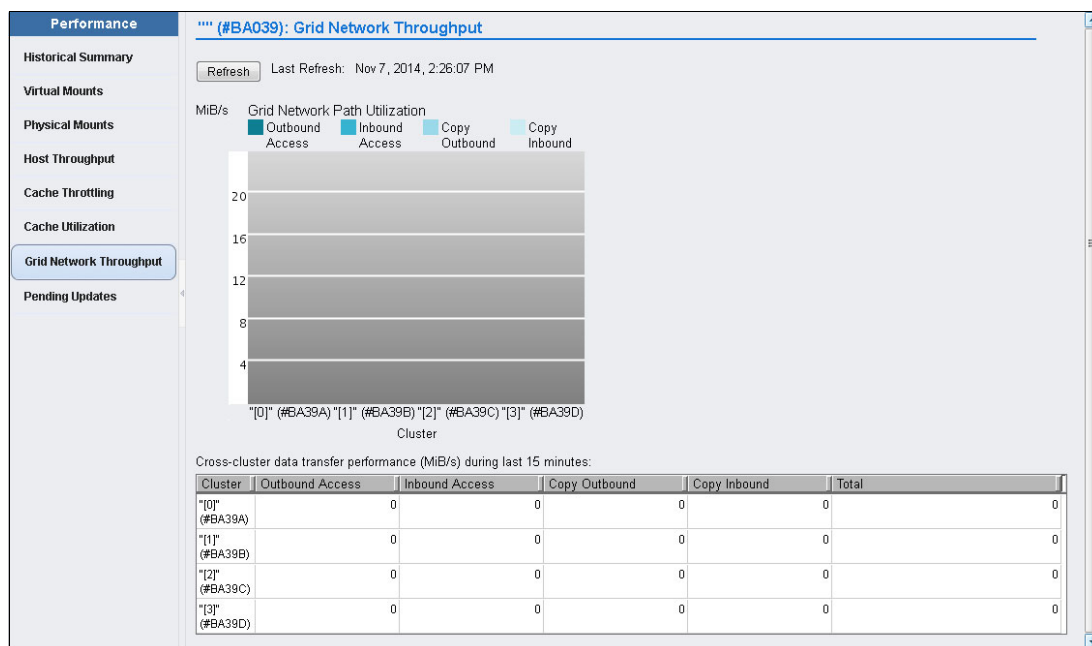


Figure 9-48 Grid Network Throughput window

For more information about this window, see the TS7700 R4.0 IBM Knowledge Center. Learn about data flow within the grid and how those numbers vary during the operation in Chapter 11, “Performance and monitoring” on page 635.

Pending Updates

The Pending Updates window is only available if the TS7700 cluster is a member of a grid. Pending updates window can be used to monitor status of outstanding updates per cluster throughout the grid. Pending updates can be caused by one cluster being offline, in service preparation or service mode while other grid peers were busy with the normal client’s production work.

A faulty grid link communication also might cause a RUN or SYNC copy to become Deferred Run or Deferred Sync. The Pending Updates window can be used to follow the progress of those copies. Figure 9-49 on page 372 shows a sample of Pending Updates window.

The Download button in the top of the window saves a comma-separated values (.csv) file that lists all volumes or grid global locks that are targeted during an ownership takeover. The volume or global pending updates are listed, along with hot tokens and stolen volumes.

Tokens are internal data structures that are used to track changes to the ownership, data, or properties of each one of the existing logical volumes in the grid. Hot tokens occur when a cluster attempts to merge its own token information with the other clusters, but the clusters are not available for the merge operation (tokens not able to merge became ‘hot’).

Stolen volume describes a volume whose ownership has been taken over during a period in which the owner cluster was in service mode or offline, or if an unexpected cluster outage occurs when the volume ownership is taken over under an operator’s direction, or by using AOTM. See Figure 9-49.

"" (#BA039): Pending Updates

Refresh Last Refresh: Nov 7, 2014, 2:39:34 PM

Takeover ownership updates required in each cluster (.csv file)

Download

Number of updates required against each cluster:

Cluster	Standard	Read/Write Takeover	Read-Only Takeover	Service Takeover
*[0] (#BA39A)	0	0	0	0
*[1] (#BA39B)	0	0	0	0
*[2] (#BA39C)	0	0	0	0
*[3] (#BA39D)	0	0	0	0

Number of **Immediate Deferred** copies for each cluster:

Cluster	Quantity	Size (GiBs)
*[0] (#BA39A)	0	0
*[1] (#BA39B)	0	0
*[2] (#BA39C)	0	0
*[3] (#BA39D)	0	0

Number of **Synchronous Deferred** copies for each cluster:

Cluster	Quantity	Size (GiBs)
*[0] (#BA39A)	0	0
*[1] (#BA39B)	0	0
*[2] (#BA39C)	0	0
*[3] (#BA39D)	0	0

Figure 9-49 Pending Updates window

For more information about copy mode and other concepts referred to in this section, see Chapter 2, “Architecture, components, and functional characteristics” on page 15. For other information about this MI function, see the TS7700 R4.0 IBM Knowledge Center on the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_ua_pending_logical_volume_updates.html

Tasks window

This window is used to monitor the status of tasks that are submitted to the TS7700. The information in this window refers to the entire grid operation if the accessing cluster is part of a grid, or only for this individual cluster if it is a stand-alone configuration. The table can be formatted by using filters, or the format can be reset to its default by using reset table preferences. Information is available in the task table for 30 days after the operation stops or the event or action becomes inactive.

Tasks are listed by starting date and time. Tasks that are still running are shown on the top of the table, and the completed tasks are listed at the bottom. Figure 9-50 shows an example of the Tasks window. Notice that the information on this page and the task status pods are of grid scope.

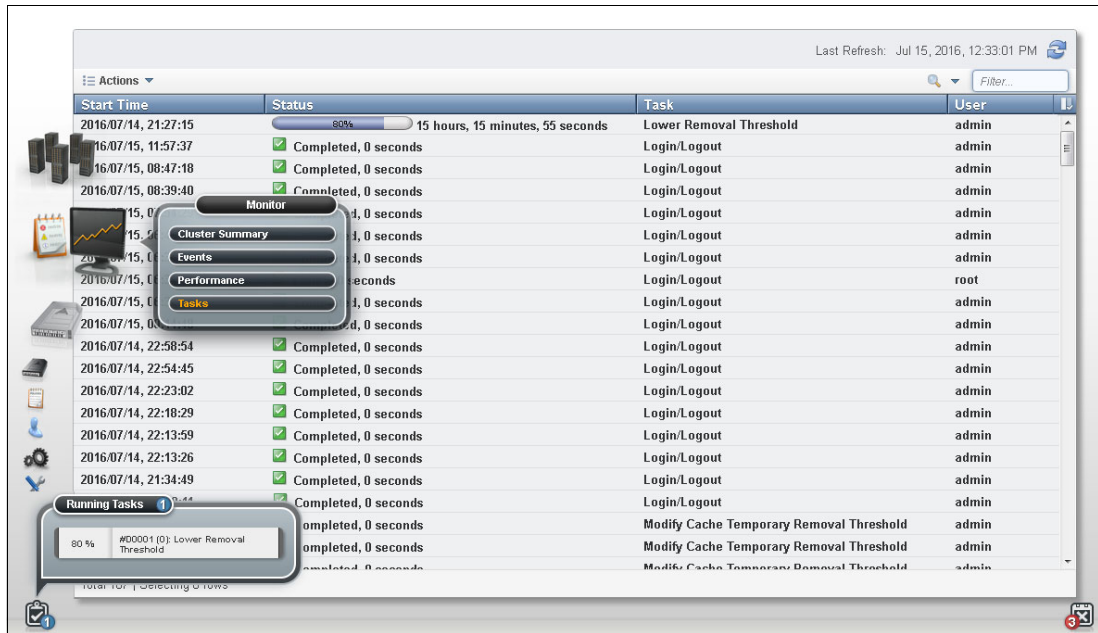


Figure 9-50 Tasks window

Note: The Start Time column refers to the time of starting a task to the local time on the computer where the MI was started. If the DATE/TIME is modified in the TS7700 from the Coordinated Universal Time during installation, the time that is shown in the *Start Times* field is offset by the same difference from the local time of the MI. Use Coordinated Universal Time in all TS7700 clusters whenever possible.

9.2.6 The Virtual icon

TS7700 MI windows that are under the **Virtual** icon can help you view or change settings that are related to virtual volumes and their queues, virtual drives, and scratch (Fast Ready) categories. For the TS7720T (Tape Attach) a new item, *Cache Partitions*, was added under the **Virtual** icon, which you can use to create, modify, or delete cache partitions.

Figure 9-51 shows the **Virtual** icon and the options available. The **Cache Partitions** item is available only for the TS7720T and TS7760T models, where the **Incoming Copy Queue** item shows only in grid configurations.



Figure 9-51 The Virtual icon and options

The available items under the Virtual icon are described in the following topics.

Cache Partitions

In the Cache Partitions window in the MI, you can create a cache partition, or reconfigure or delete an existing cache partition for the TS7720 or TS7760 tape-attached models. Also, you can use this window to monitor the cache and partitions occupancy and usage.

Figure 9-52 shows a Cache Partitions window.

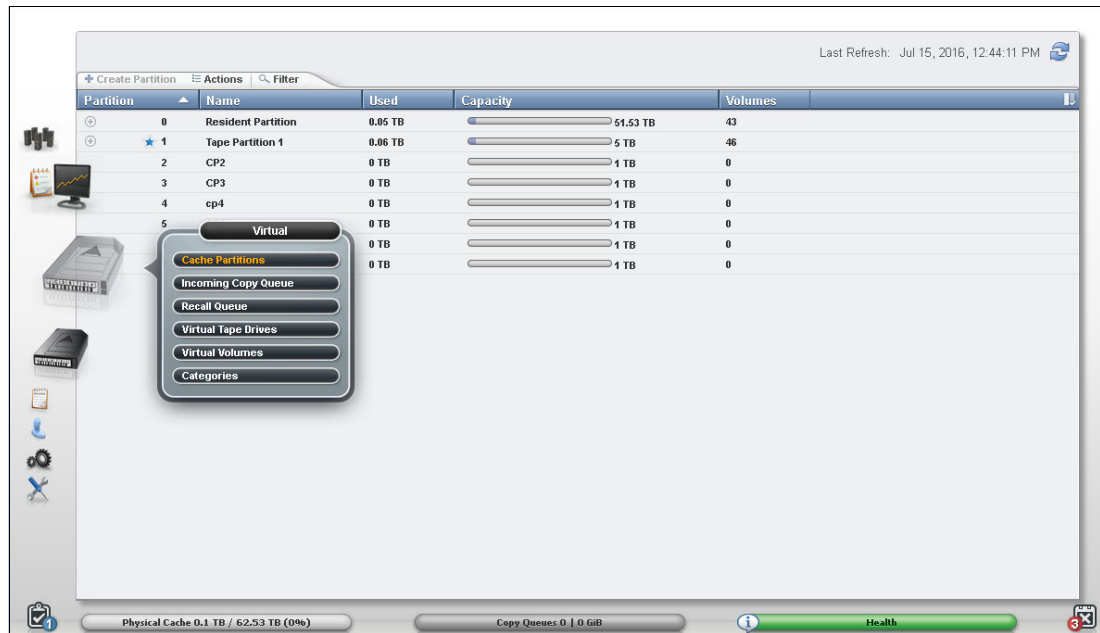


Figure 9-52 Cache Partitions window in MI

Figure 9-53 on page 376 shows a sequence for creating a new partition. There can be as many as eight partitions, from Resident partition (partition 0) to Tape Partition 7, if needed. The tape partition allocated size is subtracted from the actual resident partition capacity, if there is at least more than 2 TB of free space in the resident partition (CP0). For the complete set of rules and allowed values in effect for this window, see the TS7700 R4.0 IBM Knowledge Center. Also, learn about the tape attach TS7700, cache partitions, and usage in Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Considerations: No new partition can be created if Resident-Only (CP0) has 2 TB or less of free space. Creation of new partitions is blocked by a FlashCopy for DR in progress, or by one of the existing partitions being in an overcommitted state.

Figure 9-53 illustrates creating a new partition.

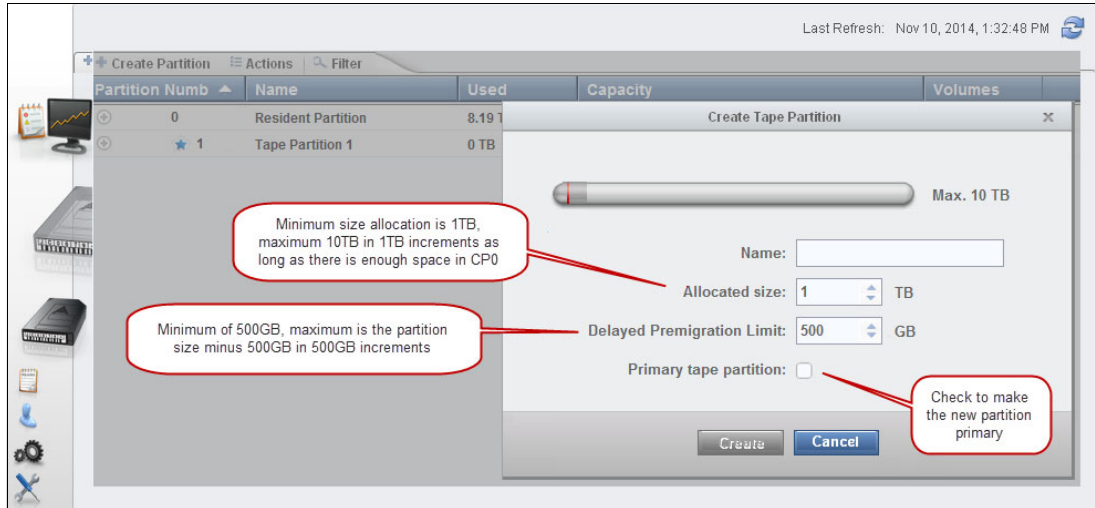


Figure 9-53 Create a tape partition

Figure 9-54 shows an example of successful creation in the upper half. The lower half shows an example where the user failed to observe the amount of free space available in CP0.

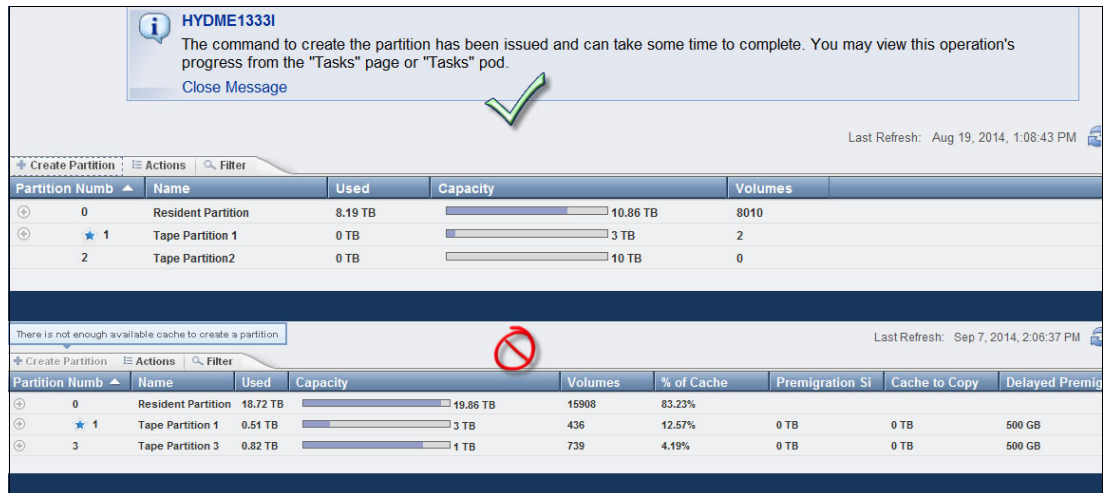


Figure 9-54 Example of a success and a failure to create a new partition

Notice that redefining the size of existing partitions in an operational TS7720T might create unexpected load peak in the overall premigration queue, causing host write throttling to be applied to the tape partitions.

For instance, consider the following example, where a tape attached partition is downsized, and become instantly overcommitted. In this example, the TS7720T premigration queue is flooded by volumes that got dislodged by the size of this cache partition becoming smaller. Partition readapts to the new size by migrating volumes in excess to physical tape.

Figure 9-55 shows the previous scenario, TS7720T operating and Tape Partition 1 operating with 12 TB cache.

Partition	Name	Used	Capacity	Volumes	% of Cache	Premigration ...	Cache to Copy	Delayed Prem...
0	Resident Partition	10.31 TB	10.86 TB	9560	45.51%			
1	Tape Partition 1	8.4 TB	12 TB	9199	50.3%	0 TB	0 TB	500 GB
	PG1 premigrated	8.4 TB		9199				
3	Tape Partition 3	0.77 TB	1 TB	695	4.19%	0 TB	0 TB	500 GB

Figure 9-55 Tape partition 1 operating with 12 TB cache

Figure 9-56 shows tape partition 1 being downsized to 8 TB. Note the initial warning and subsequent overcommit statement that shows up when resizing the tape partition results in overcommitted cache size.

Clicking **Resize** will bring up the **Overcommit State** acceptance box. Resizing will only occur after clicking **Yes**.

Overcommit State
 Decreasing Tape Partition 1 to this size will require 0.4TB of data to be migrated out of this partition's cache and on to physical tape. The partition will be in an overcommitted state until the condition clears. This action may cause performance degradation and will take some time to complete. Do you want to continue?

Figure 9-56 Downsizing Tape Partition 1, and the overcommit warning

Accepting the Overcommit statement initiates the resizing action. If this is not the best-suited time for the partition resizing (as during the peak load period), the user can click **No** and decline to take the action, and then resume it at a more appropriate time. Figure 9-57 shows the final sequence of the operation.

HYDME1340I
 The command to resize the partition has been issued and can take some time to complete. You may view this operation's progress from the "Tasks" page or "Tasks" pod.
 Close Message

Partition	Name	Used	Capacity	Volumes	% of Cache	Premigration ...	Cache to Copy	Delayed Prem...
0	Resident Partition	10.31 TB	14.86 TB	9560	62.28%			
1	Tape Partition 1	8.4 TB	8 TB	9199	33.53%	0 TB	0 TB	500 GB
3	Tape Partition 3	0.77 TB	1 TB	695	4.19%	0 TB	0 TB	500 GB

Figure 9-57 Resizing message and resulting cache partitions window

Read more about this subject in Chapter 11, "Performance and monitoring" on page 635 and Chapter 2, "Architecture, components, and functional characteristics" on page 15.

Tip: Consider limiting the MI user roles that are allowed to change the partition configurations through this window.

Incoming Copy Queue

The Incoming Copy Queue window is used for a grid-member TS7700 cluster. Use this window to view the virtual volume incoming copy queue for a TS7700 cluster. The *incoming copy queue* represents the amount of data that is waiting to be copied to a cluster. Data that is written to a cluster in one location can be copied to other clusters in a grid to achieve uninterrupted data access.

It can be specified through policies and settings on which clusters (if any) copies are, and how quickly copy operations should occur. Each cluster maintains its own list of copies to acquire, and then satisfies that list by requesting copies from other clusters in the grid according to queue priority.

Table 9-6 shows the values that are displayed in the copy queue table.

Table 9-6 Values in the copy queue table

Column type	Description
Copy Type	<p>The type of copy that is in the queue. The following values are possible:</p> <ul style="list-style-type: none"> ▶ Immediate: Volumes can be in this queue if they are assigned to a Management Class (MC) that uses the Rewind Unload (RUN) copy mode. ▶ Synchronous-deferred: Volumes can be in this queue if they are assigned to an MC that uses the Synchronous mode copy and some event (such as the secondary cluster going offline) prevented the secondary copy from occurring. ▶ Immediate-deferred: Volumes can be in this queue if they are assigned to an MC that uses the RUN copy mode and some event (such as the secondary cluster going offline) prevented the immediate copy from occurring. ▶ Deferred: Volumes can be in this queue if they are assigned to an MC that uses the Deferred copy mode. ▶ Time Delayed: Volumes can be in this queue if they are eligible to be copied based on either their creation time or last access time. ▶ Copy-refresh: Volumes can be in this queue if the MC assigned to the volumes changed and a LI REQ command was sent from the host to initiate a copy. ▶ Family-deferred: Volumes can be in this queue if they are assigned to an MC that uses RUN or Deferred copy mode and cluster families are being used.
Last TVC Cluster	<p>The name of the cluster where the copy last was in the TVC. Although this might not be the cluster from which the copy is received, most copies are typically obtained from the TVC cluster.</p> <p>This column is only shown when View by Last TVC is selected.</p>
Size	<p>Total size of the queue, which is displayed in GiB.</p> <p>See 1.6, "Data storage values" on page 12 for additional information about the use of binary prefixes.</p> <p>When Copy Type is selected, this value is per copy type. When View by Last TVC is selected, this value is per cluster.</p>
Quantity	<p>The total number of copies in queue for each type.</p>

Figure 9-58 shows the incoming copy queue window and other places in the Grid Summary and Cluster Summary that inform the user about the current copy queue for a specific cluster.

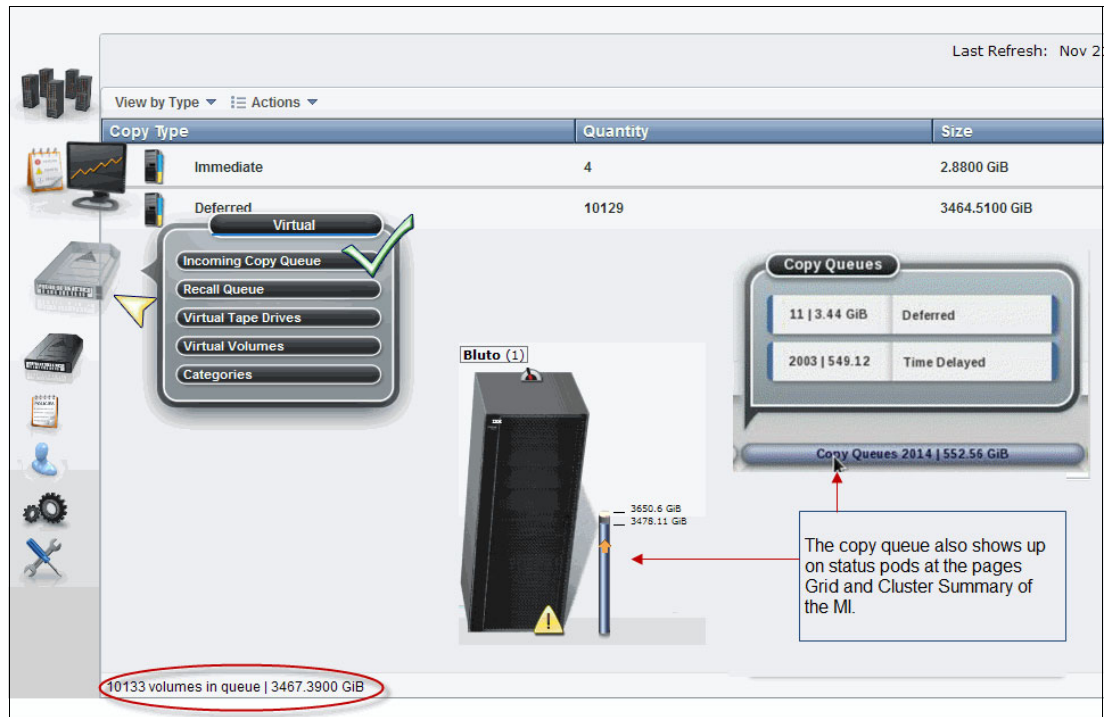


Figure 9-58 Incoming copy queue

Using the upper-left option, choose between **View by Type** and **View by Last TVC Cluster**. Use the **Actions** menu to download the Incoming Queued Volumes list.

Recall queue

The Recall Queue window of the MI displays the list of virtual volumes in the recall queue. Use this window to promote a virtual volume or filter the contents of the table. The Recall Queue item is visible but disabled on the TS7700 MI if there is no physical tape attachment to the selected cluster, but there is at least one tape attached cluster (TS7740 or TS7720T, which are connected to a TS3500 tape library) within the grid. Trying to access the Recall queue link from a cluster with no tape attachment causes the following message to display:

The cluster is not attached to a physical tape library.

Tip: This item is not visible on the TS7700 MI if there is no TS7700 tape-attached cluster in grid.

A *recall of a virtual volume* retrieves the virtual volume from a physical cartridge and places it in the cache. A *queue* is used to process these requests. Virtual volumes in the queue are classified into three groups:

- ▶ In Progress
- ▶ Scheduled
- ▶ Unscheduled

Figure 9-59 shows an example of the Recall queue window.

Position	Virtual Volume	Physical Cartridges	Time in Queue
In Progress	L0396	P0665, P0603	0 hours, 13 minutes, 1 seconds
In Progress	L0398	P0158, P0901	0 hours, 8 minutes, 50 seconds
In Progress	L0494	P0859, P093	0 hours, 14 minutes, 10 seconds
In Progress	L0353	P0786, P0811	0 hours, 10 minutes, 14 seconds
In Progress	L0601	P0404, P0726	0 hours, 10 minutes, 22 seconds
In Progress	L0203	P0733, P0509	0 hours, 12 minutes, 40 seconds
Scheduled	L0246	P0583, P0893	0 hours, 8 minutes, 45 seconds
1	L0659	P0677, P0497	0 hours, 9 minutes, 11 seconds
2	L0357	P035, P0158	0 hours, 11 minutes, 57 seconds
3	L0487	P0655, P0579	0 hours, 12 minutes, 59 seconds
4	L0871	P0420, P0954	0 hours, 8 minutes, 6 seconds
5	L0525	P0875, P0109	0 hours, 10 minutes, 58 seconds
6	L0666	P0239, P0513	0 hours, 8 minutes, 25 seconds
7	L04	P0537, P0751	0 hours, 8 minutes, 47 seconds
8	L0117	P0437, P0953	0 hours, 12 minutes, 8 seconds
25	L0321	P0432	0 hours, 9 minutes, 59 seconds

Figure 9-59 Recall queue window

Table 9-7 shows the names and the descriptions of the values that might be seen on the Recall window.

Table 9-7 Recall window values

Column name	Description
Position	The position of the virtual volume in the recall queue. The following values are possible: <ul style="list-style-type: none"> ▶ In Progress: A recall is in progress for the volume. ▶ Scheduled: The volume is scheduled to be recalled. If optimization is enabled, the TS7700 schedules recalls to be processed from the same physical cartridge. ▶ Position: A number that represents the volume's current position in the list of volumes that are not scheduled. These unscheduled volumes can be promoted by using the Actions menu.
Virtual Volume	The virtual volume to be recalled.
Physical Cartridges	The serial number of the physical cartridge on which the virtual volume is. This column can be hidden.
Time in Queue	Length of time that the virtual volume was in the queue, which is displayed in hours, minutes, and seconds as <i>HH:MM:SS</i> . This column can be hidden.

In addition to changing the recall table's appearance by hiding and showing some columns, the user can filter the data that is shown in the table by a string of text, or by the column heading. Possible selections are by Virtual Volume, Position, Physical Cartridge, or by Time in Queue. To reset the table to its original appearance, click **Reset Table Preferences**.

Another interaction now available in the Recall window is that the user can promote an unassigned volume recall to the first position in the unscheduled portion of the recall queue. This is available by checking an unassigned volume in the table, and clicking **Actions** → **Promote Volume**.

Virtual tape drives

The Virtual Tape Drives window of the MI presents the status of all virtual tape drives in a cluster. Use this window to check the status of a virtual mount, to perform a stand-alone mount or unmount, or assign host device numbers to a specific virtual drive.

The field *Cache Mount Cluster* in the virtual tape drives page identifies to which cluster TVC the volume is mounted. The user can recognize remote (crossed) or synchronous mounts simply by looking in this field. Remote mounts show other clusters that are being used by a mounted volume instead of a local cluster (the cluster to whom the Virtual Tape Drives belong to in the page currently on display), while synchronous mounts show both clusters used by the mounted volume.

The page contents can be customized by selecting specific items to display. Figure 9-60 shows the Virtual tape drives window and the available items under **Actions** menu.

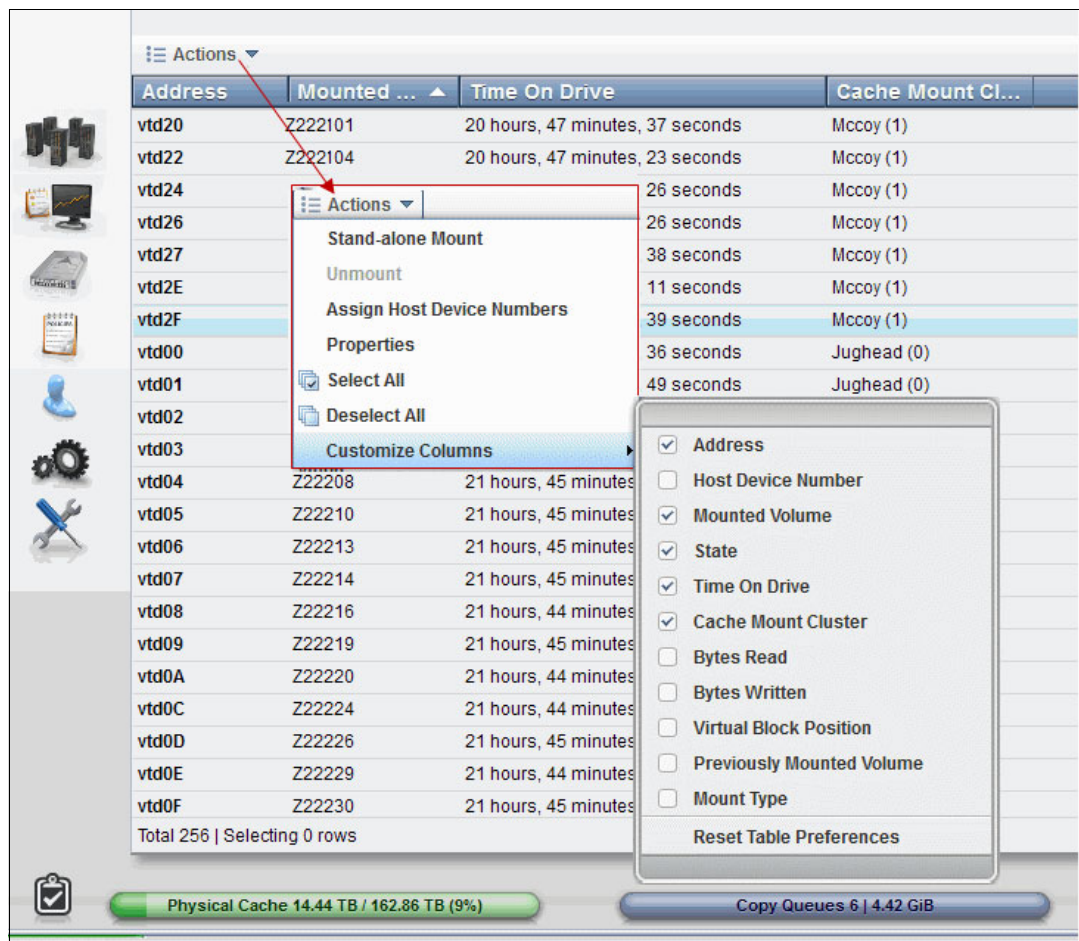


Figure 9-60 Virtual Tape Drives window

The user can perform a stand-alone mount to a logical volume against a TS7700 logical drive for special purposes, such as to perform an initial program load (IPL) of a stand-alone services core image from a virtual tape drive. Also, MI allows the user to manually unmount a logical drive that is mounted and in the idle state. The Unmount function is available not only for those volumes that have been manually mounted, but also for occasions when a logical volume has been left mounted on a virtual drive by an incomplete operation or some test rehearsal, therefore creating the need to unmount it through MI operations.

To perform a stand-alone mount, run the steps that are shown in Figure 9-61.

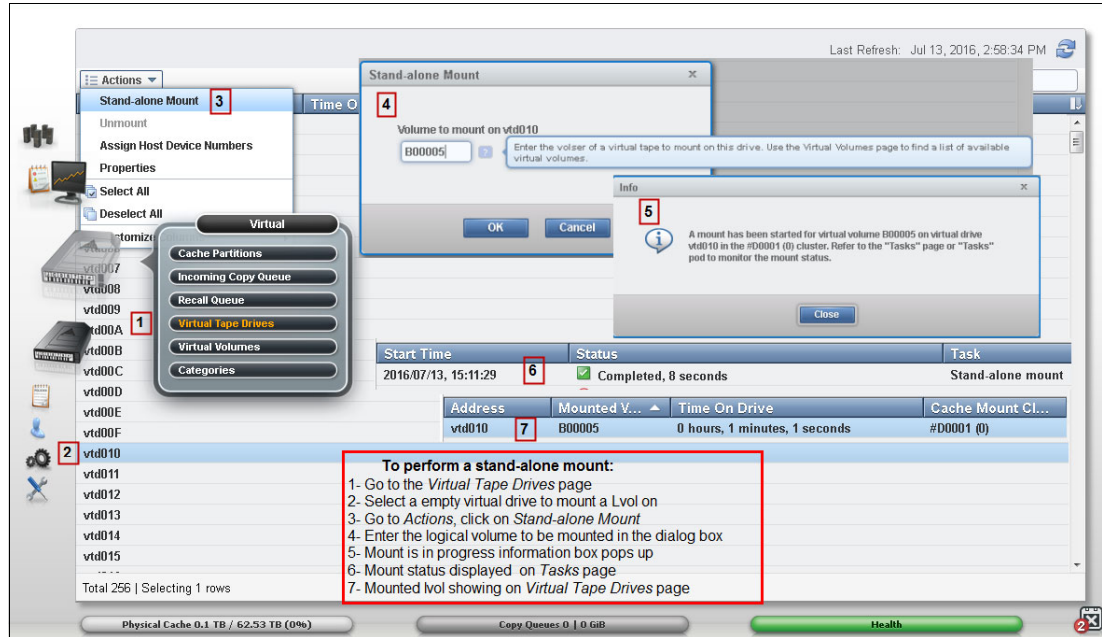


Figure 9-61 Stand-alone mount procedure

The user can mount only virtual volumes that are not already mounted, on a virtual drive that is online.

If there is a need to unmount a logical volume that is currently mounted to a virtual tape drive, follow the procedure that is shown in Figure 9-62.

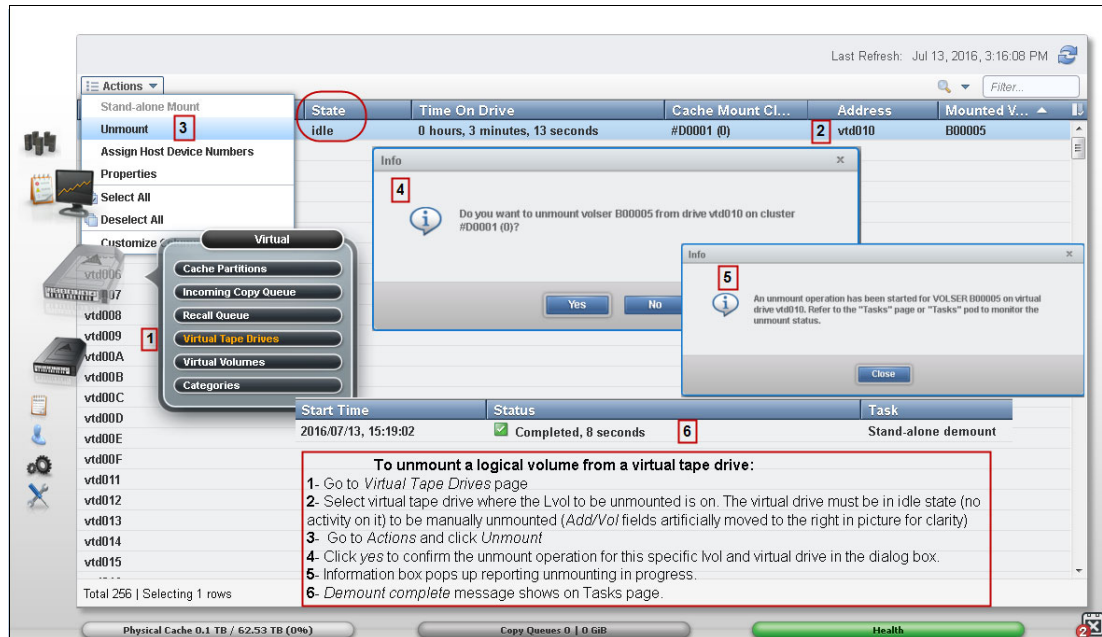


Figure 9-62 Demounting a logical volume via MI

The user can unmount only those virtual volumes that are mounted and have a status of Idle.

Table 9-8 shows the properties of virtual tape drives.

Table 9-8 Properties that are displayed for virtual tape drives

Column name	Description
Address	The virtual drive address takes the format vtdXXXX, where X is a hexadecimal number.
Host Device Number	The device identifier as defined in the attached host. The value in this field does not affect drive operations, but if the host device number is set, it is easier to compare the virtual tape drives to their associated host devices. Follow these steps to add host device numbers: 1. Select one or more virtual tape drives. 2. Select Assign Host Device Numbers from the Actions menu. 3. Enter the host device address to assign to the first virtual tape drive. Host device numbers are added to subsequent virtual tape drives incrementally.
Mounted Volume	The volume serial number (VOLSER) of the mounted virtual volume.
Previously Mounted Volume	The VOLSER of the virtual volume that is mounted on the drive before this one.
State	The role that the drive is performing. The following values are possible: Idle: The drive is not in use. Read: The drive is reading a virtual volume. Write: The drive is writing a virtual volume. This column is blank if no volume is mounted. With R3.1, this column shows <i>Offline</i> status for the virtual tape drive when appropriate. The column <i>Online</i> in previous versions of this window was removed; drive status is assumed Online unless stated Offline in this column.
Time on Drive	The elapsed time that the virtual volume was on the virtual tape drive. This column is blank if no volume is mounted.
Cache Mount Cluster	The TVC cluster that is running the mount operation. If a synchronous mount exists, this field displays two clusters. This column is blank if no volume is mounted.
Bytes Read	Amount of data that is read from the mounted virtual volume. This value is shown as Raw KiB (Compressed KiB).
Bytes Written	Amount of data that was written to the mounted virtual volume. This value is shown as Raw KiB (Compressed KiB).
Virtual Block Position	The position of the drive on the tape surface in a block number, as calculated from the beginning of the tape. This value is displayed in hexadecimal form. When a volume is not mounted, this value is 0x0.

Column name	Description
Mount Type	<p>The type of mount on the drive. This field is blank if no volume is mounted. Possible values are:</p> <ul style="list-style-type: none"> ▶ Live Copy: The mount is a live copy of the volume. This is the type that is used during normal production. ▶ FlashCopy: The mount is a FlashCopy used for disaster recovery testing. ▶ Stand-alone: The mount request was initiated by the cluster and not the host. ▶ Stand-alone FlashCopy: A user initiated a stand-alone mount for a FlashCopy. <p>To mount a virtual volume:</p> <ol style="list-style-type: none"> 1. Select a volume. 2. Select Stand-alone Mount from the Actions menu. <p>The user can mount only virtual volumes that are not already mounted, on a drive that is online.</p> <p>Follow these steps to unmount a virtual volume:</p> <ol style="list-style-type: none"> 1. Select a mounted volume. 2. Select Unmount from the Actions menu. <p>The user can unmount only those virtual volumes that are mounted and have a status of Idle.</p>

Virtual volumes

The topics in this section present information about monitoring and manipulating virtual volumes in the TS7700 MI.

Virtual volume details

Use this window to obtain detailed information about the state of a virtual volume or a FlashCopy of a virtual volume in the TS7700 Grid. Figure 9-64 on page 386 and Figure 9-66 on page 388 show an example of the resulting windows for a Virtual Volume query. The entire window can be subdivided in three parts:

1. Virtual volume summary
2. Virtual volume details
3. Cluster-specific virtual volume properties

There is a tutorial available about virtual volume display and how to interpret the windows accessible directly from the MI window. To watch it, click the **View Tutorial** link on the Virtual Volume Detail window.

Figure 9-63 shows an example of the graphical summary for a virtual volume (Z22208, in this example). The first part of the Virtual Volume Details window in the MI shows a graphical summary of the status of the virtual volume that is being displayed.

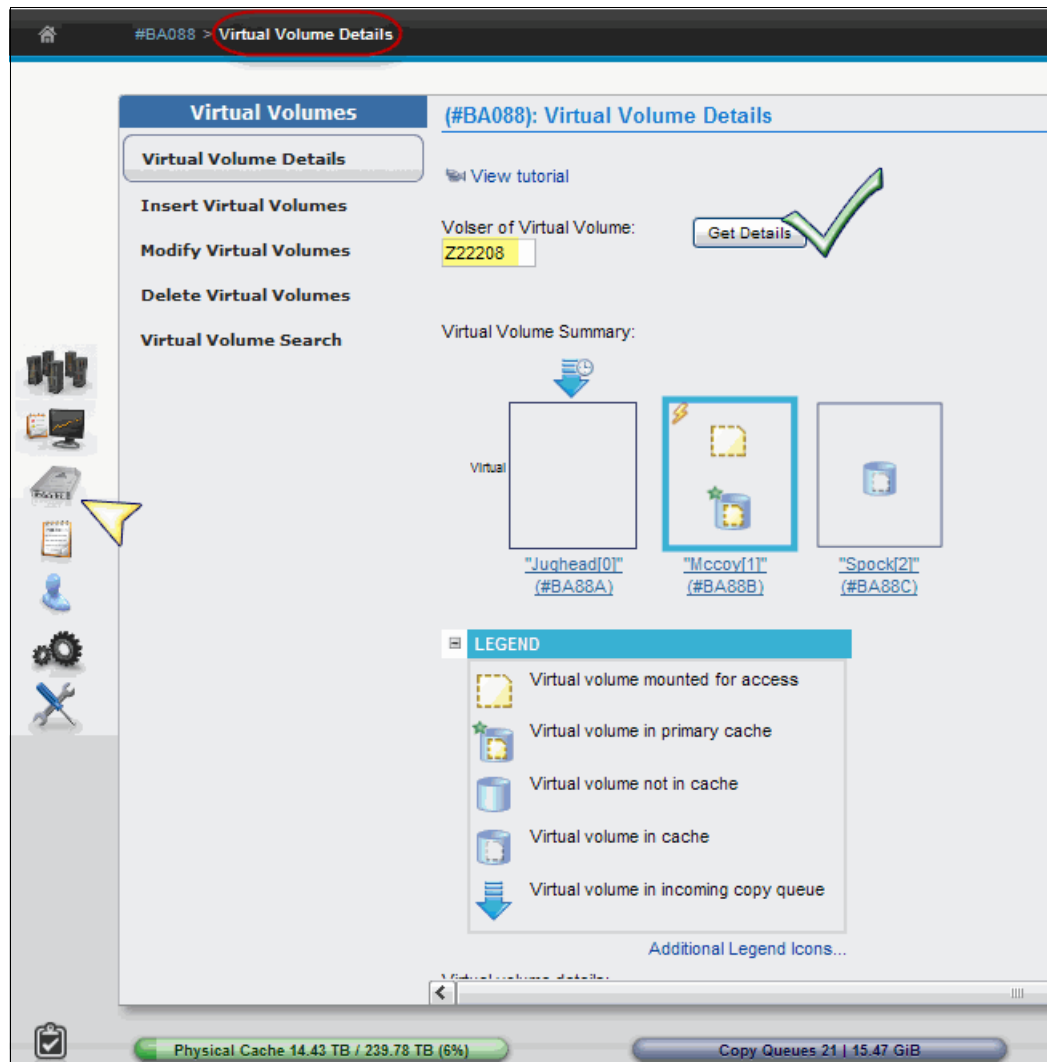


Figure 9-63 Virtual Volume Details: Graphical summary

This graphical summary brings details of the present status of the virtual volume within the grid, plus the current operations that are taking place throughout the grid concerning that volume. The graphical summary helps you understand the dynamics of a logical mount, whether the volume is in the cache at the mounting cluster, or whether it is being recalled from tape in a remote location.

Note: The physical resources are shown in the virtual volume summary, virtual volume details table, or the cluster-specific virtual volume properties table for the TS7720T and TS7740 cluster models.

The Virtual Volume Details window shows all clusters where the selected virtual volume is located within the grid. The icon that represents each individual cluster is divided in three different areas by broken lines:

- ▶ The top area relates to the logical volume status.
- ▶ The intermediate area shows actions that are currently in course or pending for the cluster.
- ▶ The bottom area reports the status of the physical components that are related to that cluster.

The cluster that owns the logical volume being displayed is identified by the blue border around it. For instance, referring to Figure 9-63 on page 385, volume Z22208 is owned by cluster 1, where volume and a FlashCopy are in cache. Volume Z22208 is not mounted or available in the primary cache at cluster 2. At the same time, Z22208 is in the deferred incoming copy queue for cluster 0.

Figure 9-64 shows how the icons are distributed through the window, and where the pending actions are represented. The blue arrow icon over the cluster represents data that is being transferred from another cluster. The icon in the center of the cluster indicates data that is being transferred within the cluster.

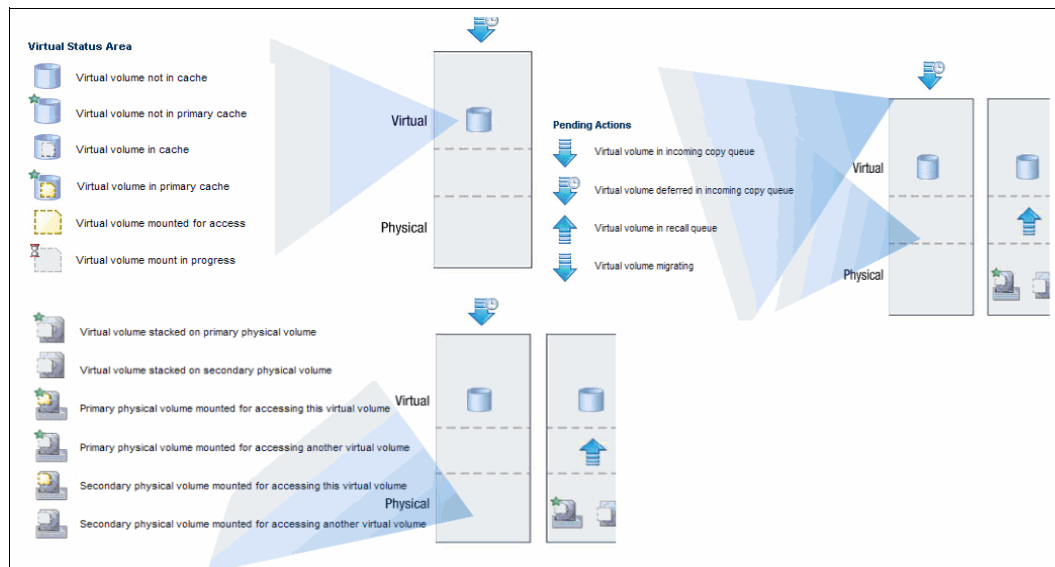


Figure 9-64 Details of the graphical summary area

Figure 9-65 shows a list of legends that can appear in the virtual volume details display, along with a brief description of the meaning of the icons.











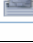


ICON	Description
	A flash copy is active on the cluster. Mind that volume properties between live copy and the flash copy may not have changed.
	Virtual volume not in primary cache. The volume was in the primary cache but has since been migrated. The current cluster may or may not be the owner. This volume should never be mounted.
	Virtual volume not in cache on the owning cluster (green star indicates the owning cluster).
	Virtual volume in cache
	Virtual volume in primary cache (owning cluster indicated by green star). The current cluster may or may not be the owner. This volume can be mounted.
	Virtual volume mounted for access.
	Virtual volume mount in progress.
	Virtual volume stacked on primary physical volume. Green star indicates physical primary volume or the first physical volume written to when a virtual volume is migrated to a physical tape.
	Virtual volume stacked on secondary physical volume.
	Primary physical volume mounted for accessing this virtual volume. Green star indicates physical primary volume or the first physical volume written to when a virtual volume is migrated to physical tape.
	Primary physical volume mounted for accessing another virtual volume. Green star indicates physical primary volume or the first physical volume written to when a virtual volume is migrated to physical tape.
	Secondary physical volume mounted for accessing this virtual volume.
	Secondary physical volume mounted for accessing another virtual volume.

Figure 9-65 Legend list for the graphical representation of the virtual volume details

Figure 9-66 shows the text section, which follows the graphical representation for the logical volume details window of the TS7700 MI, complementing the information about the logical volume. Details, such as the media type, compressed data size, current owner, and whether the volume is currently mounted and where, are presented. It displays other properties for this volume, such as copy retention, copy policy, whether an automatic removal was attempted or not, and when if so.

Virtual volume details:	
Volser	A50826
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	1,179.5 MiB
Maximum Volume Capacity (Device)	4,000 MiB
Current Owner	"[2]" (#00001)
Currently Mounted	No
vNode	-
Virtual Drive	-
Cached Copy Used for Mount	"[2]" (#00001)
Mount State	-
Cache Management Preference Group	1
Last Accessed by a Host	Aug 28, 2014, 8:24:06 AM
Last Modified	Aug 28, 2014, 8:23:53 AM
Category	002F
Storage Group	SGG00001
Management Class	MNSSN068
Storage Class	SCT0002K
Data Class	D000N004
Volume Data State	Active
Flash Copy	Not Active
Earliest Deletion On	-
Logical WORM	No

Figure 9-66 Virtual volume details: Text section

Virtual volume details: Text

The virtual volume details and status are displayed in the Virtual volume details table:

- ▶ **Volser.** The VOLSER of the virtual volume. This value is a six-character number that uniquely represents the virtual volume in the virtual library.
- ▶ **Media Type.** The media type of the virtual volume. The possible values are Cartridge System Tape (400 MiB) or Enhanced Capacity Cartridge System Tape (800 MiB).
- ▶ **Current Volume Size (device).**
- ▶ **Actual size (MiB) of the virtual volume.**
- ▶ **Maximum Volume Capacity (device).** The maximum size (MiB) of the virtual volume. This capacity is set upon insert by the DC of the volume. Changes to one logical volume size are applied only when a load-point write (scratch mount) occurs for that volume.
At volume close time, the value that is defined by DC or override is bounded to the volume and cannot be changed until the volume is reused. Any further changes to a DC override are not inherited by a volume until it is written again during a scratch mount and closed.
- ▶ **Current Owner.** The name of the cluster that currently owns the current version of the virtual volume.
- ▶ **Currently Mounted.** Indicates whether the virtual volume is mounted in a virtual drive.

- ▶ **Vnode.** The name of the vnode on which the virtual volume is mounted.
- ▶ **Virtual Drive.** The ID of the virtual drive on which the virtual volume is mounted.
- ▶ **Cache Copy Used for Mount.** The cluster name of the cache that was chosen for I/O operations for mount based on Consistency policy, volume validity, residency, performance, and cluster mode.
- ▶ **Cache Management Preference Group.** The preference level for the Storage Group (SG). It determines how soon volumes are removed from cache after their copy to tape. This information is only displayed if a physical library exists in the grid. The following values are possible:
 - **0.** Volumes in this group have preference to be removed from cache over other volumes.
 - **1.** Volumes in this group have preference to be retained in cache over other volumes. A least recently used (LRU) 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.
- ▶ **Mount State** The current mount state of the virtual volume. The following values are possible:
 - **Mounted.** The volume is mounted.
 - **Mount Pending.** A mount request is received and is in progress.
 - **Recall Queued/Requested.** A mount request is received and a recall request is queued.
 - **Recalling.** A mount request is received and the virtual volume is being staged into the TVC from physical tape.
- ▶ **Last Accessed by a Host.** The date and time that the virtual volume was last accessed by a host. The time that is recorded reflects the time zone in which the user's browser is.
- ▶ **Last Modified.** The date and time the virtual volume was last modified by a host. The time that is recorded reflects the time zone in which the user's browser is.
- ▶ **Category.** The number of the category to which the virtual volume belongs.
- ▶ **Storage Group.** The name of the SG that defines the primary pool for the premigration of the virtual volume. Only displayed for virtual volumes belonging to a private category (non-scratch).
- ▶ **Management Class.** The name of the MC applied to the volume. This policy defines the copy process for volume redundancy. Only displayed for virtual volumes belonging to a private category (non-scratch).
- ▶ **Storage Class.** The name of the Storage Class (SC) applied to the volume. This policy classifies virtual volumes to automate storage management. Only displayed for virtual volumes belonging to a private category (non-scratch).
- ▶ **Data Class.** The name of the DC applied to the volume. This policy classifies virtual volumes to automate storage management. Only displayed for virtual volumes belonging to a private category (non-scratch).
- ▶ **Volume Data State.** The state of the data on the virtual volume:
 - **New.** The virtual volume is in the insert category or a private (non-Fast Ready) category and data was never written to it.
 - **Active.** The virtual volume is located within a private category and contains data.

- **Scratched.** The virtual volume is located within a scratch (Fast Ready) category and its data is not scheduled to be automatically deleted.
 - **Pending Deletion.** The volume is within a scratch (Fast Ready) category and its contents are a candidate for automatic deletion when the earliest deletion time passes. Automatic deletion then occurs sometime later. The volume can be accessed for mount or category change before the automatic deletion and the deletion can be incomplete.
 - **Pending Deletion with Hold.** The volume is within a scratch (Fast Ready) category that is 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 passes, the volume then becomes a candidate for deletion and moves 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 scratch (Fast Ready) category or previously was in a scratch (Fast Ready) category where it became a candidate for automatic deletion and was deleted. Any mount operation to this volume is treated as a scratch (Fast Ready) mount because no data is present.
- ▶ **FlashCopyDetails of any existing FlashCopy copies.** This field is only available in the TS7720 clusters. Possible values are:
 - **Not active.** No FlashCopy is active. No FlashCopy was enabled at the host by a LI REQ operation.
 - **Active.** A FlashCopy that affects this volume was enabled at the host by a LI REQ operation. Volume properties have not changed since FlashCopy time zero.
 - **Created.** A FlashCopy that affects this volume was enabled at the host by a LI REQ operation. Volume properties between the live copy and the FlashCopy have changed. If the value **Created** is displayed under FlashCopy item, click it to start the FlashCopy details window. See the FlashCopy details window later in this chapter.
 - ▶ **Earliest Deletion On.** This is the date and time when the virtual volume is deleted. Time that is recorded reflects the time zone in which the user's browser is located. If there is no expiration date set, this value displays as a dash (-).
 - ▶ **Logical WORM.** Whether the virtual volume is formatted as a Write Once Read Many (WORM) volume. The possible values are Yes and No.

Cluster-specific Virtual Volume Properties

Figure 9-67 shows the Cluster-specific Virtual Volume Properties table that is shown in the last part of the Virtual volume details window.

Cluster-specific Virtual Volume Properties:		
	"Cluster[1]" (#01052)	"Cluster[2]" (#00001)
In Cache	No	Yes
Device Bytes Stored	1,179.5 MiB (Device)	1,179.5 MiB (Device)
Primary Physical Volume	HYDE11	A00003
Secondary Physical Volume	None	None
Copy Activity	Complete	Complete
Queue Type	-	-
Copy Mode	Synchronous Copy	Synchronous Copy
Deleted	-	-
Removal Residency	-	-
Removal Time	-	-
Partition Number	-	2
Premigration Delay Time	-	-
Storage Preference	-	-

Figure 9-67 Cluster-specific Virtual Volume Properties

The Cluster-specific Virtual Volume Properties table displays information about requesting virtual volumes on each cluster. These are properties that are specific to a cluster. Virtual volume details and the status that is displayed include the following properties:

- ▶ **Cluster.** The cluster location of the virtual volume copy. Each cluster location occurs as a separate column header.
- ▶ **In Cache.** Whether the virtual volume is in cache for this cluster.
- ▶ **Primary Physical Volume.** The physical volume that contains the specified virtual volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details window 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 a dash for those clusters that are not attached to a physical library.
- ▶ **Secondary Physical Volume.** A secondary physical volume that contains the specified virtual volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details window 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 a dash for those clusters that are not attached to a physical library.
- ▶ **Copy Activity.** Status information about the copy activity of the virtual volume copy. The following values are possible:
 - **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. The following values are possible:
 - **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.
 - **Time Delayed.** The copy occurs after the defined time delays (create / access) in hours has expired.
 - **Sync Deferred.** The copy was set to be synchronized, according to the synchronized mode copy settings, but the synchronized cluster was unable to be accessed. The copy is in the Deferred state. See "Synchronous mode copy" on page 81 for additional information about Synchronous mode copy settings and considerations.
 - **Immediate Deferred.** A RUN copy that has been moved to the Deferred state due to copy timeouts or TS7700 grid states.
- ▶ **Copy Mode.** The copy behavior of the virtual volume copy. The following values are possible:
 - **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.
 - **Time Delayed.** A copy is only made if the specified time has elapsed.
 - **No Copy.** No copy is made.

- **Sync.** The copy occurs upon any synchronization operation. See “Synchronous mode copy” on page 81 for additional information about settings and considerations.
- **Exist.** A consistent copy exists at this location although No Copy is intended. A consistent copy existed at this location at the time that the virtual 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 virtual volume on the cluster was deleted. The time that is recorded reflects the time zone in which the user’s browser is located. If the volume has not been deleted, this value is displayed as a dash.
- ▶ **Removal Residency.** The residency state of the virtual volume. This field is displayed only if the grid contains a disk-only cluster. The following values are possible:
 - “-” Removal Residency does not apply to the cluster. This value is displayed if the cluster attaches to a physical tape library.
 - **Removed.** The virtual volume has been removed from the cluster.
 - **No Removal Attempted.** The virtual volume is a candidate for removal, but the removal has not yet occurred.
 - **Retained.** An attempt to remove the virtual 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 and that copies are being replicated to other peer clusters. This copy can be removed only after enough replicas exist on other peer clusters.
 - **Deferred.** An attempt to remove the virtual 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 enough 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.
 - **Pinned.** The virtual volume is pinned by the virtual volume SC. The copy on this cluster cannot be removed until it is unpinned. When this value is present, the Removal Time value is Never.
 - **Held.** The virtual volume is held in cache on the cluster at least until the Removal Time has passed. After the removal time has passed, the virtual volume copy is a candidate for removal. The Removal Residency value becomes No Removal Attempted if the volume is not accessed before the Removal Time passes.

The copy on this cluster is moved to the Resident state if it is not accessed before the Removal Time passes. If the copy on this cluster is accessed after the Removal Time has passed, it is moved back to the Held state.
- ▶ **Removal Time.** This field is displayed only if the grid contains a disk-only cluster. Values that are displayed in this field depend on values that are displayed in the Removal Residency fields shown in Table 9-9 on page 393.

Table 9-9 Removal Time and Removal Residency value

Removal Residency state	Removal Time indicator
Removed	The date and time the virtual volume was removed from the cluster.
Held	The date and time the virtual volume becomes a candidate for removal.
Pinned	The virtual volume is never removed from the cluster.
No Removal Attempted, “-”, Retained, or Deferred	“-” The Removal Time field is not applicable.

The time that is recorded reflects the time zone in which the user’s browser is located.

Note: If the cluster contains a physical library, Removal Residency does not apply and this field displays a dash.

- ▶ **Partition number.** The partition number for a TS7720T tape attach cluster. Possible values are C0 - C7.
- ▶ **Premigration delay time.** The basis for calculating the time period that is defined in Time-Delayed Premigration Delay. The following values are possible:
 - **Volume Creation.** Calculate premigration delay, starting with the time when the virtual volume was created.
 - **Volume Last Accessed.** Calculate premigration delay, starting with the time when the virtual volume was most recently accessed.
- ▶ **Volume Copy Retention Group.** The name of the group that defines the preferred Auto Removal policy applicable to the virtual volume. The Volume Copy Retention Group provides more options to remove data from a TS7720 cluster or for a partition 0 (CP0) on a TS7720T tape attach cluster as the active data reaches full capacity.

Volumes become candidates for removal if an appropriate number of copies exist on peer clusters *and* the volume copy retention time has elapsed since the volume was last accessed. Volumes in each group are removed in order based on their least recently used access times. The volume copy retention time describes the number of hours a volume remains in cache before becoming a candidate for removal.

This field is displayed only if the cluster is a TS7720 (disk-only) part of a *hybrid* grid (one that combines TS7740 and TS7720 clusters), or for a partition 0 (CP0) on a TS7720T (tape attach) cluster. If the virtual volume is in a scratch (Fast Ready) category and is on a disk-only cluster, removal settings no longer apply to the volume, and the volume is a candidate for removal.

In this instance, the value that is displayed for the Volume Copy Retention Group is accompanied by a warning icon. The following values are possible:

- **Prefer Remove.** Removal candidates in this group are removed before removal candidates in the Prefer Keep group.
- **Prefer Keep.** Removal candidates in this group are removed after removal candidates in the Prefer Remove group.
- **Pinned.** Copies of volumes in this group are never removed from the accessing cluster. The volume copy retention time does not apply to volumes in this group. Then, volumes in this group that are later moved to scratch become priority candidates for removal.
- **“-”.** Volume Copy Retention does not apply to a TS7740 cluster. This value (a dash that indicates an empty value) is displayed if the cluster attaches to a physical tape library.

FlashCopy details

This section provides detailed information about the state of a virtual volume FlashCopy in the TS7700 grid.

This window is available only for volumes with a *created* FlashCopy of a virtual volume. In this context, created FlashCopy means an existing FlashCopy, which becomes different from the live virtual volume. The live volume has been modified after FlashCopy time zero. For the volumes with a FlashCopy *active* (meaning no difference between the FlashCopy and live volume) as in Figure 9-66 on page 388, only the Virtual Volume details window is available (FlashCopy and live volume are identical).

Figure 9-68 shows a FlashCopy details pane in the MI.

Virtual volume details:	
Volser	A08472
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	1,175.3 MiB
Maximum Volume Capacity (Device)	4,000 MiB
Current Owner	"[2]" (#00001)
Currently Mounted	No
vNode	-
Virtual Drive	-
Cached Copy Used for Mount	"[2]" (#00001)
Mount State	-
Cache Management Preference Group	1
Last Accessed by a Host	Sep 8, 2014, 9:35:05 AM
Last Modified	Sep 2, 2014, 8:30:47 AM
Category	002F
Storage Group	SGG00001
Management Class	MNNDN040
Storage Class	SCT0003K
Data Class	D000N004
Volume Data State	Active
Flash Copy	Active
Earliest Deletion On	-
Logical WORM	No

The same volume seen through LI REQ LVOL FLASH command

```
> SHOWING RESULTS FOR COMMANDS: LVOL,A08472,FLASH
"LOGICAL VOLUME INFORMATION V4 .1
" FLASH COPY VOLUME:      A08472
" MEDIA TYPE:              ECST
" COMPRESSED SIZE (MB):    1175
" MAXIMUM VOLUME CAPACITY (MB): 4000
" CURRENT OWNER:          CLUSTER2
" MOUNTED LIBRARY:
" MOUNTED VNODE:
" MOUNTED DEVICE:
" TVC LIBRARY:            CLUSTER2
" MOUNT STATE:
" CACHE PREFERENCE:      ---
" CATEGORY:              002F
```

Figure 9-68 FlashCopy details window

The virtual volume details and status are displayed in the Virtual volume details table:

- ▶ **Volser.** The VOLSER of the virtual volume, which is a six-character value that uniquely represents the virtual volume in the virtual library.
- ▶ **Media type.** The media type of the virtual volume. Possible values are:
 - **Cartridge System Tape**
 - **Enhanced Capacity Cartridge System Tape**
- ▶ **Maximum Volume Capacity.** The maximum size in MiB of the virtual volume. This capacity is set upon insert, and is based on the media type of a virtual volume.
- ▶ **Current Volume Size.** Size of the data in MiB for this virtual volume.
- ▶ **Current Owner.** The name of the cluster that currently owns the latest version of the virtual volume.
- ▶ **Currently Mounted.** Whether the virtual volume is mounted in a virtual drive. If this value is Yes, these qualifiers are also displayed:
 - **Vnode.** The name of the vnode that the virtual volume is mounted on.
 - **Virtual drive.** The ID of the virtual drive the virtual volume is mounted on.

- ▶ **Cache Copy Used for Mount.** The name of the cluster that owns the cache chosen for I/O operations for mount. This selection is based on consistency policy, volume validity, residency, performance, and cluster mode.
- ▶ **Mount State.** The mount state of the logical volume. The following values are possible:
 - **Mounted.** The volume is mounted.
 - **Mount Pending.** A mount request has been received and is in progress.
- ▶ **Last Accessed by a Host.** The date and time the virtual volume was last accessed by a host. The time that is recorded reflects the time zone in which the user's browser is located.
- ▶ **Last Modified.** The date and time the virtual volume was last accessed by a host. The time that is recorded reflects the time zone in which the user's browser is located.
- ▶ **Category.** The category to which the volume FlashCopy belongs.
- ▶ **Storage Group.** The name of the SG that defines the primary pool for the pre-migration of the virtual volume.
- ▶ **Management Class.** The name of the MC applied to the volume. This policy defines the copy process for volume redundancy.
- ▶ **Storage Class.** The name of the SC applied to the volume. This policy classifies virtual volumes to automate storage management.
- ▶ **Data Class.** The name of the DC applied to the volume.
- ▶ **Volume Data State.** The state of the data on the FlashCopy volume. The following values are possible:
 - **Active.** The virtual volume is located within a private category and contains data.
 - **Scratched.** The virtual volume is located within a scratch category and its data is not scheduled to be automatically deleted.
 - **Pending Deletion.** The volume is located within a scratch category and its contents are a candidate for automatic deletion when the earliest deletion time has passed. Automatic deletion then occurs sometime thereafter. This volume can be accessed for mount or category change before the automatic deletion, in which case the deletion can be postponed or canceled.
 - **Pending Deletion with Hold.** The volume is located within a scratch category that is 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 passes, the volume becomes a candidate for deletion and moved to the *Pending Deletion* state. While in this state, the volume is accessible by all legal host operations.
- ▶ **Earliest Deletion On.** Not applicable to FlashCopy copies (-).
- ▶ **Logical WORM.** Not applicable to FlashCopy copies (-).

Cluster-specific FlashCopy volume properties

This is the second part of the *FlashCopy details* window. Figure 9-69 on page 396 shows the cluster-specific FlashCopy volume properties in the MI. Notice that cluster location shows as a separated column header. Only clusters that are part of a disaster recovery family are shown.

Cluster-specific Virtual Volume Properties:		
	"Cluster[1]" (#01052)	"Cluster[2]" (#00001)
In Cache	Yes	Yes
Device Bytes Stored	1,175.7 MiB (Device)	1,175.7 MiB (Device)
Primary Physical Volume	None	None
Secondary Physical Volume	None	None
Copy Activity	Complete	Complete
Queue Type	-	-
Copy Mode	Rewind unload (RUN)	Rewind unload (RUN)
Deleted	-	-
Removal Residency	-	-
Removal Time	-	-
Partition Number	-	0
Premigration Delay Time	-	-
Storage Preference	-	-

Figure 9-69 Cluster-specific FlashCopy Volume Properties

The Cluster-specific FlashCopy Properties window displays cluster-related information for the FlashCopy volume that is being displayed:

- ▶ **Cluster.** The cluster location of the FlashCopy, on the header of the column. Only clusters that are part of a disaster recovery family are shown.
- ▶ **In Cache.** Whether the virtual volume is in cache for this cluster.
- ▶ **Device Bytes Stored.** The number of actual bytes (MiB) used by each cluster to store the volume. This amount can vary between clusters based on settings and configuration.
- ▶ **Copy Activity.** Status information about the copy activity of the virtual volume copy:
 - **Complete.** This cluster location completed a consistent copy of the volume.
 - **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.
 - **Time Delayed Until [time].** A copy is delayed as a result of Time Delayed Copy mode. The value for [time] is the next earliest date and time that the volume is eligible for copies.
- ▶ **Queue Type.** The type of queue as reported by the cluster. Possible values are:
 - **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.
 - **Sync Deferred.** The copy was set to be synchronized, according to the synchronized mode copy settings, but the synchronized cluster could not be accessed. The copy is in the deferred state.
 - **Immediate Deferred.** A RUN copy that has been moved to the deferred state due to copy timeouts or TS7700 Grid states.
 - **Time Delayed.** The copy occurs sometime after the delay period has been exceeded.
- ▶ **Copy Mode.** The copy behavior of the virtual volume copy. Possible values are:
 - **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.
 - **No Copy.** No copy is made.
 - **Sync.** The copy occurs upon any synchronization operation.
 - **Time Delayed.** The copy occurs sometime after the delay period has been exceeded.
 - **Exists.** A consistent copy exists at this location although *No Copy* is intended. This happens when a consistent copy existed at this location at the time the virtual 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 virtual volume on the cluster was deleted. The time that is recorded reflects the time zone in which the user's browser is located. If the volume has not been deleted, this value displays a dash.
 - ▶ **Removal Residency.** Not applicable to FlashCopy copies.
 - ▶ **Removal Time.** Not applicable to FlashCopy copies.
 - ▶ **Volume Copy Retention Group.** Not applicable to FlashCopy copies.

Insert Virtual Volumes window

Use this window to insert a range of virtual volumes in the TS7700 subsystem. Virtual volumes that are inserted in an individual cluster are available to all clusters within a grid configuration.

The Insert Virtual Volumes window is shown in Figure 9-70.

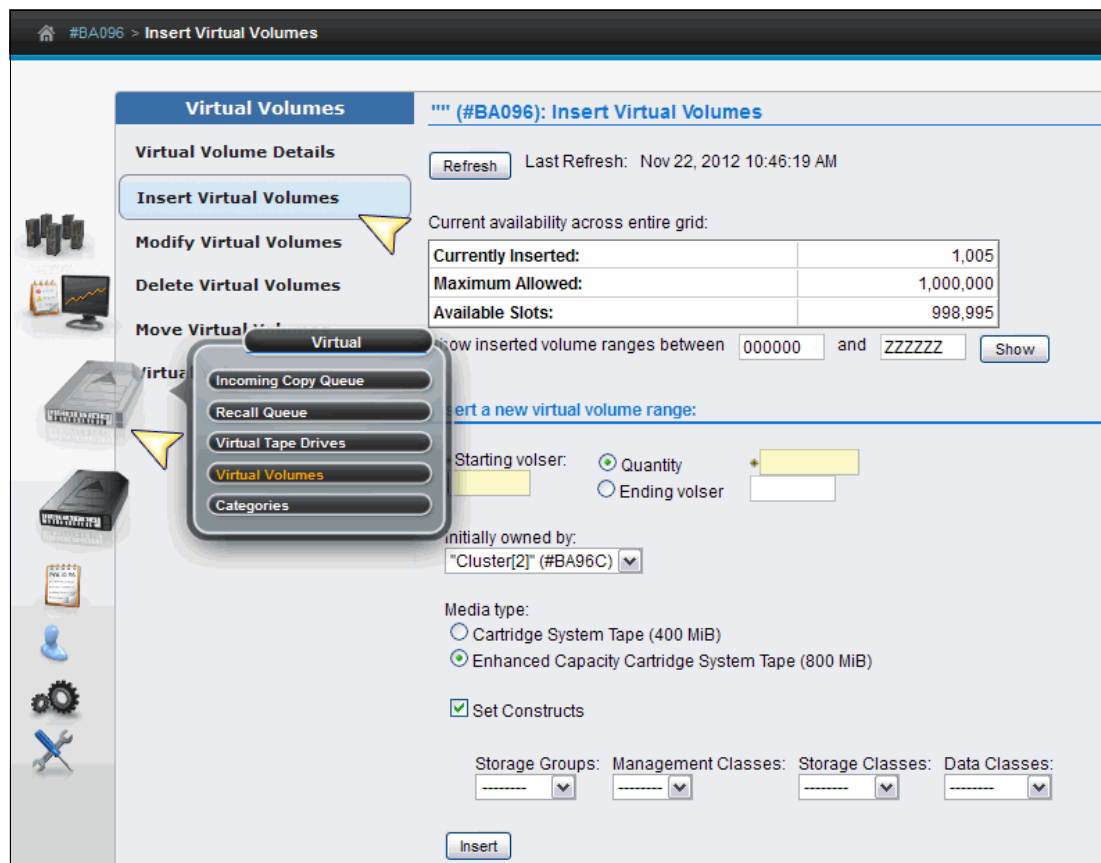


Figure 9-70 *Insert Virtual Volumes window*

The Insert Virtual Volume window shows the *Currently availability across entire grid* table. This table shows the total of the already inserted volumes, the maximum number of volumes allowed in the grid, and the available slots (the difference between the maximum allowed and the currently inserted numbers). Clicking **Show/Hide** under the table shows or hides the information box with the already inserted volume ranges, quantities, media type, and capacity. Figure 9-71 shows the inserted ranges box.

From ^	To ^	Quantity ^	Media Type ^	Capacity
111111	111115		5 ECCST	800 MiB
222222	223221		1,000 ECCST	800 MiB
Total: 2				

Figure 9-71 Show logical volume ranges

Insert a New Virtual Volume Range window

Use the following fields to insert a range of new virtual volumes:

- ▶ **Starting VOLSER.** The first virtual volume to be inserted. The range for inserting virtual volumes begins with this VOLSER number.
- ▶ **Quantity.** Select this option to insert a set number of virtual volumes beginning with the Starting VOLSER. Enter the quantity of virtual volumes to be inserted in the adjacent field. The user can insert up to 10,000 virtual volumes at one time.
- ▶ **Ending VOLSER.** Select this option to insert a range of virtual volumes. Enter the ending VOLSER number in the adjacent field.
- ▶ **Initially owned by.** The name of the cluster that will own the new virtual volumes. Select a cluster from the menu.
- ▶ **Media type.** Media type of the virtual volumes. The following values are possible:
 - Cartridge System Tape (400 MiB)
 - Enhanced Capacity Cartridge System Tape (800 MiB)

See 1.6, “Data storage values” on page 12 for additional information about the use of binary prefixes.

- ▶ **Set Constructs.** Select this check box to specify constructs for the new virtual volumes. Then, use the menu under each construct to select a predefined construct name.

Set constructs only for virtual volumes that are used by hosts that are *not* multiple virtual systems (MVS) hosts. MVS hosts automatically assign constructs for virtual volumes and overwrite any manually assigned constructs.

The user can specify any or all of the following constructs: SG, MC, SC, or DC.

Modify Virtual Volumes window

Use the Modify Virtual Volumes window that is shown in Figure 9-72 to modify the constructs that are associated with existing virtual volumes in the TS7700 composite library.

Virtual Volumes (#BA068): Modify Virtual Volumes

Virtual Volume Details Refresh Last Refresh: Nov 23, 2012 10:15:12 PM

Insert Virtual Volumes

Modify Virtual Volumes ✓ The modify virtual volumes function is only valid for virtual volumes that are used by non-MVS hosts. MVS hosts will automatically assign constructs for virtual volumes.

Delete Virtual Volumes Show inserted volume ranges between 000000 and ZZZZZZ Show

Move Virtual Volumes

Virtual Volume Search

Volume Range:

*From: *To:

Storage Groups: Management Classes: Storage Classes: Data Classes:

No Change No Change No Change No Change

Modify

Figure 9-72 Modify Virtual Volumes window

Note: You can use the Modify Virtual Volume function to manage virtual volumes that belong to a *non-MVS* host that is not aware of constructs.

MVS hosts automatically assign constructs for virtual volumes, and manual changes are not recommended. The Modify Virtual Volumes window acts on any logical volume belonging to the cluster or grid regardless of the host that owns the volume. The changes that are made on this window take effect *only* on the modified volume or range *after* a mount-demount sequence, or by using the **LI REQ COPYRFSH** command.

To display a range of existing virtual volumes, enter the starting and ending VOLSERS in the fields at the top of the window and click **Show**.

To modify constructs for a range of logical volumes, identify a Volume Range, and then, click the **Constructs** menu to select construct values and click **Modify**. The menus have these options:

- ▶ 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 menus to change one or more constructs for the identified Volume Range. From each menu, the user 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: Changes the SG for the identified Volume Range.
 - Storage Classes: Changes the SC for the identified Volume Range.
 - Data Classes: Changes the DC for the identified Volume Range.
 - Management Classes: Changes the MC for the identified Volume Range.

The user is asked to confirm the decision to modify logical volume constructs. To continue with the operation, click **OK**. To abandon the operation without modifying any logical volume constructs, click **Cancel**.

Delete Virtual Volumes window

Use the Delete Virtual Volumes window that is shown in Figure 9-73 to delete *unused* virtual volumes from the TS7700 that are in the Insert Category (FF00).

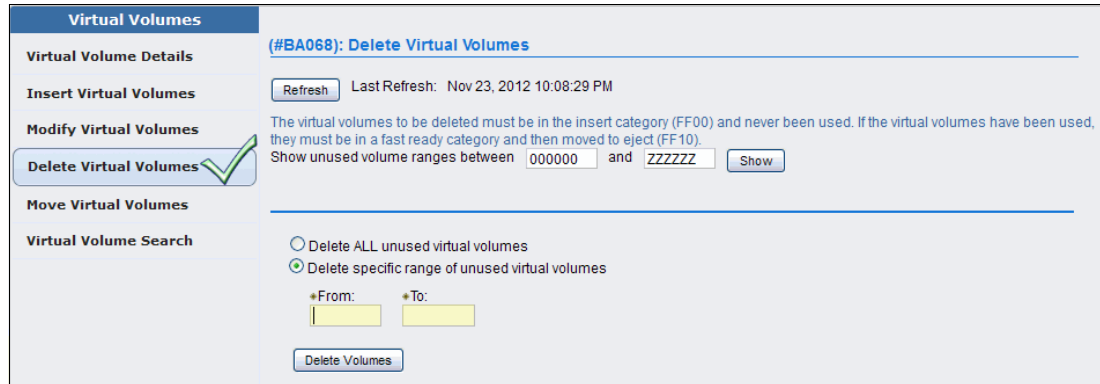


Figure 9-73 Delete Virtual Volumes window

Note: Only the unused logical volumes can be deleted through this window, meaning volumes in insert category FF00 that have never been mounted or have had their category, constructs, or attributes modified by a host. Otherwise, those logical volumes can be deleted only from the host.

The normal way to delete several virtual scratch volumes is by initiating the activities from the host. With Data Facility Storage Management Subsystem (DFSMS)/Removable Media Management (RMM) as the tape management system (TMS), it is done by using RMM commands.

To delete unused virtual volumes, select one of the options described next, and click **Delete Volumes**. A confirmation window is displayed. Click **OK** to delete or **Cancel** to cancel. To view the current list of unused virtual volume ranges in the TS7700 grid, enter a virtual volume range at the bottom of the window and click **Show**. A virtual volume range deletion can be canceled while in progress at the Cluster Operation History window.

This window has the following options:

- ▶ Delete ALL unused virtual volumes: Deletes all unused virtual volumes across all VOLSER ranges.
- ▶ Delete specific range of unused virtual volumes: All unused virtual volumes in the entered VOLSER range are deleted. Enter the VOLSER range:
 - From: The start of the VOLSER range to be deleted if **Delete specific range of unused virtual volumes** is selected.
 - To: The end of the VOLSER range to be deleted if **Delete specific range of unused virtual volumes** is selected.

Move Virtual Volumes window

Use the window that is shown in Figure 9-74 on page 401 to move a range of virtual volumes that are used by the TS7740 or TS7720T from one physical volume or physical volume range to a new target pool, or to cancel a move request already in progress. If a move operation is already in progress, a warning message displays. The user can view move operations in progress from the Events window.

Figure 9-74 MI Move Virtual Volumes

To cancel a move request, select the **Cancel Move Requests** link. The following options to cancel a move request are available:

- ▶ 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.

To move virtual volumes, 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 from where the virtual volumes must be removed. Either use this option or Existing Ranges to define the range of volumes to move, but not both.
 - From: VOLSER of the first physical volume in the range.
 - To: VOLSER of the last physical volume in the range.
- ▶ Existing Ranges: The list of existing physical volume ranges. 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 media type is in the range that is specified, no virtual volume is moved.
- ▶ Target Pool: The number (1 - 32) of the target pool to which virtual volumes are moved.
- ▶ Move Type: Used to determine when the move operation occurs. The following values are possible:
 - Deferred: Move operation will occur in the future as schedules enable.
 - Priority: Move operation occurs as soon as possible.
 - Honor Inhibit Reclaim schedule: An option of the Priority Move Type, it specifies that the move schedule occurs with the Inhibit Reclaim schedule. If this option is selected, the move operation does not occur when Reclaim is inhibited.

After defining the move operation parameters and clicking **Move**, confirm the request to move the virtual volumes from the defined physical volumes. If **Cancel** is selected, you return to the Move Virtual Volumes window.

Virtual Volumes Search window

To search for virtual volumes in a specific TS7700 cluster by VOLSER, category, media type, expiration date, or inclusion in a group or class, use the window that is shown in Figure 9-75. With the TS7720T, a new search option is available to search by Partition Number.

Virtual Volumes (#BA64A): Virtual Volume Search

Make selections below to define the database search criteria. The more search criteria used, the more restrictive the search.

[Previous Searches](#)

New Search Name:

Search Options

Volser:	Category:	Partition Number:	Media Type:
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Equal to <input type="text"/>	Expire Time: <input type="text"/>	Time Units: <input type="text"/>	Inserted Resident Partition (0) Tape Partition 1 (1)
Removal Residency: <input type="text"/>	Removal Time: <input type="text"/>		
Storage Preference: <input type="text"/>			
Storage Group: <input type="text"/>	or	--Storage Groups in Jafar--	<input type="text"/>
Management Class: <input type="text"/>	or	--Management Classes in Jafar--	<input type="text"/>
Storage Class: <input type="text"/>	or	--Storage Classes in Jafar--	<input type="text"/>
Data Class: <input type="text"/>	or	--Data Classes in Jafar--	<input type="text"/>

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 9-75 MI Virtual Volume Search entry window

You can view the results of a previous query, or create a query to search for virtual volumes.

Tip: Only one search can be run at a time. If a search is in progress, an information message displays at the top of the Virtual Volumes Search window. You can cancel a search in progress by clicking **Cancel Search**.

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 virtual volumes that match the search criteria.

Up to 10 previously named search queries can be saved. To clear the list of saved queries, select the check box next to one or more queries to be removed, select **Clear** from the menu, and click **Go**. This operation does not clear a search query already in progress.

Confirm the decision to clear the query list. Select **OK** to clear the list of saved queries, or **Cancel** 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 virtual volume search. The query name, criteria, start time, and end time are saved along with the search results.

To search for a specific VOLSER, enter parameters in the New Search Name and Volser fields and then click **Search**.

Figure 9-76 shows an example of the Virtual Volume Search Results window. The result can be printed or downloaded to a spreadsheet for post-processing.

Search Name: THAG28064301UT14
 Start Time: Aug 28, 2014, 8:43:01 AM
 End Time: Aug 28, 2014, 8:43:01 AM

Search Criteria Set:
 Partition Number: 2

Search Results:

Print report Download spreadsheet

Volser	Category	Partition N...	Media Type	Expire Time	Storage Gr...	Manageme...	Storage Cl...	Data Class
A50826	002F	2	ECCST	Not Set	SGG00001	MNSSN068	SCT0002K	D000N004
A50843	002F	2	ECCST	Not Set	SGG00001	MNSSN068	SCT0002R	D000N004

Page 1 of 1 Total: 2 Displayed: 2

Continuing here...

Mounted T...	Logical W...	Removal R...	Removal Ti...	Storage Preferen...
-	No	-	-	Defer migrate(PG1)
-	No	-	-	Defer migrate(PG1)

Figure 9-76 Virtual Volume Search Results window

When looking for the results of earlier searches, click **Previous Searches** on the Virtual Volume Search window, which is shown in Figure 9-75 on page 402.

- ▶ The following entry fields, which are shown in Figure 9-75 on page 402, can be used to format a Virtual Volume search: Volser (volume's serial number). This field can be left blank. The following wildcard characters in this field are valid:
 - Percent sign (%): Represents zero or more characters.
 - Asterisk (*): Converted to % (percent). Represents zero or more characters.
 - Period (.): Converted to _ (single underscore). Represents one character.
 - A single underscore (_): Represents one character.
 - Question mark (?): Converted to _ (single underscore). Represents one character.
- ▶ Category: The name of the category to which the virtual volume belongs. This value is a four-character hexadecimal string. For instance, 0002/0102 (scratch MEDIA2), 000E (error), 000F/001F (private), FF00 (insert) are possible values for Scratch and Specific categories. Wildcard characters can be used in this field. This field can be left blank.
- ▶ Media Type: The type of media on which the volume exists. Use the menu to select from the available media types. This field can be left blank.
- ▶ Expire Time: The amount of time in which virtual volume data expires. Enter a number. This field is qualified by the values Equal to, Less than, or Greater than in the preceding menu and defined by the succeeding menu under the heading Time Units. This field can be left blank.

- ▶ **Removal Residency:** The automatic removal residency state of the virtual volume. This field is not displayed for TS7740 clusters. In a TS7720T (tape attach) configuration, this field is displayed only when the volume is in partition 0 (CP0). The following values are possible:
 - Blank (ignore): If this field is empty (blank), the search ignores any values in the Removal Residency field. This is the default selection.
 - Removed: The search includes only virtual volumes that have been removed.
 - Removed Before: The search includes only virtual volumes that are removed before a specific date and time. If this value is selected, the Removal Time field must be complete as well.
 - Removed After: The search includes only virtual volumes that are removed after a certain date and time. If this value is selected, the Removal Time field must be complete as well.
 - In Cache: The search includes only virtual volumes in the cache.
 - Retained: The search includes only virtual volumes that are classified as retained.
 - Deferred: The search includes only virtual volumes that are classified as deferred.
 - Held: The search includes only virtual volumes that are classified as held.
 - Pinned: The search includes only virtual volumes that are classified as pinned.
 - No Removal Attempted: The search includes only virtual volumes that have not previously been subject to a removal attempt.
 - Removable Before: The search includes only virtual volumes that are candidates for removal before a specific date and time. If this value is selected, the Removal Time field must be complete as well.
 - Removable After: The search includes only virtual volumes that are candidates for removal after a specific date and time. If this value has been selected, the Removal Time field must complete as well.
- ▶ **Removal Time:** This field is displayed only if the grid contains a TS7720 or a TS7720T. Values that are displayed in this field depend on the values that are shown in the Removal Residency field. These values reflect the time zone in which the browser is:
 - Date: The calendar date according to month (M), day (D), and year (Y). It has the format *MM/DD/YYYY*. This field includes a date chooser calendar icon. The user can enter the month, day, and year manually, or can use the calendar chooser to select a specific date. The default for this field is blank.
 - Time: The Coordinated Universal Time (Coordinated Universal Time) in hours (H), minutes (M), and seconds (S). The values in this field accept the form *HH:MM:SS* only. Possible values for this field include 00:00:00 - 23:59:59. This field includes a time chooser clock icon. The user can enter hours and minutes manually by using 24-hour time designations, or can use the time chooser to select a start time based on a 12-hour (AM/PM) clock. The default for this field is midnight (00:00:00).
- ▶ **Volume Copy Retention Group:** The name of the Volume Copy Retention Group for the cluster.

The Volume Copy Retention Group provides more options to remove data from a disk-only TS7700 as the active data reaches full capacity. Volumes become candidates for removal if an appropriate number of copies exist on peer clusters *and* the volume copy retention time has elapsed since the volume was last accessed.

Volumes in each group are removed in order based on their least recently used access times. The volume copy retention time describes the number of hours a volume remains in cache before becoming a candidate for removal.

This field is only visible if the selected cluster is a TS7720 or TS7720T (for volumes in CP0). The following values are valid:

- Blank (ignore): If this field is empty (blank), the search ignores any values in the Volume Copy Retention Group field. This is the default selection.
- Prefer Remove: Removal candidates in this group are removed before removal candidates in the Prefer Keep group.
- Prefer Keep: Removal candidates in this group are removed after removal candidates in the Prefer Remove group.
- Pinned: Copies of volumes in this group are never removed from the accessing cluster. The volume copy retention time does not apply to volumes in this group. Then, volumes in this group that are moved to scratch become priority candidates for removal.

Tip: To avoid cache overruns, plan ahead when assigning volumes to this group.

- “-”: Volume Copy Retention does not apply to the TS7740 cluster and TS7720T (for volume in CP1 to CP7). This value (a dash indicating an empty value) is displayed if the cluster attaches to a physical tape library.
- ▶ Storage Group: The name of the SG in which the virtual volume is. The user can enter a name in the empty field, or select a name from the adjacent menu. This field can be left blank.
- ▶ Management Class: The name of the MC to which the virtual volume belongs. The user can enter a name in the empty field, or select a name from the adjacent menu. This field can be left blank.
- ▶ Storage Class: The name of the SC to which the virtual volume belongs. The user can enter a name in the empty field, or select a name from the adjacent menu. This field can be left blank.
- ▶ Data Class: The name of the DC to which the virtual volume belongs. The user can enter a name in the empty field, or select a name from the adjacent menu. This field can be left blank.
- ▶ Mounted: Whether the virtual volume is mounted. The following values are possible:
 - Ignore: Ignores any values in the Mounted field. This is the default selection.
 - Yes: Includes only mounted virtual volumes.
 - No: Includes only unmounted virtual volumes.
- ▶ Logical WORM: Whether the logical volume is defined as Write Once Read Many (WORM). The following values are possible:
 - Ignore: Ignores any values in the Logical WORM field. This is the default selection.
 - Yes: Includes only WORM logical volumes.
 - No: Does not include any WORM logical volumes.

Remember: The user can print or download the results of a search query by using Print Report or Download Spreadsheet on the Volumes found table at the end of the Search Results window, as shown in Figure 9-76 on page 403.

Categories

Use this window to add, modify, or delete a scratch (Fast Ready) category of virtual volumes. The user can also use this window to view the total number of volumes that are defined by custom, inserted, and damaged categories. A *category* is a grouping of virtual volumes for a predefined use. A *scratch (Fast Ready) category* groups virtual volumes for non-specific use. This grouping enables faster mount times because the TS7700 can order category mounts without recalling data from a stacked volume.

Figure 9-77 shows the Category window in the TS7700 MI.

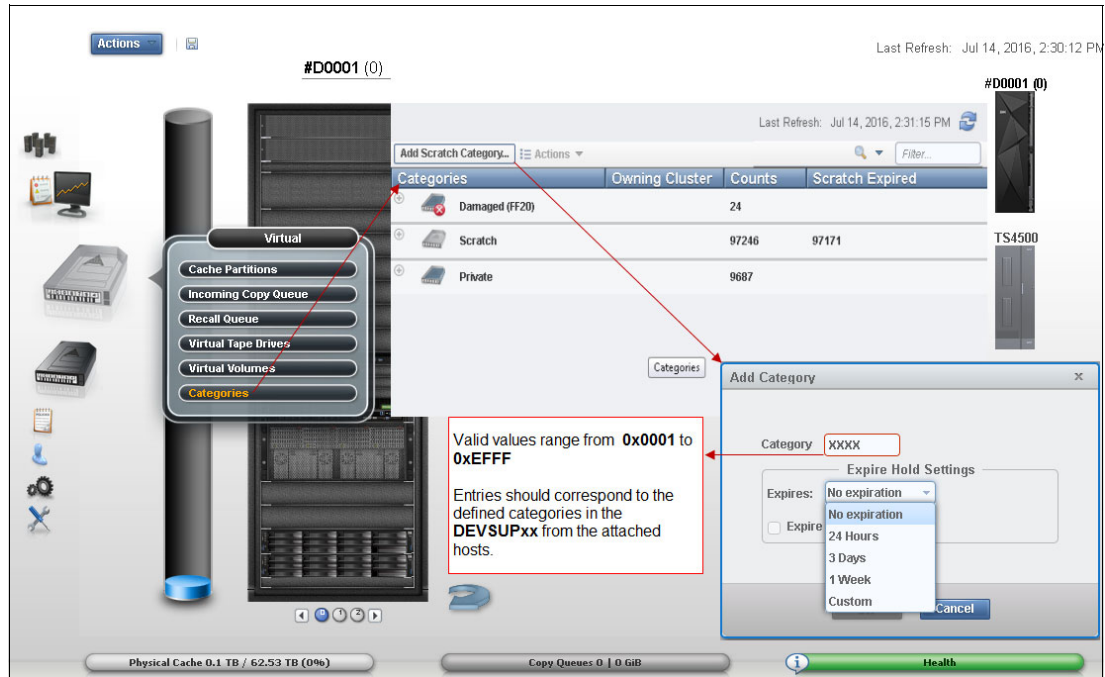


Figure 9-77 Categories window

You can display the already defined categories, as shown in the Figure 9-78.

Categories	Owning Cluster	Counts	Scratch Expired
Scratch		3179593	
Private		819979	
1111		1459	
500F		817323	
	Pesto (Cluster 0)	527311	
	Squint (Cluster 1)	0	
	Tom (Cluster 3)	286869	
	Spike (Cluster 5)	3143	
1112		1188	
111F		9	

Figure 9-78 Displaying existing categories

Table 9-10 describes values that are displayed on the Categories table, as shown in Figure 9-78 on page 406.

Table 9-10 Category values

Column name	Description
Categories	<p>The type of category that defines the virtual volume. The following values are valid:</p> <ul style="list-style-type: none"> ▶ Scratch: Categories within the user-defined private range 0x0001 through 0xEFFF that are defined as scratch (Fast Ready). Click the plus sign (+) icon to expand this heading and reveal the list of categories that are defined by this type. Expire time and hold values are shown in parentheses next to the category number. See Table 9-10 for descriptions of these values. ▶ Private: Custom categories that are established by a user, within the range of 0x0001 - 0xEFFF. Click the plus sign (+) icon to expand this heading and reveal the list of categories that are defined by this type. ▶ Damaged: A system category that is identified by the number 0xFF20. Virtual volumes in this category are considered damaged. ▶ Insert: A system category that is identified by the number 0xFF00. Inserted virtual volumes are held in this category until moved by the host into a scratch category. <p>If no defined categories exist for a certain type, that type is not displayed on the Categories table.</p>
Owning Cluster	Names of all clusters in the grid. Expand a category type or number to display. This column is visible only when the accessing cluster is part of a grid.
Counts	The total number of virtual volumes according to category type, category, or owning cluster.
Scratch Expired	The total number of scratch volumes per owning cluster that are expired. The total of all scratch expired volumes is the number of ready scratch volumes.

The user can use the Categories table to add, modify, or delete a scratch category, or to change the way information is displayed.

Tip: The total number of volumes within a grid is *not* always equal to the sum of all category counts. Volumes can change category multiple times per second, which makes the snapshot count obsolete.

Table 9-11 describes the actions that can be performed on the Categories window.

Table 9-11 Available actions on the Categories window

Action	Steps to perform action
Add a scratch category	<ol style="list-style-type: none"> 1. Select Add Scratch Category. 2. Define the following category properties: <ul style="list-style-type: none"> – Category: A four-digit hexadecimal number that identifies the category. The valid characters for this field are A - F and 0 - 9. Do 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: The amount of time after a virtual volume is returned to the scratch (Fast Ready) category before its data content is automatically delete-expired. Select an expiration time from the menu. If the user selects No Expiration, volume data never automatically delete-expires. If the user selects Custom, enter values for the following fields: <ul style="list-style-type: none"> • Time: Enter a number in the field according to these restrictions: <ul style="list-style-type: none"> 1 - 32,767 if unit is hours 1 - 1365 if unit is days 1 - 195 if unit is weeks • Time Unit: Select a corresponding unit from the menu. – Set Expire Hold: Check this box to prevent the virtual volume from being mounted or having its category and attributes changed before the expire time has elapsed. Checking this field activates the hold state for any volumes currently in the scratch (Fast Ready) category and for which the expire time has not yet elapsed. Clearing this field removes the access restrictions on all volumes currently in the hold state within this scratch (Fast Ready) category.
Modify a scratch category	<p>The user can modify a scratch category in two ways:</p> <ul style="list-style-type: none"> ▶ Select a category on the table, and then, select Actions → Modify Scratch Category. ▶ Right-click a category on the table and either hold, or select Modify Scratch Category from the menu. <p>The user can modify the following category values:</p> <ul style="list-style-type: none"> ▶ Expire ▶ Set Expire Hold <p>The user can modify one category at a time.</p>
Delete a scratch category	<p>The user can delete a scratch category in two ways:</p> <ol style="list-style-type: none"> 1. Select a category on the table, and then, select Actions → Delete Scratch Category. 2. Right-click a category on the table and select Delete Scratch Category from the menu. <p>The user can delete only one category at a time.</p>
Hide or show columns on the table	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Click the check box next to a column heading to hide or show that column in the table. Column headings that are checked display on the table.

Action	Steps to perform action
Filter the table data	<p>Follow these steps to filter by using a string of text:</p> <ol style="list-style-type: none"> 1. Click in the Filter field. 2. Enter a search string. 3. Press Enter. <p>Follow these steps to filter by column heading:</p> <ol style="list-style-type: none"> 1. Click the down arrow next to the Filter field. 2. Select the column heading to filter by. 3. Refine the selection: <ul style="list-style-type: none"> – Categories: Enter a whole or partial category number and press Enter. – Owing Cluster: Enter a cluster name or number and press Enter. Expand the category type or category to view highlighted results. – Counts: Enter a number and press Enter to search on that number string. – Scratch Expired: Enter a number and press Enter to search on that number string.
Reset the table to its default view	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Click Reset Table Preferences.

Note: There is no cross-check between defined categories in the z/OS systems and the definitions in the TS7700.

9.2.7 The Physical icon

The topics in this section present information that is related to monitoring and manipulating physical volumes in the TS7740 and TS7720T clusters. To view or modify settings for physical volume pools to manage the physical volumes that are used by the tape-attached clusters, use the window that is shown in Figure 9-79.

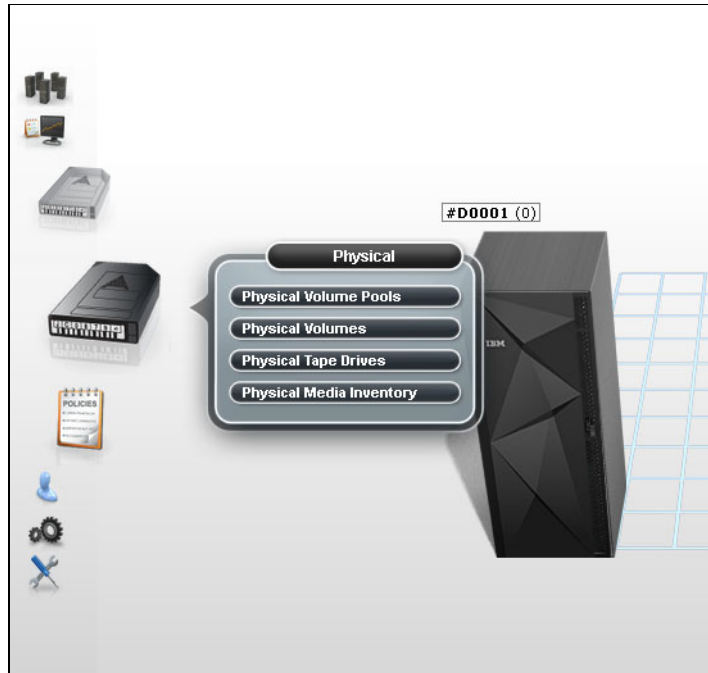


Figure 9-79 Physical icon

Physical Volume Pools

The Physical Volume Pools properties table displays the media properties and encryption settings for every physical volume pool that is defined for a specific TS7700T cluster in the grid. This table contains these tabs:

- ▶ Pool Properties
- ▶ Encryption Settings

Tip: Pools 1 - 32 are preinstalled and initially set to default attributes. Pool 1 functions as the default pool and is used if no other pool is selected.

Figure 9-80 on page 411 show an example of the Physical Volume Pools window. There is a link that is available for a tutorial showing how to modify pool encryption settings. Click the link to see the tutorial material. This window is visible but disabled on the TS7700 MI if the grid possesses a physical library, but the selected cluster does not. This message is displayed:

The cluster is not attached to a physical tape library.

You can use the window that is shown in Figure 9-80 on page 411 to view or modify settings for physical volume pools.

"Bluto[1]" (#BA67B): Physical Volume Pools

Refresh Last Refresh: Nov 22, 2012 6:21:25 PM

View tutorial

Physical volume pool properties:

Pool Properties Physical Tape Encryption Settings

Select Action

Select	Pool	Media	Second Media	Borrow Indicator	Re...	Maximum Devic...	Export Pool	Export Format	Da...	Da...	Ag...	Da...	Ma...	Reclai...
<input type="checkbox"/>	1	3592	None	Borrow, Return	1	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	2	3592	None	Borrow, Return	2	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	3	3592	None	Borrow, Return	3	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	4	3592	None	Borrow, Return	4	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	5	3592	Any 3592	Borrow, Return	5	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	6	3592	Any 3592	Borrow, Return	6	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	7	3592	Any 3592	Borrow, Return	7	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	8	3592	Any 3592	Borrow, Return	8	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	9	3592	Any 3592	Borrow, Return	9	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	10	3592	Any 3592	Borrow, Return	10	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	11	3592	Any 3592	Borrow, Return	11	All Devices	Not Defined	Default	0	0	0	0	0	35
<input type="checkbox"/>	12	3592	Any 3592	Borrow, Return	12	All Devices	Not Defined	Default	0	0	0	0	0	35
Total: 32														

Figure 9-80 Physical Volume Pools Properties table

The Physical Volume Pool Properties table displays the encryption setting and media properties for every physical volume pool that is defined in a TS7740 and TS7720T. This table contains two tabs: Pool Properties and Physical Tape Encryption Settings. The information that is displayed in a tape-attached cluster depends on the current configuration and media availability. Figure 9-81 shows an example in a TS7760T attached to a TS3500 tape library.

test1 #00186 > #D0001 (Cluster 0) - Physical Volume Pools IBM™ TS7700

Selected pools:

Pool: 2

Pool Properties:

Media Class: 3592

First Media (Primary): 2

Second Media (Secondary): 1

Borrow Indicator: 3

Reclaim Pool: 4

Maximum Devices: 5

Export Pool: 6

Export Format: 7

Days Before Secure Data Erase: 8

Days Without Access: 9

Age of Last Data Written: 10

Days Without Data Inactivation: 11

Maximum Active Data: 12

Reclaim Threshold Percentage (%): 13

Sunset Media Reclaim Threshold Percentage (%): 14

TS4500 Tape Library Expanded

TS4500 - Logical Library Hydra_VEC

Tape Library Health: Normal

Fibre Switch Health: Normal

Tape Drive Health: Normal

Tape Drive Path Health: Normal

State: Online

Operational Mode: Auto

Frame Doors: Closed

Virtual I/O Slots: Empty

Physical Cartridges: 21

Tape Drives (Available/Total): 4/4

All Compatible Devices

1

2

3

4

All Compatible Devices

Figure 9-81 Pool properties and media selection

The following information is displayed:

► Under Pool Properties:

- Pool: The pool number. This number is a whole number 1 - 32, inclusive.
- Media Class: The supported media class of the storage pool. The valid value is 3592.
- First Media (Primary): The primary media type that the pool can borrow from or return to the common scratch pool (Pool 0). The values that are displayed in this field are dependent upon the configuration of physical drives in the cluster. See Table 4-9 on page 132 for First and Second Media values based on drive configuration.

The primary media type can have the following values:

Any 3592	Any media with a 3592 format.
None	The only option available if the Primary Media type is any 3592. This option is only valid when the Borrow Indicator field value is No Borrow, Return or No Borrow, Keep.
JA	Enterprise Tape Cartridge (ETC).
JB	Extended Data Enterprise Tape Cartridge (EETC).
JC	Advanced Type C Data (ATCD)
JD	Advanced Type D Data (ATDD)
JJ	Enterprise Economy Tape Cartridge (EETC)
JK	Advanced Type K Economy (ATKE)
JL	Advanced Type L Economy (ATLE)

- Second Media (Secondary): The second choice of media type from which the pool can borrow. Options that are shown exclude the media type chosen for First Media. The following values are possible:

Any 3592	Any media with a 3592 format.
None	The only option available if the Primary Media type is any 3592. This option is only valid when the Borrow Indicator field value is No Borrow, Return or No Borrow, Keep.
JA	ETC.
JB	EETC.
JC	ATCD.
JD	ATDD.
JJ	EETC.
JK	ATKE.
JL	ATLE.

- Borrow Indicator: Defines how the pool is populated with scratch cartridges. The following values are possible:

Borrow, Return	A cartridge is borrowed from the Common Scratch Pool (CSP) and returned to the CSP when emptied.
Borrow, Keep	A cartridge is borrowed from the CSP and retained by the actual pool, even after being emptied.

- No Borrow, Return** A cartridge is not borrowed from CSP, but an emptied cartridge is placed in CSP. This setting is used for an empty pool.
- No Borrow, Keep** A cartridge is not borrowed from CSP, and an emptied cartridge is retained in the actual pool.

- Reclaim Pool: The pool to which virtual volumes are assigned when reclamation occurs for the stacked volume on the selected pool.

Important: The reclaim pool that is designated for the Copy Export pool needs to be set to the same value as the Copy Export pool. If the reclaim pool is modified, Copy Export disaster recovery capabilities can be compromised.

If there is a need to modify the reclaim pool that is designated for the Copy Export pool, the reclaim pool *cannot* be set to the same value as the primary pool or the reclaim pool that is designated for the primary pool. If the reclaim pool for the Copy Export pool is the same as either of the other two pools that are mentioned, the primary and backup copies of a virtual volume might exist on the same physical media. If the reclaim pool for the Copy Export pool is modified, it is the user's responsibility to Copy Export volumes from the reclaim pool.

- Maximum Devices: The maximum number of physical tape drives that the pool can use for premigration.
- Export Pool: The type of export that is supported if the pool is defined as an Export Pool (the pool from which physical volumes are exported). The following values are possible:

Not Defined The pool is not defined as an Export pool.

Copy Export The pool is defined as a Copy Export pool.

- Export Format: The media format used when writing volumes for export. This function can be used when the physical library recovering the volumes supports a different media format than the physical library exporting the volumes. This field is only enabled if the value in the Export Pool field is Copy Export. The following values are valid for this field:

Default The highest common format that is supported across all drives in the library. This is also the default value for the Export Format field.

E06 Format of a 3592-E06 Tape Drive.

E07 Format of a 3592-E07 Tape Drive.

E08 Format of a 3592-E08 Tape Drive

- 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 virtual volume on the stacked physical volume is accessed. Secure Data Erase occurs later, 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 as a candidate for reclamation. Each physical stacked volume has a timer for this purpose, which is reset when a virtual volume is accessed. The reclamation occurs later, 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 virtual volume is accessed. The reclamation occurs later, 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 that the data ratio of the pool has been higher than the Maximum Active Data used 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 later, 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 is not 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 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 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 is performed on the stacked volume. The valid range of possible values is 0 - 95% and can be selected in 5% increments; 35% is the default value.
- ▶ Sunset Media Reclaim Threshold: Lists the percentage that is used to determine when to reclaim sunset media. This is a new option of Release 3.3. To modify pool properties, select the check box next to one or more pools that are shown on the Pool Properties tab, select **Modify Pool Properties** from the menu, and click **Go**.
- ▶ Physical Tape Encryption Settings: The Physical Tape Encryption Settings tab displays the encryption settings for physical volume pools. The following encryption information is displayed on this tab:
 - Pool: The pool number. This number is a whole number 1 - 32, inclusive.
 - Encryption: The encryption state of the pool. The following values are possible:

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.
 - Key Mode 1: Encryption mode that is used with Key Label 1. The following values are available:

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: The current encryption key (EK) 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. Therefore, key labels are reported by using uppercase characters.

Note: You can use identical values in Key Label 1 and Key Label 2, but you must define each label for each key.

If the encryption state is Disabled, this field is blank. If the default key is used, the value in this field is default key.

- Key Mode 2: Encryption mode that is used with Key Label 2. The following values are valid:

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.
“-”	The default key is in use.

- Key Label 2: The current EK 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. Therefore, key labels are reported by 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, complete these steps:

1. Select one or more pools that are shown on the Physical Tape Encryption Settings tab.
2. Select **Modify Encryption Settings** from the menu and click **Go**.

Physical volumes

The topics in this section present information that is related to monitoring and manipulating physical volumes in the TS7700T and TS7740. This window is visible but disabled on the TS7700 MI if the grid possesses a physical library, but the selected cluster does not.

The following message is displayed:

The cluster is not attached to a physical tape library.

Tip: This window is not visible on the TS7700 MI if the grid does not possess a physical library.

TS7700 MI windows that are under the Physical icon can help you view or change settings or actions that are related to the physical volumes and pools, physical drives, media inventory, TVC, and a physical library. Figure 9-82 shows the navigation and the Physical Volumes window. Physical Volumes page is available for all TS7700 tape-attached models. Picture shows a TS7760T model attached to a TS4500 tape library.

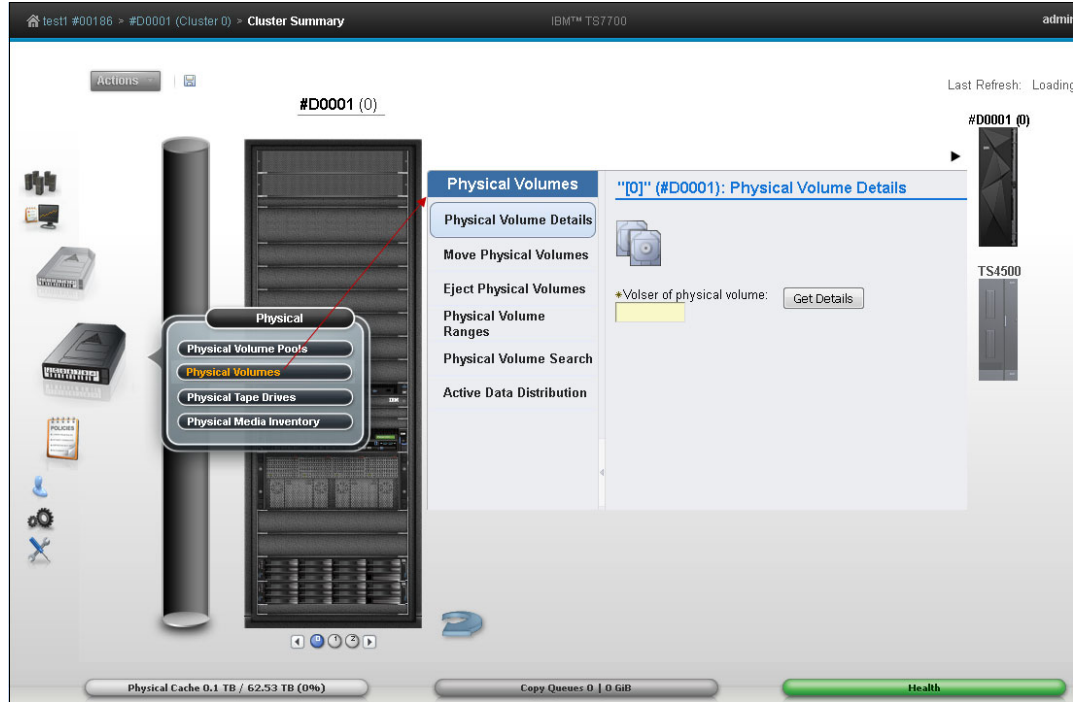


Figure 9-82 Physical Volumes navigation and options

The following options are available selections under Physical Volumes:

- ▶ Physical Volume Details
- ▶ Move Physical Volumes window
- ▶ Eject Physical Volumes window
- ▶ Physical Volume Ranges window
- ▶ Physical Volume Search window
- ▶ Active Data Distribution

Physical Volume Details

Use this window to obtain detailed information about a physical stacked volume in the TS7740 and TS7700T clusters. Figure 9-83 on page 417 shows a sample of the Physical Volume Details window.

You can download the list of virtual volumes in the physical stacked volume being displayed by clicking **Download List of Virtual Volumes** under the table.

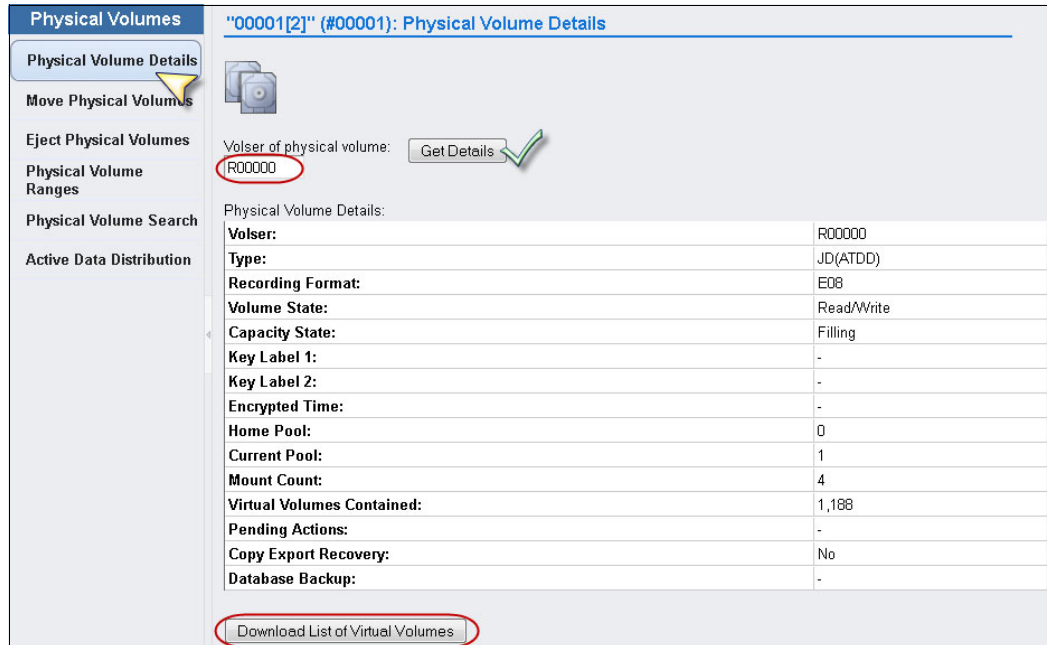


Figure 9-83 Physical Volume Details window

The following information is displayed when details for a physical stacked volume are retrieved:

- ▶ **VOLSER.** Six-character VOLSER number of the physical stacked volume.
- ▶ **Type.** The media type of the physical stacked volume. The following values are possible:
 - **JA (ETC).** ETC.
 - **JB (ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC (ATCD).** ATCD.
 - **JD (ATDD).** ATDD.
 - **JJ (EETC).** EETC.
 - **JK (ATKE).** ATKE.
 - **JL(ATLE).** ATLE.

Note: JD (ATDD) and JL (ATLE) media types are only available if the highest common format (HCF) is set to E08 or higher.

- ▶ **Recording Format.** The format that is used to write the media. The following values are possible:
 - **Undefined.** The recording format that is used by the volume is not recognized as a supported format.
 - **J1A**
 - **E05**
 - **E05E.** E05 with encryption.
 - **E06**
 - **E06E.** E06 with encryption.

- **E07**
- **E07E.** E07 with encryption.
- **E08**
- **E08E.** E08 with encryption.
- ▶ **Volume State** The following values are possible:
 - **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.
 - **Copy Export Pending.** The volume is in a pool that is being exported as part of an in-progress Copy Export.
 - **Copy Exported.** The volume has been ejected from the library and removed to offsite storage.
 - **Copy Export Reclaim.** The host can send a Host Console Query request to reclaim a physical volume currently marked Copy Exported. The data mover then reclaims the virtual volumes from the primary copies.
 - **Copy Export No Files Good.** The physical volume has been ejected from the library and removed to offsite storage. The virtual volumes on that physical volume are obsolete.
 - **Misplaced.** The library cannot locate the specified volume.
 - **Inaccessible.** The volume exists in the library inventory but is in a location that the cartridge accessor cannot access.
 - **Manually Ejected.** The volume was previously present in the library inventory, but cannot currently be located.
- ▶ **Capacity State.** Possible values are empty, filling, and full.
- ▶ **Key Label 1/Key Label 2.** The EK label that is associated with a physical volume. Up to two key labels can be present. If there are no labels present, the volume is not encrypted. If the EK used is the default key, the value in this field is default key.
- ▶ **Encrypted Time.** The date the physical volume was first encrypted using the new EK. If the volume is not encrypted, the value in this field is “-”.
- ▶ **Home Pool.** The pool number to which the physical volume was assigned when it was inserted into the library, or the pool to which it was moved through the library manager Move/Eject Stacked Volumes function.
- ▶ **Current Pool.** The current storage pool in which the physical volume is.
- ▶ **Mount Count.** The number of times the physical volume has been mounted since being inserted into the library.
- ▶ **Virtual Volumes Contained.** Number of virtual volumes that are contained on this physical stacked volume.
- ▶ **Pending Actions.** Whether a move or eject operation is pending. The following values are possible:
 - **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.
- ▶ **Copy Export Recovery.** Whether the database backup name is valid and can be used for recovery. Possible values are **Yes** and **No**.
- ▶ **Database Backup.** The time stamp portion of the database backup name.

Move Physical Volumes window

To move a range or quantity of physical volumes that is used by the TS7720T or TS7740 to a target pool, or cancel a previous move request, use the window that is shown in Figure 9-84.



Figure 9-84 Move Physical Volumes options

The Select Move Action menu provides options for moving physical volumes to a target pool. The following options are available to move physical volumes to a target pool:

- ▶ **Move Range of Physical Volumes.** Moves physical volumes to the target pool physical volumes in the specified range. This option requires you to select a Volume Range, Target Pool, and Move Type. The user can also select a Media Type.
- ▶ **Move Range of Scratch Only Volumes.** Moves physical volumes to the target pool scratch volumes in the specified range. This option requires you to select a Volume Range and Target Pool. The user can also select a Media Type.
- ▶ **Move Quantity of Scratch Only Volumes.** Moves a specified quantity of physical volumes from the source pool to the target pool. This option requires to select Number of Volumes, Source Pool, and Target Pool. The user can also select a Media Type.
- ▶ **Move Export Hold to Private.** Moves all Copy Export volumes in a source pool back to a private category if the volumes are in the Export/Hold category but are not selected to be ejected from the library. This option requires to select a Source Pool.
- ▶ **Cancel Move Requests.** Cancels any previous move request.

If the user selects **Move Range of Physical Volumes** or **Move Range of Scratch Only Volumes** from the **Select Move Action** menu, the user must define a volume range or select an existing range, select a target pool, and identify a move type. A media type can be selected as well.

If the user selects **Move Quantity of Scratch Only Volumes** from the **Select Move Action** menu, the user must define the number of volumes to be moved, identify a source pool, and identify a target pool. A media type can be selected as well.

If the user selects **Move Export Hold to Private** from the **Select Move Action** menu, the user must identify a source pool.

The following move operation parameters are available:

- ▶ **Volume Range.** The range of physical volumes to move. The user can use either this option or the Existing Ranges option to define the range of volumes to move, but not both. Specify the range:
 - **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. The user can use either this option or the Volume Range option to define the range of volumes to move, but not both.
- ▶ **Source Pool.** The number (0 - 32) of the source pool from which physical volumes are moved. If the user is selecting a source pool for a Move Export Hold to Private operation, the range of volumes that is displayed is 1 - 32.
- ▶ **Target Pool.** The number (0 - 32) of the target pool to which physical volumes are moved.
- ▶ **Move Type.** Used to determine when the move operation occurs. The following values are possible:
 - **Deferred Move.** The move operation occurs based on the first Reclamation policy that is triggered for the applied source pool. This operation depends on reclaim policies for the source pool and might take some time to complete.
 - **Priority Move.** The move operation occurs as soon as possible. Use this option to complete the operation sooner.
 - **Honor Inhibit Reclaim schedule.** An option of the Priority Move Type, it specifies that the move schedule occurs with the Inhibit Reclaim schedule. If this option is selected, the move operation does not occur when Reclaim is inhibited.
- ▶ **Number of Volumes.** The number of physical volumes to be moved.
- ▶ **Media Type.** Specifies the media type of the physical volumes in the range to be moved. The physical volumes in the range that is specified to be moved must be of the media type that is designated by this field, or the move operation fails.

After the user defines move operation parameters and clicks **Move**, the user confirms the request to move physical volumes. If the user selects **Cancel**, the user returns to the Move Physical Volumes window. To cancel a previous move request, select **Cancel Move Requests** from the **Select Move Action** menu. The following options are available to cancel a move request:

- ▶ **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

To eject a range or quantity of physical volumes that is used by the TS7720T or TS7740, or to cancel a previous eject request, use the window that is shown in Figure 9-85.

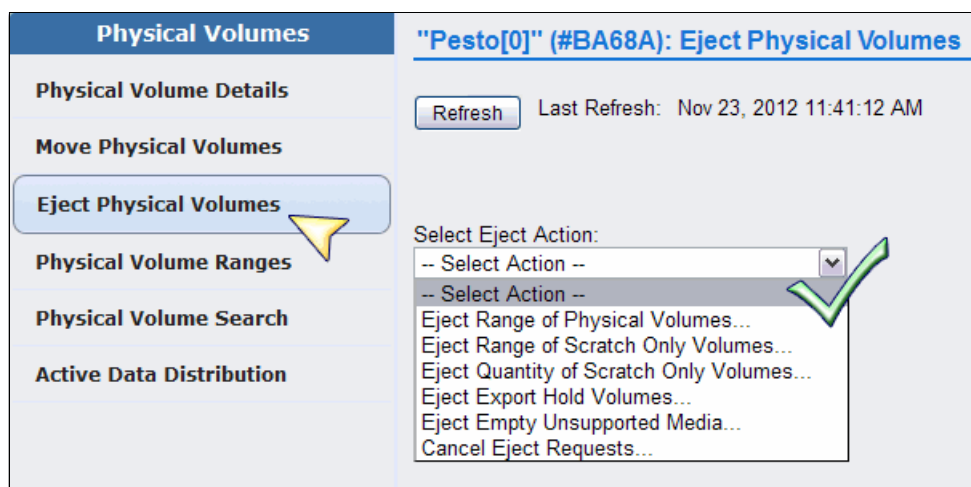


Figure 9-85 Eject Physical Volumes window

The Select Eject Action menu provides options for ejecting physical volumes.

Note: Before a stacked volume with active virtual volumes can be ejected, all active logical volumes in it are copied to a different stacked volume.

The following options are available to eject physical volumes:

- ▶ **Eject Range of Physical Volumes.** Ejects physical volumes in the range that is specified. This option requires you to select a volume range and eject type. A media type can be selected as well.
- ▶ **Eject Range of Scratch Only Volumes.** Ejects scratch volumes in the range that is specified. This option requires you to select a volume range. A media type can be selected as well.
- ▶ **Eject Quantity of Scratch Only Volumes.** Ejects a specified quantity of physical volumes. This option requires you to select several volumes and a source pool. A media type can be selected as well.
- ▶ **Eject Export Hold Volumes.** Ejects a subset of the volumes in the Export/Hold Category.
- ▶ **Eject Empty Unsupported Media.** Ejects physical volumes on unsupported media after the existing read-only data is migrated to new media.
- ▶ **Cancel Eject Requests.** Cancels any previous eject request.

If the user selects **Eject Range of Physical Volumes** or **Eject Range of Scratch Only Volumes** from the **Select Eject Action** menu, the user must define a volume range or select an existing range and identify an eject type. A media type can be selected as well.

If the user selects **Eject Quantity of Scratch Only Volumes** from the **Select Eject Action** menu, the user must define the number of volumes to be ejected, and to identify a source pool. A media type can be selected as well.

If the user selects **Eject Export Hold Volumes** from the **Select Eject Action** menu, the user must select the VOLSERS of the volumes to be ejected. To select all VOLSERS in the Export Hold category, select **Select All** from the menu. The eject operation parameters include these parameters:

- ▶ **Volume Range.** The range of physical volumes to eject. The user can use either this option or the Existing Ranges option to define the range of volumes to eject, but not both. Define the range:
 - **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. The user can use either this option or the Volume Range option to define the range of volumes to eject, but not both.
- ▶ **Eject Type.** Used to determine when the eject operation will occur. The following values are possible:
 - ▶ **Deferred Eject.** The eject operation occurs based on the first Reclamation policy that is triggered for the applied source pool. This operation depends on reclaim policies for the source pool and can take some time to complete.
 - ▶ **Priority Eject.** The eject operation occurs as soon as possible. Use this option to complete the operation sooner.
 - **Honor Inhibit Reclaim schedule.** An option of the Priority Eject Type, it specifies that the eject schedule occurs with the Inhibit Reclaim schedule. If this option is selected, the eject operation does not occur when Reclaim is inhibited.
- ▶ **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 that are specified to eject must be of the media type designated by this field, or else the eject operation fails.

After the user defines the eject operation parameters and clicks **Eject**, the user must confirm the request to eject physical volumes. If the user selects **Cancel**, the user returns to the Eject Physical Volumes window.

To cancel a previous eject request, select **Cancel Eject Requests** from the **Select Eject Action** menu. The following options are available to cancel an eject request:

- ▶ **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

To view physical volume ranges or unassigned physical volumes in a library that is attached to an TS7700T cluster, use this window.

Figure 9-86 on page 423 shows the Physical Volume Ranges window and options. When working with volumes recently added to the attached TS3500 tape library that are not showing in the Physical Volume Ranges window, click **Inventory Upload**. This action requests the physical inventory from the defined logical library in the TS3500 to be uploaded to the TS7720T or TS7740, repopulating the Physical Volume Ranges window.

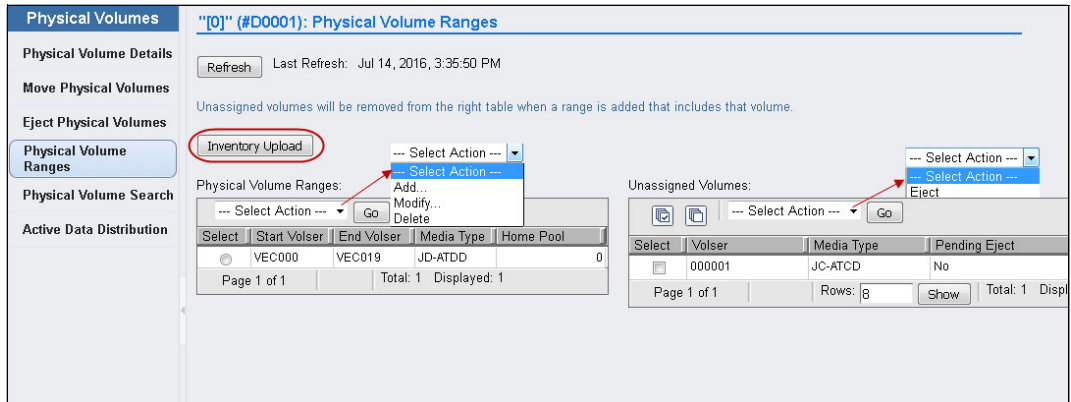


Figure 9-86 Physical Volume Ranges window

Tip: When inserting a VOLSER that belongs to a defined tape attach TS7700 range, it is presented and inventoried according to the setup in place. If the newly inserted VOLSER does not belong to any defined range in the TS7700T, an intervention-required message is generated, requiring the user to correct the assignment for this VOLSER.

Important: If a physical volume range contains virtual volumes with active data, those virtual volumes must be moved or deleted before the physical volume range can be moved or deleted.

The following information is displayed in the Physical Volume Ranges table:

- ▶ **Start VOLSER.** The first VOLSER in a defined range.
- ▶ **End VOLSER.** The last VOLSER in a defined range.
- ▶ **Media Type.** The media type for all volumes in a VOLSER range. The following values are possible:
 - **JA-ETC.** ETC.
 - **JB(ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD).** ATCD.
 - **JD(ATDD).** ATDD.
 - **JJ(EETC).** EETC.
 - **JK(ATKE).** ATKE.
 - **JL(ATLE).** ATLE.

Note: JA and JJ media are supported only for read-only operations with 3592 E07 tape drives. 3592-E08 does not support JA, JJ, or JB media.

- ▶ **Home Pool.** The home pool to which the VOLSER range is assigned.

Use the menu on the Physical Volume Ranges table to add a VOLSER range, or to modify or delete a predefined range.

Unassigned Volumes

The Unassigned Volumes table displays the list of unassigned physical volumes that are pending ejection for a cluster. A VOLSER is removed from this table when a new range that

contains the VOLSER is added. The following status information is displayed in the Unassigned Volumes table:

- ▶ **VOLSER.** The VOLSER associated with a given physical volume.
- ▶ **Media Type.** The media type for all volumes in a VOLSER range. The following values are possible:
 - **JA-ETC.** ETC.
 - **JB(ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD).** ATCD.
 - **JD(ATDD).** ATDD.
 - **JJ(EETC).** EETC.
 - **JK(ATKE).** ATKE.
 - **JL(ATLE).** ATLE.

Note: JA and JJ media are supported only for read-only operations with 3592 E07 tape drives. 3592-E08 does not support JA, JJ, or JB media.

- ▶ **Pending Eject.** Whether the physical volume associated with the VOLSER is awaiting ejection.

Use the Unassigned Volumes table to eject one or more physical volumes from a library that is attached to a TS7720T or TS7740.

Physical Volume Search window

To search for physical volumes in a TS7720T or TS7740 cluster according to one or more identifying features, use this window.

Figure 9-87 on page 425 shows the Physical Volume Search window. Click the **Previous Searches** hyperlink to view the results of a previous query on the Previous Physical Volumes Searches window.

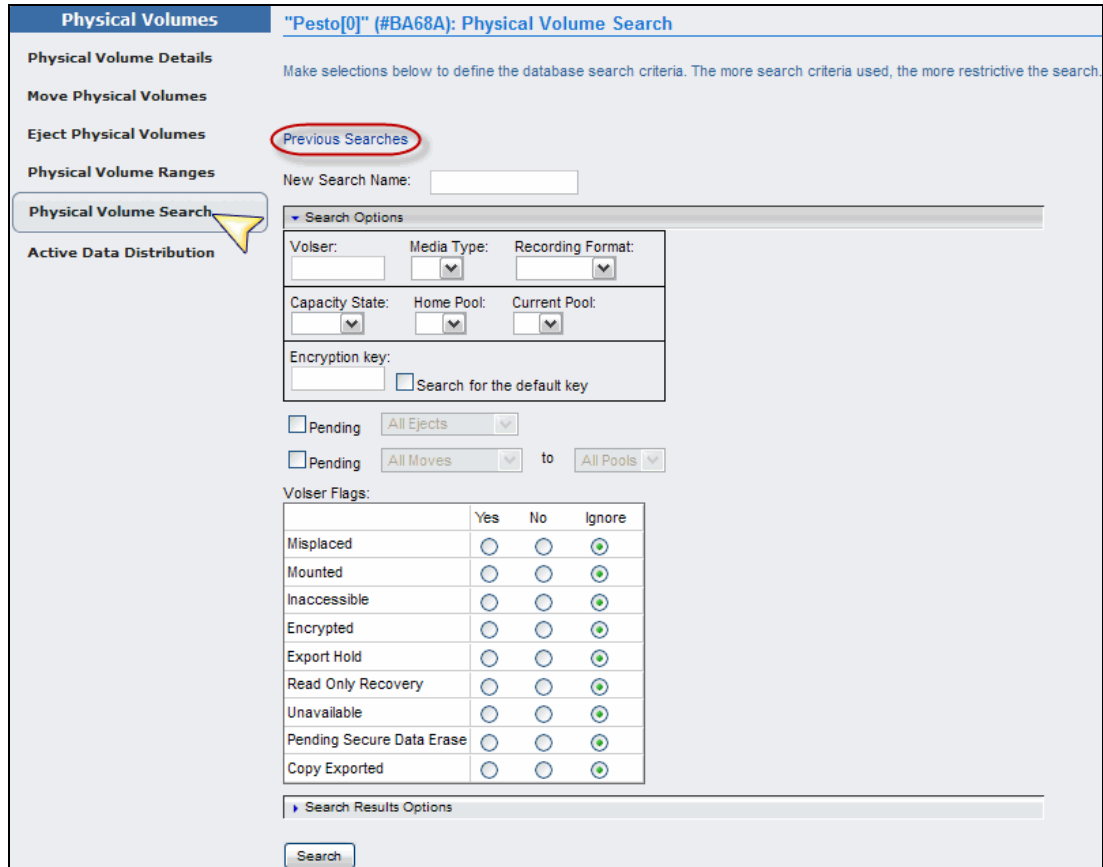


Figure 9-87 Physical Volume Search window

The following information can be seen and requested on the Physical Volume Search window:

- ▶ **New Search Name.** Use this field to create a new search query.
Enter a name for the new query in the New Search Name field.
Enter values for any of the search parameters that are defined in the Search Options table.
- ▶ **Search Options.** Use this table to define the parameters for a new search query.
Click the down arrow next to Search Options to open the Search Options table.

Note: Only one search can be run at a time. If a search is in progress, an information message displays at the top of the Physical Volume Search window. The user can cancel a search in progress by clicking **Cancel Search** within this message.

Define one or more of the following search parameters:

- ▶ **VOLSER.** The volume serial number. This field can be left blank. The following wildcard characters can be used in this field:
 - **% (percent).** Represents zero or more characters.
 - *** (asterisk).** Converted to % (percent). Represents zero or more characters.
 - **. (period).** Represents one character.
 - **_ (single underscore).** Converted to period (.). Represents one character.
 - **? (question mark).** Converted to period (.). Represents one character.

- ▶ **Media Type.** The type of media on which the volume exists. Use the menu to select from available media types. This field can be left blank. The following other values are possible:
 - **JA-ETC.** ETC.
 - **JB(ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD).** ATCD.
 - **JD(ATDD).** ATDD.
 - **JJ(EETC).** EETC.
 - **JK(ATKE).** ATKE.
 - **JL(ATLE).** ATLE.
- ▶ **Recording Format.** The format that is used to write the media. Use the menu to select from the available media types. This field can be left blank. The following other values are possible:
 - **Undefined.** The recording format that is used by the volume is not recognized as a supported format.
 - **J1A.**
 - **E05.**
 - **E05E.** E05 with encryption.
 - **E06.**
 - **E06E.** E06 with encryption.
 - **E07.**
 - **E07E.** E07 with encryption.
 - **E08.**
 - **E08E.** E08 with encryption.
- ▶ **Capacity State.** Whether any active data exists on the physical volume and the status of that data in relation to the volume's capacity. This field can be left blank. The following other values are valid:
 - **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 extra data.
 - **Full.** The volume contains valid data. At some point, it was marked as full and extra data cannot be added to it. In some cases, a volume can be marked full and yet be short of the volume capacity limit.

Enter a name for the new query in the New Search Name field. Enter values for any of the search parameters that are defined in the Search Options table.

Search Options table

Use this table to define the parameters for a new search query. Click the down arrow next to Search Options to open the Search Options table.

Note: Only one search can be run at a time. If a search is in progress, an information message displays at the top of the Physical Volume Search window. The user can cancel a search in progress by clicking **Cancel Search** within this message.

Define one or more of the following search parameters:

- ▶ **VOLSER.** The volume serial number. This field can be left blank. The user can also use the following wildcard characters in this field:
 - % (**percent**). Represents zero or more characters.
 - * (**asterisk**). Is converted to % (percent). Represents zero or more characters.
 - . (**period**). Represents one character.
 - _ (**single underscore**). Converted to . (period). Represents one character.
 - ? (**question mark**). Is converted to . (period). Represents one character.
- ▶ **Media Type.** The type of media on which the volume is. Use the menu to select from available media types. This field can be left blank. The following other values are possible:
 - **JA-ETC.** ETC.
 - **JB(ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD).** ATCD.
 - **JD(ATDD).** ATDD.
 - **JJ(EETC).** EETC.
 - **JK(ATKE).** ATKE.
 - **JL(ATLE).** ATLE.
- ▶ **Recording Format.** The format that is used to write the media. Use the menu to select from available media types. This field can be left blank. The following other values are possible:
 - **Undefined.** The recording format that is used by the volume is not recognized as a supported format.
 - **J1A.**
 - **E05.**
 - **E05E.** E05 with encryption.
 - **E06.**
 - **E06E.** E06 with encryption.
 - **E07.**
 - **E07E.** E07 with encryption
 - **E08.**
 - **E08E.** E08 with encryption.
- ▶ **Capacity State.** Whether any active data exists on the physical volume and the status of that data in relation to the volume's capacity. This field can be left blank. The following other values are possible:
 - **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 more data.
 - **Full.** The volume contains valid data. At some point, it was marked as full and more data cannot be added to it. In some cases, a volume can be marked full and yet be short of the volume capacity limit.
- ▶ **Home Pool.** The pool number (0 - 32) to which the physical volume was assigned when it was inserted into the library, or the pool to which it was moved through the library manager Move/Eject Stacked Volumes function. This field can be left blank.

- ▶ **Current Pool.** The number of the storage pool (0 - 32) in which the physical volume currently exists. This field can be left blank.
- ▶ **Encryption Key.** The EK label that is designated when the volume was encrypted. This is a text field. The following values are valid:
 - 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.** Whether to include physical volumes pending an eject in the search query. The following values are valid:
 - **All Ejects.** All physical volumes pending eject are included in the search.
 - **Priority Ejects.** Only physical volumes that are classified as priority eject are included in the search.
 - **Deferred Ejects.** Only physical volumes that are 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. The following values are possible:
 - **All Moves.** All physical volumes pending a move are included in the search.
 - **Priority Moves.** Only physical volumes that are classified as priority move are included in the search.
 - **Deferred Moves.** Only physical volumes that are classified as deferred move are included in the search.

Any of the previous values can be modified by using the adjacent menu. Use the adjacent menu to narrow the search down to a specific pool set to receive physical volumes. The following values are possible:

 - **All Pools.** All pools are included in the search.
 - **0 - 32.** The number of the pool to which the selected physical volumes are moved.
- ▶ **VOLSER flags.** Whether to include, exclude, or ignore any of the following VOLSER flags in the search query. Select only one:
 - **Yes** to include.
 - **No** to exclude.
 - **Ignore** to ignore the following VOLSER types during the search:
 - Misplaced
 - Mounted
 - Inaccessible
 - Encrypted
 - Export Hold
 - Read Only Recovery
 - Unavailable
 - Pending Secure Data Erase
 - Copy Exported

Search Results Options

Use this table to select the properties that are displayed on the Physical Volume Search Results window.

Click the down arrow next to **Search Results Options** to open the Search Results Options table. Select the check box next to each property that should display on the Physical Volume Search Results window.

Review the property definitions from the Search Options table section. The following properties can be displayed on the Physical Volume Search Results window:

- ▶ Media Type
- ▶ Recording Format
- ▶ Home Pool
- ▶ Current Pool
- ▶ Pending Actions
- ▶ Volume State
- ▶ Mounted Tape Drive
- ▶ Encryption Key Labels
- ▶ Export Hold
- ▶ Read Only Recovery
- ▶ Copy Export Recovery
- ▶ Database Backup

Click **Search** to initiate a new physical volume search. After the search is initiated but before it completes, the Physical Volume Search window displays the following information message:

The search is currently in progress. The user can check the progress of the search on the Previous Search Results window.

Note: The search-in-progress message is displayed on the Physical Volume Search window until the in-progress search completes or is canceled.

Figure 9-88 shows the result of a search.

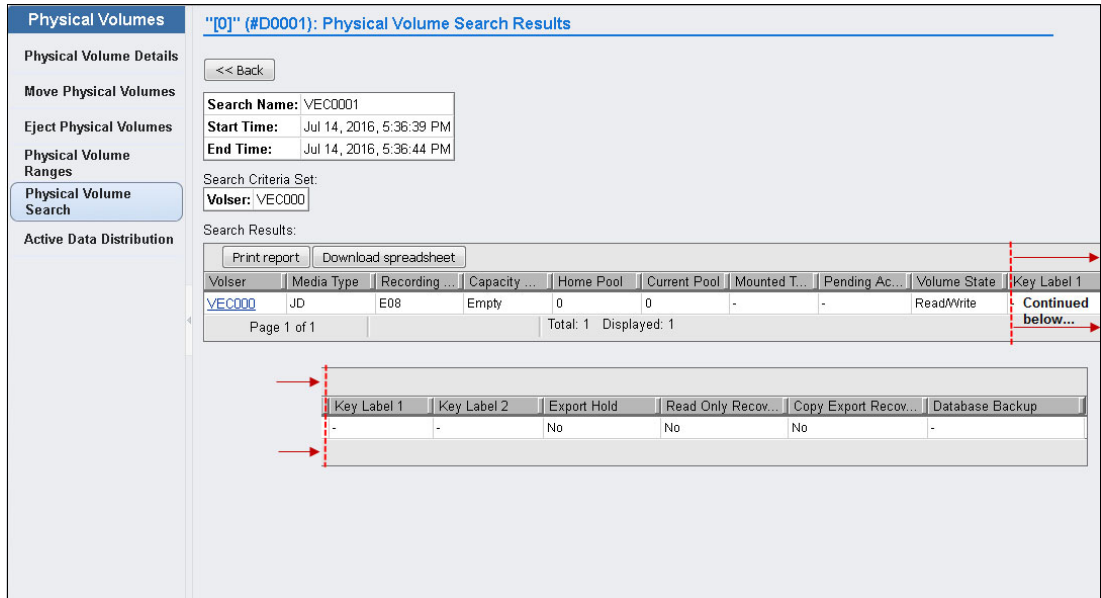


Figure 9-88 Physical Volume Search Results window

To check the progress of the search being run, click the **Previous Search Results** hyperlink in the information message. To cancel a search in progress, click Cancel Search. When the search completes, the results are displayed in the Physical Volume Search Results window. The query name, criteria, start time, and end time are saved along with the search results. A maximum of 10 search queries can be saved.

Active Data Distribution

To view the distribution of data on physical volumes that are marked full on an TS7700T cluster, use this window. The distribution can be used to select an appropriate reclaim threshold. The Active Data Distribution window displays the utilization percentages of physical volumes in increments of 10%.

Figure 9-89 shows the Active Data Distribution window.




Figure 9-89 Active Data Distribution window

Number of Full Volumes at Utilization Percentages window

The tables in Figure 9-90 show the number of physical volumes that are marked as full in each physical volume pool, according to % of volume used. The following fields are displayed:

- **Pool.** The physical volume pool number. This number is a hyperlink; click it to display a graphical representation of the number of physical volumes per utilization increment in a pool. If the user clicks the pool number hyperlink, the Active Data Distribution subwindow opens.

Figure 9-90 shows a sample of the window that is reached by clicking the Pool 4 hyperlink.

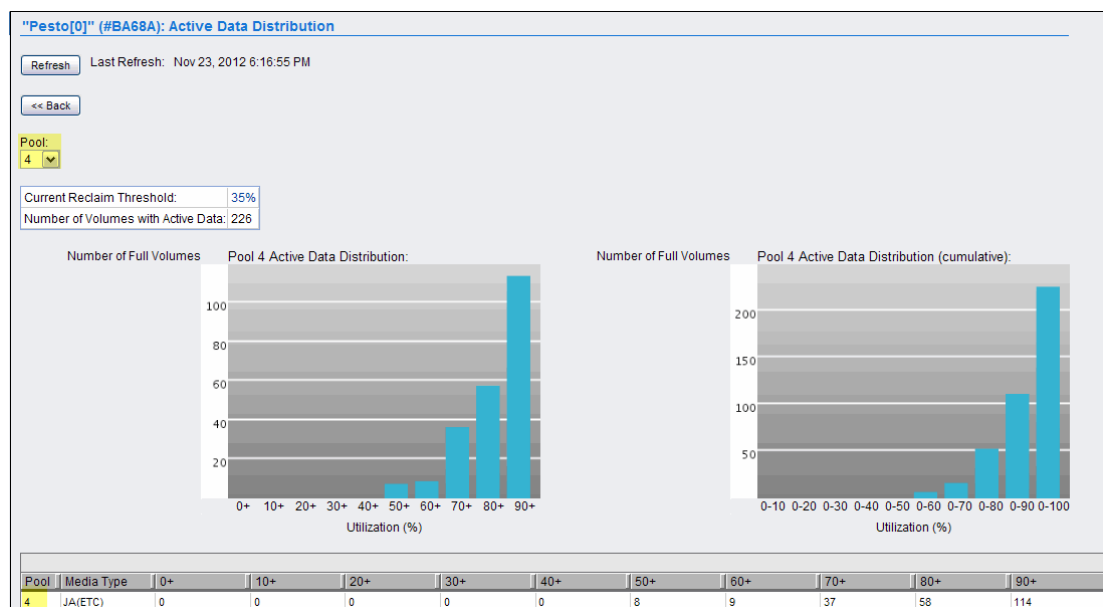


Figure 9-90 Active Data Distribution for a specific pool

This subwindow contains the following fields and information:

- ▶ **Pool.** To view graphical information for another pool, select the target pool from this menu.
- ▶ **Current Reclaim Threshold.** 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 under this percentage, a reclaim operation is performed on the stacked volume. The valid range of possible values is 0 - 95% and can be selected in 5% increments; 35% is the default value.

Tip: This percentage is a hyperlink; click it to open the Modify Pool Properties window, where the user can modify the percentage that is used for this threshold.

- ▶ **Number of Volumes with Active Data.** The number of physical volumes that contain active data.
- ▶ **Pool *n* Active Data Distribution.** This graph displays the number of volumes that contain active data per volume utilization increment for the selected pool. On this graph, utilization increments (x axis) do not overlap.
- ▶ **Pool *n* Active Data Distribution (cumulative).** This graph displays the cumulative number of volumes that contain active data per volume utilization increment for the selected pool. On this graph, utilization increments (x axis) overlap, accumulating as they increase.

The Active Data Distribution subwindow also displays utilization percentages for the selected pool, excerpted from the Number of Full Volumes at Utilization Percentages table.

- ▶ **Media Type.** The type of cartridges that are contained in the physical volume pool. If more than one media type exists in the pool, each type is displayed, separated by commas. The following values are possible:
 - **Any 3592.** Any media with a 3592 format.
 - **JA-ETC.** ETC.
 - **JB(ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD).** ATCD.
 - **JD(ATDD).** ATDD.
 - **JJ(EETC).** EETC.
 - **JK(ATKE).** ATKE.
 - **JL(ATLE).** ATLE.
- ▶ **Percentage of Volume Used (0+, 10+, 20+, and so on).** Each of the last 10 columns in the table represents a 10% increment of total physical volume space used. For instance, the column heading 20+ represents the 20% - 29% range of a physical volume used. For each pool, the total number of physical volumes that occur in each range is listed.

Physical Tape Drives window

To view a summary of the state of all physical drives that are accessible to the TS7700T cluster, use this window.

This window is visible but disabled on the TS7700 MI if the grid possesses a physical library, but the selected cluster does not. The following message is displayed:

The cluster is not attached to a physical tape library.

Tip: This window is not visible on the TS7700 MI if the grid does not possess a physical library.

Figure 9-91 shows the Physical Tape Drives window.

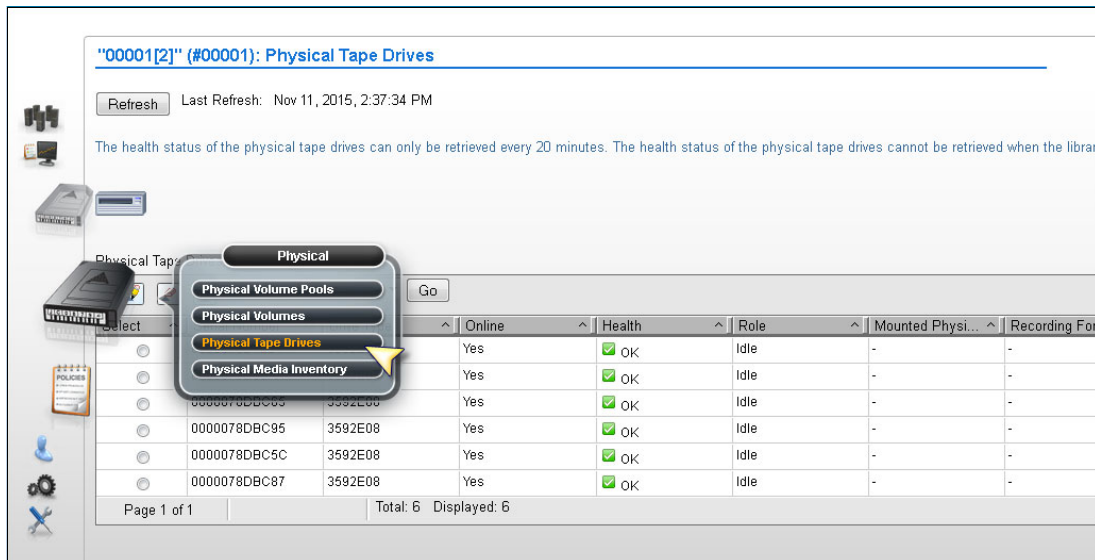


Figure 9-91 Physical Tape Drives window

The Physical Tape Drives table displays status information for all physical drives accessible by the cluster, including the following information:

- ▶ **Serial Number.** The serial number of the physical drive.
- ▶ **Drive Type.** The machine type and model number of the drive. The following values are possible:
 - **3592J1A.**
 - **3592E05.**
 - **3592E05E.** A 3592 E05 drive that is Encryption Capable.
 - **3592E06.**
 - **3952E07.**
 - **3952E08.**
- ▶ **Online.** Whether the drive is online.
- ▶ **Health.** The health of the physical drive. This value is obtained automatically at times that are determined by the TS7700. The following values are possible:
 - **OK.** The drive is fully functioning.
 - **WARNING.** The drive is functioning but reporting errors. Action needs to be taken to correct the errors.
 - **DEGRADED.** The drive is operational but has lost some redundancy resource and performance.
 - **FAILURE.** The drive is not functioning and immediate action must be taken to correct it.
 - **OFFLINE/TIMEOUT.** The drive is out of service or unreachable within a certain time frame.

- ▶ **Role.** The current role the drive is performing. The following values are possible:
 - **IDLE.** The drive is not in use.
 - **MIGRATION.** The drive is being used to copy a virtual volume from the TVC to the physical volume.
 - **RECALL.** The drive is being used to recall a virtual volume from a physical volume to the TVC.
 - **RECLAIM SOURCE.** The drive is being used as the source of a reclaim operation.
 - **RECLAIM TARGET.** The drive is being used as the target of a reclaim operation.
 - **EXPORT.** The drive is being used to export a volume.
 - **SECURE ERASE.** The drive is being used to erase expired volumes from the physical volume securely and permanently.
- ▶ **Mounted Physical Volume.** VOLSER of the physical volume that is mounted by the drive.
- ▶ **Recording Format.** The format in which the drive operates. The following values are possible:
 - **J1A.** The drive is operating with J1A data.
 - **E05.** The drive is operating with E05 data.
 - **E05E.** The drive is operating with E05E encrypted data.
 - **E06.** The drive is operating with E06 data.
 - **E06E.** The drive is operating with E06E encrypted data.
 - **E07.** The drive is operating with E07 data.
 - **E07E.** The drive is operating with E07E encrypted data.
 - **E08.** The drive is operating with E08 data.
 - **E08E.** The drive is operating with E08 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 window to determine whether the format state has changed. If the Unknown state persists for 1 hour or longer, contact your IBM SSR.
- ▶ **Requested Physical Volume.** The VOLSER of the physical volume that is requested for mount. If no physical volume is requested, this field is blank.

To view additional information for a specific, selected drive, see the Physical Drives Details table on the Physical Tape Drive Details window:

1. Select the radio button next to the serial number of the physical drive in question.
2. Select **Details** from the Select Action menu.
3. Click **Go** to open the Physical Tape Drives Details window.

Figure 9-92 shows a Physical Tape Drive Details window.

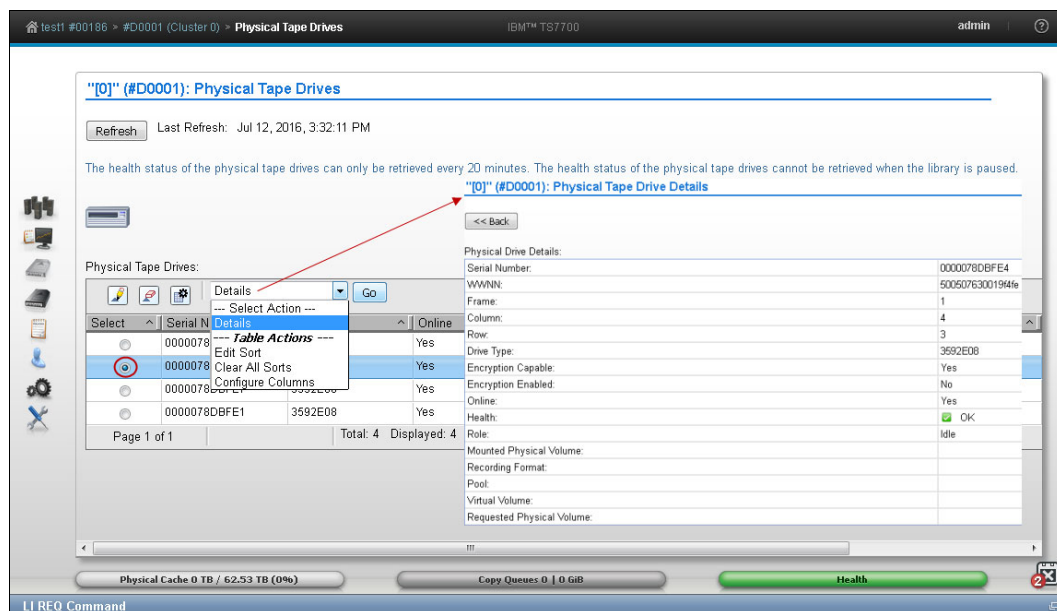


Figure 9-92 Physical Tape Drive Details window

The Physical Drives Details table displays detailed information for a specific physical tape drive:

- **Serial Number.** The serial number of the physical drive.
- **Drive Type.** The machine type and model number of the drive. The following values are possible:
 - **3592J1A.**
 - **3592E05.**
 - **3592E05E.** A 3592 E05 drive that is Encryption Capable.
 - **3592E06.**
 - **3952E07.**
 - **3592E08.**
- **Online.** Whether the drive is online.
- **Health.** The health of the physical drive. This value is obtained automatically at times that are determined by the TS7740. The following values are possible:
 - **OK.** The drive is fully functioning.
 - **WARNING.** The drive is functioning but reporting errors. Action needs to 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 must be taken to correct it.
 - **OFFLINE/TIMEOUT.** The drive is out of service or cannot be reached within a certain time frame.
- **Role.** The current role that the drive is performing. The following values are possible:
 - **IDLE.** The drive is not in use.

- **MIGRATION.** The drive is being used to copy a virtual volume from the TVC to the physical volume.
 - **RECALL.** The drive is being used to recall a virtual volume from a physical volume to the TVC.
 - **RECLAIM SOURCE.** The drive is being used as the source of a reclaim operation.
 - **RECLAIM TARGET.** The drive is being used as the target of a reclaim operation.
 - **EXPORT.** The drive is being used to export a volume.
 - **SECURE ERASE.** The drive is being used to erase expired volumes from the physical volume securely and permanently.
- **Mounted Physical Volume.** VOLSER of the physical volume mounted by the drive.
 - **Recording Format.** The format in which the drive operates. The following values are possible:
 - **J1A.** The drive is operating with J1A data.
 - **E05.** The drive is operating with E05 data.
 - **E05E.** The drive is operating with E05E encrypted data.
 - **E06.** The drive is operating with E06 data.
 - **E06E.** The drive is operating with E06E encrypted data.
 - **E07.** The drive is operating with E07 data.
 - **E07E.** The drive is operating with E07E encrypted data.
 - **E08.** The drive is operating with E08 data.
 - **E08E.** The drive is operating with E08 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 window to determine whether the format state has changed. If the Unknown state persists for 1 hour or longer, contact your IBM SSR.
 - **Requested Physical Volume.** The VOLSER of the physical volume that is requested for mount. If no physical volume is requested, this field is blank.
 - **WWNN.** The worldwide node name that is used to locate the drive.
 - **Frame.** The frame in which the drive is.
 - **Row.** The row in which the drive is.
 - **Encryption Enabled.** Whether encryption is enabled on the drive.

Note: If the user is monitoring this field while changing the encryption status of a drive, the new status does not display until you bring the TS7700 Cluster offline and then back online.

- **Encryption Capable.** Whether the drive is capable of encryption.
 - **Physical Volume.** VOLSER of the physical volume that is mounted by the drive.
 - **Pool** The pool name of the physical volume that is mounted by the drive.
 - **Virtual Volume.** VOLSER of the virtual volume being processed by the drive.
4. Click **Back** to return to the Physical Tape Drives window.

Physical Media Inventory window

To view physical media counts for media types in storage pools in the TS7700, use this window.

This window is visible but disabled on the TS7700 MI if the grid possesses a physical library, but the selected cluster does not. The following message is displayed:

The cluster is not attached to a physical tape library.

Tip: This window is not visible on the TS7700 MI if the grid does not possess a physical library.

Figure 9-93 shows the Physical Media Inventory window.

The screenshot shows the Physical Media Inventory window for pool "00001[2]" (#00001). It includes a Refresh button and a last refresh timestamp of Nov 11, 2015, 2:50:21 PM. Below the button, a note states: "The table below shows the number of physical media for each media type, for each pool." The table is titled "Inventory of Physical Media Pools:" and contains the following data:

Pool	Media Type	Empty	Filling	Full	Queued for ...	ROR	Unavailable	Unsupported
0	JB	31	-	-	-	0	1	30
0	JD	12	-	-	-	0	0	0
1	JB	2	0	2	0	0	0	2
1	JD	2	3	0	0	0	0	0
3	JB	0	0	1	0	0	0	0
9	JB	2	0	1	0	0	0	2
9	JD	2	1	0	0	0	0	0

Figure 9-93 Physical Media Inventory window

The following physical media counts are displayed for each media type in each storage pool:

- ▶ **Pool.** The storage pool number.
- ▶ **Media Type.** The media type defined for the pool. A storage pool can have multiple media types and each media type is displayed separately. The following values are possible:
 - **JA-ETC.** ETC.
 - **JB(ETCL).** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD).** ATCD.
 - **JD(ATDD).** ATDD.
 - **JJ(EETC).** EETC.
 - **JK(ATKE).** ATKE.
 - **JL(ATLE).** ATLE.
- ▶ **Empty.** The count of physical volumes that are empty for the pool.
- ▶ **Filling.** The count of physical volumes that are 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.

Tip: A value in the Full field is displayed as a hyperlink; click it to open the Active Data Distribution subwindow. The Active Data Distribution subwindow displays a graphical representation of the number of physical volumes per utilization increment in a pool. If no full volumes exist, the hyperlink is disabled.

- ▶ **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.** The count of physical volumes in the Read Only Recovery (ROR) state that are damaged or corrupted.
- ▶ **Unavailable.** The count of physical volumes that are in the unavailable or destroyed state.
- ▶ **Unsupported.** Unsupported media (for example: JA and JJ) type present in tape library and inserted for the TS7740 and TS7720T. Based on the drive configuration, the TS7700 cannot use one or more of the specified media, which can result in the out-of-scratch condition.

9.2.8 The Constructs icon

The topics in this section present information that is related to TS7700 storage constructs. Figure 9-94 shows the Constructs icon and the options under it.

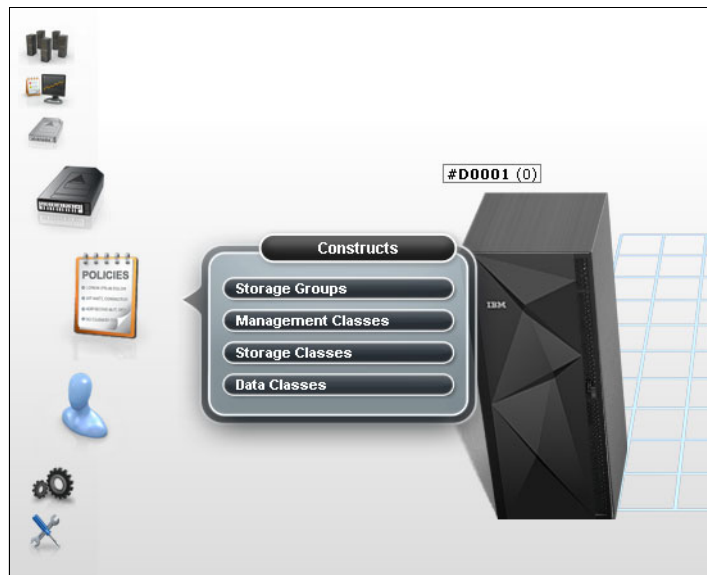


Figure 9-94 The Constructs icon

Storage Groups window

Use the window that is shown in Figure 9-95 to add, modify, or delete an SG.

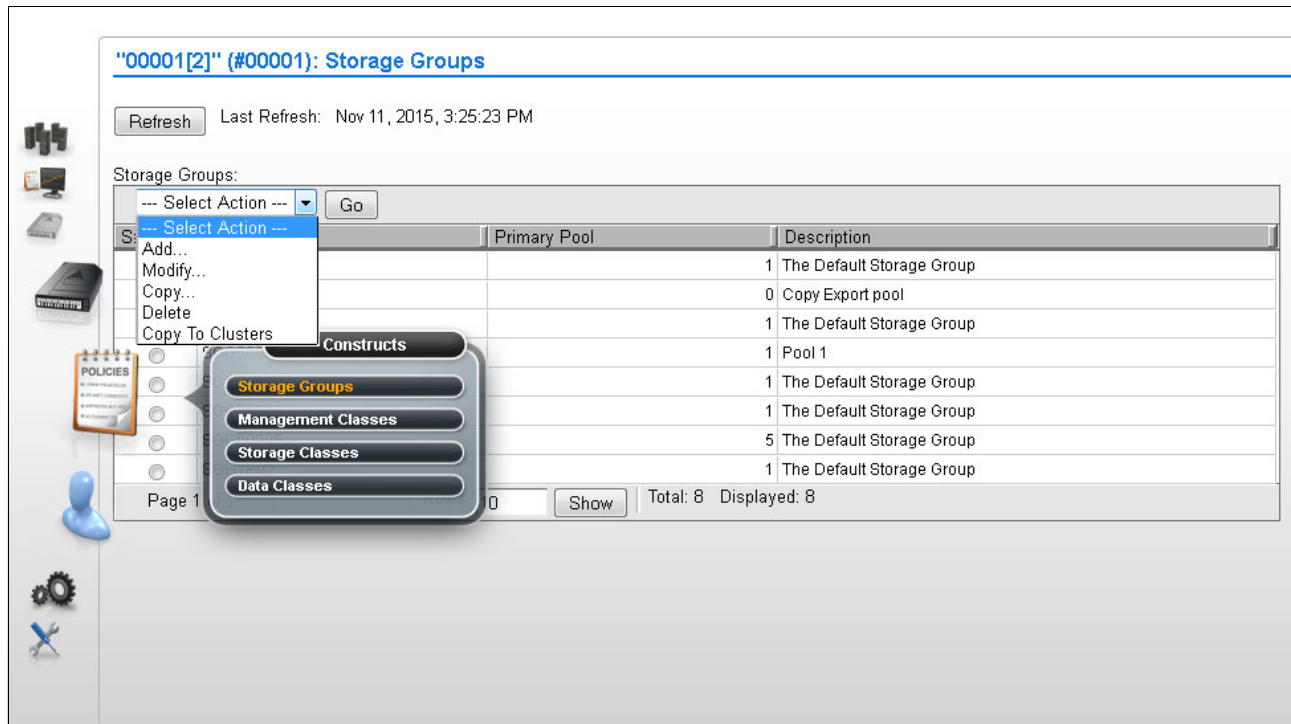


Figure 9-95 MI Storage Groups window

The SGs table displays all existing SGs available for a cluster.

The user can use the SGs table to create an SG, modify an existing SG, or delete an SG. Also, the user can copy selected SGs to the other clusters in this grid by using the Copy to Clusters action available in the menu.

The SGs table shows the following status information:

- ▶ **Name.** The name of the SG. Each SG 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 migration. Only validated physical primary pools can be selected. If the cluster does not possess a physical library, this column is not visible, and the MI categorizes newly created SGs by using pool 1.
- ▶ **Description.** A description of the SG.

Use the menu in the SGs table to add an SG, or modify or delete an existing SG.

To add an SG, select **Add** from the menu. Complete the fields for information that will be displayed in the SGs table.

Consideration: If the cluster does not possess a physical library, the Primary Pool field is not available in the Add or Modify options.

To modify an existing SG, select the radio button from the Select column that appears next to the name of the SG that needs to be modified. Select **Modify** from the menu. Complete the fields for information that must be displayed in the SGs table.

To delete an existing SG, select the radio button from the Select column that appears next to the name of the SG to delete. Select **Delete** from the menu. Confirm the decision to delete an SG. If you select **OK**, the SG is deleted. If you select **No**, the request to delete is canceled.

Management Classes window

To define, modify, copy, or delete the MC that defines the TS7700 copy policy for volume redundancy, use this window (Figure 9-96). The table displays the copy policy that is in force for each component of the grid.

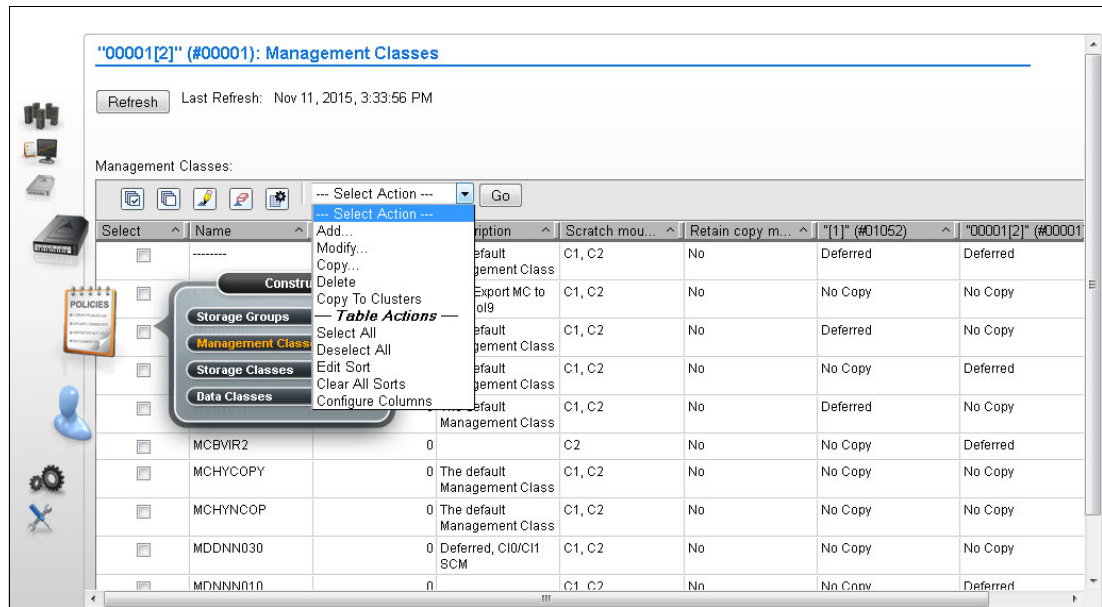


Figure 9-96 MI Management Classes window on a grid

The secondary copy pool column shows only in a tape attach TS7700 cluster. This is a requirement for using the Copy Export function.

Figure 9-97 shows the MCs options, including the Time Delayed option.

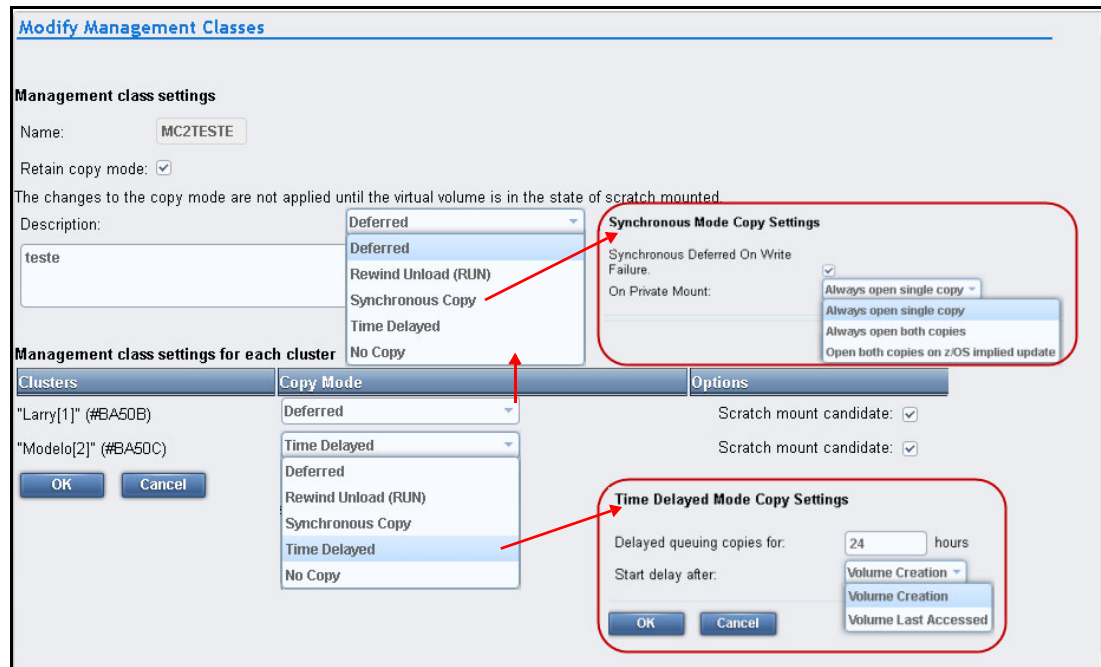


Figure 9-97 Modify Management Classes options

The user can use the MCs table to create, modify, and delete MCs. The default MC can be modified, but cannot be deleted. The default MC uses dashes (-----) for the symbolic name.

The MCs table shows the following status information:

- ▶ **Name.** The name of the MC. Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field cannot be a number. This is the only field that cannot 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 is not visible and the MI categorizes newly created SGs by using pool 0.
- ▶ **Description.** A description of the MC definition. The value in this field must be 1 - 70 characters in length.
- ▶ **Retain Copy Mode.** Whether previous copy policy settings on private (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 MI, previously assigned copy modes are retained and subsequent read or modify access does not refresh, update, or merge copy modes. This enables the original number of copies to be maintained.
- ▶ **Scratch Mount Candidate.** Clusters that are listed under *Scratch Mount Candidate* are selected first for scratch mounts of the volumes that are associated with the MC. If no cluster is displayed, the scratch mount process selects among the available clusters in a random mode.

To add an MC, complete the following steps:

1. Select **Add** from the Management Class menu that is shown in Figure 9-96 on page 439 and click **Go**.
2. Complete the fields for information that will be displayed in the MCs table. Up to 256 MCs can be created per TS7700 grid.

Remember: If the cluster does not possess a physical library, the Secondary Pool field is not available in the **Add** option.

You can use the Copy Action option to copy any existing MC to each cluster in the TS7700 Grid.

The following options are available in the MC:

- ▶ **No Copy.** No volume duplication occurs if this action is selected.
- ▶ **RUN.** Volume duplication occurs when the **Rewind Unload** command is received. The command returns only after the volume duplication completes successfully.
- ▶ **Deferred.** Volume duplication occurs later based on the internal schedule of the copy engine.
- ▶ **Synchronous Copy.** Volume duplication is treated as host I/O and takes place before control is returned to the application issuing the I/O. Only two clusters in the grid can have the Synchronous mode copy defined.
- ▶ **Time Delayed.** Volume duplication occurs only after the delay time that is specified by the user elapses. This option is only available if all clusters in the grid are running R3.1 or higher level of code. Selecting **Time Delayed Mode** for any cluster opens another option menu:
 - **Delay Queueing Copy for [X] Hours.** Number of hours that queuing the copies will be delayed if **Time Delayed Mode Copy** is selected. Can be set for 1 - 65,535 hours.
 - **Start Delay After:**
 - **Volume Create.** Delay time is clocked from the volume creation.
 - **Volume Last Accessed.** Delay time is clocked from the last access. Every time a volume is accessed, elapsed time is zeroed for that volume and countdown starts again from the delay value set by user.

To modify an existing MC, complete the following steps:

1. Select the check box from the Select column that appears in the same row as the name of the MC to modify.

The user can modify only one MC at a time.

2. Select **Modify** from the menu and click **Go**.

Of the fields that were listed previously in the MCs table, the user can change all of them except the MC name.

To delete one or more existing MCs, complete the following steps:

1. Select the check box from the Select column that appears in the same row as the name of the MC to delete.
2. Select multiple check boxes to delete multiple MCs.
3. Select **Delete** from the menu.
4. Click **Go**.

Note: The default MC cannot be deleted.

Storage Classes window

To define, modify, or delete an SC that is used by the TS7700 to automate storage management through classification of data sets and objects within a cluster, use the window that is shown in Figure 9-98. Also, you can use this window to copy an existing SC to the same cluster being accessed, or to another cluster in the grid.

You can view SCs from any TS7700 in the grid, but TVC preferences can be altered only from a tape-attached cluster. Figure 9-98 shows the window in a TS7720T model.

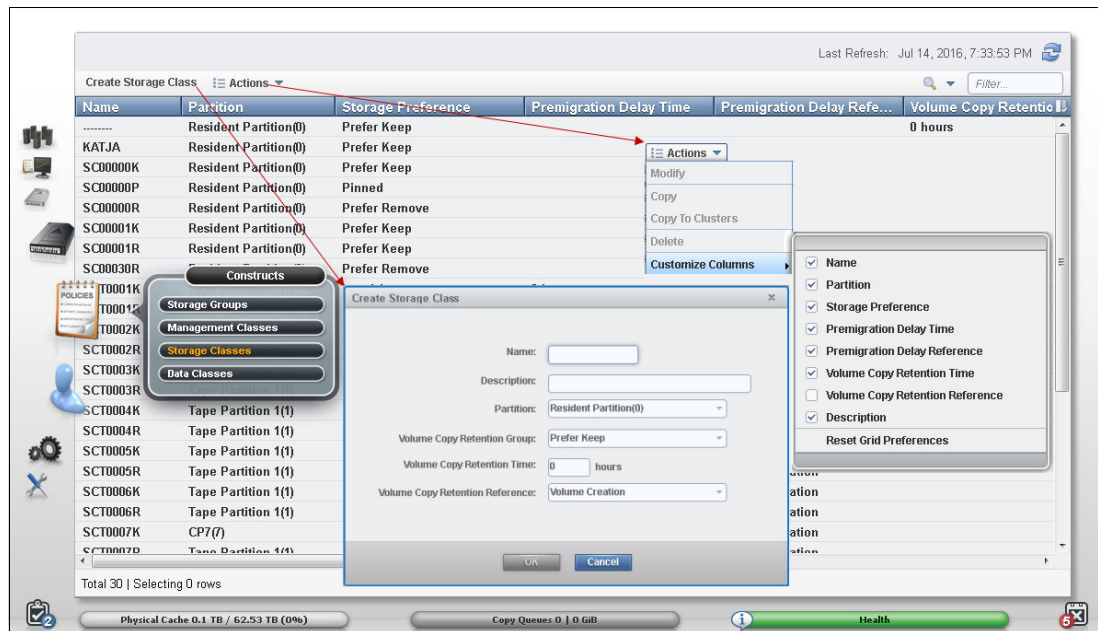


Figure 9-98 MI Storage Classes window on a TS7700T

The SCs table lists defined SCs that are available to control data sets (CDSs) and objects within a cluster.

The **Create Storage Class** box is slightly different depending on the TS7700 cluster model being accessed by the MI. Figure 9-99 shows appearance for the Create Storage Class box in different TS7700 models.

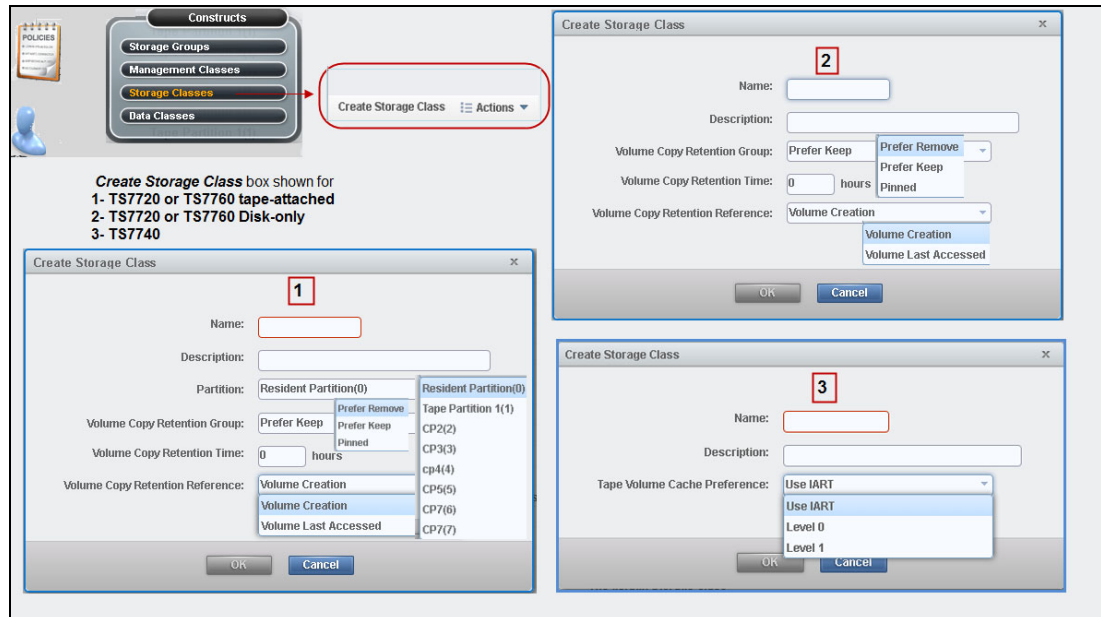


Figure 9-99 The Create Storage Class box

The default SC can be modified, but cannot be deleted. The default SC has dashes (-----) as the symbolic name.

The SCs table displays the following status information:

- ▶ **Name.** The name of the SC. The value in this field must be 1 - 8 characters. Each SC within a cluster must have a unique name. Valid characters for this field are A-Z, 0-9, \$, @, *, #, and %. The first character of this field cannot be a number. This is the only field that cannot be modified after it is added.
- ▶ **Description.** An optional description of the SC. The value in this field must be 0 - 70 characters.
- ▶ **Partition.** The name of the partition that is associated with the SC. A partition must be active before it can be selected as a value for this field. This field is displayed only if the cluster is a TS7720 that is attached to a physical library. A dash (-) indicates that the SC contains a partition that was deleted. Any volumes that are assigned to go to the deleted partition are redirected to the primary tape partition.
- ▶ **Tape Volume Cache Preference.** The preference level for the SC. It determines how soon volumes are removed from cache after their copy to tape. This field is visible only if the TS7700 Cluster attaches to a physical library. If the selected cluster does not possess a physical library, volumes in that cluster's cache display a Level 1 preference. The following values are possible:
 - **Use IART.** Volumes are removed according to the TS7700's Initial Access Response Time (IART).
 - **Level 0.** Volumes are removed from the TVC as soon as they are copied to tape.

- **Level 1.** Copied volumes remain in the TVC until more space is required, then the first volumes are removed to free space in the cache. This is the default preference level that is assigned to new preference groups.
- ▶ **Premigration Delay Time.** The number of hours until premigration can begin for volumes in the SC, based on the volume time stamp designated by Premigration Delay Reference. Possible values are 0 - 65535. If 0 is selected, premigration delay is disabled. This field is visible only if the TS7700 cluster attaches to a physical library.
- ▶ **Premigration Delay Reference.** The volume operation that establishes the time stamp from which Premigration Delay Time is calculated. This field is visible only if the TS7700 cluster attaches to a physical library. The following values are possible:
 - **Volume Creation.** The time at which the volume was created by a scratch mount or write operation from beginning of tape.
 - **Volume Last Accessed.** The time at which the volume was last accessed.
- ▶ **Volume Copy Retention Group.** The name of the group that defines the preferred auto removal policy applicable to the virtual volume. The Volume Copy Retention Group provides more options to remove data from a disk-only TS7700 or resident-only (CP0) partition in the TS7700T as the active data reaches full capacity. Volumes become candidates for removal if an appropriate number of copies exist on peer clusters *and* the volume copy retention time has elapsed since the volume was last accessed. Volumes in each group are removed in order based on their least recently used access times.

The volume copy retention time describes the number of hours a volume remains in cache before becoming a candidate for removal. This field is only displayed for disk-only clusters when they are part of a hybrid grid (one that combines TS7700 clusters that both *do* and *do not* attach to a physical library).

If the virtual volume is in a scratch category and is on a disk-only cluster, removal settings no longer apply to the volume, and it is a candidate for removal. In this instance, the value that is displayed for the Volume Copy Retention Group is accompanied by a warning icon:

- **Prefer Remove.** Removal candidates in this group are removed before removal candidates in the Prefer Keep group.
- **Prefer Keep.** Removal candidates in this group are removed after removal candidates in the Prefer Remove group.
- **Pinned.** Copies of volumes in this group are never removed from the accessing cluster. The volume copy retention time does not apply to volumes in this group. Volumes in this group that are later moved to scratch become priority candidates for removal.
- ▶ **Volume Copy Retention Time.** The minimum amount of time (in hours) that a volume remains temporarily pinned in cache (counting from the volume creation or last access time) before transitioning either to *Prefer Keep* or *Remove* groups. When the amount of retention time elapses, the copy then becomes a candidate for removal. Possible values include 0 - 65,536; the default is 0.

This field is only visible if the selected cluster is a TS7720 *and* all of the clusters in the grid operate at Licensed Internal Code level 8.7.0.xx or later. If the Volume Copy Retention Group displays a value of Pinned, this field is disabled.
- ▶ **Volume Copy Retention Reference.** The volume operation that establishes the time stamp from which Volume Copy Retention Time is calculated. The following list describes the possible values:
 - **Volume Creation.** The time at which the volume was created by a scratch mount or write operation from beginning of tape.
 - **Volume Last Accessed.** The time at which the volume was last accessed.

If the Volume Copy Retention Group displays a value of Pinned, this field is disabled.

Data Classes window

To define, modify, copy, or delete a TS7700 DC that is used to automate storage management through the classification of data sets, use the window that is shown in Figure 9-100, defining volume sizes and LWORM policy assignment.

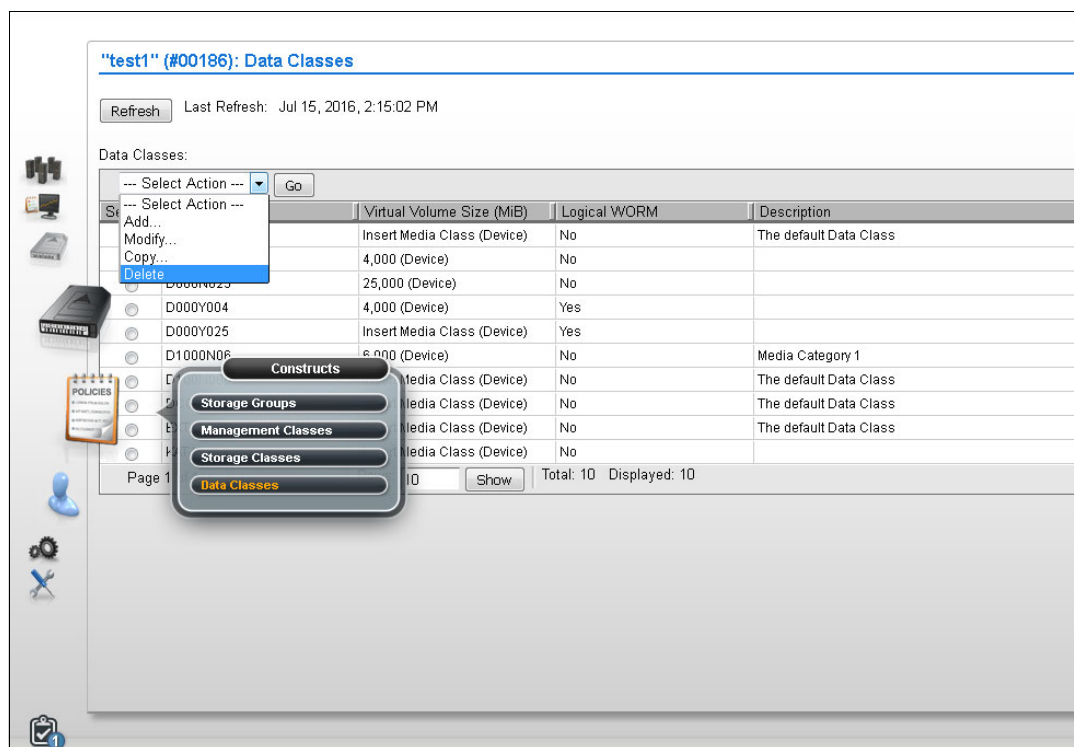


Figure 9-100 MI Data Classes window

Important: Scratch (Fast Ready) categories and DCs work at the system level and are unique for all clusters in a grid. Therefore, if they are modified in one cluster, they are applied to all clusters in the grid.

The DC table (Figure 9-100) displays the list of DCs defined for each cluster of the grid.

The user can use the DCs table to create a DC or modify, copy, or delete an existing DC. The default DC can be modified, but cannot be deleted. The default DC has dashes (-----) as the symbolic name.

The DCs table lists the following status information:

- ▶ **Name.** The name of the DC. Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field cannot be a number. The value in this field must be 1 - 8 characters in length. The first character of this field cannot be a number. This field is the only field that cannot be modified after it is added.
- ▶ **Virtual Volume Size (MiB).** The logical volume size of the DC, which determines the maximum number of MiB for each logical volume in a defined class. One possible value is Insert Media Class, where the logical volume size is not defined (the DC is not defined by a maximum logical volume size). Other possible values are 1,000 MiB, 2,000 MiB, 4,000 MiB, 6,000 MiB, or 25,000 MiB.

A maximum size of 25,000 MiB for logical volumes is allowed without any restriction if all clusters in a grid operate at R3.2 or higher level of Licensed Internal Code.

Otherwise, the 25,000 MiB is not supported whenever one or more TS7740 clusters are present in grid, and at least one cluster in the grid operates at a Licensed Internal Code level earlier than R3.2.

Also, being the grid that is formed exclusively by TS7720 clusters that are not attached to a physical library, Feature Code 0001 is required on every cluster operating at a LIC level earlier than R3.2.

- ▶ **Description.** A description of the DC definition. The value in this field must be 0 - 70 characters.
- ▶ **Logical WORM.** Whether LWORM is set for the DC. LWORM 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. The following values are valid:
 - **Yes.** LWORM is set for the DC. Volumes belonging to the DC are defined as LWORM.
 - **No.** LWORM is not set. Volumes belonging to the DC are not defined as LWORM. This is the default value for a new DC.

Use the menu in the DCs table to add a DC, or modify or delete an existing DC.

Tip: The user can create up to 256 DCs per TS7700 grid.

9.2.9 The Access icon

The topics in this section present information that is related to managing user access in a TS7700 subsystem. A series of enhancements in user access management has been introduced, from Release 1.6 to the current level. All the user authentication modalities that are introduced by the different levels of code are supported in the current level. For example:

- ▶ The TS7700 supports a role-based access control (RBAC) policy through the System Storage Productivity Center or by direct LDAP by using Microsoft Active Directory (MSAD) or IBM Resource Access Control Facility (RACF).

Note: SSPC is no longer offered by IBM. Use direct LDAP.

- ▶ RBAC policies are applied globally to all users through all ports, not only to the MI users.
- ▶ Provides an option to exclude service personnel (local and remote) from RBAC policies if convenient.
- ▶ All user access to TS7700 is secured and controlled by RACF.

Although RACF is intended to address all of the secure access needs for the z Systems environment, RACF does not provide a direct interface to the external storage devices, such as the TS7700 cluster. The current implementation uses the IBM Security Directory Server (formerly Tivoli LDAP server) for z Systems as a bridge between the TS7700 clusters and RACF. The authentication policies can be classified into two categories:

- ▶ Local, which replicates users and respective assigned roles across the grid.
- ▶ External, which keeps the list of users and group data on a separated server, mapping the relationship between users, group, and authorization roles when a user logs in to a cluster.

External authentication policies include Storage Authentication Service policies and Direct LDAP (lightweight directory access protocol) policies.

Important: Before enabling *External* policies, create an account that can be used by service personnel (local and remote), or select the boxes in MI to exclude your IBM SSR (local and remote) from RBAC policies.

Local Authentication Policy is managed and applied within the cluster or clusters participating in a grid. In a multi-cluster grid configuration, user IDs and their associated roles are defined through the MI on one of the clusters. The user IDs and roles are then propagated through the grid to all participating clusters. You can use Storage Authentication Service and Direct LDAP policies to manage centrally user IDs and roles:

- ▶ The Storage Authentication Service policy stores user and group data on a separate server and maps relationships among users, groups, and authorization roles when a user signs in to a cluster. Network connectivity to an external System Storage Productivity Center (SSPC) is required. Each cluster in a grid can operate its own Storage Authentication Service policy.
- ▶ Direct LDAP policy is an RBAC policy that authenticates and authorizes users through direct communication with an LDAP (MSAD platform) server or Tivoli LDAP server for z Systems. Only one authentication policy can be enabled per cluster at one time.

RACF authentication uses the Secure Database Manager (SDBM) in IBM Security Directory Server for z Systems as a front end to the RACF.

The user can access the following options through the User Access (blue man icon) link:

- ▶ **Security Settings.** Use this window to view security settings for a TS7700 grid. From this window, the user can also access windows 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 grid.
- ▶ **SSL Certificates.** Use this window to view, import, or delete Secure Sockets Layer (SSL) certificates to support connection to a Storage Authentication Service server from a TS7700 cluster.
- ▶ **InfoCenter Settings.** Use this window to upload a new TS7700 IBM Knowledge Center to the cluster's MI.

Note: Although the term "InfoCenter" is still used in the interface of the product, the term "IBM Knowledge Center" is the current correct term.

The options for User Access Management in the MI for TS7700 are shown in Figure 9-101.

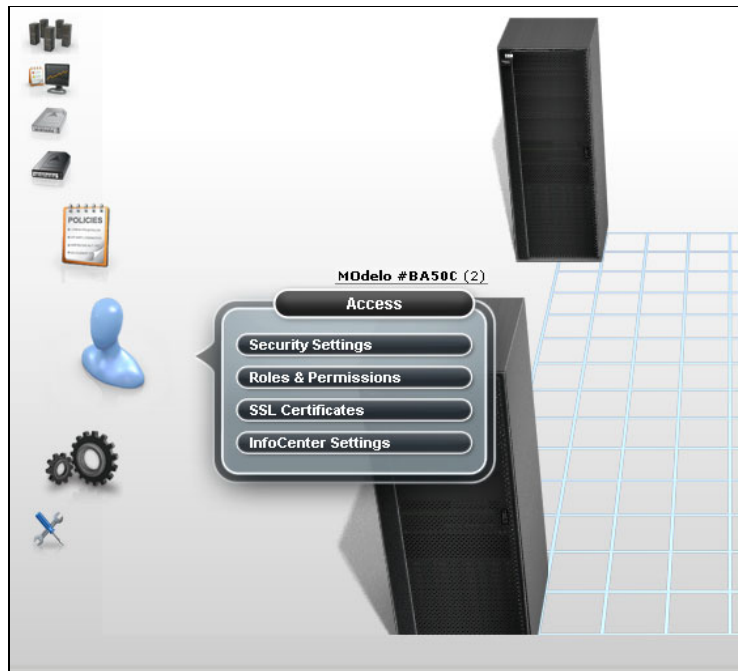


Figure 9-101 The Access icon and options

Security Settings

Figure 9-102 shows the Security Settings window, which is the entry point to enabling security policies.

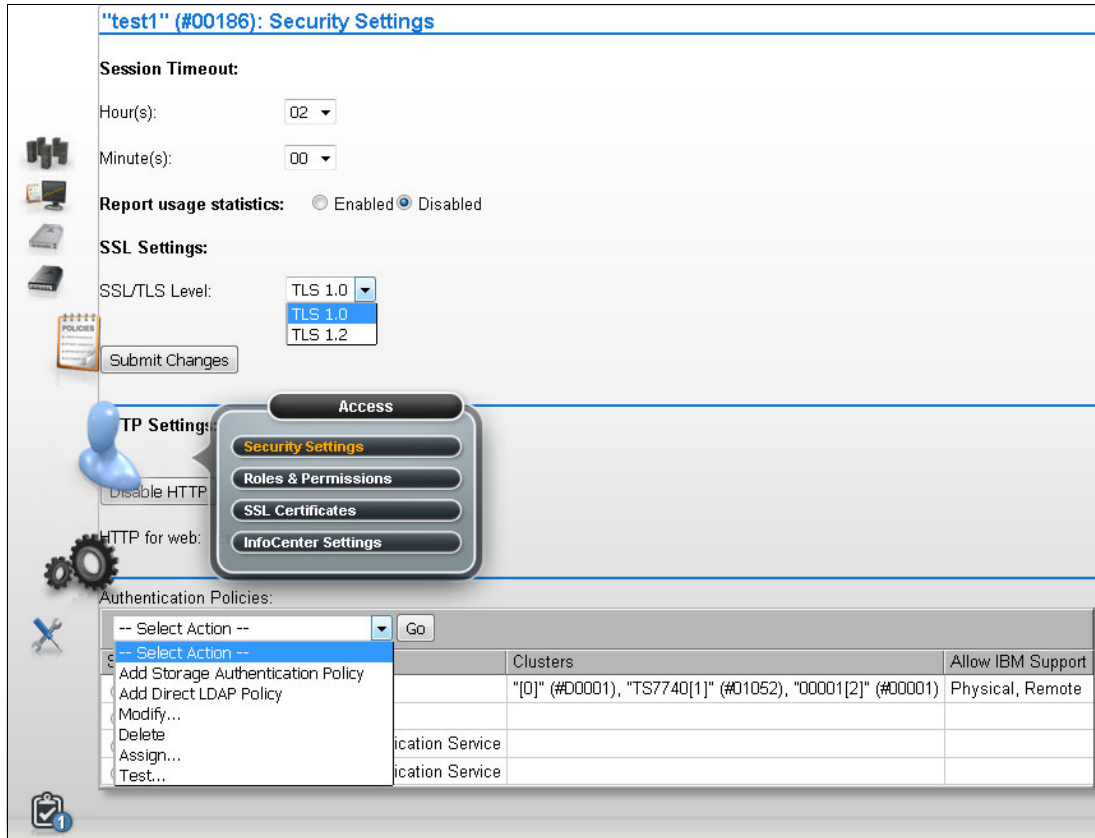


Figure 9-102 TS7700 Security Settings

Use the Session Timeout policy to specify the number of hours and minutes that the MI can be idle before the current session expires and the user is redirected to the login window. This setting is valid for all users in the grid.

To modify the maximum idle time, select values from the Hours and Minutes menus and click **Submit Changes**. The following parameters are valid for Hours and Minutes:

- ▶ **Hours.** The number of hours the MI can be idle before the current session expires. Possible values for this field are 00 - 23.
- ▶ **Minutes.** The number of minutes the MI can be idle before the current session expires. Possible values for this field are 00 - 55, selected in 5-minute increments.

The Authentication Policies table lists the following information:

- ▶ **Policy Name.** The name of the policy that defines the authentication settings. The policy name is a unique value that is composed of 1 - 50 Unicode characters. Heading and trailing blank spaces are trimmed, although internal blank spaces are retained. After a new authentication policy is created, its policy name cannot be modified.

Tip: 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 using 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 be deleted only 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 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. This section looks at the various windows that are required to manage the Local Authentication Policy.

To add a user to the Local Authentication Policy for a TS7700 Grid, complete the following steps:

1. On the TS7700 MI, click **Access** → **Security Settings** from the left navigation window.
2. Click **Select** next to the Local policy name on the Authentication Policies table.
3. Select **Modify** from the Select Action menu and click **Go**.
4. On the Local Accounts table, select **Add** from the Select Action menu and click **Go**.
5. In the Add User window, enter values for the following required fields:
 - **User name:** The new user's login name. This value must be 1 - 128 characters and composed of Unicode characters. Spaces and tabs are not allowed.
 - **Role:** The role that is assigned to the user account. The role can be a predefined role or a user-defined role. The following values are possible:
 - **Operator:** The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, and 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 cannot change network configuration, feature licenses, user accounts, or custom roles.
 - **Administrator:** The administrator has the highest level of authority, and can view all windows and perform any action, including the addition and removal of user accounts. The administrator has access to all service functions and TS7700 resources.

- **Manager:** The manager has access to monitoring information, and performance data and functions, and can perform actions for users, including adding, modifying, and deleting user accounts. The manager is restricted from changing most other settings, including those for logical volume management, network configuration, feature licenses, and custom roles.
- **Custom roles:** The administrator can name and define two custom roles by selecting the individual tasks that are permitted to each custom role. Tasks can 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 window, click **Cancel**.

Figure 9-103 shows the sequence of creating a new user. It is used for managing users with the Local Authentication Policy method.

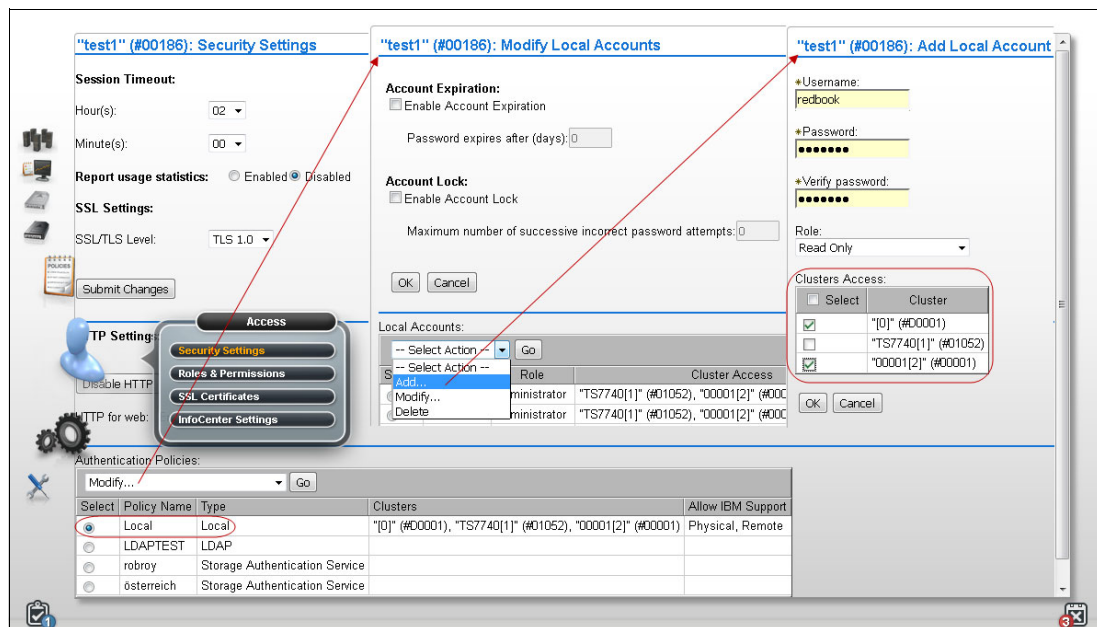


Figure 9-103 Creating a new local user

Modifying the user or group of the Local Authentication Policy

To modify a user or group property for a TS7700 grid, use this window.

Tip: Passwords for the users are changed from this window also.

To modify a user account belonging to the Local Authentication Policy, complete these steps:

1. On the TS7700 MI, click **Access** (blue man icon) → **Security Settings** from the left navigation window.
2. Click **Select** next to the Local policy name on the Authentication Policies table.
3. Select **Modify** from the Select Action menu and click **Go**.
4. On the Local Accounts table, click **Select** next to the user name of the policy to modify.

5. Select **Modify** from the Select Action menu and click **Go**. See Figure 9-103 for the Modify Local Accounts options.
6. Modify the values for any of the following fields:
 - Role: The role that is assigned to the user account. The following values are possible:
 - Operator: The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, and 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 cannot change network configuration, feature licenses, user accounts, and custom roles.
 - Administrator: The administrator has the highest level of authority, and can view all windows and perform any action, including the addition and removal of user accounts. The administrator has access to all service functions and TS7700 resources.
 - Manager: The manager has access to monitoring information and performance data and functions, and can perform actions for users, including adding, modifying, and deleting user accounts. The manager is restricted from changing most other settings, including those for logical volume management, network configuration, feature licenses, and custom roles.
 - Custom roles: The administrator can name and define two custom roles by selecting the individual tasks that are permitted for each custom role. Tasks can 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 window, click **Cancel**.

Note: The user 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 window is shown in Figure 9-106 on page 456. In the Cluster Access table, select the **Select** check box to toggle all the cluster check boxes on and off.

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). This section highlights the various windows that are required to manage the Storage Authentication Service Policy.

Important: When a Storage Authentication Service policy is enabled for a cluster, service personnel are required to log in with the setup user or group. Before enabling storage authentication, create an account that can be used by service personnel.

To add a Storage Authentication Service Policy for a TS7700 Grid, complete the following steps:

1. On the TS7700 MI, click **Access** (blue man icon) → **Security Settings** from the left navigation window.
2. On the Authentication Policies table, select **Add Storage Authentication Service Policy** from the Select Action menu.
3. Click **Go** to open the Add Storage Authentication Service Policy window that is shown in Figure 9-104 on page 454. 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 that is composed of 1 - 50 Unicode characters. Heading and trailing blank spaces are trimmed, although internal blank spaces are retained. After a new authentication policy is created, its policy name cannot be modified.
 - b. **Primary Server URL:** The primary URL for the Storage Authentication Service. The value in this field consists of 1 - 256 Unicode characters and takes the following format:
`https://<server_IP_address>:secure_port/TokenService/services/Trust`
 - c. **Alternative Server URL:** The alternative URL for the Storage Authentication Service if the primary URL cannot be accessed. The value in this field consists of 1 - 256 Unicode characters and takes the following format:
`https://<server_IP_address>:secure_port/TokenService/services/Trust`
4. To complete the operation, click **OK**. To abandon the operation and return to the Security Settings window, click **Cancel**.

Remember: If the Primary or alternative Server URL uses the HTTPS protocol, a certificate for that address must be defined on the SSL Certificates window.

- d. **Server Authentication:** Values in the following fields are required if IBM 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:
 - **User ID:** The user name that is used with HTTP basic authentication for authenticating to the Storage Authentication Service.
 - **Password:** The password that is used with HTTP basic authentication for authenticating to the Storage Authentication Service.

Figure 9-104 shows an example of adding a Storage Authentication Service Policy. We suggest marking the highlighted boxes to allow IBM service personnel access to the TS7700.

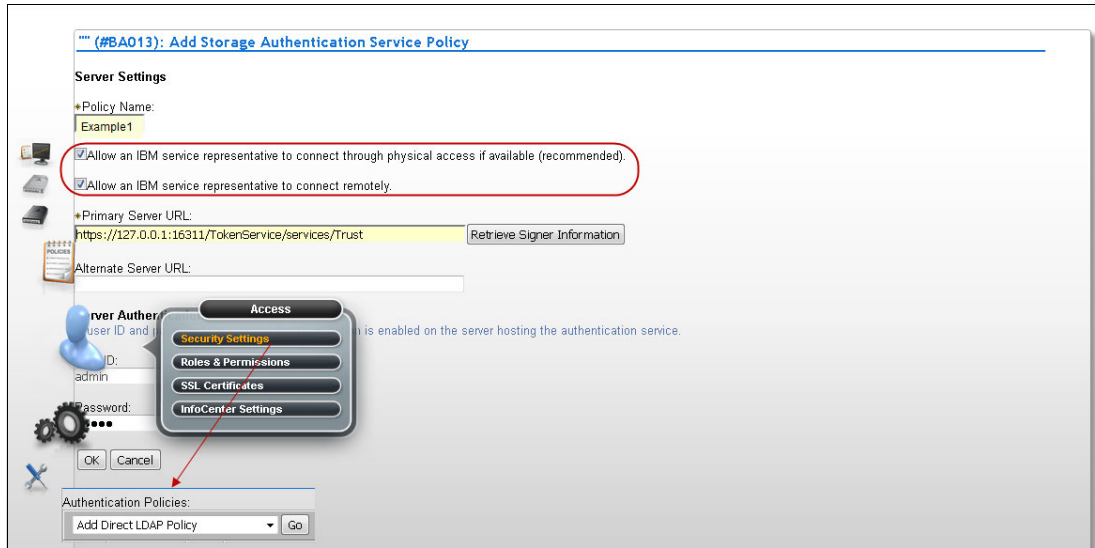


Figure 9-104 Add Storage Authentication Service Policy

Click **OK** to confirm the creation of the Storage Authentication Policy. The window that is shown in Figure 9-105 opens. In the Authentication Policies table, no clusters are assigned to the newly created policy, so the Local Authentication Policy is enforced. When the newly created policy is in this state, it can be deleted because it is not applied to any of the clusters.

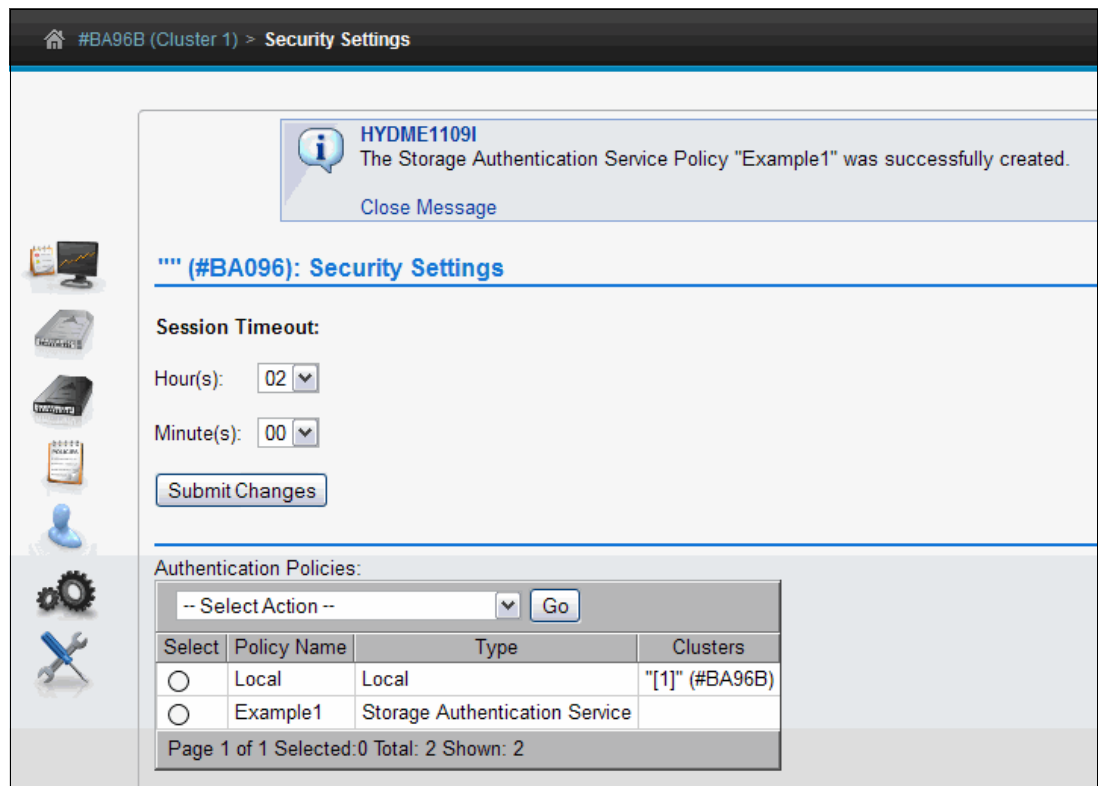


Figure 9-105 Addition of Storage Authentication Service Policy completed

Adding a user to a Storage Authentication Policy

To add a user to a Storage Authentication Service Policy for a TS7700 Grid, complete the following steps:

1. On the TS7700 MI, click **Access** → **Security Settings** from the left navigation window.
2. Select the policy to be modified.
3. In the Authentication Policies table, select **Modify** from the **Select Action** menu.
4. Click **Go** to open the Modify Storage Authentication Service Policy window.
5. In the Modify Storage Authentication Service Policy page, go to the **Storage Authentication Service Users/Groups** table at the bottom.
6. Select **Add User** from the Select Action menu.
7. Click **Go** to open the Add External Policy User window.
8. In the Add External Policy User window, enter values for the following required fields:
 - User name: 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 that is assigned to the user account. The role can be a predefined role or a user-defined role. The following values are valid:
 - Operator: The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, and 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 as the administrator, but cannot change network configuration, feature licenses, user accounts, and custom roles.
 - Administrator: The administrator has the highest level of authority, and can view all windows and perform any action, including the addition and removal of user accounts. The administrator has access to all service functions and TS7700 resources.
 - Manager: The manager has access to monitoring information and performance data and functions, and can perform actions for users, including adding, modifying, and deleting user accounts. The manager is restricted from changing most other settings, including those for logical volume management, network configuration, feature licenses, and custom roles.
 - Custom roles: The administrator can name and define two custom roles by selecting the individual tasks that are permitted for each custom role. Tasks can be assigned to a custom role in the Roles and assigned permissions table in the Roles & Permissions Properties window.
 - Cluster Access: The clusters (can be multiple) to which the user has access.
9. To complete the operation, click **OK**. To abandon the operation and return to the Modify Local Accounts window, click **Cancel**.

10. Click **OK** after the fields are complete. Figure 9-106 shows the sequence to add a new user to an existing Storage authentication Service Policy.

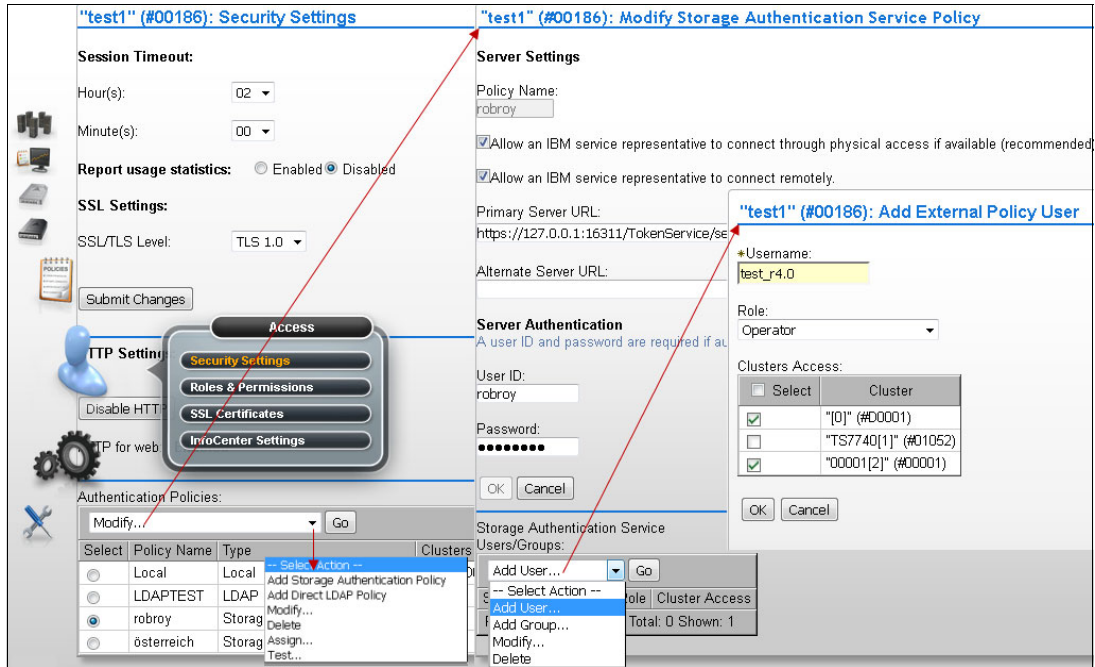


Figure 9-106 Adding a new user to an existent Storage Authentication Service Policy

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 it is necessary to have sufficient privileges with the new policy, you must enter a user name and password that is recognized by the new authentication policy.

To add a user to a Storage Authentication Service Policy for a TS7700 grid, complete the following steps:

1. On the TS7700 MI, click **Access** → **Security Settings** from the left navigation window.
2. In the Authentication Policies table, select **Assign** from the Select Action menu as shown in Figure 9-107.

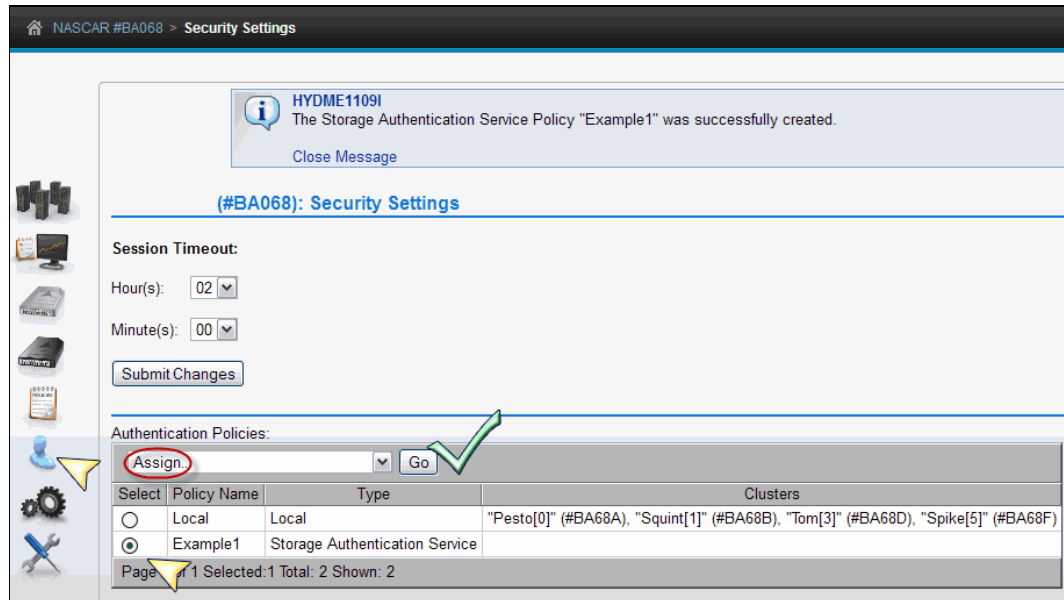


Figure 9-107 Assign Storage Authentication Service Policy to grid resources

3. Click **Go** to open the Assign Authentication Policy window that is shown in Figure 9-108.

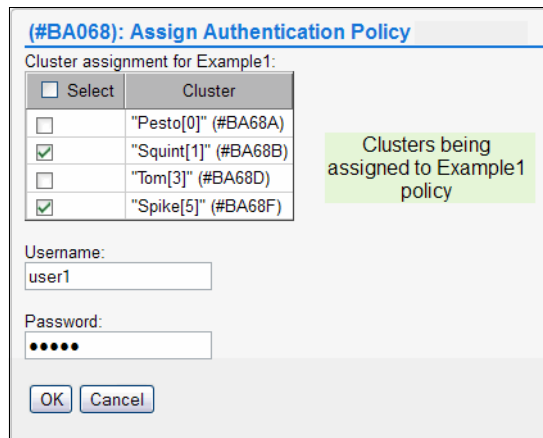


Figure 9-108 Cluster assignment selection

4. To apply the authentication policy to a cluster, select the check box next to the cluster's name.

Enter values for the following fields:

- User name: User name for the TS7700 MI.
- Password: Password for this TS7700 MI user.

5. To complete the operation, click **OK**. To abandon the operation and return to the Security Settings window, click **Cancel**.

Deleting a Storage Authentication Policy

The user can delete a Storage Authentication Service policy if it is not in effect on any cluster. The Local policy cannot be deleted. In the Authentication Policies table in Figure 9-109, no clusters are assigned to the policy, so it can be deleted. If clusters are assigned to the policy, use **Modify** from the **Select Action** menu to remove the assigned clusters.

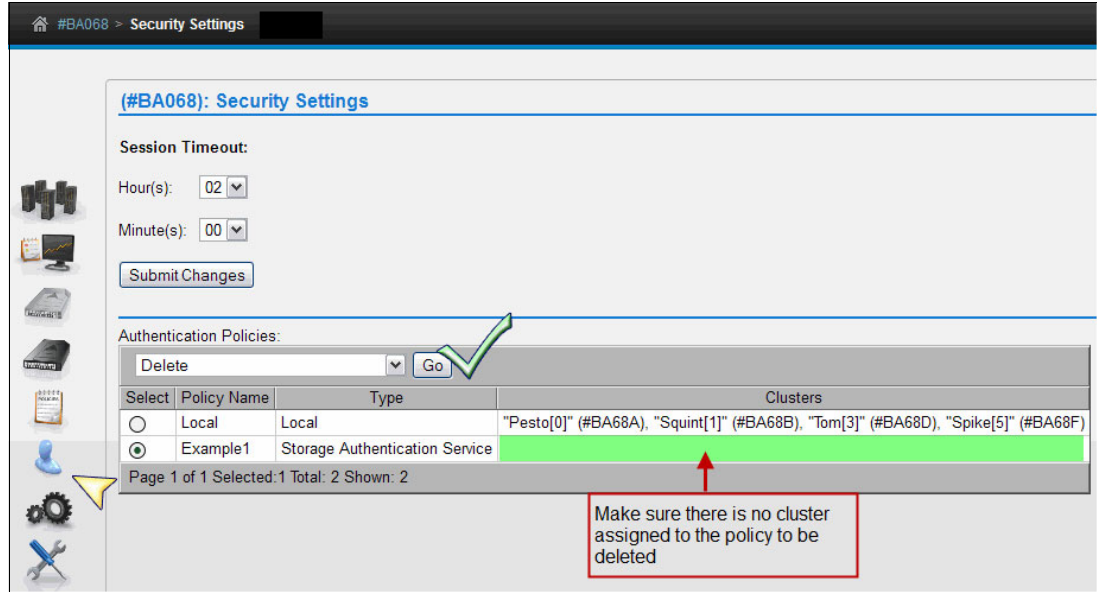


Figure 9-109 Delete a Storage Authentication Service Policy

To delete a Storage Authentication Service Policy from a TS7700 grid, complete the following steps:

1. On the TS7700 MI, click **Access** → **Security Settings** from the left navigation window.
2. In the Authentication Policies table, select **Delete** go to the **Select Action** menu as shown in Figure 9-109. From the Security Settings window, go to the Authentication Policies table and complete the following steps:
 - a. Select the radio button next to the policy that must be deleted.
 - b. Select **Delete** from the **Select Action** menu.
 - c. Click **Go** to open the Confirm Delete Storage Authentication Service policy window.
 - d. Click **OK** to delete the policy and return to the Security Settings window, or click **Cancel** to abandon the delete operation and return to the Security Settings window.

3. Confirm the policy deletion, as shown in Figure 9-110. Click **OK** to delete the policy.

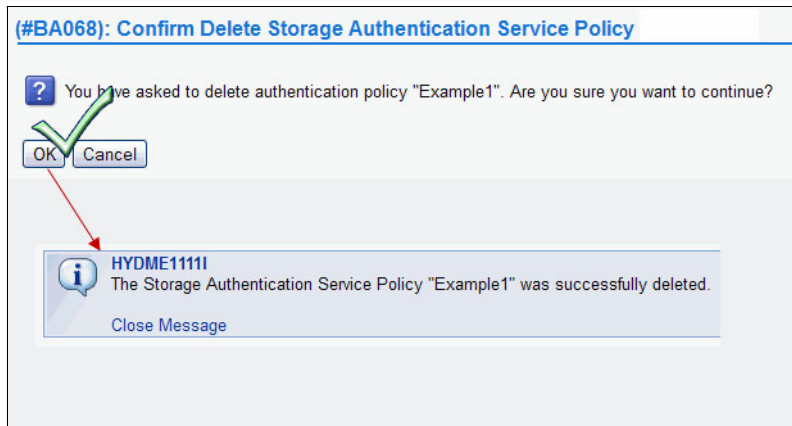


Figure 9-110 Confirm delete and delete action successful messages

Testing an Authentication Policy

Before a new Authentication Policy can be used, it must be tested. The test validates the login credentials (user ID and password) in all clusters for which this user ID and role are authorized. Also, access to the external resources needed by an external authentication policy, such as an SSPC or an LDAP server, is tested. The credentials that are entered in the test window (User ID and Password) are authenticated and validated by the LDAP server, for an external policy.

Tip: The policy needs to be configured to an LDAP server before being added in the TS7700 MI. External users and groups to be mapped by the new policy are checked in LDAP before being added.

To test the security settings for the TS7700 grid, complete the following steps. Use these steps to test the roles that are assigned to the user name by an existing policy.

1. From the Security Settings window, go to the Authentication Policies table:
 - Select the radio button next to the policy to test.
 - Select **Test** from the Select Action menu.
 - Click **Go** to open the Test Authentication Policy window.
2. Check the check box next to the name of each cluster on which to conduct the policy test.
3. Enter values for the following fields:
 - User name: The user name for the TS7700 MI. This value consists of 1 - 16 Unicode characters.
 - Password: The password for the TS7700 MI. This value consists of 1- 16 Unicode characters.

Note: If the user name entered belongs to a user not included on the policy, test results show success, but the result comments show a null value for the role and access fields. Additionally, the user name that entered cannot be used to log in to the MI.

4. Click **OK** to complete the operation. If you must abandon the operation, click **Cancel** to return to the Security Settings window.

When the authentication policy test completes, the Test Authentication Policy results window opens to display results for each selected cluster. See Figure 9-111 for an example.

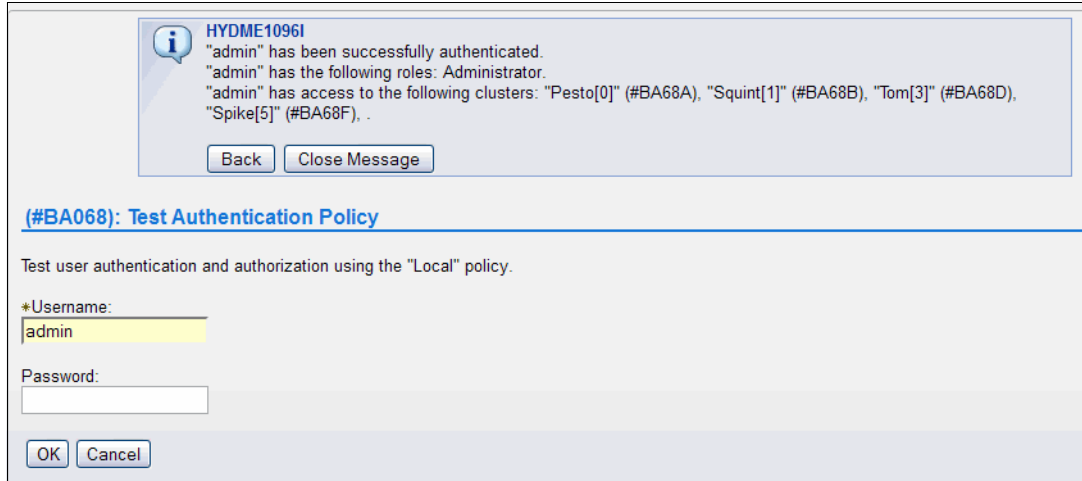


Figure 9-111 Test Authentication Policy results

The results include a statement indicating whether the test succeeded or failed, and if it failed, the reason for the failure. The Test Authentication Policy results window also displays the Policy Users table. Information that is shown on that table includes the following fields:

- ▶ **Username.** The name of a user who is authorized by the selected authentication policy.
- ▶ **Role.** The role that is assigned to the user under the selected authentication policy.
- ▶ **Cluster Access.** A list of all the clusters in the grid for which the user and user role are authorized by the selected authentication policy. Check Figure 9-112 for an example of a failure in the Test Authentication Policy.

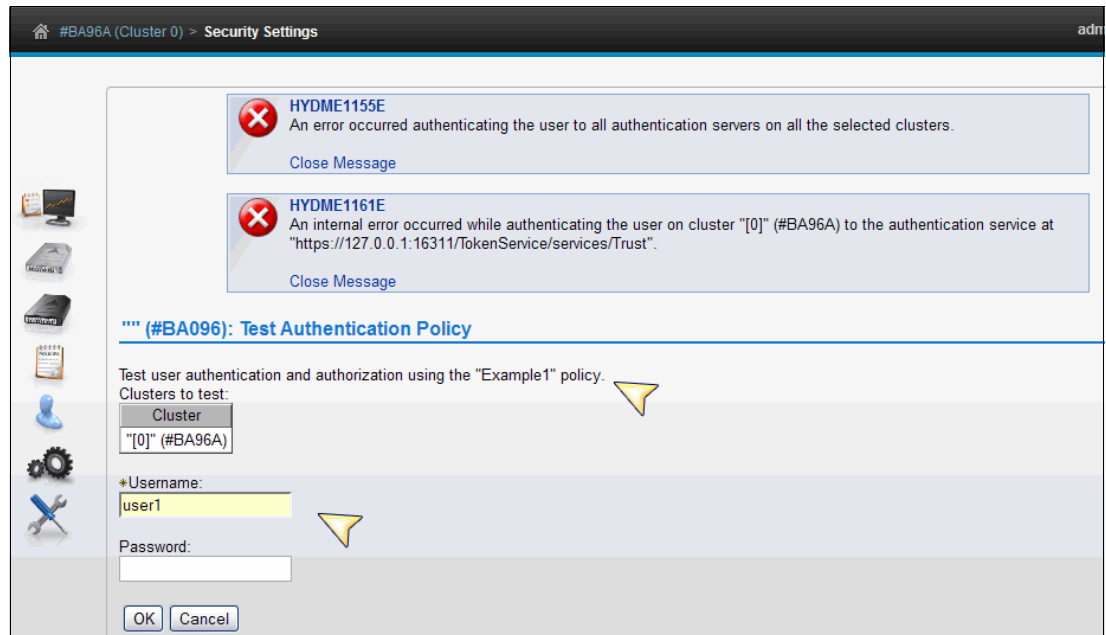


Figure 9-112 Failure in the Test Authentication Policy

To return to the Test Authentication Policy window, click **Close Window**. To return to the Security Settings window, click **Back** at the top of the Test Authentication Policy results window.

Adding a Direct LDAP policy

A *Direct LDAP Policy* is an external policy that maps user, group, and role relationships. Users are authenticated and authorized through a direct communication with an LDAP server. This section highlights the various windows that are required to manage a Direct LDAP policy.

Important: When a Direct LDAP policy is enabled for a cluster, service personnel are required to log in with the setup user or group. Before enabling LDAP authentication, create an account that can be used by service personnel. Also, the user can enable an IBM SSR to connect to the TS7700 through physical access or remotely by selecting those options in the DIRECT LDAP POLICY window.

To add a Direct LDAP Policy for a TS7700 grid, complete the following steps:

1. On the TS7700 MI, click **Access** → **Security Settings** from the left navigation window.
2. From the menu, select **Add Direct LDAP Policy** and click **GO**. Use Figure 9-113 for reference.

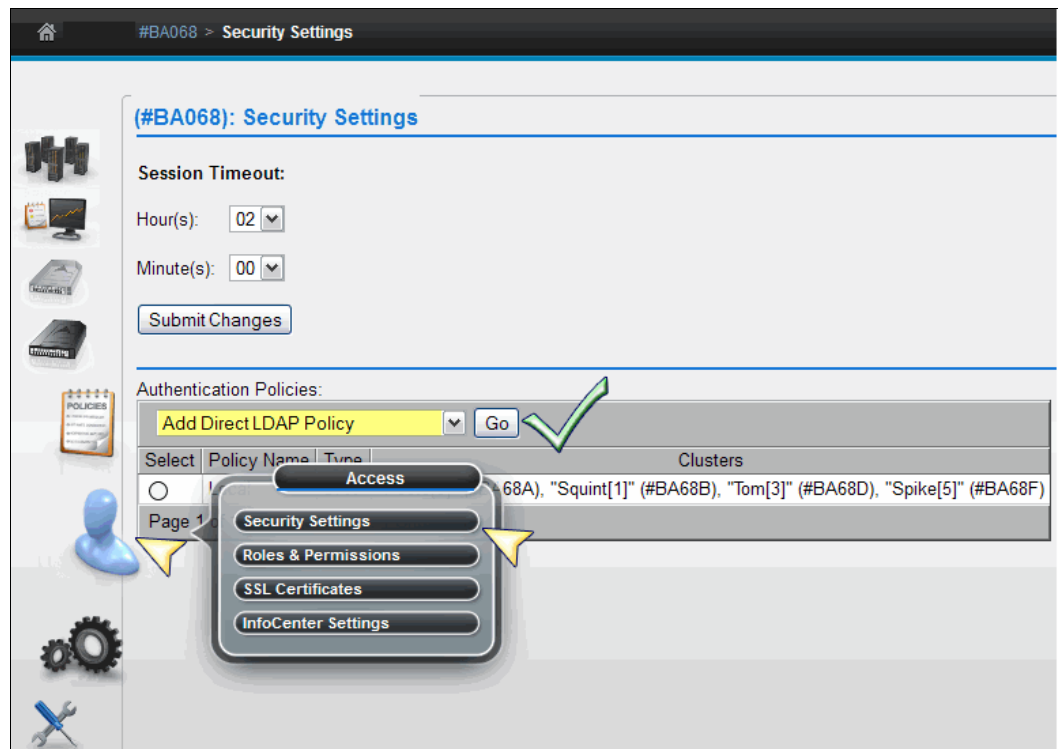


Figure 9-113 LDAP Policy selection

- Use the options in Figure 9-114 to grant to the IBM SSR a local or remote connection for service support.

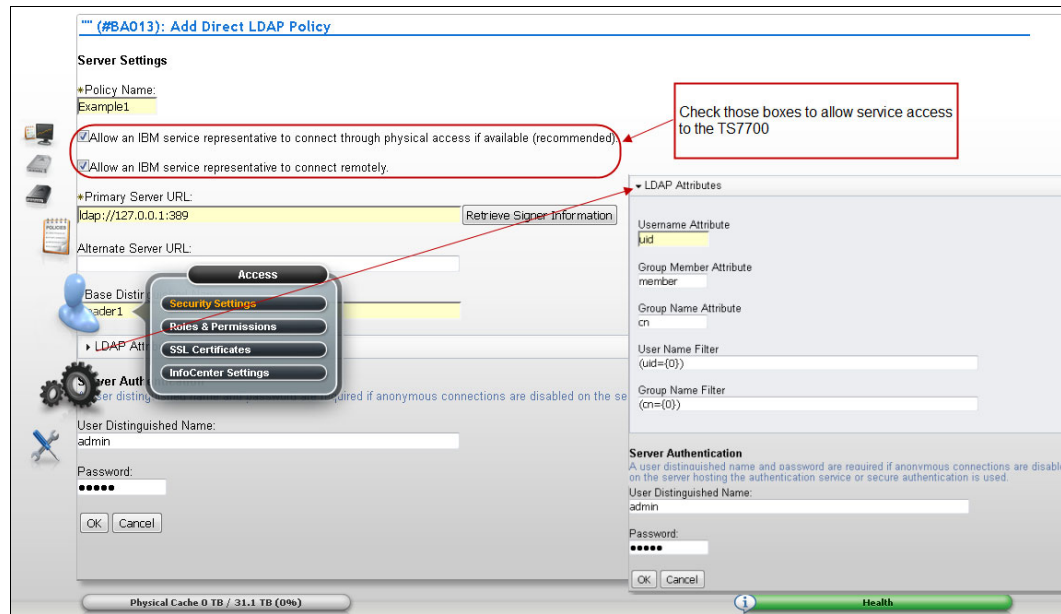


Figure 9-114 Adding a Direct LDAP Policy

Note: LDAP external authentication policies are not available for backup or recovery through the backup or restore settings operations. Record it, keep it safe, and have it available for a manual recovery as dictated by the security standards.

The values in the following fields are required if secure authentication is used or anonymous connections are disabled on the LDAP server:

- ▶ **User Distinguished Name:** The user distinguished name is used to authenticate to the LDAP authentication service. This field supports a maximum length of 254 Unicode characters, for example:
CN=Administrator,CN=users,DC=mycompany,DC=com
- ▶ **Password:** The password is used to authenticate to the LDAP authentication service. This field supports a maximum length of 254 Unicode characters.

When modifying an LDAP Policy, the LDAP attributes fields also can be changed:

- ▶ **Base Distinguish Name:** The LDAP distinguished name (DN) that uniquely identifies a set of entries in a realm. This field is required but blank by default. The value in this field consists of 1 - 254 Unicode characters.
- ▶ **User Name Attribute:** The attribute name that is used for the user name during authentication. This field is required and contains the value uid by default. The value in this field consists of 1 - 61 Unicode characters.
- ▶ **Password:** The attribute name that is used for the password during authentication. This field is required and contains the value userPassword by default. The value in this field consists of 1 - 61 Unicode characters.
- ▶ **Group Member Attribute:** The attribute name that is used to identify group members. This field is optional and contains the value member by default. This field can contain up to 61 Unicode characters.

- ▶ **Group Name Attribute:** The attribute name that is used to identify the group during authorization. This field is optional and contains the value `cn` by default. This field can contain up to 61 Unicode characters.
- ▶ **User Name filter:** Used to filter and verify the validity of an entered user name. This field is optional and contains the value `(uid={0})` by default. This field can contain up to 254 Unicode characters.
- ▶ **Group Name filter:** Used to filter and verify the validity of an entered group name. This field is optional and contains the value `(cn={0})` by default. This field can contain up to 254 Unicode characters.

Click **OK** to complete the operation. Click **Cancel** to abandon the operation and return to the Security Settings window.

Creating a RACF based LDAP Policy

The process is similar to the previous item “Adding a Direct LDAP policy” on page 461. There are some required configurations on the host side regarding the RACF, SDBM, and IBM Security Directory Server that should be performed in advance before this capability can be made operational. See Chapter 10, “Host Console operations” on page 601 for the description of the parameters and configurations. When those configurations are ready, the RACF based LDAP Policy can be created and activated. See Figure 9-115 to add a RACF policy.

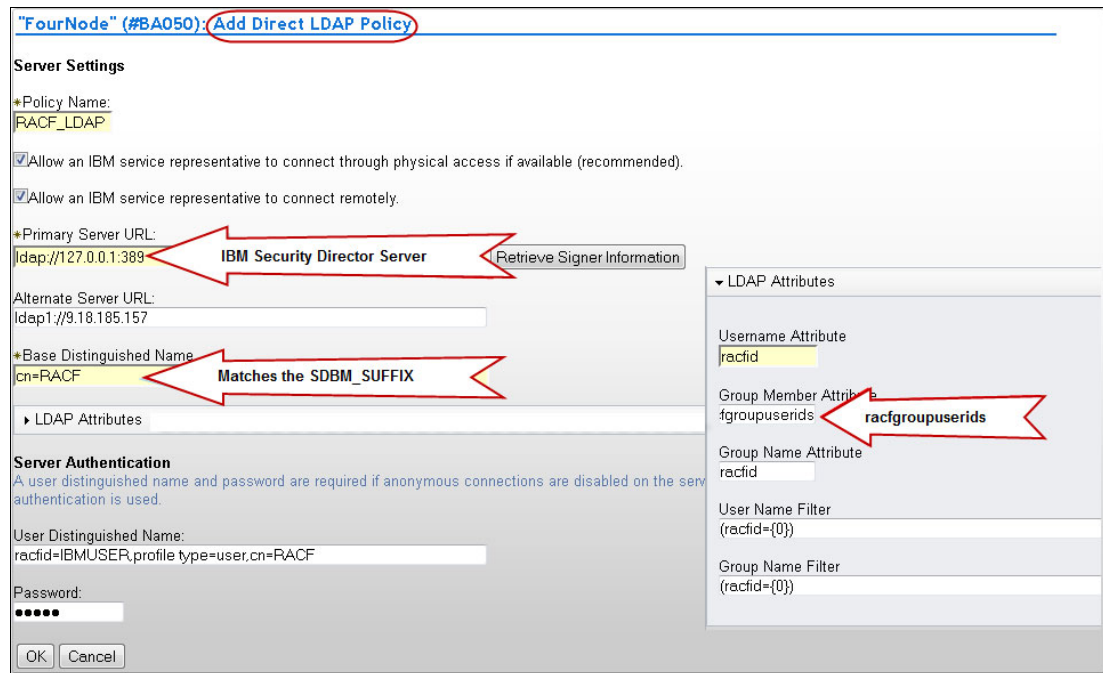


Figure 9-115 Adding a RACF based Direct LDAP Policy

As shown in Figure 9-115, a new policy is created, which is called `RACF_LDAP`. The Primary Server URL is that of the IBM Security Directory Server, the same way any regular LDAP server is configured.

The Base Distinguished Name matches the `SDBM_SUFFIX`.

In the screen capture, the **Group Member Attribute** was set to `racfgroupuserids` (it shows truncated in MI's text box).

The item **User Distinguished Name** should be specified with all of the following parameters:

- ▶ racfid
- ▶ profilename
- ▶ cn

When the previous setup is complete, more users can be added to the policy, or clusters can be assigned to it, as described in the topics to follow. There are no specific restrictions for these RACF/LDAP user IDs, and they can be used to secure the MI, or the IBM service login (for the IBM SSR) just as any other LDAP user ID.

See the TS7700 3.3 IBM Knowledge Center, available locally on the MI window by clicking the question mark on the upper right upper bar and selecting **Help**, or on the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_3.3.0/ts7700_security_ldap_racf.html

Adding users to a Direct LDAP Policy

See the process that is described in the “Adding a user to a Storage Authentication Policy” on page 455. The same steps apply when adding users to a Direct LDAP Policy.

Assigning a Direct LDAP Policy to a cluster or clusters

See the procedure that is described in “Assigning clusters to a Storage Authentication Policy” on page 456. The same steps apply when working with a Direct LDAP Policy.

Deleting a Direct LDAP Policy

See the procedure that is described in “Deleting a Storage Authentication Policy” on page 458. The same steps apply when deleting a Direct LDAP Policy.

Roles & Permissions window

You can use the window that is shown in Figure 9-116 to set and control user roles and permissions for a TS7700 Grid.



Figure 9-116 TS7700 MI Roles & Permissions window

Figure 9-117 shows the Roles & Permissions window, listing the user roles and a summary of each role.

"" (#BA096): Roles & Permissions

Refresh Last Refresh: Nov 15, 2012 9:26:12 PM

View tutorial

Manage users and assign user roles

Select	Roles	Descriptions
<input type="checkbox"/>	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 virtual volumes.
<input type="checkbox"/>	Lead Operator	The lead operator has almost all of the same permissions as does the administrator, but may not change network configuration, feature licenses, user accounts or custom roles.
<input type="checkbox"/>	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.
<input type="checkbox"/>	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 virtual volume management, network configuration, or feature licenses.
<input type="checkbox"/>	Custom Role 1	The administrator can name and define this custom role by selecting the individual tasks permitted to this custom role.
<input type="checkbox"/>	Custom Role 2	The administrator can name and define this custom role by selecting the individual tasks permitted to this custom role.
<input type="checkbox"/>	Custom Role 3	The administrator can name and define this custom role by selecting the individual tasks permitted to this custom role.
<input type="checkbox"/>	Custom Role 4	The administrator can name and define this custom role by selecting the individual tasks permitted to this custom role.
<input type="checkbox"/>	Custom Role 5	The administrator can name and define this custom role by selecting the individual tasks permitted to this custom role.
<input type="checkbox"/>	Custom Role 6	The administrator can name and define this custom role by selecting the individual tasks permitted to this custom role.

Page 1 of 2 1 Go Total: 14 Displayed: 10

Figure 9-117 Roles & Permissions window

Each role is described in the following list:

- ▶ **Operator:** The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, and 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 cannot change network configuration, feature licenses, user accounts, and custom roles.
- ▶ **Administrator:** The administrator has the highest level of authority, and can view all windows and perform any action, including the addition and removal of user accounts. The administrator has access to all service functions and TS7700 resources.
- ▶ **Manager:** The manager has access to monitoring information, performance data, and functions, and can perform actions for users. The manager is restricted from changing most settings, including those for logical volume management, network configuration, feature licenses, user accounts, and custom roles.
- ▶ **Custom roles:** The administrator can name and define 10 custom roles by selecting the individual tasks that are permitted to each custom role. Tasks can be assigned to a custom role in the Roles and Assigned Permissions window.

Roles and Assigned Permissions table

The Roles and Assigned Permissions table is a dynamic table that displays the complete list of TS7700 grid tasks and the permissions that are assigned to selected user roles.

To view the Roles and Assigned Permissions table, complete the following steps:

1. Select the check box to the left of the role to be displayed. The user can 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. Subsequent columns show the assigned permissions for selected role (or roles). A check mark denotes permitted tasks for a user role. A null dash (-) denotes prohibited tasks for a user role.

Permissions for predefined user roles cannot be modified. The user can name and define up to 10 different custom roles, if necessary. The user can modify permissions for custom roles in the Roles and Assigned Permissions table. The user can modify only one custom role at a time.

To modify a custom role, complete 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 the requirements by selecting (permitting) or clearing (prohibiting) tasks. Selecting or clearing a parent task affects any child tasks. However, a child task can be selected or cleared independently of a parent task. The user can apply the permissions of a predefined role to a custom role by selecting a role from the Role Template menu and clicking **Apply**. The user can then customize the permissions by selecting or clearing tasks.
3. After all tasks for the custom role are selected, click **Submit Changes** to activate the new custom role.

Remember: The user can apply the permissions of a predefined role to a custom role by selecting a role from the Role Template menu and clicking **Apply**. The user can then customize the permissions by selecting or clearing tasks.

SSL Certificates window

Use the window that is shown in Figure 9-118 to view, import, or delete SSL certificates that support secure connections to a Storage Authentication Service server from a TS7700 cluster. If a Primary or alternative Server URL, defined by a Storage Authentication Service Policy, uses the HTTPS protocol, a certificate for that address must be defined in this window.

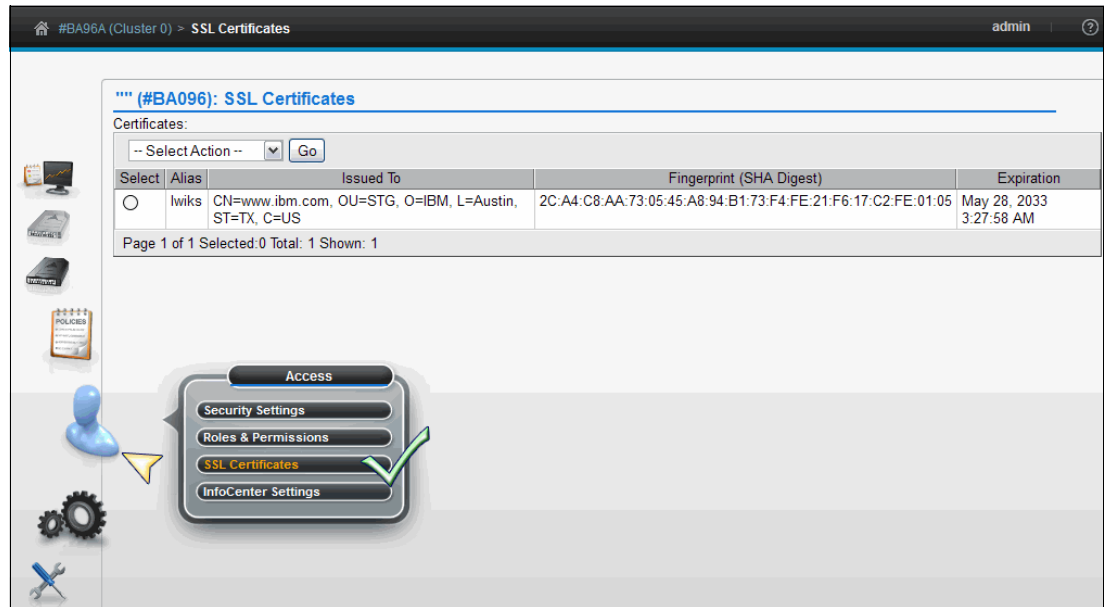


Figure 9-118 SSL Certificates window

The Certificates table displays the following identifying information for SSL certificates on the cluster:

- ▶ **Alias:** A unique name to identify the certificate on the system.
- ▶ **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, complete the following steps:

1. Select **Retrieve from port** from the Select Action 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, complete the following steps:

1. Select the radio button next to the certificate to delete, select **Delete** from the Select Action menu, and click **Go**. The Confirm Delete SSL Certificate window opens and prompts to confirm the 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

To upload a new TS7700 IBM Knowledge Center to the cluster's MI, use the window that is shown in Figure 9-119.

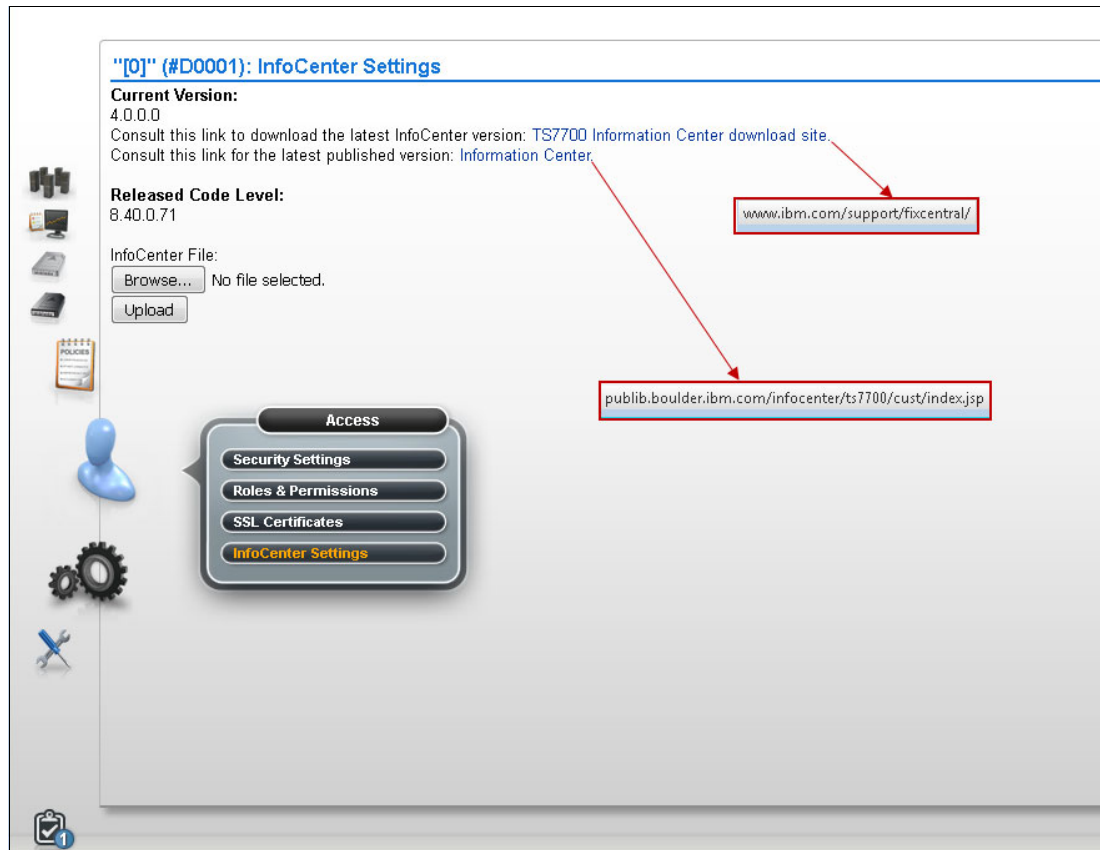


Figure 9-119 InfoCenter Settings window

This window has the following items:

- ▶ Current Version section, where the following items can be identified or accessed:
 - Identify the version level and date of IBM Knowledge Center that is installed on the cluster.
 - Access a product database to download a JAR file containing a newer version of IBM Knowledge Center.
 - Access an external site displaying the most recently published version of IBM Knowledge Center.

- ▶ The TS7700 IBM Knowledge Center download site link.
Click this link to open the Fix Central product database so that the user can download a new version of the TS7700 IBM Knowledge Center as a .jar file (if available):
 - a. Select **System Storage** from the Product Group menu.
 - b. Select **Tape Systems** from the Product Family menu.
 - c. Select **TS7700** from the Product menu.
 - d. Click **Continue**.
 - e. On the **Select Fixes** window, check the box next to the wanted **InfoCenter Update file** (if available).
 - f. Click **Continue**.
 - g. On the Download Options window, select **Download using Download Director**.
 - h. Select the check box next to **Include prerequisites and co-requisite fixes**.
 - i. Click **Continue**.
 - j. On the Download files using Download Director window, ensure that the check box next to the correct InfoCenter Update version is checked and click **Download now**. The Download Director applet opens. The downloaded file is saved at C:\DownloadDirector\.
- ▶ With the new .jar file that contains the updated IBM Knowledge Center (either from the Fix Central database or from an IBM SSR), save the .jar file to a local directory.
To upload and install the new IBM Knowledge Center, complete the following steps:
 - a. Click **Browse** to open the File Upload window.
 - b. Go to the folder that contains the new .jar file.
 - c. Highlight the new .jar file name and click **Open**.
 - d. Click **Upload** to install the new IBM Knowledge Center on the cluster's MI.

9.2.10 The Settings icon

The TS7700 MI windows that are collected under the Settings icon can help you view or change cluster network settings, feature licenses, SNMP, and library port access groups.

Figure 9-120 shows the Settings icon and options.



Figure 9-120 The Setting icon and options

Cluster network settings

To set or modify IP addresses for the selected TS7700 cluster, use this window.

Figure 9-121 shows the Cluster Network Setting navigation and the Customer IP addresses tab.

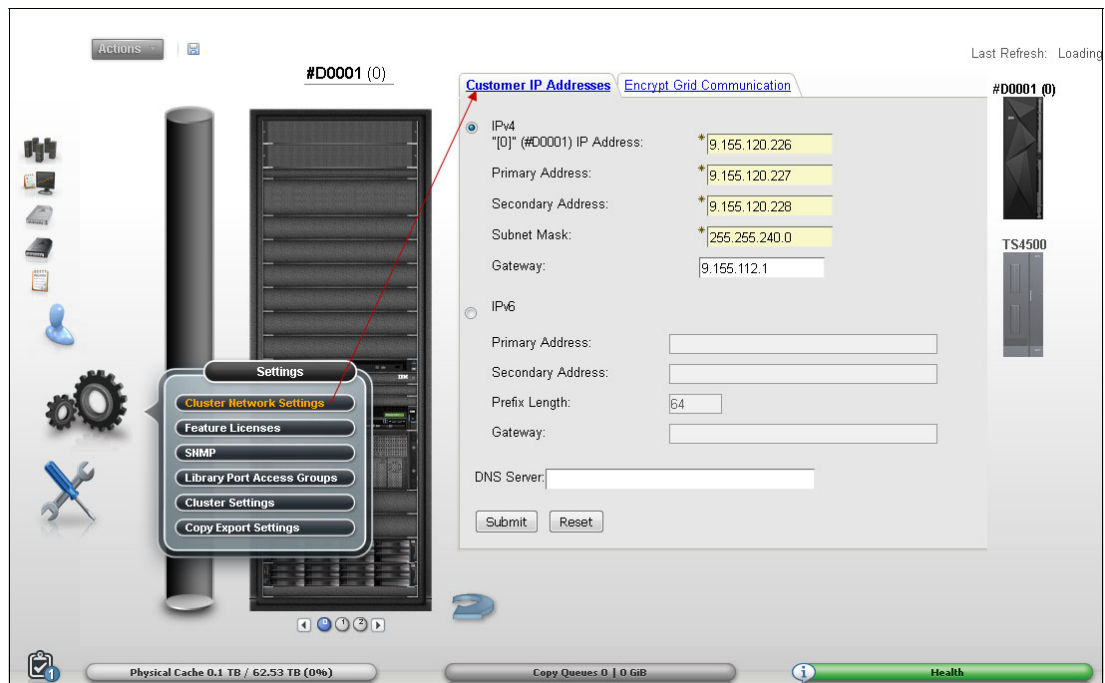


Figure 9-121 Customer Networking settings page and the IP tab

Customer IP addresses tab

Use this tab to set or modify the MI IP addresses for the selected cluster. Each cluster is associated with two routers or switches. Each router or switch is assigned an IP address and one virtual IP address is shared between routers or switches.

Note: Any modifications to IP addresses on the accessing cluster interrupt access to that cluster for all current users. If the accessing cluster IP addresses are modified, the current users are redirected to the new virtual address.

The following fields show on this tab:

- ▶ IPv4: Select this radio button if the cluster can be accessed by an IPv4 address. If this option is disabled, all incoming IPv4 traffic is blocked, although loop-back traffic is still permitted.

If this option is enabled, specify the following addresses:

- *<Cluster Name>* IP address: An AIX virtual IPv4 address that receives traffic on both customer networks. This field cannot be blank if IPv4 is enabled.
- Primary Address: The IPv4 address for the primary customer network. This field cannot be blank if IPv4 is enabled.
- Secondary Address: The IPv4 address for the secondary customer network. This field cannot be blank if IPv4 is enabled.
- Subnet Mask: The IPv4 subnet mask that is used to determine the addresses present on the local network. This field cannot be blank if IPv4 is enabled.
- Gateway: The IPv4 address that is used to access systems outside the local network.

A valid IPv4 address is 32 bits long, consists of four decimal numbers, each 0 - 255, separated by periods, such as 98.104.120.12.

- ▶ IPv6: Select this radio button if the cluster can be accessed by an IPv6 address. If this option is disabled, all incoming IPv6 traffic is blocked, although loop-back traffic is still permitted. If the user enables this option and does not designate any additional IPv6 information, the minimum required local addresses for each customer network interface will automatically be enabled and configured by using neighbor discovery. If this option is enabled, the user can specify the following addresses:

- Primary Address: The IPv6 address for the primary network. This field cannot be blank if IPv6 is enabled.
- Secondary Address: The IPv6 address for the secondary network. This field cannot be blank if IPv6 is enabled.
- Prefix Length: The IPv6 prefix length that is used to determine the addresses present on the local network. The value in this field is an integer 1 - 128. This field cannot be blank if IPv6 is enabled.
- Gateway: The IPv6 address that is used to access systems outside the local network.

A valid IPv6 address is a 128-bit long hexadecimal value that is separated into 16-bit fields by colons, such as 3afa:1910:2535:3:110:e8ef:ef41:91cf.

Leading zeros can be omitted in each field so that :0003: can be written as :3:. A double colon (::) can be used once per address to replace multiple fields of zeros, for example:

3afa:0:0:0:200:2535:e8ef:91cf

can be written as:

3afa::200:2535:e8ef:91cf

- ▶ DNS Server: The IP addresses of any domain name server (DNS), separated by commas. DNS addresses are needed only when specifying a symbolic domain name rather than a numeric IP address for one or more of the following types of information:
 - Primary Server URL on the Add External policy window
 - Encryption Key Server (EKS) address
 - SNMP server address
 - Security server address

If this field is left blank, the DNS server address is populated by Dynamic Host Configuration Protocol (DHCP).

The address values can be in IPv4 or IPv6 format. A maximum of three DNS servers can be added. Any spaces that are entered in this field are removed.

To submit changes, click **Submit**. If the user changes apply to the accessing cluster, a warning message is displayed that indicates that the current user access will be interrupted. To accept changes to the accessing cluster, click **OK**. To reject changes to the accessing cluster and return to the IP addresses tab, click **Cancel**.

To reject the changes that are made to the IP addresses fields and reinstate the last submitted values, select **Reset**. The user can also refresh the window to reinstate the last submitted values for each field.

Encrypt Grid Communication tab

Use this tab to encrypt grid communication between specific clusters.

Important: Enabling grid encryption significantly affects the performance of the TS7700. System performance can be reduced by 70% or more when grid encryption is enabled.

Figure 9-122 shows the Encrypt Grid Communication tab. In the example, the option was to encrypt grid communications between Cluster 3 and Cluster 5, and also between Cluster 0 and Cluster 3. The remaining paths in the example are not encrypted.

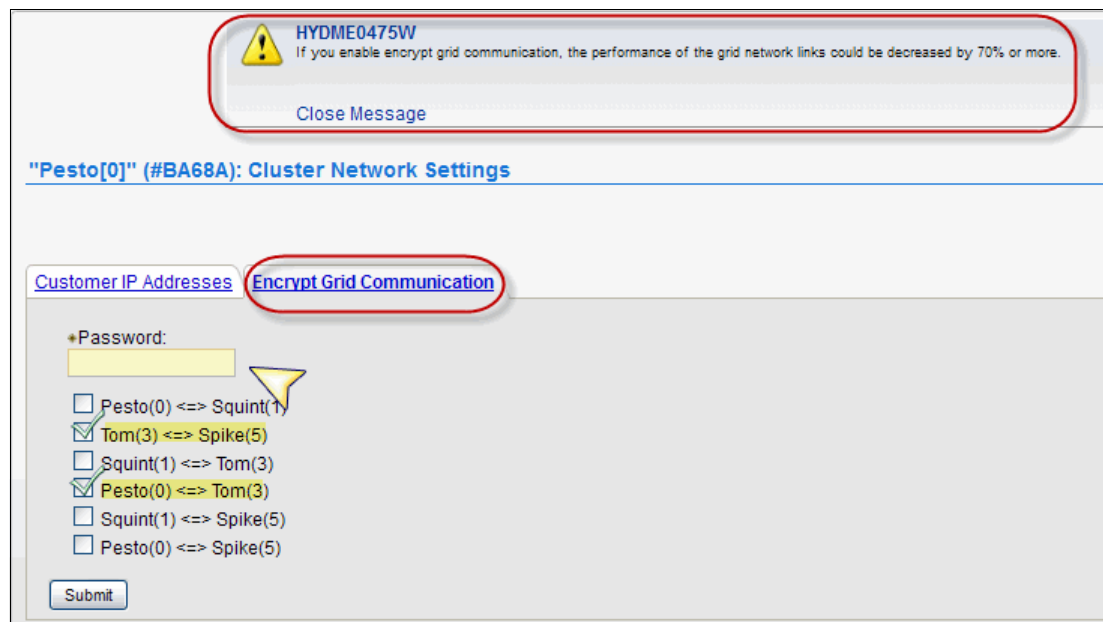


Figure 9-122 Encrypt Grid Communication tab

This tab includes the following fields:

- ▶ Password: This password is used as an EK to protect grid communication. This value has a 255 ASCII character limit, and is required.
- ▶ Cluster communication paths: Select the box next to each cluster communication path to be encrypted.

The user can select a communication path between two clusters only if both clusters meet all the following conditions:

- ▶ Are online
- ▶ Operate at a Licensed Internal Code level of 8.30.0.x or higher
- ▶ Operate by using IPv6-capable servers (3957-V07/VEB)

To submit changes, click **Submit**.

Feature licenses

Use this window to view information about feature licenses, or to activate or remove feature licenses from the TS7700 cluster.

Figure 9-123 shows the Feature Licenses window.

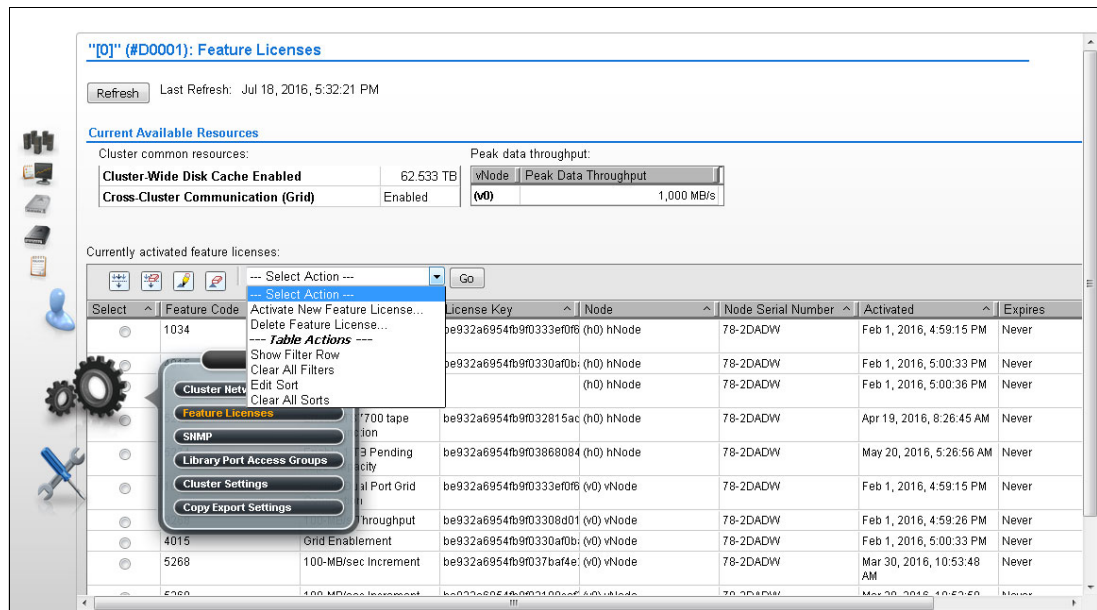


Figure 9-123 Feature Licenses window

The fields on the Feature Licenses window in the MI are described:

- ▶ **Cluster common resources.** The Cluster common resources table displays a summary of resources that are affected by activated features. The following information is displayed:
 - **Cluster-Wide Disk Cache Enabled.** The amount of disk cache that is enabled for the entire cluster, in terabytes (TB). If the selected cluster does not possess a physical library, the value in this field displays the total amount of cache that is installed on the cluster. Access to cache by a cluster without a physical library is not controlled by feature codes.
 - **Cross-Cluster Communication (Grid).** Whether cross-cluster communication is enabled on the grid. If this option is enabled, multiple clusters can form a grid. The possible values are Enabled and Disabled.

- **Peak data throughput.** The Peak data throughput table displays for each vnode the peak data throughput in megabytes per second (MBps). The following information is displayed:
 - **Vnode.** Name of the vnode.
 - **Peak data throughput.** The upper limit of the data transfer speed between the vnode and the host, which is displayed in MBps.
- ▶ **Currently activated feature licenses.** The Currently activated feature licenses table displays a summary of features that are installed on each cluster:
 - **Feature Code.** The feature code number of the installed feature.
 - **Feature Description.** A description of the feature that was installed by the feature license.
 - **License Key.** The 32-character license key for the feature.
 - **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. The following values are possible:
 - **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.

Note: The user can back up these settings as part of the `ts7700_cluster<cluster ID>.xmi` file and restore them for later use. When the backup settings are restored, new settings are added but no settings are deleted. The user cannot restore feature license settings to a cluster that is different from the cluster that created the `ts7700_cluster<cluster ID>.xmi` backup file. After restoring feature license settings on a cluster, log out and then log in to refresh the system.

Use the menu on the Currently activated feature licenses table to activate or remove a feature license. The user can also use this menu to sort and filter feature license details.

Simple Network Management Protocol

To view or modify the SNMP configured on a TS7700 Cluster, use this window on the TS7700 MI. Figure 9-124 shows the SNMP window in the MI.

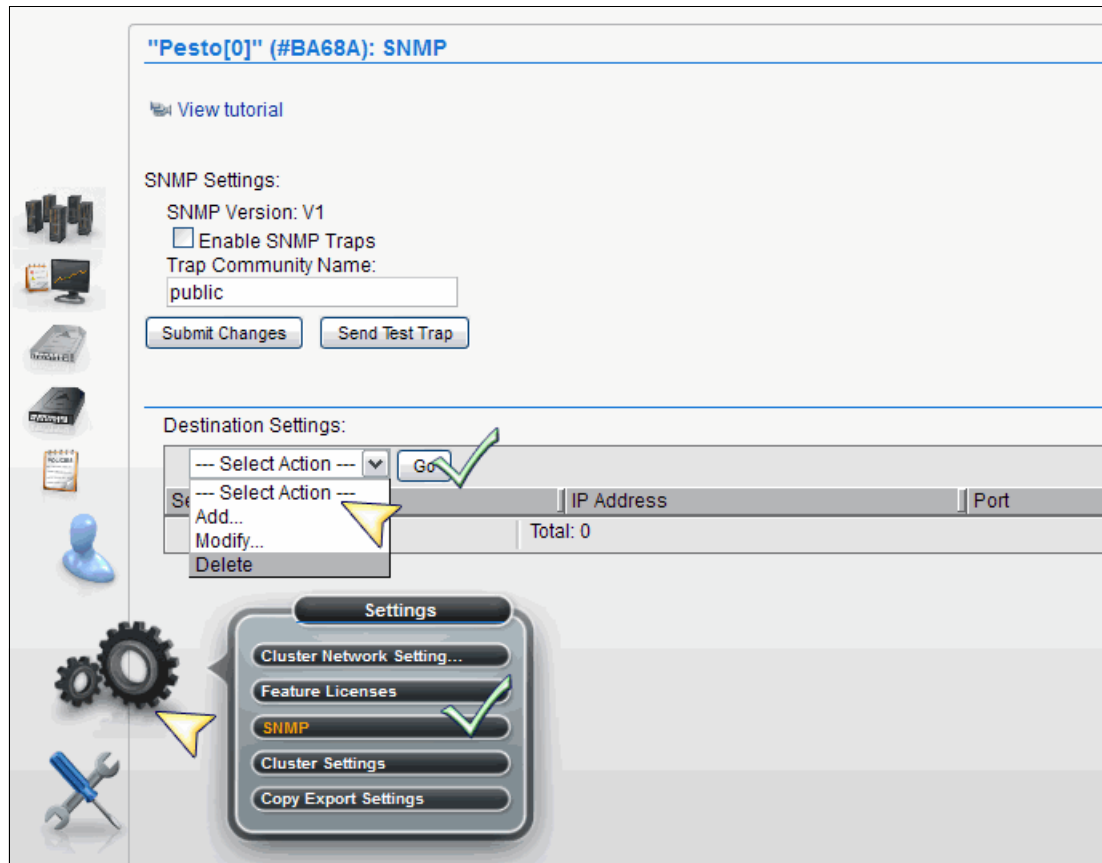


Figure 9-124 SNMP window and options

You use this window to configure SNMP traps that log events, such as logins, configuration changes, status changes (vary on, vary off, or service prep), shutdown, and code updates. SNMP is a networking protocol that enables a TS7700 to gather and transmit automatically information about alerts and status to other entities in the network.

When adding or modifying SNMP destinations, follow this advice:

- ▶ Use IPv4 or IPv6 addresses as destinations rather than a fully qualified domain name (FQDN).
- ▶ Verify that any FQDN used correctly addresses its IP address.

Test only *one* destination at a time when testing SNMP configuration to ensure that FQDN destinations are working properly.

SNMP settings

Use this section to configure global settings that apply to SNMP traps on an entire cluster. The following settings are configurable:

- ▶ **SNMP Version.** The SNMP version. It defines the protocol that is used in sending SNMP requests and is determined by the tool that is used to monitor SNMP traps. Different versions of SNMP traps work with different management applications. The following values are possible:
 - **V1.** The suggested trap version that is compatible with the greatest number of management applications. No alternative version is supported.
 - **Enable SNMP Traps.** A check box that enables or disables SNMP traps on a cluster. A checked box enables SNMP traps on the cluster; a cleared box disables SNMP traps on the cluster. The check box is cleared, by default.
 - **Trap Community Name.** The name that 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. The default value for this field is `public`.
- ▶ **Send Test Trap.** Select this button to send a test SNMP trap to all destinations listed in the Destination Settings table by using the current SNMP trap values. The Enable SNMP Traps check box does not need to be checked to send a test trap. If the SNMP test trap is received successfully and the information is correct, click **Submit Changes**.
- ▶ **Submit Changes.** Select this button to submit changes to any of the global settings, including the fields SNMP Version, Enable SNMP Traps, and Trap Community Name.
- ▶ **Destination Settings.** Use the Destination Settings table to add, modify, or delete a destination for SNMP trap logs. The user can add, modify, or delete a maximum of 16 destination settings at one time.

Note: A user with read-only permissions cannot modify the contents of the Destination Settings table.

The following settings are configurable:

- ▶ **IP address.** The IP address of the SNMP server. This value can take any of the following formats: IPv4, IPv6, a *host name* that is resolved by the system (such as `localhost`), or an FQDN if a domain name server (DNS) is provided. A value in this field is required.

Tip: A valid IPv4 address is 32 bits long, consists of four decimal numbers, each 0 - 255, separated by periods, such as `98.104.120.12`.

A valid IPv6 address is a 128-bit long hexadecimal value that is separated into 16-bit fields by colons, such as `3afa:1910:2535:3:110:e8ef:ef41:91cf`. Leading zeros can be omitted in each field so that `:0003:` can be written as `:3:.` A double colon (`::`) can be used once per address to replace multiple fields of zeros.

For example, `3afa:0:0:0:200:2535:e8ef:91cf` is also `3afa::200:2535:e8ef:91cf`.

- ▶ **Port.** The port to which the SNMP trap logs are sent. This value must be 0 - 65535. A value in this field is required.

Use the Select Action menu on the Destination Settings table to add, modify, or delete an SNMP trap destination. Destinations are changed in the vital product data (VPD) as soon as they are added, modified, or deleted. These updates do not depend on clicking **Submit Changes**.

Note: Any change to SNMP settings is logged on the Tasks window.

Library Port Access Groups window

To view information about library port access groups that are used by the TS7700, use this window. You use *library port access groups* to segment resources and authorization by controlling access to library data ports. Figure 9-125 shows the library port access group link.

Tip: This window is visible only if at least one instance of FC5271 (selective device access control (SDAC)) is installed on all clusters in the grid.

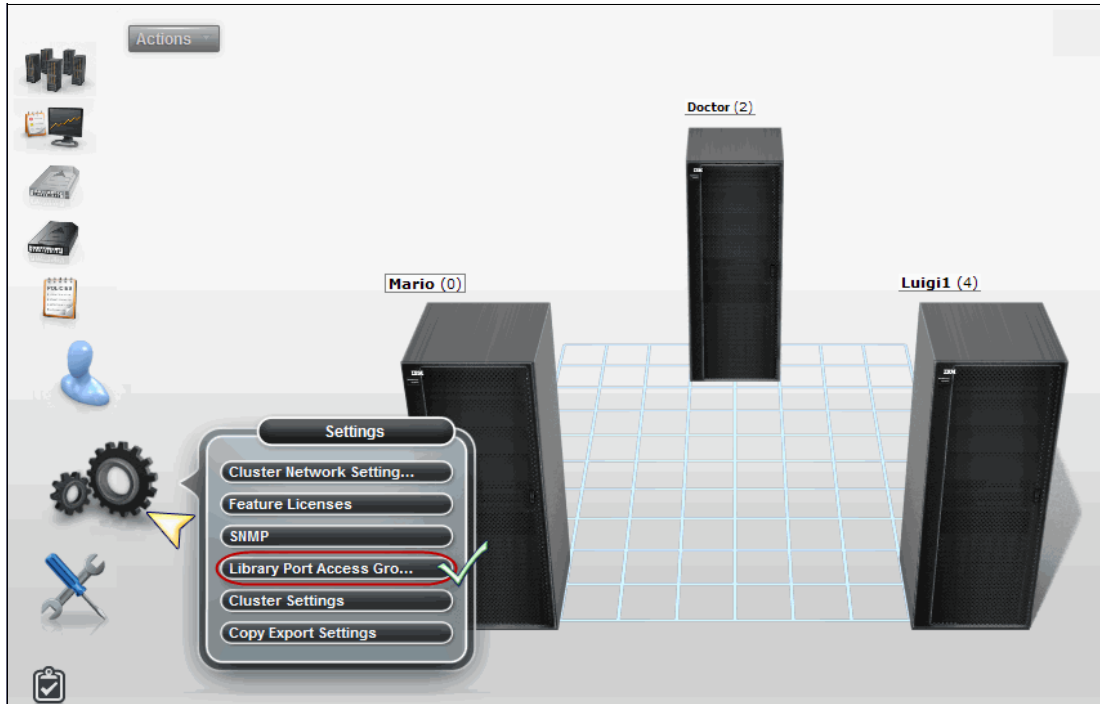


Figure 9-125 Library Port Access Group link

Access Groups table

The Access Groups table displays information about existing library port access groups. Figure 9-126 on page 478 shows the Library Port Access Groups window.

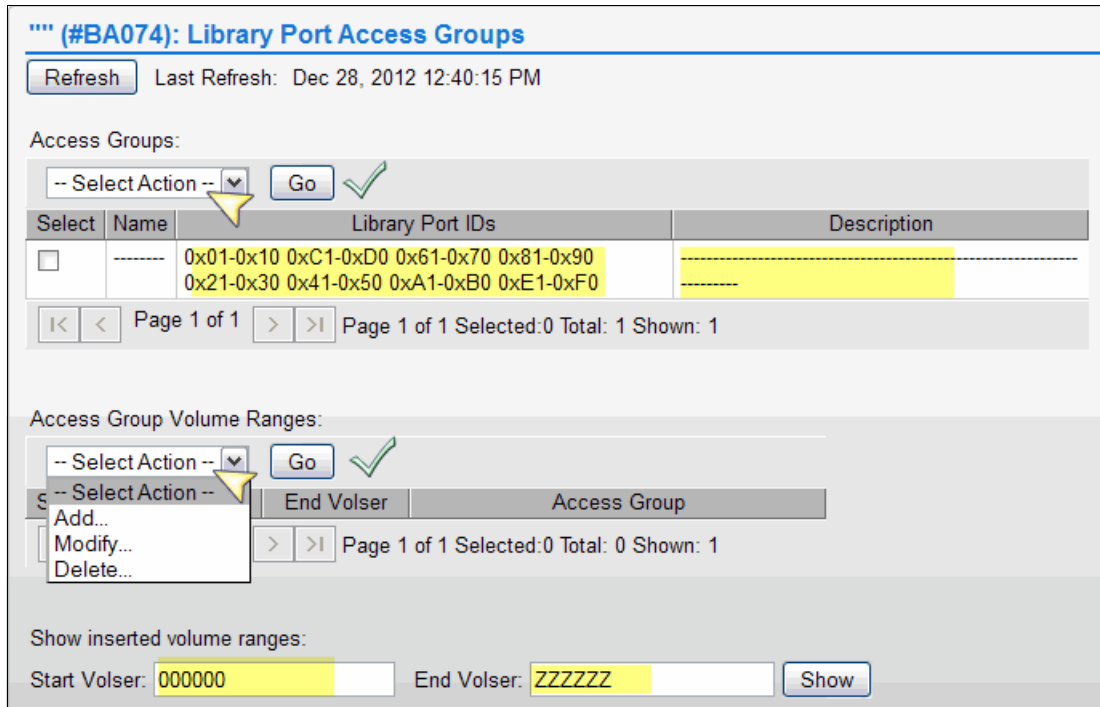


Figure 9-126 Library Port Access Groups window

The user can use the Access Groups table to create a library port access group. Also, the user can modify or delete an existing access group, as shown in Figure 9-127.

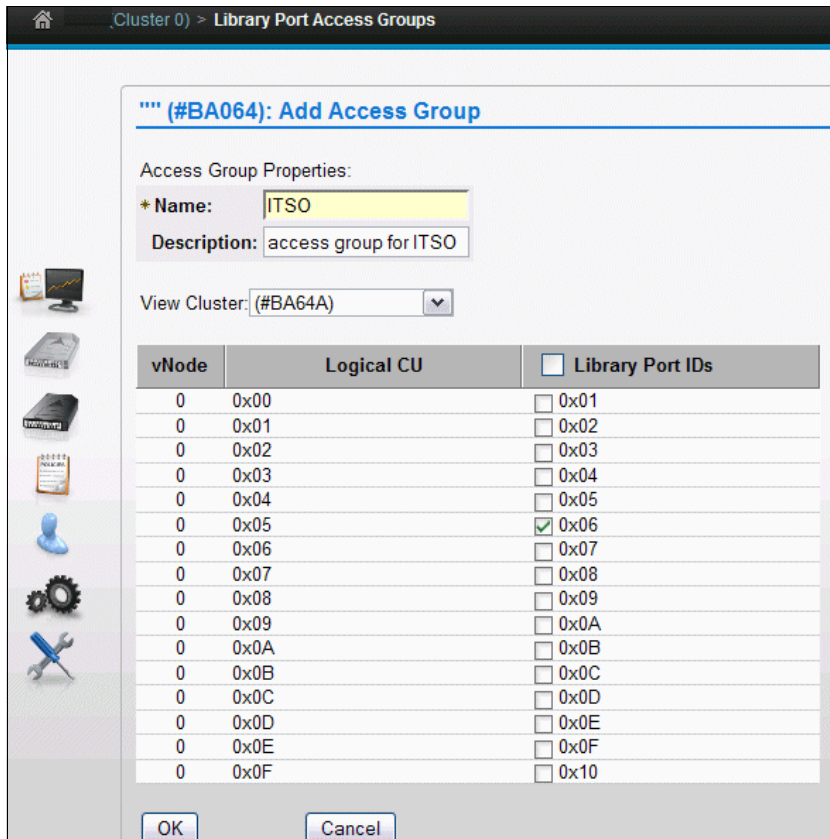


Figure 9-127 Add Access group

The following status information is displayed in the Access Groups table:

- ▶ **Name.** The identifying name of the access group. This name must be unique and cannot be modified after it is created. It must contain 1 - 8 characters, and the first character in this field cannot be a number. Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The default access group is identified by the name “-----”. This group can be modified but cannot be deleted.
- ▶ **Library Port IDs.** A list of Library Port IDs accessible using the defined access group. This field contains a maximum of 750 characters, or 31 Library Port IDs separated by commas or spaces. A range of Library Port IDs is signified by using a hyphen (-). This field can be left blank.

The default access group has a value in this field that is 0x01-0xFF. Initially, all port IDs are shown by default. However, after modification, this field can change to show only the IDs corresponding to the existing vnodes.

Important: VOLSERS that are not found in the SDAC VOLSER range table use this default group to determine access. The user can modify this group to remove any or all default Library Port IDs. However, if all default Library Port ID values are removed, no access is granted to any volumes not in a defined range.

Use the Select Action menu on the Access Groups table to add, modify, or delete a library port access group.

- ▶ **Description.** A description of the access group (a maximum of 70 characters).
- ▶ **Access Groups Volume Ranges.** The Access Groups Volume Ranges table displays VOLSER range information for existing library port access groups. The user can also use the Select Action menu on this table to add, modify, or delete a VOLSER range that is defined by a library port access group.
- ▶ **Start VOLSER.** The first VOLSER in the range that is defined by an access group.
- ▶ **End VOLSER.** The last VOLSER in the range that is defined by an access group.
- ▶ **Access Group.** The identifying name of the access group, which is defined by the Name field in the Access Groups table.

Use the Select Action menu on the Access Group Volume Ranges table to add, modify, or delete a VOLSER range that is associated with a library port access group. The user can show the inserted volume ranges. To view the current list of virtual volume ranges in the TS7700 cluster, enter the start and end VOLSERS and click **Show**.

Note: Access groups and access group ranges are backed up and restored together. For additional information, see “Backup settings” on page 489 and “Restore Settings window” on page 492.

Cluster Settings

You can use the Cluster Settings to view or change settings that determine how a cluster runs copy policy overrides, applies Inhibit Reclaim schedules, uses an EKS, implements write protect mode, and runs backup and restore operations.

For an evaluation of different scenarios and examples where those overrides benefit the overall performance, see 4.2, “Planning for a grid operation” on page 150.

Copy Policy Override

Figure 9-128 shows the Cluster Settings window being used to open the Copy Policy Override window.

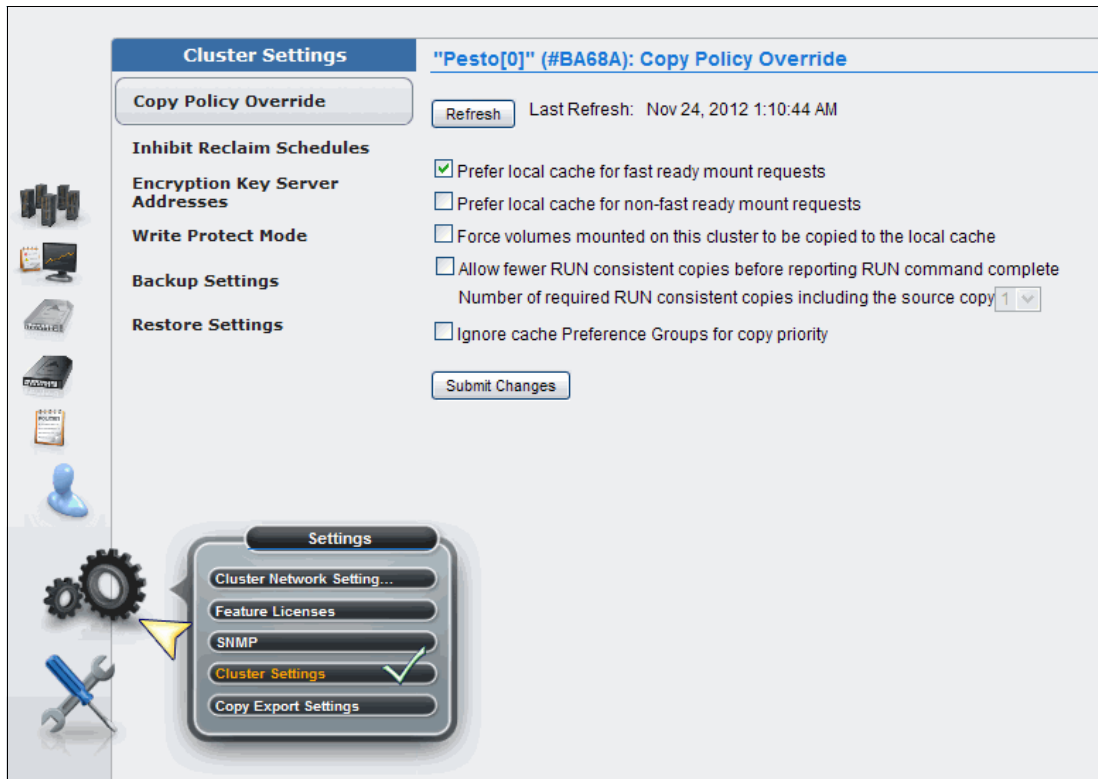


Figure 9-128 Cluster Settings and Copy Policy Override

To override local copy and I/O policies for a TS7700 cluster, use this window.

Reminder: The items on this window can modify the cluster behavior regarding local copy and certain I/O operations. Some **LI REQUEST** commands also can do this action.

For the selected cluster, the user can tailor copy policies to override certain copy or I/O operations. Select the check box next to one or more of the following settings to specify a policy override:

► Prefer local cache for Fast Ready mount requests

When this setting is selected, a scratch (Fast Ready) mount selects the local TVC in the following conditions:

- The Copy Mode field that is defined by the MC for the mount has a value other than No Copy defined for the local cluster.
- The local cluster is not in a degraded state. The following examples are degraded states:
 - Out of cache resources
 - Out of physical scratch

Note: This override can be enabled independently of the status of the copies in the cluster.

- ▶ Prefer local cache for non-Fast Ready mount requests

This override causes the local cluster to satisfy the mount request if both of the following conditions are true:

- The cluster is available.
- The local 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, the default cluster selection criteria applies.

- ▶ Force volumes that are mounted on this cluster to be copied to the local cache

When this setting is selected for a private (non-Fast Ready) mount, a copy operation is performed on the local cluster as part of the mount processing. When this setting is selected for a scratch (Fast Ready) mount, the Copy Consistency Point on the specified MC is overridden for the cluster with a value of Rewind Unload. This override does not change the definition of the MC, but influences the replication policy.

- ▶ Enable fewer RUN consistent copies before reporting RUN command complete

When this setting is selected, the maximum number of RUN copies, including the source, is determined by the value that is entered at **Number of required RUN consistent copies including the source copy**. This value must be consistent before the RUN operation completes. If this option is not selected, the MC definitions are used explicitly. Therefore, the number of RUN copies can be from one to the number of clusters in the grid configuration or the total number of clusters that are configured with a RUN Copy Consistency Point.

- ▶ Ignore cache preference groups for copy priority

If this option is selected, copy operations ignore the cache preference group when determining the priority of volumes that are copied to other clusters.

Note: These settings override the default TS7700 behavior and can be different for every cluster in a grid.

To change any of the settings on this window, complete the following steps:

1. Select or clear the box next to the setting that must be changed. If the user enables **Enable fewer RUN consistent copies before reporting RUN command complete**, the user can alter the value for **Number of required RUN consistent copies including the source copy**.
2. Click **Submit Changes**.

Inhibit Reclaim Schedules window

To add, modify, or delete Inhibit Reclaim schedules that are used to postpone tape reclamation in a TS7740 or TS7720T cluster, use this window.

This window is visible but disabled on the TS7700 MI if the grid possesses a physical library, but the selected cluster does not. The following message is displayed:

The cluster is not attached to a physical tape library.

Tip: This window is not visible on the TS7700 MI if the grid does not possess a physical library.

Reclamation can improve tape usage by consolidating data on some physical volumes, but it uses 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. Figure 9-129 shows an Inhibit Reclamation Schedules window.

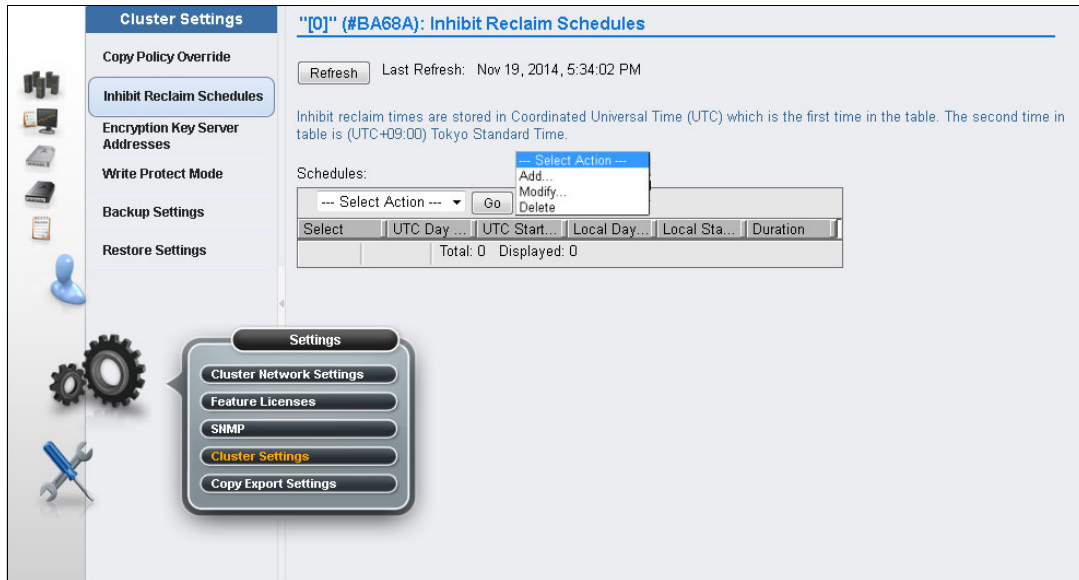


Figure 9-129 Inhibit Reclaim Schedules window

The following fields on this window are described:

- ▶ **Schedules.** The Schedules table displays the list of Inhibit Reclaim schedules that are 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 (Coordinated Universal Time) and then in local time. The status information is displayed in the Schedules table:
 - **Coordinated Universal Time Day of Week.** The Coordinated Universal Time day of the week on which the reclamation will be inhibited. The following values are possible:
 - **Every Day.** Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
 - **Coordinated Universal Time Start Time.** The Coordinated Universal Time time in hours (H) and minutes (M) at which reclamation is 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 includes a time chooser clock icon. The user can enter hours and minutes manually by using 24-hour time designations, or can use the time chooser to select a start time based on a 12-hour (AM/PM) clock.
 - **Local Day of Week.** The day of the week in local time on which the reclamation is inhibited. The day that is recorded reflects the time zone in which the browser is. The following values are possible:
 - **Every Day.** Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
 - **Local Start Time.** The local time in hours (H) and minutes (M) at which reclamation is inhibited. The values in this field must take the form HH:MM. The time that is recorded reflects the time zone in which the web browser is. Possible values for this field include 00:00 - 23:59. The Start Time field includes a time chooser clock icon. The user can enter hours and minutes manually by using 24-hour time designations, or can use the time chooser to select a start time based on a 12-hour (AM/PM) clock.

- **Duration.** The number of days (D) hours (H) and minutes (M) that the reclamation is 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 minutes 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 minutes.

Note: Inhibit Reclaim schedules cannot overlap.

Use the menu on the Schedules table to add a new Inhibit Reclaim schedule or to modify or delete an existing schedule. Figure 9-130 shows the Add Inhibit Reclaim Schedule window.

Figure 9-130 Add Inhibit Reclaim Schedule window

To modify an Inhibit Reclaim schedule, follow these steps:

1. From the Inhibit Reclaim Schedules window, go to the Schedules table.
2. Select the radio button next to the Inhibit Reclaim schedule to be modified.
3. Select **Modify** from the Select Action menu.
4. Click **Go** to open the Modify Inhibit Reclaim Schedule window.

The values are the same as for the Add Inhibit Reclaim Schedule, which is listed in Figure 9-130.

To delete an Inhibit Reclaim schedule, follow these steps:

1. From the Inhibit Reclaim Schedules window, go to the Schedules table.
2. Select the radio button next to the Inhibit Reclaim schedule that must be deleted.
3. Select **Delete** from the **Select Action** menu.
4. Click **Go** to open the Confirm Delete Inhibit Reclaim Schedule window.
5. Click **OK** to delete the Inhibit Reclaim schedule and return to the Inhibit Reclaim Schedules window, or click **Cancel** to abandon the delete operation and return to the Inhibit Reclaim Schedules window.

Note: Plan the Inhibit Reclaim schedules carefully. Running the reclaims during peak times can affect production, and not having enough reclaim schedules influences the media consumption.

Encryption Key Server Addresses window

To set the EKS addresses in the TS7700 cluster, use this window.

In the TS7700 subsystem, user data can be encrypted on tape cartridges by the encryption capable tape drives that are available to the TS7700 tape-attached clusters. Also, data can be encrypted by the full data encryption (FDE) DDMs in TS7700 TVC cache.

With R4.0, the TVC cache encryption can be configured either for local or external encryption key management with 3956-CC9, 3956-CS9, and 3956-CSA cache types. Tape encryption uses an out-of-band external key management. For more information, see Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Note: Only IBM Security Key Lifecycle Manager or IBM Security Key Lifecycle Manager for z/OS are supported for external disk encryption or TS1140 and TS1150 tape drives. The settings for Encryption Server are shared for both tape and external disk encryption.

For both tape and disk encryption, an EKS is required in the network that is accessible by the TS7700 cluster. For more information, see Chapter 4, “Preinstallation planning and sizing” on page 125 and Chapter 7, “Hardware configurations and upgrade considerations” on page 227.

There is a tutorial available on MI that shows the properties of the EKS. To watch it, click the **View tutorial** link in the MI window. Figure 9-131 shows the Encryption Key Server Addresses setup window.

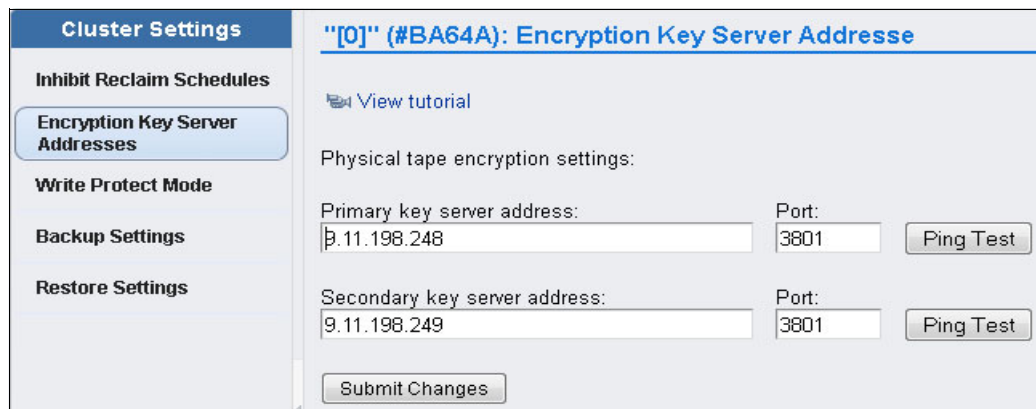


Figure 9-131 Encryption Key Server Addresses window

If the cluster has the feature code for disk or tape encryption enabled, this window is visible on the TS7700 MI. Otherwise, it can be visible but disabled on the TS7700 MI (when the grid possesses a physical library, but the selected cluster does not) or not be visible (no clusters in the grid possess either a physical library or a cache encryption feature enabled).

Note: IP addresses 10.x.x.x, 192.168.x.x, and 172.16.x.x shall not be routed through the Internet, and therefore must be confined to a local customer network. If TS7700 and key manager reside in different networks, separated by the Internet, IP addresses outside of those listed ranges must be used.

The EKS assists encryption-enabled tape drives in generating, protecting, storing, and maintaining EKs that are used to encrypt information that is being written to, and to decrypt information that is being read from, tape media (tape and cartridge formats). Also, EKS manages the EK for the cache disk subsystem, with the External key management disk encryption feature installed, which removes the responsibility of managing the key from the 3957 V07 or VEB engine, and from the disk subsystem controllers. To read more about data encryption with TS7700, see Chapter 2, “Architecture, components, and functional characteristics” on page 15.

The following settings are used to configure the TS7700 connection to an EKS.

► Primary key server address

The key server name or IP address that is primarily used to access the EKS. This address can be a fully qualified host name or an IP address in IPv4 or IPv6 format. This field is not required if the user does not want to connect to an EKS.

Tip: A valid IPv4 address is 32 bits long, consists of four decimal numbers, each 0 - 255, separated by periods, such as 98.104.120.12.

A valid IPv6 address is a 128-bit long hexadecimal value separated into 16-bit fields by colons, such as 3afa:1910:2535:3:110:e8ef:ef41:91cf. Leading zeros can be omitted in each field so that :0003: can be written as :3:. A double colon (::) can be used once per address to replace multiple fields of zeros. For example, 3afa:0:0:0:200:2535:e8ef:91cf can be written as: 3afa::200:2535:e8ef:91cf.

A fully qualified host name is a domain name that uniquely and absolutely names a computer. It consists of the host name and the domain name. The domain name is one or more domain labels that place the computer in the DNS naming hierarchy. The host name and the domain name labels are separated by periods and the total length of the *host name* cannot exceed 255 characters.

► Primary key server port

The port number of the primary key server. Valid values are any whole number 0 - 65535; the default value is 3801. This field is only required if a primary key address is used.

► Secondary key server address

The key server name or IP address that is used to access the EKS when the primary key server is unavailable. This address can be a fully qualified host name or an IP address in IPv4 or IPv6 format. This field is not required if the user does not want to connect to an EKS. See the primary key server address description for IPv4, IPv6, and fully qualified host name value parameters.

► Secondary key server port

The port number of the secondary key server. Valid values are any whole number 0 - 65535; the default value is 3801. This field is only required if a secondary key address is used.

► Using the Ping Test

Use the Ping Test buttons to check the cluster network connection to a key server after changing a cluster's address or port. If the user changes a key server address or port and do not submit the change before using the Ping Test button, the user receives the following message:

To perform a ping test you must first submit your address and/or port changes.

After the ping test starts, one of the following two messages will occur:

- The ping test against the address “<address>” on port “<port>” was successful.
- The ping test against the address “<address>” on port “<port>” from “<cluster>” has failed. The error returned was: <error text>.

Click **Submit Changes** to save changes to any of these settings.

Tip: The user can back up these settings as part of the `ts7700_cluster<cluster ID>.xmi` file and restore them for later use or use with another cluster. If a key server address is empty at the time that the backup is run, when it is restored, the port settings are the same as the default values.

Write Protect Mode window

To view Write Protect Mode settings in a TS7700 cluster, use this window. This window also is displayed if the Write Protect Mode is enabled due to FlashCopy being enabled (the Current State field shows Write protect for FlashCopy enabled).

With FlashCopy in progress, no modifications are allowed on the Write Protect Mode window until the FlashCopy testing is completed. When Write Protect Mode is enabled on a cluster, host commands fail if they are sent to virtual devices in that cluster and attempt to modify a volume's data or attributes.

Note: FlashCopy is enabled from LI REQ (Library Request Host Console) command.

Meanwhile, host commands that are sent to virtual 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 client-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 the mode is set, the mode is retained through intentional and unintentional outages and can be disabled only through the same MI window that is used to enable the function. When a cluster within a grid configuration has Write Protect Mode enabled, standard grid functions, such as virtual volume replication and virtual volume ownership transfer, are unaffected.

Virtual volume categories can be excluded from Write Protect Mode. Up to 32 categories can be identified and set to include or exclude from Write Protect Mode by using the Category Write Protect Properties table. Additionally, write-protected volumes in any scratch (Fast Ready) category can be mounted as private volumes if the Ignore Fast Ready characteristics of write-protected categories check box is selected.

Figure 9-132 shows the Write Protect Mode window.

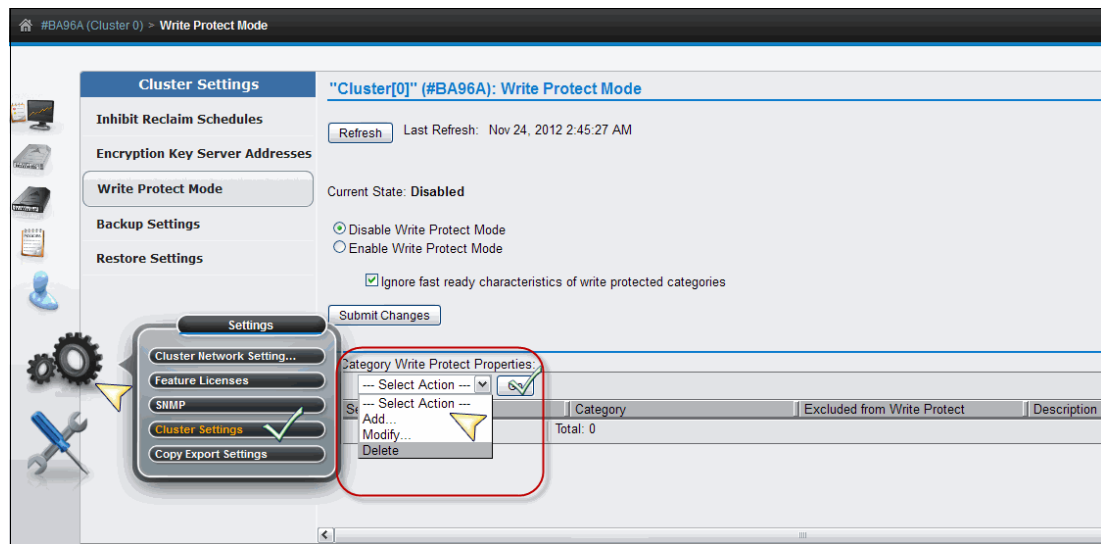


Figure 9-132 Write Protect Mode window

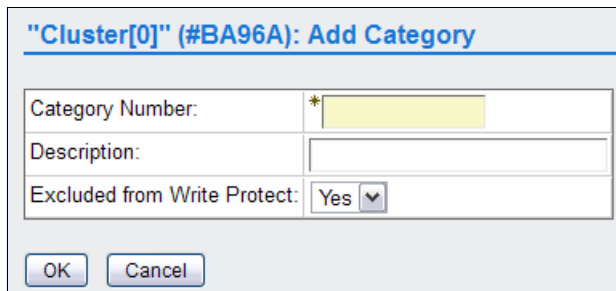
The page that is shown in Figure 9-132 has the following information:

- ▶ **Current State.** The status of Write Protect Mode on the active cluster. The following values are possible:
 - **Disabled.** Write Protect Mode is disabled on the cluster. No Write Protect settings are in effect for the cluster.
 - **Enabled.** Write Protect Mode is enabled on the cluster. Any attempt by an attached host to modify a volume or its attributes fails, subject to any defined category exclusions.
- ▶ **Disable Write Protect Mode.** Select this radio button to disable Write Protect Mode on the active cluster. If the user select this option, no volumes on the cluster are write-protected; all Write Protect settings are disabled.
- ▶ **Enable Write Protect Mode.** Select this radio button to enable Write Protect Mode on the active cluster. This option prevents hosts that are attached to this cluster from modifying volumes or their attributes.
- ▶ **Ignore Fast Ready characteristics of write protected categories.** If this check box is selected, write-protected volumes that have been returned to a scratch or a Fast Ready category continue to be viewed as private volumes. This enables a disaster recovery test host to mount production volumes as private volumes even though the production environment has since returned them to scratch.

Peer clusters, such as production clusters, continue to view these volumes as scratch volumes. This setting does not override the scratch (Fast Ready) characteristics of the excluded categories.
- ▶ **Category Write Protect Properties.** Use the Category Write Protect Properties table to add, modify, or delete categories to be selectively excluded from Write Protect Mode.

Disaster recovery test hosts or locally connected production partitions can continue to read and write to local volumes if their volume categories are excluded from write protect. These hosts must use a set of categories different from those primary production categories that are write protected.

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. Figure 9-133 shows the Add Category window.



The screenshot shows a dialog box titled '"Cluster[0]" (#BA96A): Add Category'. It contains three input fields: 'Category Number:' with an asterisk and a yellow highlight, 'Description:', and 'Excluded from Write Protect:' with a dropdown menu set to 'Yes'. There are 'OK' and 'Cancel' buttons at the bottom.

Figure 9-133 Add Category window

The following category fields are displayed in the Category Write Protect Properties table:

- ▶ **Category Number:** The identifier for a defined category. This is an alphanumeric hexadecimal value between 0x0001 and 0xFEFF (0x0000 and 0xFFxx cannot be used). Values that are entered do not include the 0x prefix, although this prefix is displayed on the Cluster Summary window. Values that are entered are padded up to four places. Letters that are used in the category value must be capitalized.
- ▶ **Excluded from Write Protect:** Whether the category is excluded from Write Protect Mode. The following values are possible:
 - **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. This description must contain 0 - 63 Unicode characters.

Use the menu on the Category Write Protect Properties table to add a category, or modify or delete an existing category.

The user must click **Submit Changes** to save any changes that were made to the Write Protect Mode settings.

Backup settings

To back up the settings from a TS7700 cluster, use this window. Figure 9-134 shows an example of a Backup Settings window.

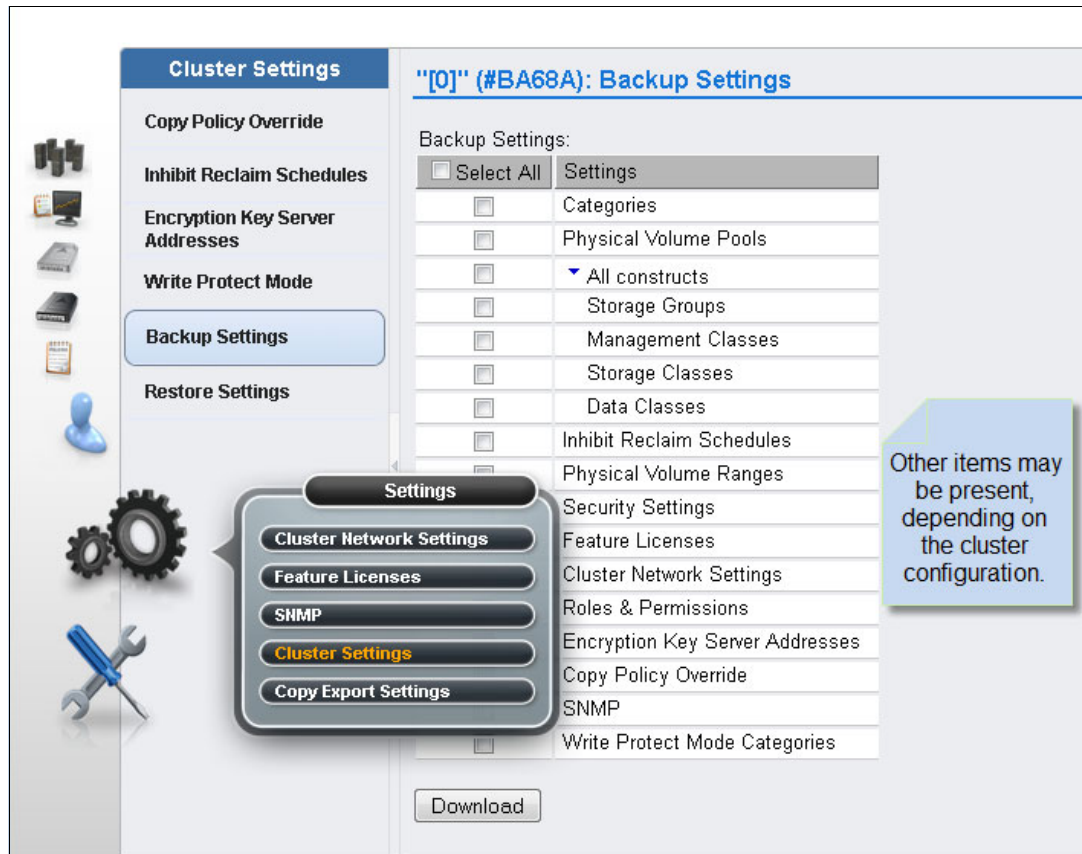


Figure 9-134 Backup Settings window

Important: Backup and restore functions are not supported between clusters operating at different code levels. Only clusters operating at the same code level as the accessing cluster (the one addressed by the web browser) can be selected for Backup or Restore. Clusters operating different code levels are visible, but options are disabled.

The Backup Settings table lists the cluster settings that are available for backup:

- ▶ **Categories:** Select this check box to back up scratch (Fast Ready) categories that are used to group virtual volumes.
- ▶ **Physical Volume Pools:** Select this check box to back up physical volume pool definitions.

Note: If the cluster does not possess a physical library, physical volume pools are not available.

- ▶ **All Constructs:** Select this check box to select all of the following constructs for backup. Alternatively, the user can select a specific construct by checking the box for the option:
 - **Storage Groups:** Select this check box to back up defined SGs.
 - **Management Classes:** Select this check box to back up defined MCs.

- **Storage Classes:** Select this check box to back up defined SCs.
- **Data Classes:** Select this check box to back up defined DCs.
- ▶ **Inhibit Reclaim Schedule:** Select this check box to back up the Inhibit Reclaim schedules that are used to postpone tape reclamation.

Note: If the cluster does not possess a physical library, the Inhibit Reclaim Schedules option is not available.

- ▶ **Library Port Access Groups:** Select this check box to back up defined library port access groups.

Note: This setting is only available if all clusters in the grid are operating with Licensed Internal Code levels of 8.20.0.xx or higher and the SDAC feature is installed.

Library port access groups and access group ranges are backed up and restored together.

- ▶ **Physical Volume Ranges:** Select this check box to back up defined physical volume ranges. If the cluster does not possess a physical library, physical volume ranges are not available.
- ▶ **Security Settings:** Select this check box to back up defined security settings:
 - **Session Timeout**
 - **Account Expiration**
 - **Account Lock**
 - **Encryption Key Server Addresses**
 - **Primary key server address**
 - **Primary key server port**
 - **Secondary key server address**
 - **Secondary key server port**
- ▶ **Cluster Network Settings:** Select this box to back up the defined cluster network settings.
- ▶ **Roles & Permissions:** Select this check box to back up defined custom user roles.

Important: A restore operation after a backup of cluster settings does *not* restore or otherwise modify any user, role, or password settings defined by a security policy.

- ▶ **Feature Licenses:** Select this check box to back up the settings for currently activated feature licenses.

Note: The user can back up these settings as part of the `ts7700_cluster<cluster ID>.xmi` file and restore them for later use on the same cluster. However, the user cannot restore feature license settings to a cluster different from the cluster that created the `ts7700_cluster<cluster ID>.xmi` backup file.

The following feature license information is available for backup:

- **Feature Code.** The feature code number of the installed feature.
- **Feature Description.** A description of the feature that was installed by the feature license.
- **License Key.** The 32-character license key for the feature.
- **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 that the feature license was activated.
- **Expires.** The expiration status of the feature license. The following values are possible:
 - **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.
- ▶ **Encryption Key Server Addresses:** Select this check box to back up defined EKS addresses:
 - **Primary key server address**
 - **Primary key server port**
 - **Secondary key server address**
 - **Secondary key server port**

Important: A restore operation after a backup of cluster settings does *not* restore or otherwise modify any user, role, or password settings defined by a security policy.

- ▶ **Copy Policy Override:** Select this check box to back up the settings to override local copy and I/O policies.
- ▶ **SNMP:** Select this check box to back up the settings for SNMP.
- ▶ **Write Protect Mode Categories:** Select this check box to back up the settings for write protect mode categories.

To back up cluster settings, click a check box next to any of the previous settings and then click **Download**. A window opens to show that the backup is in progress.

Important: If the user navigates away from this window while the backup is in progress, the backup operation is stopped and the operation must be restarted.

When the backup operation is complete, the backup file `ts7700_cluster<cluster ID>.xmi` is created. This file is an XML Meta Interchange file. The user is prompted to open the backup file or save it to a directory. Save the file. When prompted to open or save the file to a directory, save the file without changing the `.xmi` file extension or the file contents.

Any changes to the file contents or extension can cause the restore operation to fail. The user can modify the file name before saving it, if the user wants to retain this backup file after subsequent backup operations. *If the user chooses to open the file, do not use Microsoft Excel to view or save it.* Microsoft Excel changes the encoding of an XML Meta Interchange file, and the changed file is corrupted when used during a restore operation.

The following settings are not available for backup or recovery:

- ▶ User accounts
- ▶ Security policies
- ▶ Grid identification policies
- ▶ Cluster identification policies
- ▶ Grid communication encryption (IPSec)
- ▶ SSL certificates

Record these settings in a safe place and recover them manually if necessary.

Restore Settings window

To restore the settings from a TS7700 cluster to a recovered or new cluster, use this window.

Note: Backup and restore functions are not supported between clusters operating at different code levels. Only clusters operating at the same code level as the current cluster can be selected from the Current Cluster Selected graphic. Clusters operating at different code levels are visible, but not available, in the graphic.

Figure 9-135 shows the Restore Settings window.

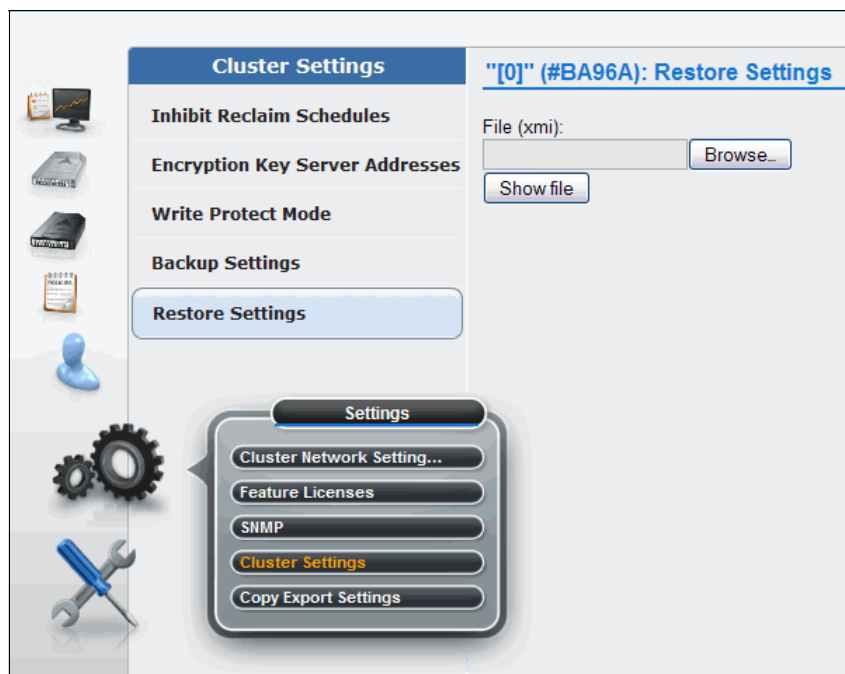


Figure 9-135 Restore Settings window

Follow these steps to restore cluster settings:

1. On the Restore Settings window, click **Browse** to open the File Upload window.
2. Go to the backup file used to restore the cluster settings. This file has an .xmi extension.
3. Add the file name to the File name field.
4. Click **Open** or press Enter from the keyboard.
5. Click **Show file** to review the cluster settings that are 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 are shown:

- ▶ **Categories:** Select this check box to restore scratch (Fast Ready) categories that are used to group virtual volumes.
- ▶ **Physical Volume Pools:** Select this check box to restore physical volume pool definitions.

Important: If the backup file was created by a cluster that did not possess a physical library, physical volume pool settings are reset to default.

- ▶ **All Constructs:** Select this check box to restore all of the displayed constructs.
- ▶ **Storage Groups:** Select this check box to restore defined SGs.

- ▶ **Management Classes:** Select this check box to restore defined MCs.

MC settings are related to the number and order of clusters in a grid. Take special care when restoring this setting. If an MC is restored to a grid that has more clusters than the grid had when the backup was run, the copy policy for the new cluster or clusters are set as No Copy.

If an MC is restored to a grid that has fewer clusters than the grid had when the backup was run, the copy policy for the now-nonexistent clusters is changed to No Copy. The copy policy for the first cluster is changed to RUN to ensure that one copy exists in the cluster.

If cluster IDs in the grid differ from cluster IDs present in the restore file, MC copy policies on the cluster are overwritten with those from the restore file. MC copy policies can be modified after the restore operation completes.

If the backup file was created by a cluster that did not define one or more scratch mount candidates, the default scratch mount process is restored. The default scratch mount process is a random selection routine that includes all available clusters. MC scratch mount settings can be modified after the restore operation completes.

- ▶ **Storage Classes:** Select this check box to restore defined SCs.
- ▶ **Data Classes:** Select this check box to restore defined DCs.

If this setting is selected and the cluster does not support logical Write Once Read Many (LWORM), the Logical WORM setting is disabled for all DCs on the cluster.

- ▶ **Inhibit Reclaim Schedule:** Select this check box to restore Inhibit Reclaim schedules that are used to postpone tape reclamation.

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.

Note: If the backup file was created by a cluster that did not possess a physical library, the Inhibit Reclaim schedules settings are reset to default.

- ▶ **Library Port Access Groups:** Select this check box to restore defined library port access groups.

This setting is only available if all clusters in the grid are operating with Licensed Internal Code levels of 8.20.0.xx or higher.

Library port access groups and access group ranges are backed up and restored together.

- ▶ **Physical Volume Ranges:** Select this check box to restore defined physical volume ranges.

If the backup file was created by a cluster that did not possess a physical library, physical volume range settings are reset to default.

- ▶ **Roles & Permissions:** Select this check box to restore defined custom user roles.

A restore operation after a backup of cluster settings does *not* restore or otherwise modify any user, role, or password settings defined by a security policy.

- ▶ **Security Settings:** Select this check box to restore defined security settings, for example:
 - **Session Timeout**
 - **Account Expiration**
 - **Account Lock**

- ▶ **Encryption Key Server Addresses:** Select this check box to restore defined EKS addresses. If a key server address is empty at the time that the backup is performed, when restored, the port settings are the same as the default values. The following EKS address settings can be restored:
 - **Primary key server address:** The key server name or IP address that is primarily used to access the EKS.
 - **Primary key server port:** The port number of the primary key server.
 - **Secondary key server address:** The key server name or IP address that is used to access the EKS when the primary key server is unavailable.
 - **Secondary key server port:** The port number of the secondary key server.
- ▶ **Cluster Network Settings:** Select this check box to restore the defined cluster network settings.

Important: Changes to network settings affect access to the TS7700 MI. When these settings are restored, routers that access the TS7700 MI are reset. No TS7700 grid communications or jobs are affected, but any current users are required to log back on to the TS7700 MI by using the new IP address.

- ▶ **Feature Licenses:** Select this check box to restore the settings for currently activated feature licenses. When the backup settings are restored, new settings are added but no settings are deleted. After restoring feature license settings on a cluster, log out and then log in to refresh the system.

Note: The user cannot restore feature license settings to a cluster that is different from the cluster that created the `ts7700_cluster<cluster ID>.xmi` backup file.

The following feature license information is available for backup:

- **Feature Code.** The feature code number of the installed feature.
- **Feature Description.** A description of the feature that was installed by the feature license.
- **License Key.** The 32-character license key for the feature.
- **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. The following values are possible:
 - **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.

After selecting **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 that the backup occurred.

Select the box next to each setting to be restored. Click **Restore**.

Note: The restore operation overwrites existing settings on the cluster.

A warning window opens and prompts you to confirm the decision to restore settings. Click **OK** to restore settings or **Cancel** to cancel the restore operation.

The Confirm Restore Settings window opens.

Important: If the user navigates away from this window while the restore is in progress, the restore operation is stopped and the operation must be restarted.

The restore cluster settings operation can take 5 minutes or longer. During this step, the MI is communicating the commands to update settings. If the user navigates away from this window, the restore settings operation is canceled.

Copy Export Settings window

Use this window to change the maximum number of physical volumes that can be exported by the TS7700. Figure 9-136 shows the Copy Export Settings window.



Figure 9-136 Copy Export Settings window

The Number of physical volumes to export is the maximum number of physical volumes that can be exported. This value is an integer 1 - 10,000. The default value is 2000. To change the number of physical volumes to export, enter an integer in the described field and click **Submit**.

Note: The user can modify this field even if a Copy Export operation is running, but the changed value does not take effect until the next Copy Export operation starts.

For more information about Copy Export, see Chapter 12, “Copy Export” on page 757.

9.2.11 The Service icon

The following sections present information about running service operations and troubleshooting problems for the TS7700. Figure 9-137 shows the Service icon options for a stand-alone TS7720T cluster compared to a grid of TS7720 clusters.

Ownership Takeover Mode shows only when a cluster is member of a grid, while *Copy Export Recover* and *Copy Export Recovery Status* options appear for a single TS7700T configuration (that is connected to a physical library).

Note: *Copy Export Recover* or *Copy Export Recover Status* are only available in a single cluster configuration for TS7740 or TS7720T.



Figure 9-137 Service Icon options

To enable or disable *Ownership Takeover Mode* for a failed cluster in a TS7700, use the window that is shown in Figure 9-138 on page 497. The Ownership Takeover Mode must be started from any surviving cluster in the grid when a cluster becomes inaccessible. Enabling Ownership Takeover Mode from the failed cluster is not possible.

Note: Keep the IP addresses for the clusters in the configuration available for use in the event of a failure of a cluster. Thus, the MI can be accessed from a surviving cluster to initiate the ownership takeover actions.

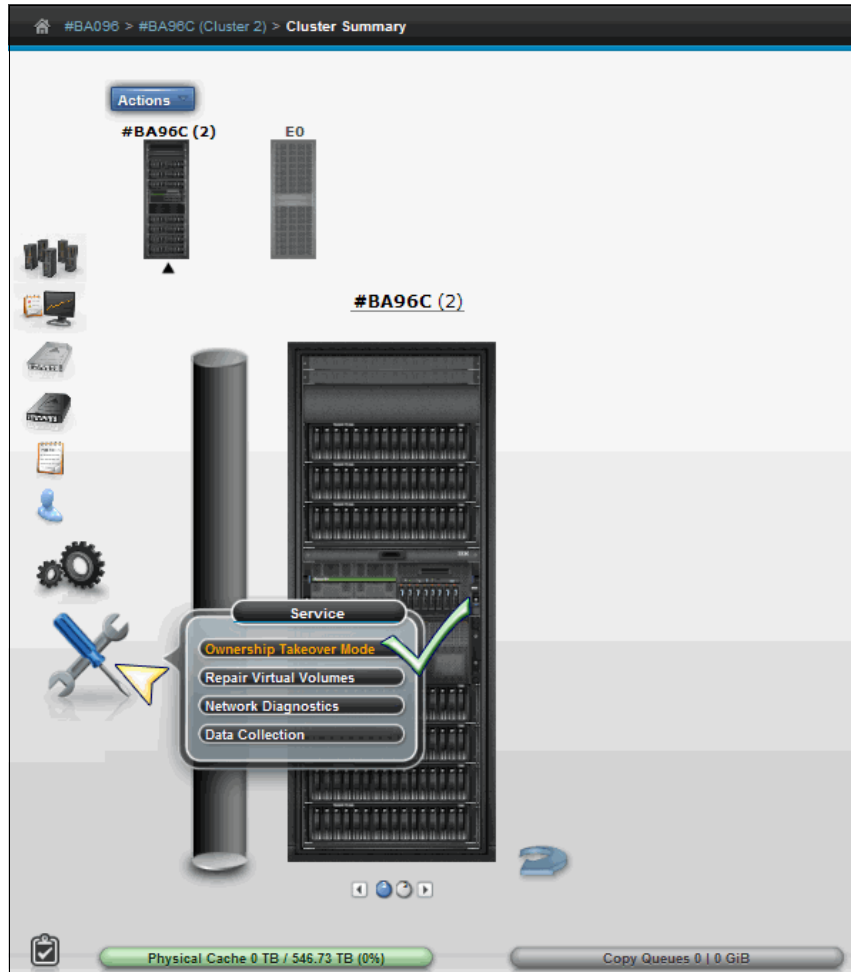


Figure 9-138 MI window to set the Ownership Takeover Mode

When a cluster enters a failed state, enabling Ownership Takeover Mode enables 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 links between clusters fail, the normal means of transferring ownership is not available.

Read/write or read-only takeover should not be enabled when only the communication path between the clusters has failed, and the isolated cluster remains operational. The integrity of logical volumes in the grid can be compromised if a takeover mode is enabled for a cluster that is only isolated from the rest of the grid (not failed) and there is active host access to it.

A takeover decision should be made only for a cluster that is indeed no longer operational.

AOTM, when available and configured, verifies the real status of the non-responsive cluster by using an alternative communication path other than the usual connection between clusters. AOTM uses the TSSC associated with each cluster to determine whether the cluster is alive or failed, enabling the ownership takeover only in case the unresponsive cluster has indeed failed. If the cluster is still alive, AOTM does not initiate a takeover, and the decision is up to the human operator.

If one or more clusters become isolated from one or more peers, those volumes that are owned by the inaccessible peers cannot be mounted without first enabling an ownership takeover mode. Volumes that are owned by one of the accessible clusters can be successfully mounted and modified. For those mounts that cannot obtain ownership from the inaccessible peers, the operation fails. In z/OS, the failure for this error code is not permanent, which makes it possible for the user to enable ownership takeover and retry the operation.

When Read Only takeover mode is enabled, those volumes requiring takeover are read-only, and fail any operation that attempts to modify the volume attributes or data. Read/write takeover enables full read/write access of attributes and data.

If an ownership takeover mode is enabled when only a WAN/LAN failure is present, read/write takeover should not be used because it can compromise the integrity of the volumes that are accessed by both isolated groups of clusters.

Read-only takeover mode should be used instead.

If full read/write access is required, one of the isolated groups should be taken offline to prevent any use case where both groups attempt to modify the same volume. Figure 9-139 shows the Ownership Takeover Mode window when navigating from the window that is shown in Figure 9-138 on page 497.

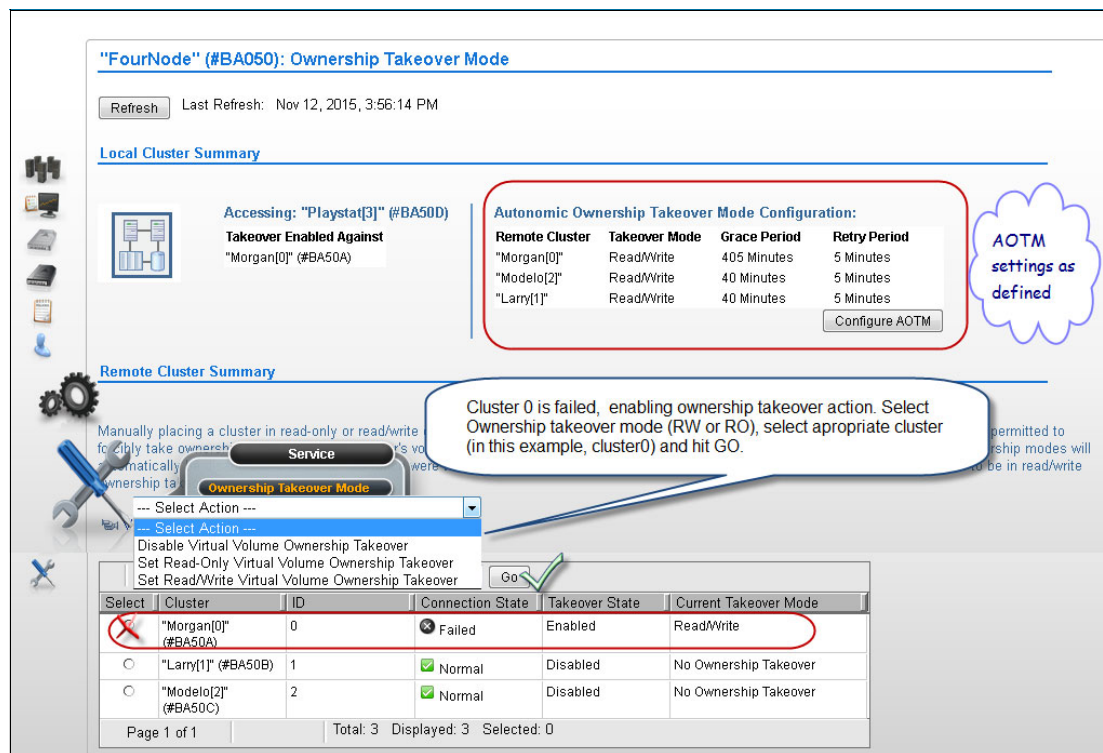


Figure 9-139 Ownership Takeover Mode

Figure 9-139 shows the local cluster summary, the list of available clusters in the grid, the connection state between local (accessing) cluster and its peers. It also shows the current takeover state for the peer clusters (if enabled or disabled by the accessing cluster) and the current takeover mode.

In the example that is shown in Figure 9-139 on page 498, cluster zero has a failed connection status. A mount request by the host for a volume that is owned by cluster zero that is issued to any of the peer clusters causes an operator intervention, reporting that the ownership request for that volume was not granted by cluster zero. The decision to takeover the ownership must be made (either by human operator or AOTM).

Here is the manual procedure to start an ownership takeover against a failed cluster:

1. Authenticate the MI to the surviving cluster that has the takeover intervention.
2. Go to the Ownership Takeover Mode by clicking **Service** → **Ownership Takeover Mode**.
3. Select the failed cluster (the one to be taken over).
4. In the Select Action box, select the appropriate Ownership takeover mode (RW or RO).
5. Click **Go**, and retry the host operation that failed.

Figure 9-139 on page 498 shows that AOTM was previously configured in this grid (for Read/Write, with a grace period of 405 minutes for Cluster 0). In this case, automatic ownership takeover takes place at the end of that period (6 hours and 45 minutes). Human operation can override that setting manually by taking the actions that are described previously, or by changing the AOTM settings to more suitable values. You can use the **Configure AOTM** button to configure the values that are displayed in the previous AOTM Configuration table.

Important: An IBM SSR must configure the TSSC IP addresses for each cluster in the grid *before* AOTM can be enabled and configured for any cluster in the grid.

Table 9-12 compares the operation of read/write and read-only ownership takeover modes.

Table 9-12 Comparing read/write and read-only ownership takeover modes

Read/write ownership takeover mode	Read-only ownership takeover mode
<p>Operational clusters in the grid can run these tasks:</p> <ul style="list-style-type: none"> ▶ Perform read and write operations on the virtual volumes that are owned by the failed cluster. ▶ Change virtual volumes that are owned by the failed cluster to private or SCRATCH status. 	<p>Operational clusters in the grid can run these tasks:</p> <ul style="list-style-type: none"> ▶ Perform read operations on the virtual volumes that are owned by the failed cluster. <p>Operational clusters in the grid cannot run these tasks:</p> <ul style="list-style-type: none"> ▶ Change the status of a volume to private or scratch. ▶ Perform read and write operations on the virtual volumes that are owned by the failed cluster.
<p>A consistent copy of the virtual volume must be available on the grid or the virtual volume must exist in a scratch category. If no cluster failure occurred (grid links down) and the ownership takeover was started by mistake, the possibility exists for two sites to write data to the same virtual volume.</p>	<p>If no cluster failure occurred, it is possible that a virtual volume that is accessed by another cluster in read-only takeover mode contains older data than the one on the owning cluster. This situation can occur if the virtual volume was modified on the owning cluster while the communication path between the clusters was down. When the links are reestablished, those volumes are marked in error.</p>

See the TS7700 IBM Knowledge Center available locally by clicking the question mark icon in the upper right corner of the MI window, or online at the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7700/cust/4.0/hydra_c_ic_home.html

Repair Virtual Volumes window

Damaged volumes typically occur due to a user intervention, such as enabling ownership takeover against a live cluster and ending up with two different versions of the same volume, or in a cluster removal scenario where the removed cluster had the only instance of a volume.

In these cases, the volume is moved to the FF20 (damaged) category by the TS7700 subsystem, and the host cannot access it. If access is attempted, messages like CBR4125I Valid copy of volume volser in library library-name inaccessible are displayed.

To repair virtual volumes in the damaged category for the TS7700 Grid, use this window. Figure 9-140 shows the Repair Virtual Volumes window of the TS7700 MI.

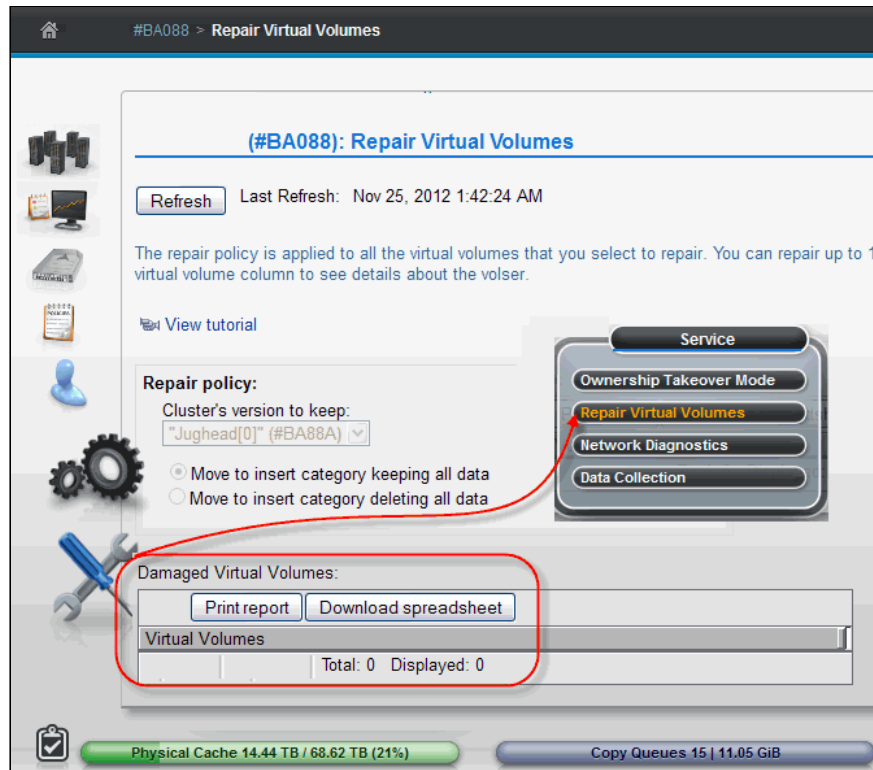


Figure 9-140 Repair Virtual Volumes window

The user can print the table data by clicking **Print report**, which is shown in Figure 9-140. A comma-separated value (.csv) file of the table data can be downloaded by clicking **Download spreadsheet**. The following information is displayed on this window:

- ▶ **Repair policy.** The Repair policy section defines the repair policy criteria for damaged virtual volumes in a cluster. The following criteria is shown:
 - **Cluster's version to keep.** The selected cluster obtains ownership of the virtual volume when the repair is complete. This version of the virtual volume is the basis for repair if the **Move to insert category keeping all data** option is selected.
 - **Move to insert category keeping all data.** This option is used if the data on the virtual volume is intact and still relevant. If data has been lost, do not use this option. If the cluster that is chosen in the repair policy has no data for the virtual volume to be repaired, choosing this option is the same as choosing **Move to insert category deleting all data**.

- **Move to insert category deleting all data.** The repaired virtual volumes are moved to the insert category and all data is erased. Use this option if the volume is returned to scratch or if data loss has rendered the volume obsolete. If the volume has been returned to scratch, the data on the volume is no longer needed. If data loss has occurred on the volume, data integrity issues can occur if the data on the volume is not erased.
- **Damaged Virtual Volumes.** The Damaged Virtual Volumes table displays all the damaged virtual volumes in a grid. The following information is shown:
 - **Virtual Volume.** The VOLSER of the damaged virtual volume. This field is also a hyperlink that opens the Damaged Virtual Volumes Details window, where more information is available.Damaged virtual volumes cannot be accessed; repair all damaged virtual volumes that appear on this table. The user can repair up to 10 virtual volumes at a time.

Follow these steps to repair damaged virtual volumes:

1. Define the repair policy criteria in the Repair policy section.
2. Select a cluster name from the Cluster's version to keep menu.
3. Click the radio button next to *either* **Move to insert category keeping all data** *or* **Move to insert category deleting all data**.
4. In the Damaged Virtual Volumes table, select the check box next to one or more (up to 10) damaged virtual volumes to be repaired by using the repair policy criteria.
5. Select **Repair** from the Select Action menu.
6. A confirmation message appears at the top of the window to confirm the repair operation. Click **View Task History** to open the Tasks window to monitor the repair progress. Click **Close Message** to close the confirmation message.

Network Diagnostics window

The Network Diagnostics window can be used to initiate ping or trace route commands to any IP address or host name from this TS7700 cluster. The user can use these commands to test the efficiency of grid links and the network system.

Figure 9-141 shows the navigation to the Network Diagnostics window and a ping test example.

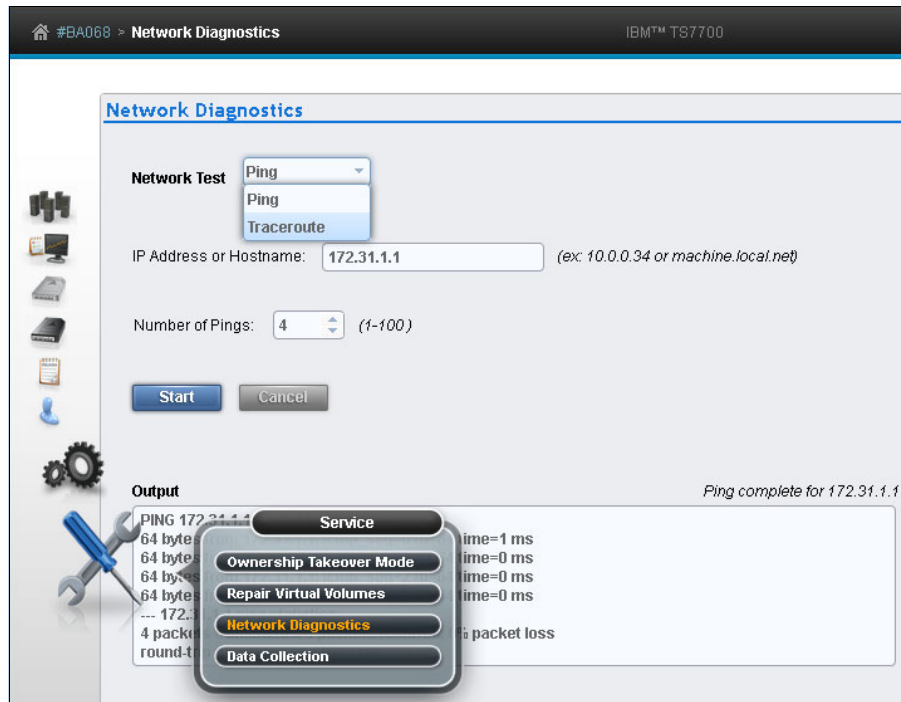


Figure 9-141 Network Diagnostics window

The following information is shown on this window:

- ▶ **Network Test:** The type of test to be run from the accessing cluster. The following values are available:
 - **Ping:** Select this option to initiate a ping test against the IP address or host name entered in the IP Address/Hostname field. This option tests the length of time that is required for a packet of data to travel from the computer to a specified host, and back again. This option can test whether a connection to the target IP address or host name is enabled, the speed of the connection, and the distance to the target.
 - **Traceroute:** Select this option to initiate a trace route test against the IP address or host name that is entered in the IP Address/Hostname field. This option traces the path that a packet follows from the accessing cluster to a target address and displays the number of times packets are rebroadcasted by other servers before reaching their destination.

Important: The **Traceroute** command is intended for network testing, measurement, and management. It imposes a heavy load on the network and should not be used during normal operations.

- **IP Address/Hostname:** The target IP address or host name for the selected network test. The value in this field can be an IP address in IPv4 or IPv6 format or a fully qualified host name.
- **Number of Pings:** Use this field to select the number of pings that are sent by the **Ping** command. The range of available pings is 1 - 100. The default value is 4. This field is only displayed if the value in the Network Test field is Ping.

- ▶ **Start:** Click this button to begin the selected network test. This button is disabled if required information is not yet entered on the window or if the network test is in progress.
- ▶ **Cancel:** Click this button to cancel a network test in progress. This button is disabled unless a network test is in progress.
- ▶ **Output:** This field displays the progress output that results from the network test command. Information that is retrieved by the web interface is displayed in this field as it is received. The user can scroll within this field to view output that exceeds the space provided.

The status of the network command is displayed in line with the Output field label and right-aligned over the Output field. The format for the information that is displayed is shown:

```
Pinging 98.104.120.12...
Ping complete for 98.104.120.12
Tracing route to 98.104.120.12...
Trace complete to 98.104.120.12
```

Data Collection window

To collect a snapshot of data or a detailed log to help check system performance or troubleshoot a problem during the operation of the TS7700, use this window.

If the user is experiencing a performance issue on a TS7700, the user has two options to collect system data for later troubleshooting. The first option, System Snapshot, collects a summary of system data that includes the performance state. This option is useful for intermittently checking the system performance. This file is built in approximately 5 minutes.

The second option, TS7700 Log Collection, enables you to collect historical system information for a time period up to the past 12 hours. This option is useful for collecting data during or soon after experiencing a problem. Based on the number of specified hours, this file can become large and require over an hour to build.

Figure 9-142 shows the Data Collection window in the MI.

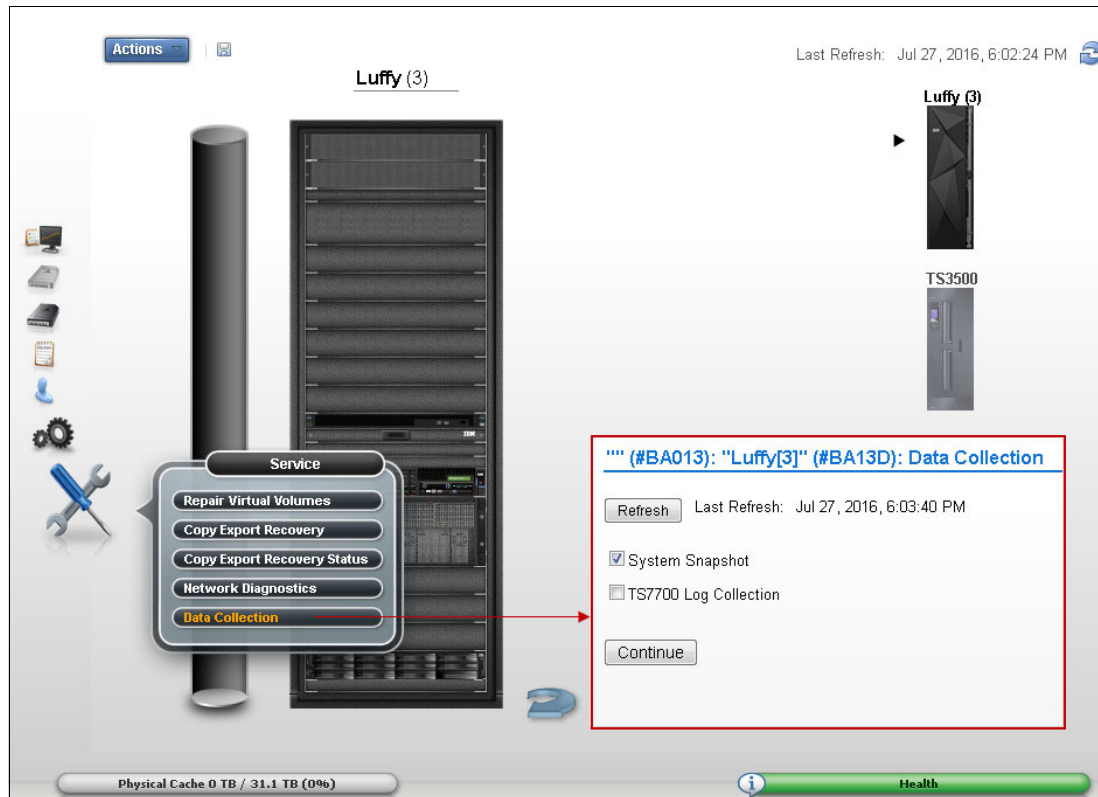


Figure 9-142 Data Collection window

The following information is shown on the Data Collection window:

- ▶ **System Snapshot.** Select this box to collect a summary of system health and performance from the preceding 15-minute period. The user can collect and store up to 24 System Snapshot files at the same time.
- ▶ **TS7700 Log Collection.** Check this box to collect and package all logs from the time period that is designated by the value in the Hours of Logs field. The user can collect and store up to two TS7700 Log Collection files at the same time.
- ▶ **Hours of Logs.** Use this menu to select the number of preceding hours from which system logs are collected. Possible values are 1 - 12, with a default of 2 hours. The time stamp next to the hours field displays the earliest time from which logs are collected. This time stamp is automatically calculated based on the number that is displayed in the hours field.

Note: Periods that are covered by TS7700 Log Collection files cannot overlap. If the user attempts to generate a log file that includes a period that is covered by an existing log file, a message prompts the user to select a different value for the hours field.

- ▶ **Continue.** Click this button to initiate the data collection operation. This operation cannot be canceled after the data collection begins.

Note: Data that is collected during this operation is not automatically forwarded to IBM. The user must contact IBM and open a problem management report (PMR) to move manually the collected data off the system.

When data collection is started, a message is displayed that contains a button linking to the Tasks window. The user can click this button to view the progress of data collection.

Important: If data collection is started on a cluster that is in service mode, the user might not be able to check the progress of data collection. The Tasks window is not available for clusters in service mode, so there is no link to it in the message.

- ▶ **Data Collection Limit Reached.** This dialog box opens if the maximum number of System Snapshot or TS7700 Log Collection files exists. The user can save a maximum number of 24 System Snapshot files or two TS7700 Log Collection files. If the user attempted to save more than the maximum of either type, the user is prompted to delete the oldest existing version before continuing. The name of any file to be deleted is displayed.
Click **Continue** to delete the oldest files and proceed. Click **Cancel** to abandon the data collection operation.
- ▶ **Problem Description.** Optional: Enter a detailed description of the conditions or problem that was experienced before any data collection has been initiated, in this field. Include symptoms and any information that can assist IBM Support in the analysis process, including the description of the preceding operation, VOLSER ID, device ID, any host error codes, any preceding messages or events, time and time zone of incident, and any PMR number (if available). The number of characters in this description cannot exceed 1000.

Copy Export Recovery window

To test a Copy Export recovery, or to run an actual Copy Export recovery on the TS7700 cluster, use this window.

Tip: This window is only visible in a single TS7740 or TS7720T configuration (both of which are connected to a physical tape library).

Figure 9-143 shows the Copy Export Recovery window.

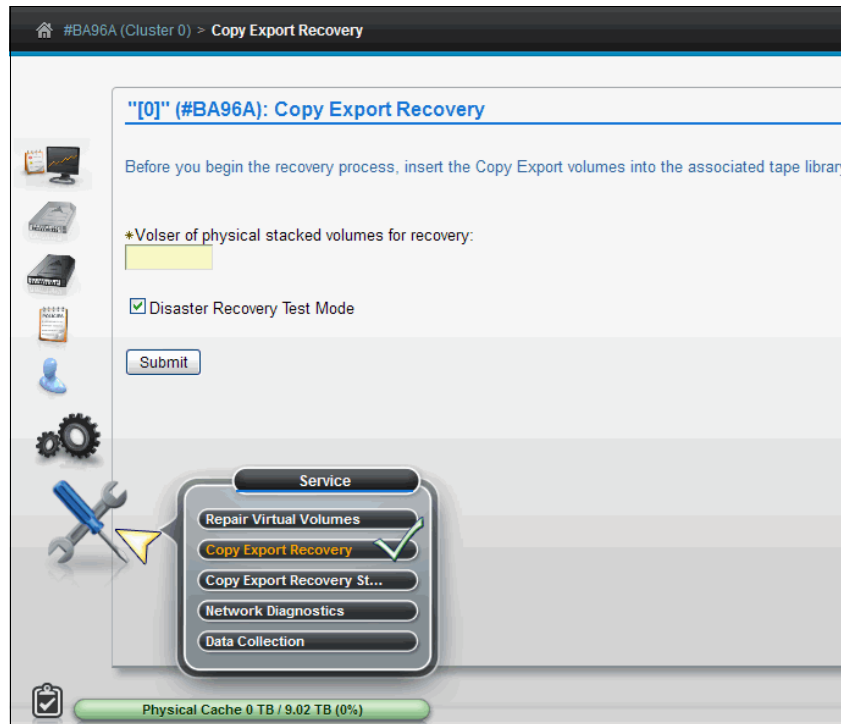


Figure 9-143 Copy Export recovery window

Copy Export enables the export of all virtual volumes and the virtual volume database to physical volumes, which can then be ejected and saved as part of a data retention policy for disaster recovery. The user can also use this function to test system recovery.

For a detailed explanation of the Copy Export function, see Chapter 12, “Copy Export” on page 757.

Reminder: The recovery cluster needs tape drives that are compatible with the exported media. Also, if encrypted tapes are used for export, access to the EKs must be provided.

Before the user attempts a Copy Export, ensure that all physical media that is used in the recovery is inserted. During a Copy Export recovery, all current virtual and physical volumes are erased from the database and virtual volumes are erased from the cache. Do not attempt a Copy Export operation on a cluster where current data is to be saved.

Important: In a grid configuration, each TS7700 is considered a separate source. Therefore, only the physical volume that is exported from a source TS7700 can be used for the recovery of that source. Physical volumes that are exported from more than one source TS7700 in a grid configuration cannot be combined to use in recovery. Recovery can occur only to a single cluster configuration; the TS7700 that is used for recovery must be configured as Cluster 0.

Secondary Copies window

If the user creates a new secondary copy, the original secondary copy is deleted because it becomes inactive data. For example, if the user modifies constructs for virtual volumes that already were exported and the virtual volumes are remounted, a new secondary physical volume is created.

The original physical volume copy is deleted without overwriting the virtual volumes. When the Copy Export operation is rerun, the new, active version of the data is used.

The following fields and options are presented to the user to help testing recovery or running a recovery:

- ▶ **Volser of physical stacked volume for Recovery Test.** The physical volume from which the Copy Export recovery attempts to recover the database.
- ▶ **Disaster Recovery Test Mode.** This option determines whether a Copy Export Recovery is run as a test or to recover a system that has suffered a disaster. If this box contains a check mark (default status), the Copy Export Recovery runs as a test. If the box is cleared, 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 that are written to the system are written to newly added scratch tapes, and do 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 in normal mode (box cleared) rewrites virtual volumes to physical storage if the constructs change so that the virtual volume's data can be put in the correct pools. Also, in this type of recovery, reclaim processing remains enabled and primary physical copies are restored, requiring the addition of scratch physical volumes.

A recovery that is run in this mode enables the data on the Copy Export tapes to expire in the normal manner and those physical volumes to be reclaimed.

Note: The number of virtual volumes that can be recovered depends on the number of FC5270 licenses that are installed on the TS7700 that is used for recovery. Additionally, a recovery of more than 2 million virtual volumes must be run by a TS7740 operating with a 3957-V07 and a code level of 8.30.0.xx or higher.

- ▶ **Erase all existing virtual volumes during recovery.** This check box is shown if virtual 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 running the recovery. The user can check this check box to proceed with the Copy Export Recovery operation.
- ▶ **Submit.** Click this button to initiate the Copy Export Recovery operation.
- ▶ **Confirm Submission of Copy Export Recovery.** The user is asked to confirm the decision to initiate a Copy Export Recovery option. Click **OK** to continue with the Copy Export Recovery operation. Click **Cancel** to abandon the Copy Export Recovery operation and return to the Copy Export Recovery window.
- ▶ **Password.** The user password. If the user selected the **Erase all existing virtual volumes during recovery** check box, the confirmation message includes the Password field. The user must provide a password to erase all current data and proceed with the operation.
- ▶ **Canceling a Copy Export Recovery operation in progress.** The user can cancel a Copy Export Recovery operation that is in progress from the Copy Export Recovery Status window.

Copy Export Recovery Status window

Use this window to view information about or to cancel a currently running Copy Export recovery operation on a TS7700 cluster.

Figure 9-144 shows the Copy Export Recovery Status window in the MI.

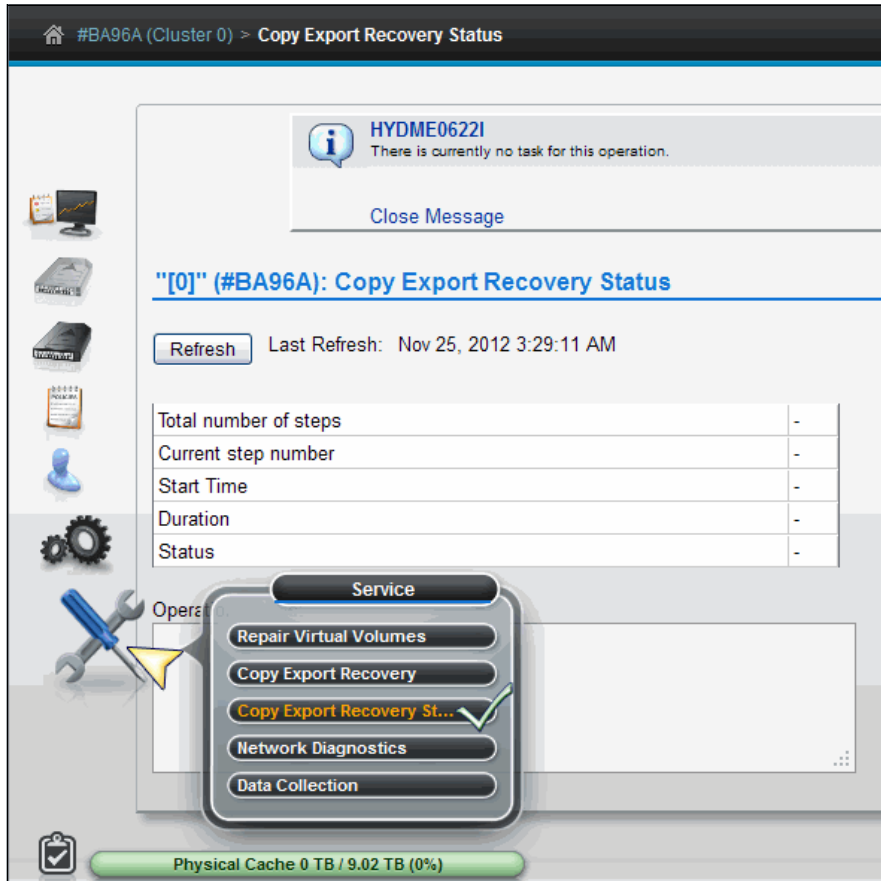


Figure 9-144 Copy Export Recovery Status window

Important: The Copy Export recovery status is only available for a stand-alone TS7700T cluster.

The table on this window displays the progress of the current Copy Export recovery operation. This window includes the following information:

- ▶ **Total number of steps.** The total number of steps that are required to complete the Copy Export recovery operation.
- ▶ **Current step number.** The number of steps completed. This value is a fraction of the total number of steps that are required to complete, not a fraction of the total time that is required to complete.
- ▶ **Start time.** The time stamp for the start of the operation.
- ▶ **Duration.** The amount of time the operation has been in progress, in hours, minutes, and seconds.

- ▶ **Status.** The status of the Copy Export recovery operation. The following values are possible:
 - **No task.** No Copy Export operation is in progress.
 - **In progress.** The Copy Export operation is in progress.
 - **Complete with success.** The Copy Export operation completed successfully.
 - **Canceled.** The Copy Export operation was canceled.
 - **Complete with failure.** The Copy Export operation failed.
 - **Canceling.** The Copy Export operation is in the process of cancellation.
- ▶ **Operation details.** This field displays informative status about the progress of the Copy Export recovery operation.
- ▶ **Cancel Recovery.** Click the Cancel Recovery button to end a Copy Export recovery operation that is in progress and erase all virtual and physical data. The Confirm Cancel Operation dialog box opens to confirm the decision to cancel the operation. Click **OK** to cancel the Copy Export recovery operation in progress. Click **Cancel** to resume the Copy Export recovery operation.

9.3 Common procedures

This section describes how to run some tasks that are necessary during the implementation stage of the TS7700, whether in stand-alone or grid mode. Some procedures that are described here might also be useful later during the lifecycle of the TS7700, when a change in configuration or operational parameter is necessary for the operation of the subsystem to meet the new requirements.

The tasks are grouped by these criteria:

- ▶ Procedures that are related to the tape library connected to a TS7700 tape-attached model
- ▶ Procedures that are used with all TS7700 cluster models

9.3.1 The tape library with the TS7700T cluster

The following sections describe the steps necessary to configure a TS7700 tape-attached cluster with a tape library.

Defining a logical library

The tape library GUI is required to define a logical library and run the following tasks. Therefore, ensure that it is set up correctly and working. For access through a standard-based web browser, an IP address must be configured in the tape library, which is done initially by the IBM SSR during the hardware installation at the TS3500 or TS4500.

Important:

- ▶ Each tape attach TS7000 cluster requires its own logical library in the tape library.
- ▶ The ALMS feature must be installed and enabled to define a logical library partition in both TS3500 and TS4500 tape libraries.

Ensuring that ALMS is enabled

Before enabling ALMS, the ALMS license key must be entered through the TS3500 tape library operator window because ALMS is a chargeable feature.

You can check the status of ALMS with the TS3500 tape library GUI by clicking **Library** → **ALMS**, as shown in Figure 9-145.

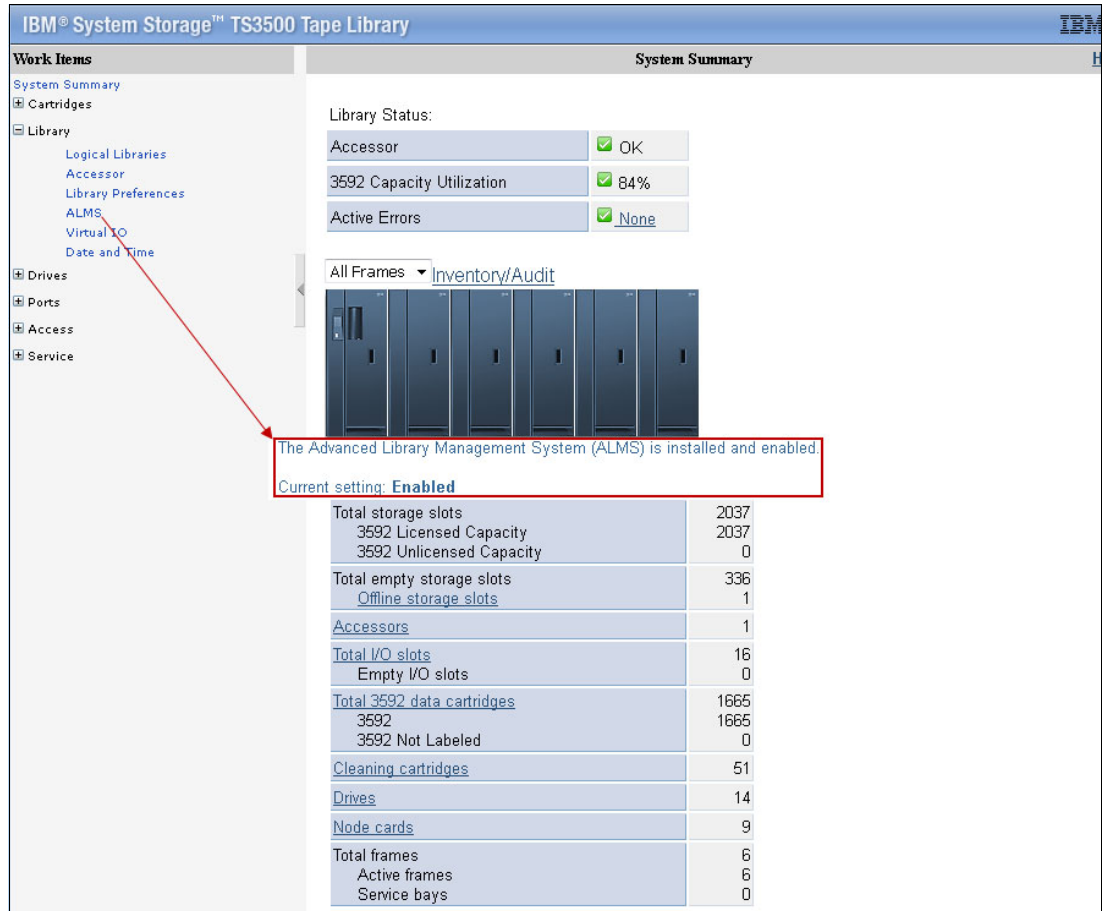


Figure 9-145 TS3500 tape library GUI Summary and ALMS window

When ALMS is enabled for the first time in a partitioned TS3500 tape library, the contents of each partition are migrated to ALMS logical libraries. When enabling ALMS in a non-partitioned TS3500 tape library, cartridges that are already in the library are migrated to the new ALMS single logical library.

Figure 9-146 shows how to check the ALMS status with the TS4500 tape library. If necessary, the license key for the ALMS feature can be entered and activated in the same page.




Figure 9-146 ALMS installed and enabled on TS4500

Creating a logical library with TS4500

See IBM Knowledge Center for more complete information about TS4500 operations and configuration. IBM Knowledge Center is available locally at TS4500 GUI by clicking in question mark icon, or online at:

<http://www.ibm.com/support/knowledgecenter/en/STQRQ9/com.ibm.storage.ts4500.doc>

Complete these steps for a TS4500 tape library:

1. From the initial page of the TS4500 GUI, select Library Icon and click logical library. See Figure 9-147 on page 512 for a visual reference.

2. Work with the *Create Logical Library* option as detailed in Figure 9-147 to complete the task.

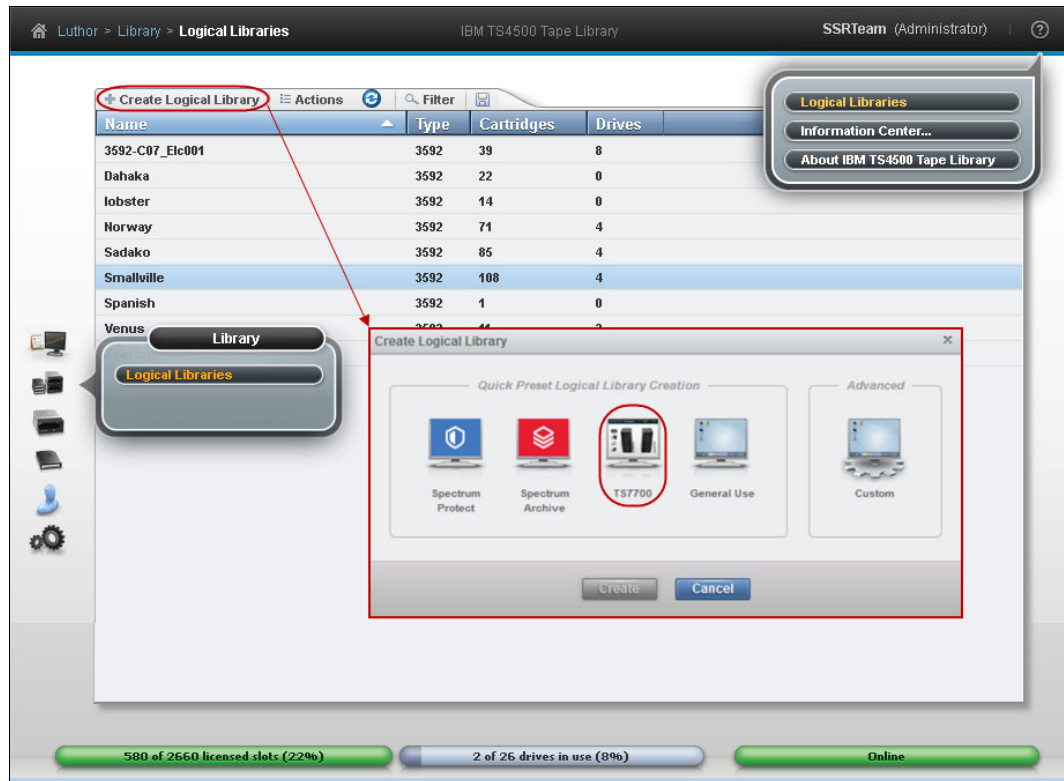


Figure 9-147 TS4500 create logical library page

Notice that the TS4500 GUI features selected presets, which helps in the setup of the new logical library. For the TS7700 library, use the TS7700 option that is highlighted in Figure 9-147. This option uses the 3592 tape drives that are not assigned to any existent logical library within the TS4500 tape library. Also, it selects up to four drives as control paths, distributing them in two separate frames, when this is possible.

Note: The TS7700 preset is disabled when less than four unassigned tape drives are available to create a new logical library.

Figure 9-148 on page 513 shows how to display which tape drives are available (unassigned) to be configured in the new logical library. Always work with your IBM service representative when defining drives for the TS7700 during installation or any further change in the environment. Those drives must be correctly cabled to the fibre switches dedicated to the TS7700, and the new back-end resources need to be configured (or reconfigured) within the cluster for proper operation.




Figure 9-148 Display available tapes

The preset also indicates the *System Managed* encryption method for the new TS7700 logical library.

See Figure 9-149 for an example of logical library definition.

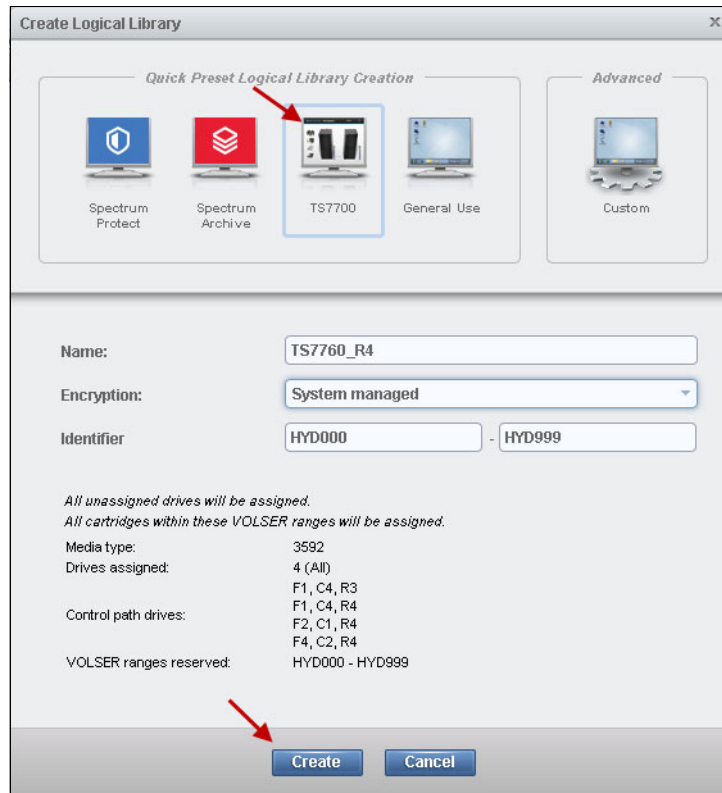


Figure 9-149 Defining the new Logical Library for the TS7700T

After the configuration of the logical library is completed, your IBM service representative can complete the TS770 tape-attached cluster installation, and the tape cartridges can be inserted in the TS4500 tape library.

Creating a logical library with ALMS on the TS3500 tape library

Complete these steps for a TS3500 tape library:

1. From the main section of the TS3500 tape library GUI Welcome window, go to the work items on the left side of the window and click **Library** → **Logical Libraries**
2. From the **Select Action** menu, select **Create** and click **Go**.
An extra window, named Create Logical Library pops up.
3. Type the logical library name (up to 15 characters), select the media type (**3592** for TS7740 or TS7720T), and then click **Apply**. The new logical library is created and is displayed in the logical library list when the window is refreshed.

4. After the logical library is created, you can display its characteristics by selecting **Library** → **Logical Libraries** under work items on the left side of the window. Figure 9-150 shows a summary of the screens in the Create logical library sequence.

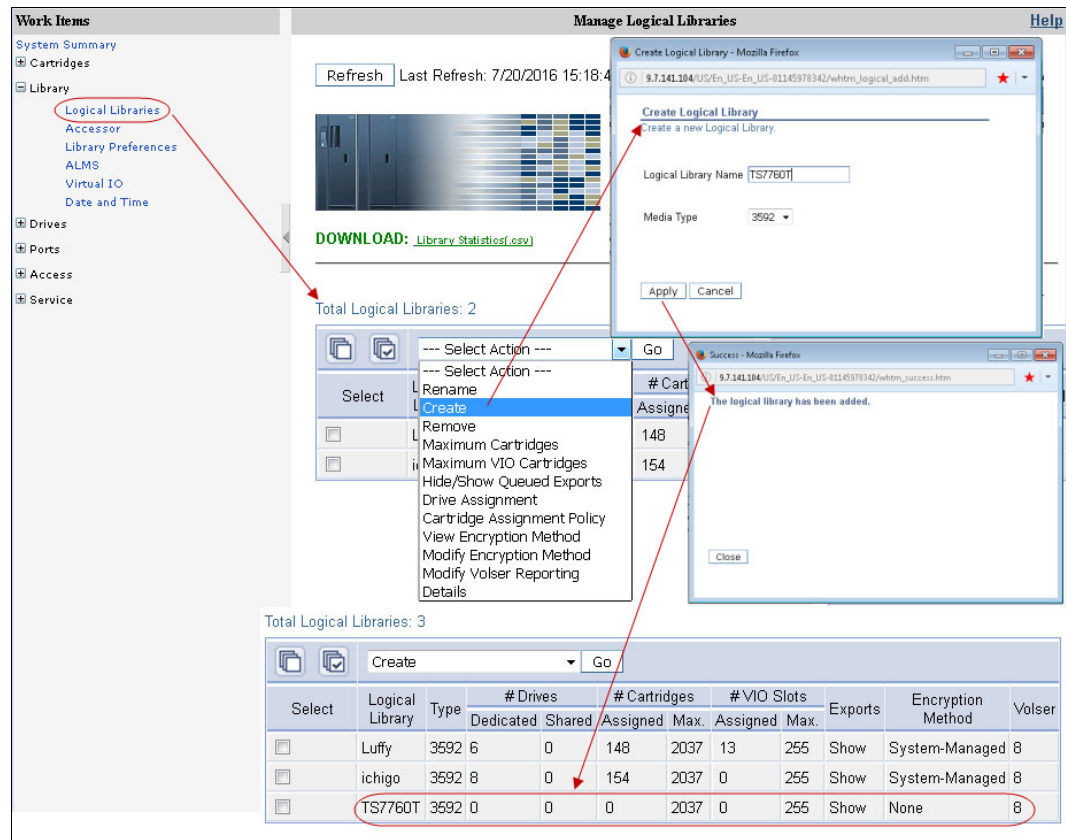


Figure 9-150 Creating a Logical Library with the TS3500 tape library

Maximum number of slots, 8-character volser, and VIO

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 tape library cartridge slots for each logical library. This enables a logical library to grow without changing the configuration each time you want to add empty slots.

Make sure that the new logical library has the **eight-character Volser** reporting option set. Another item to consider is the VIO usage - if VIO is enabled and, if so, how many cells should be defined. For more information, see the documentation regarding the TS3500 Tape Library available on virtual I/O slots and applicability:

http://www.ibm.com/support/knowledgecenter/STCMML8/com.ibm.storage.ts3500.doc/ipg_3584_a69p0vio.html

All those items can be checked and defined at the *Manage Logical Libraries* page in the TS3500 tape library, as shown in Figure 9-151.

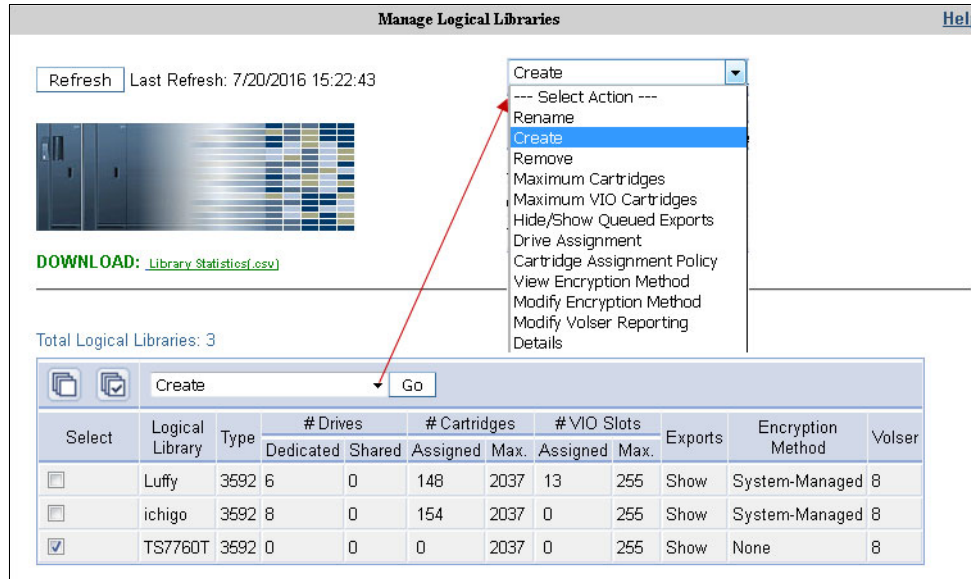


Figure 9-151 Manage Logical Libraries in the TS3500

Assigning drives

Now, the TS7700T tape drives should be added to the logical library.

From the Logical Libraries window that is shown in Figure 9-152 on page 517, use the work items on the left side of the window to go to the requested web window by clicking **Drives** → **Drive Assignment**. This link takes you to a filtering window where you can select to have the drives displayed by drive element or by logical library.

Note: For the 3592 J1A, E05, E06, and E07 drives, an intermix of tape drive models is *not* supported by TS7720T or TS7740, except for 3592-E05 tape drives working in J1A emulation mode and 3592-J1A tape drives (the first and second generation of the 3592 tape drives).

Upon selection, a window opens so that a drive can be added to or removed from a library configuration. Also, you can use this window to share a drive between Logical Libraries and define a drive as a control path.

Figure 9-152 on page 517 shows the drive assignment window of a logical library that has all drives assigned.

Unassigned drives appear in the Unassigned column with the box checked. To assign them, select the appropriate drive box under the logical library name and click **Apply**.

Note: Do not share drives belonging to a TS7700T. They must be exclusive.

Click the **Help** link at the upper-right corner of the window that is shown in Figure 9-152 to see extended help information, such as detailed explanations of all the fields and functions of the window. The other TS3500 tape library GUI windows provide similar help support.




Figure 9-152 Drive Assignment window

TS7700 R4.0 works with the TS1150 tape drives in a homogeneous or heterogeneous configuration. Heterogeneous configuration of the tape drives means a mix of TS1150 (3592 E08) and one previous generation of the 3592 tape drives to facilitate data migration from legacy media. Tape drives from previous generation are only used to read legacy media (JA/JB) while the TS1150 will read/write to the newer media types. There will be no writes to the legacy media type, so the support for heterogeneous configuration of the tape drives is deemed limited.

You can read more about heterogeneous drive support in Chapter 2, “Architecture, components, and functional characteristics” on page 15 and Chapter 8, “Migration” on page 283.

Note: For the 3592 J1A, E05, E06, and E07 drives, an intermix of tape drive models is *not* supported by TS7720T or TS7740, except for 3592-E05 tape drives working in J1A emulation mode and 3592-J1A tape drives (the first and second generation of the 3592 tape drives).

In a multi-platform environment, logical libraries show up as in Figure 9-152. Physical tape drives can be reassigned from one logical library to another. This can be easily done for the Open Systems environment, where the tape drives attach directly to the host systems without a tape controller or VTS/TS7700.

Note: Do not change drive assignments if they belong to an operating TS774000T, or tape controller. Work with your IBM SSR, if necessary.

In a z Systems environment, a tape drive always attaches to one tape control unit (CU) only. If it is necessary to change the assignment of a tape drive from a TS7720T or TS7740, the CU must be reconfigured to reflect the change. Otherwise, the missing resource is reported as defective to the MI and hosts. Work with your IBM SSRs to perform these tasks in the correct way, avoiding unplanned outages.

Important: In a z Systems environment, use the Drive Assignment window *only* for these functions:

- ▶ *Initially* assign the tape drives from the TS3500 tape library GUI to a logical partition (LPAR).
- ▶ Assign more tape drives after they are attached to the TS7740 or a tape controller.
- ▶ Remove physical tape drives from the configuration *after* they are physically detached from the TS7740 or tape controller.

In addition, never disable ALMS at the TS3500 tape library after it is enabled for z Systems host support and z Systems tape drive attachment.

Defining control paths

Each TS7740 requires four control path drives defined. If possible, distribute the control path drives over more than one TS3500 tape library frame to avoid single points of failure.

Defining the encryption method for the new logical library

After adding tape drives to the new logical library, the encryption method for the new logical library (if applicable) needs to be defined.

Reminders:

- ▶ When using encryption, tape drives must be set to Native mode.
- ▶ To activate encryption, FC9900 must have been ordered for the TS7400 or the TS7720T, and the license key must be installed. In addition, the associated tape drives must be Encryption Capable 3592-E05, 3592-E06, 3592-E07, or 3592-E08.

Complete the following steps:

1. Check the drive mode by opening the Drives summary window in the TS3500 tape library GUI, as shown in Figure 9-153, and look in the Mode column. This column is displayed only if drives in the tape library are emulation-capable.

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

Figure 9-153 Checking drive mode

- If necessary, change the drive mode to Native mode (3592-E05 only). In the Drives summary window, select a drive and select **Change Emulation Mode**, as shown in Figure 9-154.




Figure 9-154 Change the drive emulation

- In the next window that opens, select the native mode for the drive. After the drives are in the wanted mode, proceed with the Encryption Method definition.
- In the TS3500 MI, click **Library** → **Logical Libraries**, select the logical library with which you are working, select **Modify Encryption Method**, and then click **Go**. See Figure 9-155.

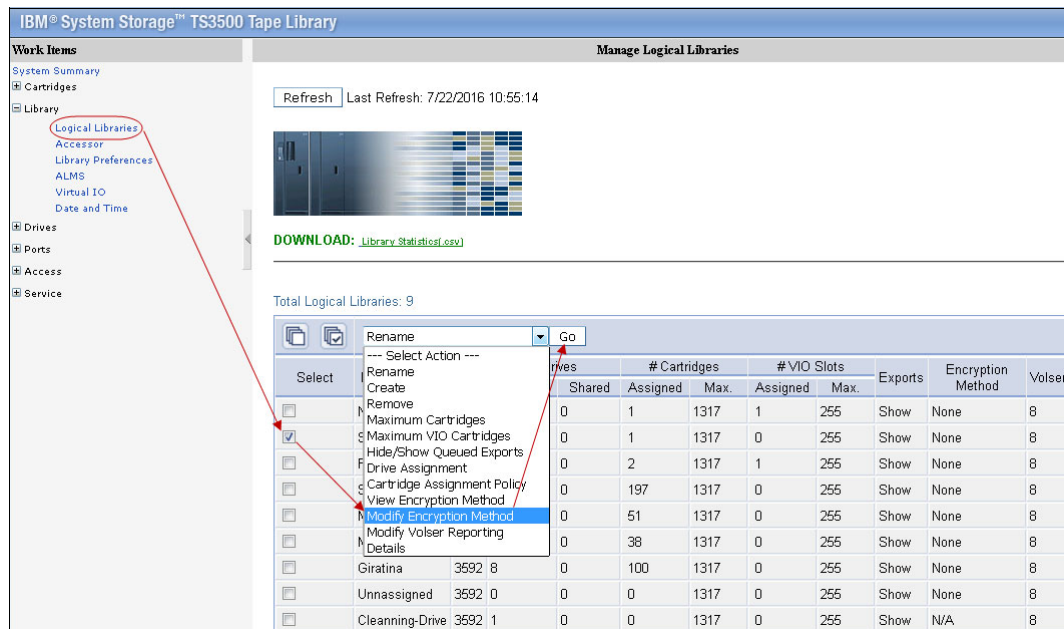


Figure 9-155 Select the encryption method

- In the window that opens, select **System-Managed** for the chosen method, and select all drives for this partition. See Figure 9-156.

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 9-156 Set the encryption method

To make encryption fully operational in the TS7740 configuration, more steps are necessary. Work with your IBM SSR to configure the Encryption parameters in the TS7740 during the installation process.

Important: Keep the Advanced Encryption Settings as *NO ADVANCED SETTING*, unless set otherwise by IBM Engineering.

Defining Cartridge Assignment Policies

The *Cartridge Assignment Policy* (CAP) of the TS3500 tape library is where the ranges of physical cartridge volume serial numbers are assigned to specific logical libraries. With CAP correctly defined, when a cartridge is inserted with a VOLSER that matches that range into the I/O station, the library automatically assigns that cartridge to the appropriate logical library.

To add, change, and remove policies, select **Cartridge Assignment Policy** from the Cartridges work items. The maximum quantity of CAPs for the entire TS3500 tape library must not exceed 300 policies.

Figure 9-157 shows the VOLSER ranges defined for logical libraries.

Select	Logical Library	Volume Serial Number Ranges
<input checked="" type="radio"/>	Newman	X00110 - X00119
<input type="radio"/>	Newman	X00140 - X00159
<input type="radio"/>	Newman	000100 - 000119
<input type="radio"/>	Newman	JA0165 - JA0166
<input type="radio"/>	Kramer	J1G000 - J1G999
<input type="radio"/>	Kramer	JJC242 - JJC262
<input type="radio"/>	Kramer	JBR075 - JBR084
<input type="radio"/>	Kramer	JA0280 - JA0299
<input type="radio"/>	Kramer	JJH040 - JJH059
<input type="radio"/>	Kramer	JJH160 - JJH199
<input type="radio"/>	Kramer	JA0800 - JA0899
<input type="radio"/>	Kramer	F00140 - F00199
<input type="radio"/>	Kramer	JJY390 - JJY394
<input type="radio"/>	Zero	310650 - 310999
<input type="radio"/>	Zero	J1M600 - J1M699
<input type="radio"/>	Zero	J1M750 - J1M799

Figure 9-157 TS3500 Tape Library Cartridge Assignment Policy

The TS3500 tape library enables duplicate VOLSER ranges for different media types only. For example, Logical Library 1 and Logical Library 2 contain Linear Tape-Open (LTO) media, and Logical Library 3 contains IBM 3592 media. Logical Library 1 has a CAP of ABC100-ABC200. The library rejects an attempt to add a CAP of ABC000-ABC300 to Logical Library 2 because the media type is the same (both LTO). However, the library does enable an attempt to add a CAP of ABC000-ABC300 to Logical Library 3 because the media (3592) is different.

In a storage management subsystem (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 sharing a TS3500 tape library between z and Open Systems hosts or not.

Tip: The CAP does not reassign an already assigned tape cartridge. If needed, you must first unassign it, then manually reassign it.

Inserting TS7700T physical volumes

The tape attach TS7700 subsystem manages both logical and physical volumes. The CAP of the TS3500 tape library or the associate volume ranges at TS4500 affects only the physical volumes that are associated with this TS7740 or TS7720T logical library. Logical Volumes are managed exclusively from the TS7700 MI.

To add physical cartridges, complete the following steps:

1. Define CAPs at the IBM TS3500 or apply the volser ranges at TS4500 tape library through the GUI. This process ensures that all TS7700 ranges are recognized and assigned to the correct logical library partition (the logical library that is created for this specific TS7700) before you begin any TS7700 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 tape attach TS7700 logical library partitions according to the definitions.

Important: Before inserting physical volumes belonging to a TS7700T into the tape library, ensure that the VOLSER ranges are defined correctly at the TS7700 MI. For more information, see “Defining VOLSER ranges for physical volumes” on page 529.

These procedures ensure that TS7700 back-end cartridges are never assigned to a host by accident. Figure 9-158 shows the flow of physical cartridge insertion and assignment to logical libraries for TS7740 or TS7720T.

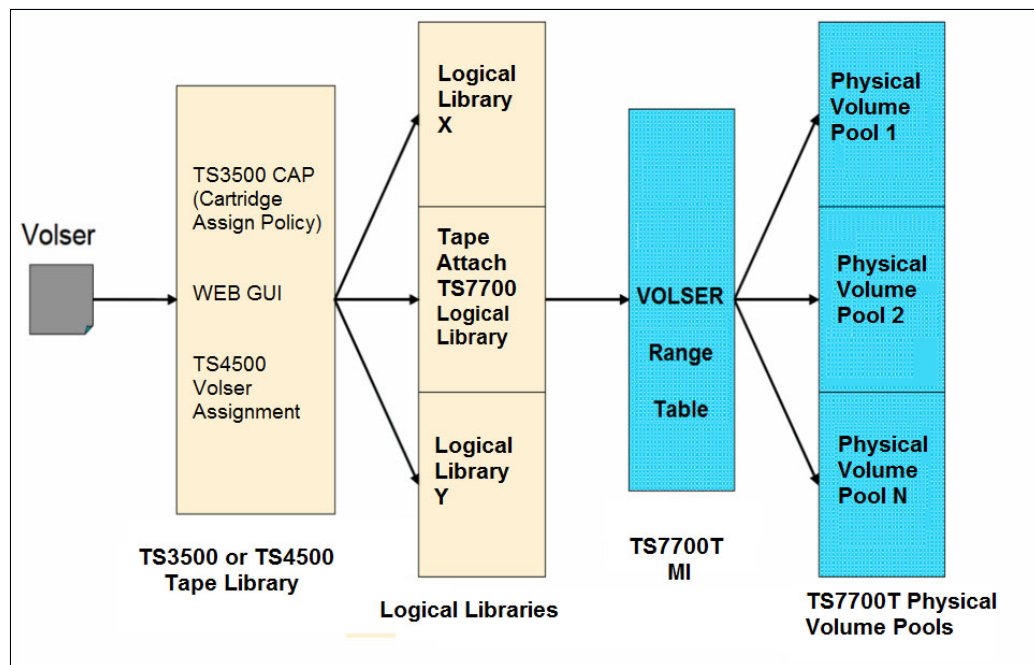


Figure 9-158 Physical volume assignment

Inserting physical volumes into the tape library

Two methods are available for inserting physical volumes into the tape library:

- ▶ Opening the library doors and inserting the volumes directly into the tape library storage empty cells (bulk loading)
- ▶ Using the tape library I/O station

Insertion directly into storage cells

Use the operator pane of the tape library to pause it. Open the door and insert the cartridges into any empty slot, except those slots that are reserved for diagnostic cartridges, which are Frame 1, Column 1 in the first Row (F01, C01, and R01) in a single media-type library. Also do not insert cartridges in the shuffle locations in the high-density frames (top two first rows in the HD frame). Always use empty slots in the same frame whose front door was opened, otherwise the cartridges will not be inventoried.

Important: Cartridges that are not in a CAP range (TS3500) or associated to any logical library (TS4500) are not assigned to any logical library.

After completing the new media insertion, close the doors. After approximately 15 seconds, the tape library automatically inventories the frame or frames of the door you opened.

When the tape library finishes the physical inventory, the TS7700T uploads the inventory from its associate logical library. At the end of the inventory upload, the tape library comes to the *Auto* status to the tape attach TS7700 cluster.

Tip: Only place cartridges in a frame whose front door is open. Do not add or remove cartridges from an adjacent frame.

Insertion by using the I/O station

The tape library can be operating with or without virtual I/O (VIO) being enabled.

Basically, with VIO enabled, the tape library moves 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 that is mapped as a VIO - SCSI element between 769 (X'301') and 1023 (X'3FF') for the logical library that is designated by the Volser association or CAP.

Each logical library has its own set of up to 256 VIO slots, as defined during logical library creation or later.

With VIO disabled, the tape library does not move cartridges from the physical I/O station unless it receives a command from the TS7700T or any other host in control.

For both cases, the tape library detects the presence of cartridges in the I/O station when it transitions from open to close, and scans all I/O cells by using the bar code reader. The CAP or volser assignment decides to which logical library those cartridges belong and then runs *one* of the following tasks:

- ▶ Moves them to the VIO slots of the designated logical library, with VIO enabled.
- ▶ Waits for a host command in this logical library. The cartridges stay in the I/O station after the bar code scan.

The volumes being inserted should belong to the range of volumes that are defined in the tape library (CAP or volser range) for the TS7700 logical library, and those ranges also should be defined in the TS7700 Physical Volume Range as described in “Defining VOLSER ranges for physical volumes” on page 529. Both conditions should be met to a physical cartridge be successfully inserted to the TS7700T.

If any VOLSER is not in the range that is defined by the policies, the cartridges need to be assigned to the correct logical library manually by operator.

Note: Make sure that CAP ranges are correctly defined. Insert Notification is not supported on a high-density library. If a cartridge outside the CAP-defined ranges is inserted, it remains unassigned without any notification, and it might be checked in by any logical library of the same media type.

Verify that the cartridges were correctly assigned - or not left unassigned by using the tape library GUI. Check Figure 9-159 for the MI page, with TS4500 and TS3500.




Figure 9-159 Check the volume assignment

When volumes that belong to a logical library are found unassigned, correct the CAP or volser assignment definitions, and reinsert them again. Optionally, cartridges could be manually assigned to the correct logical library by the operator through the GUI.

We strongly suggest having correctly defined CAP or volser assignment policies in the tape library for the best operation of the tape system.

Unassigned volumes in the tape attach TS7700

A physical volume goes to the Unassigned category in the TS7740 or TS7720T if it does not fit in any defined range of physical volumes for this TS7700 cluster. Defined Ranges and Unassigned Volumes can be checked in the TS7700 MI Physical Volume Ranges window that is shown in Figure 9-160.

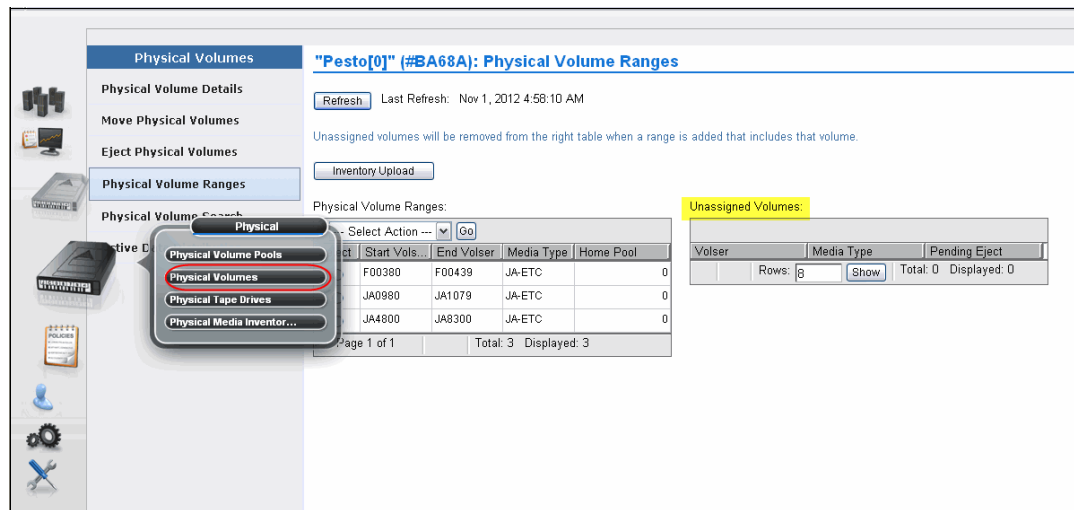


Figure 9-160 TS7740 unassigned physical volumes

If an unassigned volume should be assigned to this TS7700T, a new range that includes this volume must be created, as described in “Defining VOLSER ranges for physical volumes” on page 529. If this volume was incorrectly assigned to the TS7700 cluster, it should be ejected and reassigned to the correct logical library in the tape library. Also, make sure that CAP or volser assignments are correct in the tape library.

Assigning cartridges in the tape library to the logical library partition

This procedure is necessary only if a cartridge was inserted, without CAP or volser assignment being provided in advance (not recommended). To use this procedure, you must assign the cartridge manually to a logical library in the tape library.

Clarifications:

- ▶ Insert Notification is not supported in a high-density library for TS3500. The CAP must be correctly configured to provide automated assignment of all the inserted cartridges.
- ▶ A cartridge that has been manually assigned to the TS7700 logical library does not display automatically in the TS7700T inventory. An Inventory Upload is needed to refresh the TS7700 cluster inventory. The Inventory Upload function is available on the Physical Volume Ranges menu as shown in Figure 9-160.
- ▶ Cartridge assignment to a logical library is available only through the tape library GUI.

Assigning a data cartridge

To assign a data cartridge to a logical library in the TS3500 tape library, complete these steps:

1. Click in the *Cartridge* Icon on the TS4500 GUI, and select *Cartridges* there.
2. Find the cartridge that you want to assign (should show as unassigned at that point), and select it by clicking the line.
3. Right-click it, and select Assign. Choose the correct logical library in the list available.
4. Complete the assignment insertion by clicking *Assign* button.
5. For a TS7700T cluster, click **Physical** → **Physical Volumes** → **Physical Volume Ranges** and click **Inventory Upload**, as shown in Figure 9-164 on page 529

To assign a data cartridge to a logical library in the TS3500 tape library, complete these steps:

1. Open the TS3500 tape library GUI (go to the library’s Ethernet IP address or the library URL by using a standard browser). The Welcome window opens.
2. Click **Cartridges** → **Data Cartridges**. The Data Cartridges window opens.
3. Select the logical library to which the cartridge is assigned and select a sort view of the cartridge range. The library can sort the cartridge by volume serial number, SCSI element address, or frame, column, and row location. Click **Search**. The Cartridges window opens and shows all the ranges for the specified logical library.
4. Select the range that contains the data cartridge that should be assigned.
5. Select the data cartridge and then click **Assign**.
6. Select the logical library partition to which the data cartridge should be assigned to.
7. Click **Next** to complete the function.
8. For a TS7700T cluster, click **Physical** → **Physical Volumes** → **Physical Volume Ranges** and click **Inventory Upload**, as shown in Figure 9-164 on page 529.

Inserting a cleaning cartridge

Each drive in the tape library requires cleaning from time to time. Tape drives that are used by the TS7700 subsystem can request a cleaning action when necessary. This cleaning is carried out by the tape library automatically. However, the necessary cleaning cartridges must be provided.

Remember: Cleaning action is performed automatically by the tape libraries when necessary. A cleaning cartridge is good for 50 cleaning actions.

Use the cartridge magazine to insert cleaning cartridges into the I/O station, and then into the TS4500 tape library. The TS4500 can be set to move expired cleaning cartridges to the I/O station automatically. Figure 9-161 shows how to set it.

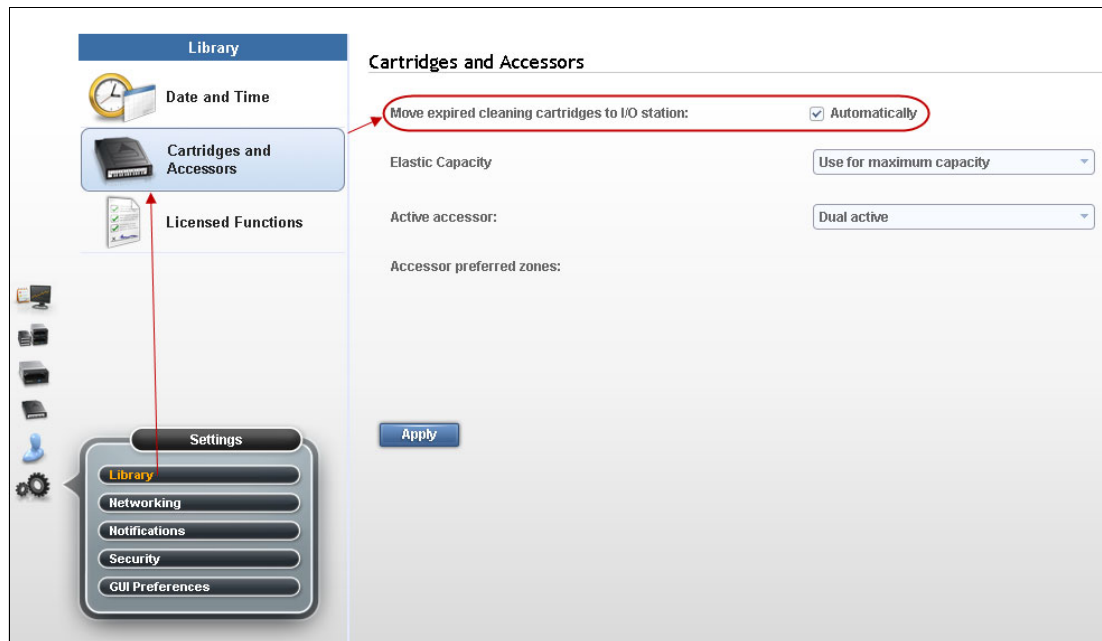


Figure 9-161 TS4500 tape library moves expired cleaning cartridge to I/O station automatically

The GUI page Cartridges by Logical Library under icon Cartridge shows how many cleaning cartridges are globally available to the TS4500 tape library, as shown in Figure 9-162.

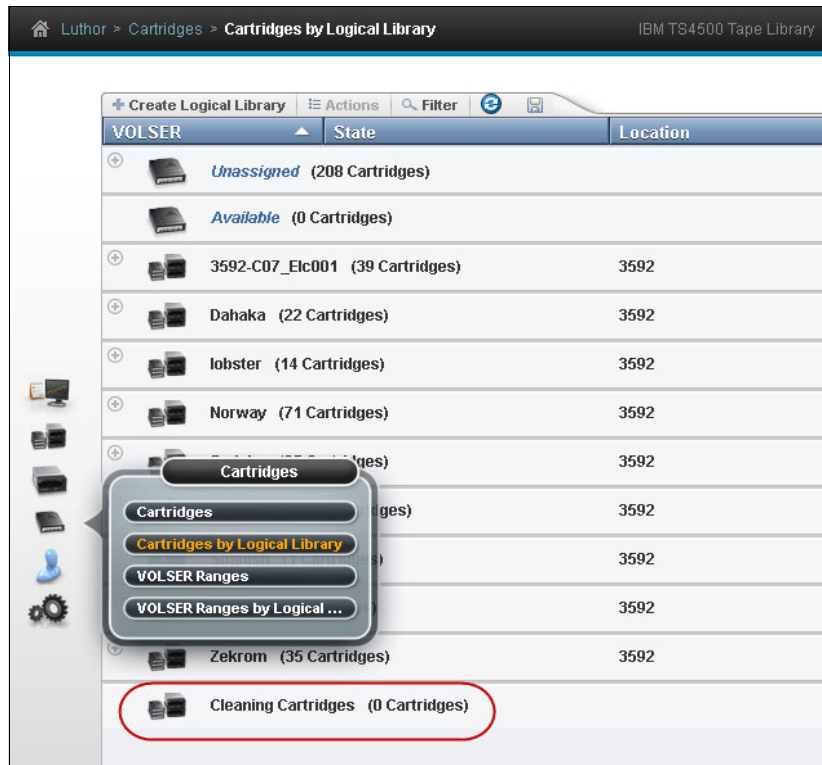


Figure 9-162 Displaying cleaning cartridges with the TS4500.

Also, there are TS4500 tape library command-line interface commands that can be used to check the status of the cleaning cartridges or alter settings in the tape library. Read about this in the documentation for TS4500, available locally by clicking the question mark icon at the top bar in the GUI, or online at:

https://www.ibm.com/support/knowledgecenter/STQRQ9/com.ibm.storage.ts4500.doc/ts4500_ichome.html

With TS3500 tape library, the process to insert cleaning cartridges varies depending on the setup of the tape library. A cleaning cartridge can be inserted by using the web interface or from the operator window. As many as 100 cleaning cartridges can be inserted in a TS3500 tape library.

To insert a cleaning cartridge by using the TS3500 tape library GUI, complete the following steps:

1. Open the door of the I/O station and insert the cleaning cartridge.
2. Close the door of the I/O station.
3. Enter the Ethernet IP address on the URL line of the browser. The Welcome Window opens.
4. Click **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 window, complete 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. If the tape is inserted incorrectly (for instance, upside down), the I/O station does not close properly. Do not force it.
3. Close the I/O station and click **Yes**. The tape library scans the I/O station for the cartridges and moves 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, and Back to return to the Activity window.

Tip: Cleaning cartridge are not assigned to specific logical libraries.

Removing cleaning cartridges from a TS3500 tape library

This section describes how to remove a cleaning cartridge by using the TS3500 tape library GUI. The operator panel in the tape library also can be used for this operation. For more information, see *IBM TS3500 Tape Library with ALMS Operator Guide, GA32-0594*.

To use the TS3500 tape library GUI to remove a cleaning cartridge from the tape library, complete 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 Cartridges window opens, as shown in Figure 9-163.
3. Select a cleaning cartridge. From the Select Action menu, select **Remove**, and then click **Go**.
4. Look at the Activity pane in the operator window to determine whether the I/O station to be used is locked or unlocked. If the station is locked, unlock it by using application software.
5. Open the door of the I/O station and remove the cleaning cartridge.
6. Close the door of the I/O station.

Determining the cleaning cartridge usage in the TS3500 tape library

The usage of the cleaning cartridge can be determined by using the same window that is used for the removal of the cleaning cartridges. See the Cleans Remaining column shown in Figure 9-163.

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 9-163 TS3500 tape library cleaning cartridges

9.3.2 TS7700T definitions

This section provides information about the following definitions:

- ▶ Defining VOLSER ranges for physical volumes
- ▶ Defining physical volume pools in the TS7700 tape-attached cluster
- ▶ Defining Encryption Key Server addresses

Defining VOLSER ranges for physical volumes

After a cartridge is assigned to a logical library that is associated to a TS7700T by CAPs or volser ranges, it is presented to the TS7700 tape attached cluster. The TS7700T uses the VOLSER ranges that are defined in its VOLSER Ranges table to set it to a proper category. Define the proper policies in the VOLSER Ranges table *before* inserting the cartridges into the tape library.

Note: VOLSER Ranges (or CAP) should be correctly assigned at tape library before using the tape library with IBM System z hosts. Native physical volume ranges must fall within ranges that are assigned to System z host logical libraries.

Use the window that is shown in Figure 9-164 to add, modify, and delete physical volume ranges. Unassigned physical volumes are listed in this window. If a volume is listed as unassigned volume, and this volume belongs to this TS7700, a new range should be added including that volume to fix it. If an unassigned volume does not belong to this TS7700 cluster, it should be ejected and reassigned to the proper logical library in the physical tape library.

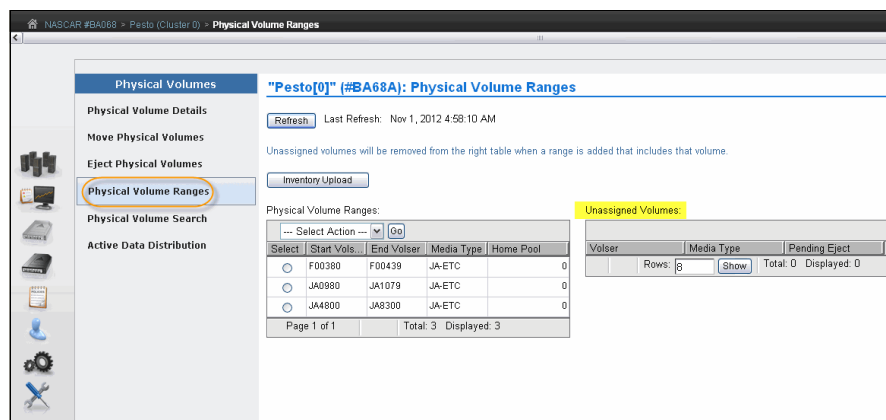


Figure 9-164 Physical Volume Ranges window

Click **Inventory Upload** to upload the inventory from the TS3500 tape library 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 specific component. The VOLSER Ranges table can be used to create a new VOLSER range, or to modify or delete a predefined VOLSER range.

Important: Operator intervention is required to resolve unassigned volumes.

Figure 9-164 on page 529 shows the status information that is displayed in the VOLSER Ranges table:

- ▶ **Start Volser:** The first VOLSER in a defined range.
- ▶ **End Volser:** The last VOLSER in a defined range.
- ▶ **Media Type:** The media type for all volumes in a certain VOLSER range. The following values are valid:
 - **JA-ETC:** ETC.
 - **JB(ETCL):** Enterprise Extended-Length Tape Cartridge.
 - **JC(ATCD):** ATCD.
 - **JD(ATDD):** ATDD.
 - **JJ(EETC):** EETC.
 - **JK(ATKE):** ATKE.
 - **JL(ATLE):** ATLE.

Use the menu in the VOLSER Ranges table to add a VOLSER range, or to modify or delete a predefined range:

- ▶ To add a VOLSER range, select **Add** from the menu. Complete the fields for the information of volume range to be added.
- ▶ To modify a predefined VOLSER range, click the radio button from the **Select** column in the same row as the name of the VOLSER range to be modified. Select **Modify** from the menu and make the changes.

Important: Modifying a predefined VOLSER range does not affect physical volumes that are already inserted and assigned to the TS7700T. Only physical volumes that are inserted after the VOLSER range modification are changed.

The VOLSER entry fields must contain 6 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 in the following manner. The VOLSER range is increased so that alphabetic characters are increased alphabetically, and numeric characters are increased numerically. For example, VOLSER range ABC000 - ABD999 results in a range of 2,000 VOLSERs (ABC000 - ABC999 and ABD000 - ABD999).

Note: VOLSER ranges that are defined at the physical tape library user interface refer exclusively to physical cartridges. *Logical Volumes* are defined only through the TS7700 MI. For more information, see “Inserting virtual volumes” on page 546.

For the TS7700, no additional definitions are required at the hardware level other than setting up the correct VOLSER ranges at the tape library.

Although now the cartridges can be inserted into the tape library, complete the required definitions at the host before inserting any physical cartridges.

Defining physical volume pools in the TS7700T

Pooling physical volume allows for enables data to be placed into separate sets of physical media, treating each media group in a specific way. For instance, there might be a need to segregate production data from test data, or encrypt part of the data. All of this can be accomplished by defining physical volume pools. Also, the reclaim parameters can be defined for each specific pool to best suit specific needs. The TS7700 MI is used for pool property definitions.

Items under Physical Volumes in the MI apply only to tape attach clusters. Trying to access those windows from a TS7720 results in the following HYDME0995E message:

This cluster is not attached to a physical tape library.

Use the window that is shown in Figure 9-165 to view or modify settings for physical volume pools, which manage the physical volumes that are used by the TS7700.

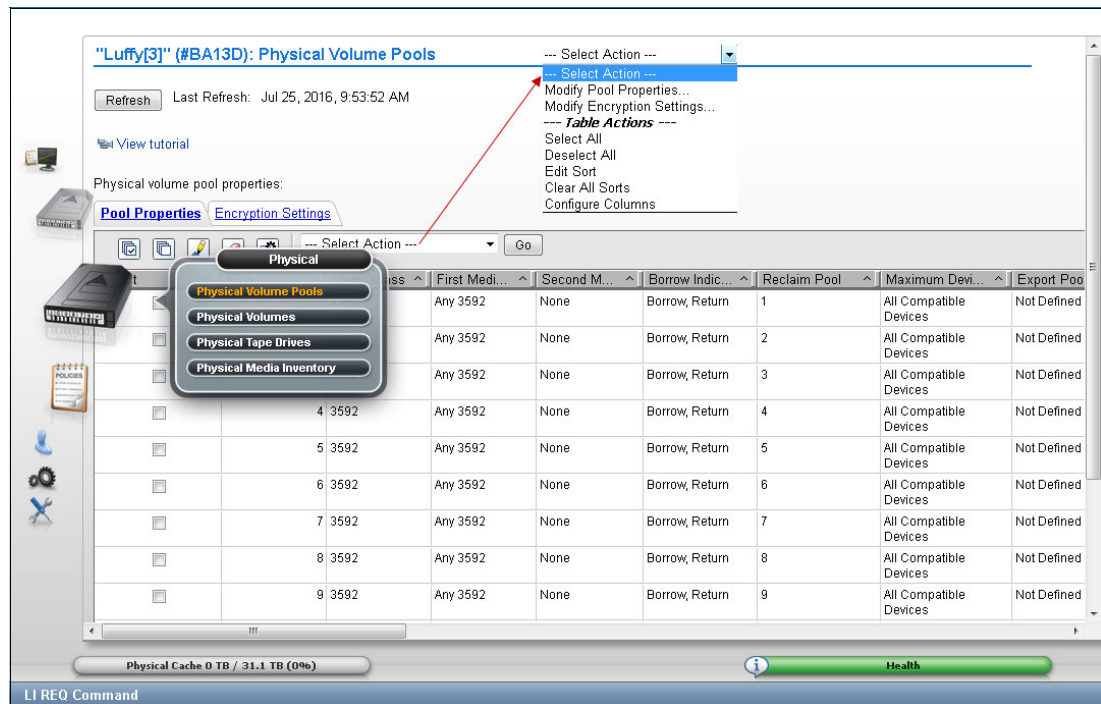


Figure 9-165 Physical Volume Pools

The Physical Volume Pool Properties table displays the encryption setting and media properties for every physical volume pool that is defined for TS7700T clusters in the grid.

Use the Physical Volume Pool Properties table to view encryption and media settings for all installed physical volume pools. To view and modify more pool properties, select a pool or pools from this table and then select either **Modify Pool Properties** or **Modify Encryption Settings** from the menu.

Tip: Pools 1 - 32 are preinstalled. Pool 1 functions as the default pool and is used if no other pool is selected.

The Physical Volume Pool Properties table displays the media properties and encryption settings for every physical volume pool that is defined for each cluster in the grid. This table contains two tabs: Pool Properties and Physical Tape Encryption Settings.

These two tabs contain the following information:

- ▶ The following information is under the **Pool Properties** tab:
 - **Pool**: Lists the pool number, which is a whole number 1 - 32, inclusive.
 - **Media Class**: Lists that 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). The following values are valid:
 - **Any 3592**. Any 3592.
 - **None**. Indicates that the pool cannot borrow or return any media to the common scratch pool. This option is valid only when the Borrow Indicator field value is no borrow, return, or no borrow, keep.
 - **JA**. ETC.
 - **JB**. EDETC.
 - **JC**. ATCD.
 - **JJ**. EETC.
 - **JK**. ATKE.
 - **JD**. ATDD.
 - **JL**. ATLE.

To modify pool properties, select the check box next to one or more pools that are listed in the Physical Volume Pool Properties table and select **Properties** from the menu. The Pool Properties table is displayed.

You can modify the fields Media Class and First Media, defined previously, and the following fields:

- **Second Media (Secondary)**: Lists the second choice of media type from which the pool can borrow. The options that are listed exclude the media type that is selected for the First Media. The following values are valid:
 - **Any 3592**. Any 3592.
 - **JA**. ETC.
 - **JB**. EDETC.
 - **JC**. ATCD.
 - **JJ**. EETC.
 - **JK**. ATKE.
 - **None**. The only option that is available if the Primary Media type is Any 3592. This option is valid only when the Borrow Indicator field value is no borrow, return, or no borrow, keep.
 - **JD**. ATDD.
 - **JL**. ATLE.
- **Borrow Indicator**: Defines how the pool is populated with scratch cartridges. The following values are valid:
 - **Borrow, Return**. A cartridge is borrowed from the common scratch pool and returned when emptied.
 - **Borrow, Keep**. A cartridge is borrowed from the common scratch pool and retained, even after being emptied.

- **No Borrow, Return.** A cartridge cannot be borrowed from the common scratch pool, but an emptied cartridge is placed in the common scratch pool. This setting is used for an empty pool.
- **No Borrow, Keep.** A cartridge cannot be borrowed from the common scratch pool, and an emptied cartridge is 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:
 - i. From the Physical Volume Pools window, click the **Pool Properties** tab.
 - ii. Select the check box next to each pool to be modified.
 - iii. Select **Modify Pool Properties** from the Physical volume pools menu.
 - iv. Click **Go** to open the Modify Pool Properties window.

The following values are valid:

- **Not Defined.** The pool is not defined as an Export pool.
- **Copy Export.** The pool is defined as a Copy Export pool.
- **Export Format:** The media format used when writing volumes for export. This function can be used when the physical library that is recovering the volumes supports a different media format than the physical library that is exporting the volumes. This field is only enabled if the value in the Export Pool field is Copy Export. The following values are valid:
 - **Default.** The highest common format that is supported across all drives in the library. It is also the default value for the Export Format field.
 - **E06.** Format of a 3592-E06 tape drive.
 - **E07.** Format of a 3592-E07 tape drive.
 - **E08.** Format of a 3592 E08 tape drive.
- **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 later, 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. Clicking to clear the check box deactivates this function.
- **Days Without Access:** Lists the number of days that the pool can persist without access to 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 later, based on an internal schedule. The valid range of possible values is 1 - 365. Clicking to clear 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 later, based on an internal schedule. The valid range of possible values is 1 - 365. Clicking to clear 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 is deactivated. The reclamation occurs later, based on an internal schedule. The valid range of possible values is 1-365. Clicking to clear the check box deactivates 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 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 reclaim free storage on a stacked volume. When the amount of active data on a physical stacked volume drops below this percentage, a reclaim operation is run on the stacked volume. The valid range of possible values is 5% - 95%. The default value is 10%. Clicking to clear the check box deactivates this function.
- ▶ The following information is under the **Physical Tape 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. The possible values are Enabled and Disabled.
 - **Key Mode 1:** Lists the encryption mode that is used with Key Label 1. The following values are valid for this field:
 - **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 EK 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, so key labels are reported by 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 that is used with Key Label 2. The following values are valid for this field:
 - **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 EK 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, so key labels are reported by 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, complete the following steps:

1. Open the Physical Volume Pools window (Figure 9-166).

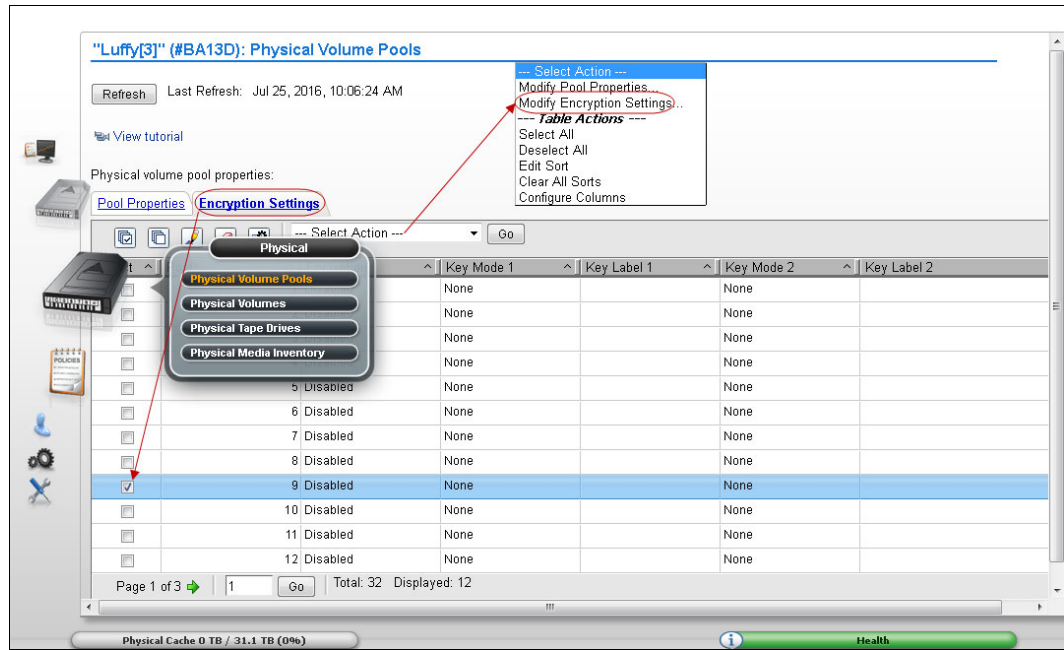


Figure 9-166 Modify encryption parameters for a pool

Tip: A tutorial is available in the Physical Volume Pools window to show how to modify encryption properties.

2. Click the **Physical Tape Encryption Settings** tab.
3. Select the check box next to each pool to be modified.
4. Click **Select Action** → **Modify Encryption Settings**.

5. Click **Go** to open the Modify Encryption Settings window (Figure 9-167).

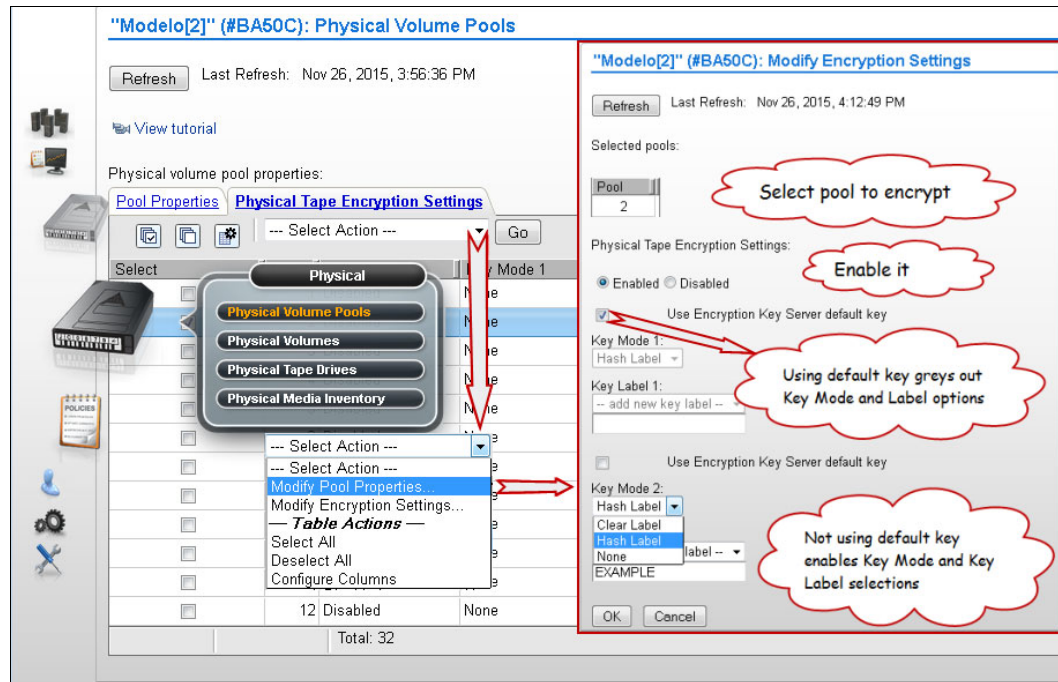


Figure 9-167 Modify encryption settings parameters

In this window, the values for any of the following controls can be modified:

– **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 that is provided by the encryption key manager.

Consideration: Your encryption key manager software must support default keys to use this option.

This check box occurs before both Key Label 1 and Key Label 2 fields. Select this check box for each label to be defined by 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. The following values are valid:

- **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 EK 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, so key labels are reported by using uppercase characters.

– **Key Mode 2**

This field is the encryption mode that is used with Key Label 2. The following values are valid:

- **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 EK 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, so key labels are reported by using uppercase characters.

6. To complete the operation, click **OK**. To abandon the operation and return to the Physical Volume Pools window, click **Cancel**.

Defining reclamation settings in a TS7700T

To optimize the use of the subsystem resources, such as processor cycles and tape drive usage, space reclamation can be inhibited during predictable busy periods and reclamation thresholds can be adjusted to the optimum point in the TS7700T through the MI. The *reclaim threshold* is the percentage that is used to determine when to run the reclamation of free space in a stacked volume.

When the amount of active data on a physical stacked volume drops below this percentage, the volume becomes eligible for reclamation. Reclamation values can be in the range of 0% - 95%, with a default value of 35%. Selecting 0% deactivates this function.

Note: Subroutines of the Automated Read-Only Recovery (ROR) process are started to reclaim space in the physical volumes. Those cartridges are made read-only momentarily during the reclaim process, returning to normal status at the end of the process.

Throughout the data lifecycle, 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 volume becomes obsolete, that space becomes a waste of capacity in that physical tape. Therefore, the active data level of that volume is decreasing over time.

TS7700T actively monitors 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 TS7700 places that volume in a candidate list for reclamation.

Reclamation copies active data from that volume to another stacked volume in the same pool. When the copy finishes and the volume becomes empty, the volume is returned to available SCRATCH status. This cartridge is now available for use and is returned to the common scratch pool or directed to the specified reclaim pool, according to the Physical Volume Pool definition.

Clarification: Each reclamation task uses two tape drives (source and target) in a tape-to-tape copy function. The TS7700 TVC is not used for reclamation.

Multiple reclamation processes can run in parallel. The maximum number of reclaim tasks is limited by the TS7700T, based on the number of available drives as shown in Table 9-13.

Table 9-13 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 by using the Physical Volume Pools window in the TS7700 MI, as shown in Figure 9-168.

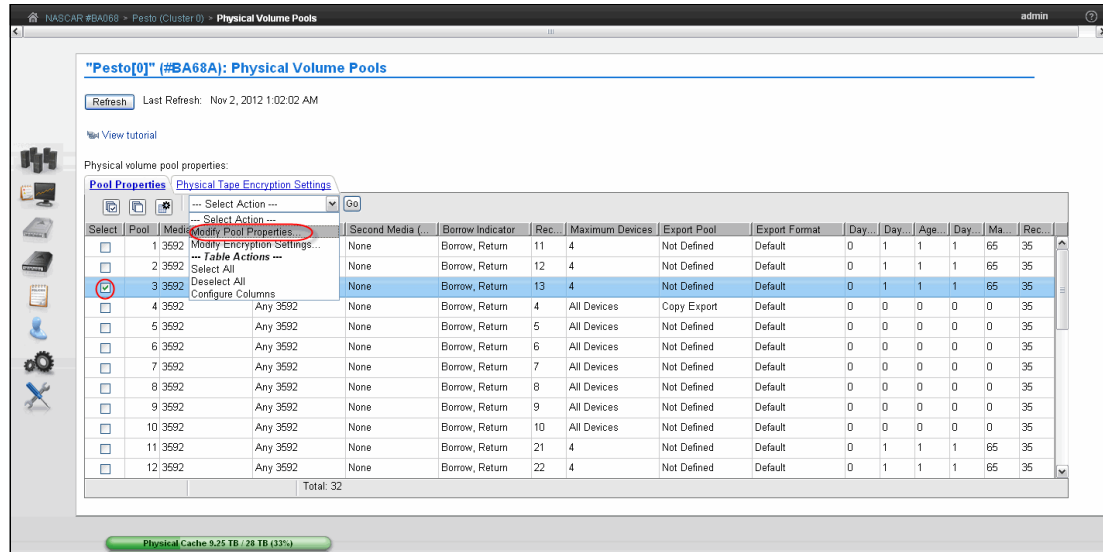


Figure 9-168 Physical Volume Pools

Select a pool and click **Modify Pool Properties** in the menu to set the reclamation level and other policies for that pool. Figure 9-169 shows the Modify Pool Properties window.

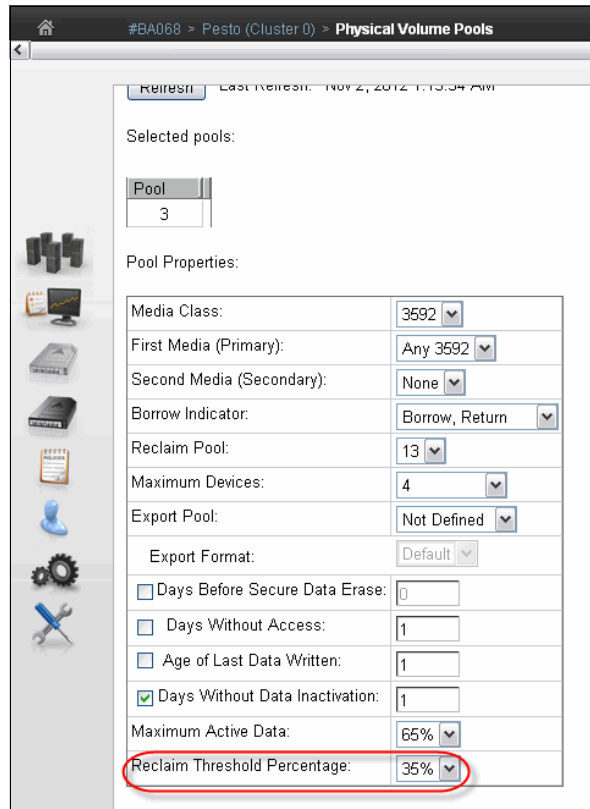


Figure 9-169 Pool properties

The example shows the borrow-return policy in effect for Pool 3, meaning that cartridges can be borrowed from the common scratch pool (and, those cartridges are returned to the CSP upon reclamation). Also, the user has defined that volumes belonging to pool 3 should reclaim into pool 13.

No more than four drives can be used for premigration in pool 3. The reclaiming threshold percentage has been set to 35%, meaning that when a physical volume in pool 3 crosses down the threshold of 35% of occupancy with active data, the stacked cartridge became candidate for reclamation. The other way to trigger a reclamation in this example is *Days Without Data Inactivation* for tape cartridges with up to 65% of occupancy level.

Reclamation enablement

To minimize any effect on TS7700 activity, the storage management software monitors resource use in the TS7700, and enables or disables reclamation. Optionally, reclamation activity can be prevented at specific times by specifying an Inhibit Reclaim Schedule in the TS7700 MI (Figure 9-170 on page 543 shows an example).

However, the TS7700T determines whether reclamation is enabled or disabled once an hour, depending on the number of available scratch cartridges. It disregards the Inhibit Reclaim Schedule if the TS7700T goes below a minimum number of scratch cartridges that are available. Now, reclamation is enforced by the tape attach TS7700 cluster.

Tip: The maximum number of Inhibit Reclaim Schedules is 14.

Using the Bulk Volume Information Retrieval (BVIR) process, the amount of active data on stacked volumes can be monitored on PHYSICAL MEDIA POOLS, helping to plan for a reasonable and effective reclaim threshold percentage. Also, Host Console Request function can be used to obtain the physical volume counts.

Although reclamation is enabled, stacked volumes 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 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 that is in the reclaim process. If this happens, reclamation is stopped for this physical tape after the current logical volume move is complete.

Pooling is enabled as a standard feature of the TS7700, even if only one pool is used. Reclamation can occur on multiple volume pools at the same time, and process 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 selects all the stacked volumes from 0% - 29% of the remaining active data. The reclaim tasks then process the volumes from least full (0%) to most full (29%) up to the defined reclaim threshold of 30%.

Individual pools can have separate reclaim policies set. The number of pools can also influence the reclamation process because the TS7740 or TS7720 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 in the following manner:

1. A pool enters a *Low scratch state* when it has access to less than 50 and two or more empty cartridges (scratch tape volumes).
2. A pool enters a *Panic scratch state* when it has access to fewer than two empty cartridges (scratch tape 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, the pool does not enter the *Low scratch state*.

Whether borrowing is configured or not, if each pool has two scratch cartridges, the Panic Reclamation mode is not entered. Panic Reclamation mode is entered when a pool has fewer than two scratch cartridges and no more scratch cartridges can be borrowed from any other pool that is defined for borrowing. Borrowing is described in “Using physical volume pools” on page 48.

Important: A physical volume pool that is running out of scratch cartridges might stop mounts in the TS7740 or TS7720T tape attach partitions, affecting host tape operations. Mistakes in pool configuration (media type, borrow and return, home pool, and so on) or operating with an empty common scratch pool might lead to this situation.

Consider that one reclaim task uses two drives for the data move, and processor cycles. When a reclamation starts, these drives are busy until the volume that is being reclaimed is empty. If the reclamation threshold level is raised too high, the result is larger amounts of data to be moved, with a resultant penalty in resources that are needed for recalls and premigration. The default setting for the reclamation threshold level is 35%.

Ideally, reclaim threshold level should be 10% - 35%. Read more about how to fine-tune this function and about the available host functions in 4.4.4, “Physical volumes for TS7740, TS7720T, and TS7760T” on page 171. Pools in either scratch state (Low or Panic state) get priority for reclamation.

Table 9-14 summarizes the thresholds.

Table 9-14 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	At least one, regardless of idle drives. If more idle drives are available, more reclaims are started, up to the maximum limit.	

Priority	Condition	Reclaim schedule honored	Active data threshold% honored	Number of concurrent reclaims	Comments
2	Priority move	Yes or No	No	At least one, regardless of idle drives. If more idle drives are available, more reclaims are started, up to the maximum limit.	<p>If a volume is within 10 days of a Secure Data Erasure and still has active data on it, it is reclaimed at this priority. An SDE priority move accepts the inhibit reclaim schedule.</p> <p>For a TS7700 MI-initiated priority move, the option to accept the inhibit reclaim schedule is given to the operator.</p>
3	Pool in Low scratch state	Yes	Yes	At least one, regardless of idle drives. If more idle drives are available, more reclaims are started, up to the maximum limit.	Volumes that are subject to reclaim because of Maximum Active Data, Days Without Access, Age of Last Data Written, and Days Without Data Inactivation 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 use priority 3 or 4 reclamation.

Tips:

- ▶ A physical drive is considered *idle* when no activity has occurred for the previous 10 minutes.
- ▶ The Inhibit Reclaim schedule is not accepted 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 must 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, the Inhibit Reclaim schedule can be used to inhibit reclaim operations for the heavy mount activity periods.

To define the Inhibit Reclaim schedule, click **Management Interface** → **Settings** → **Cluster Settings**, which opens the window that is shown in Figure 9-170.

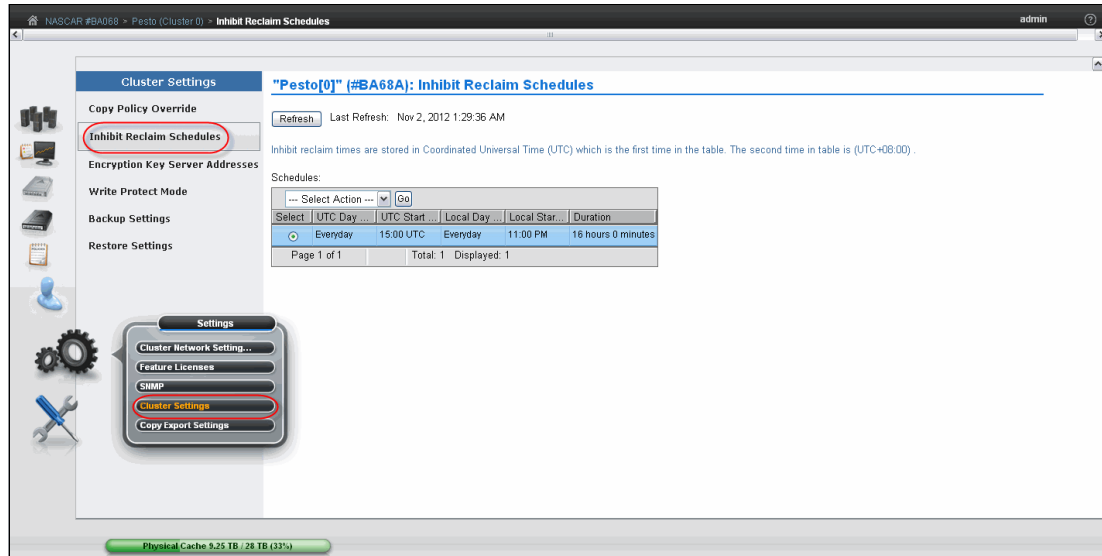


Figure 9-170 Inhibit Reclaim schedules

The Schedules table (Figure 9-171) displays the day, time, and duration of any scheduled reclamation interruption. All inhibit reclaim dates and times are first displayed in Coordinated Universal Time and then in local time. Use the menu on the Schedules table to add a new Reclaim Inhibit Schedule, or modify or delete an existing schedule, as shown in Figure 9-171.

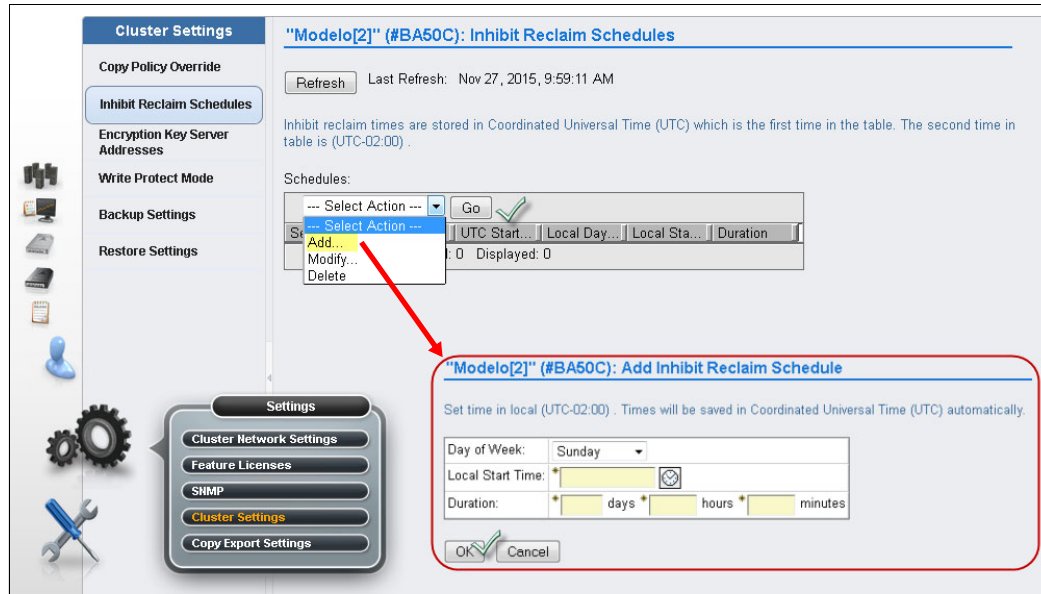


Figure 9-171 Add Inhibit Reclaim schedule

Defining Encryption Key Server addresses

Set the EKS addresses in the TS7700 cluster(Figure 9-172).



Figure 9-172 Encryption Key Server Addresses

To watch a tutorial that shows the properties of encryption key management, click the **View tutorial** link.

The EKS assists encryption-enabled tape drives in generating, protecting, storing, and maintaining EKs that are used to encrypt information being written to and decrypt information being read from tape media (tape and cartridge formats). Also, EKS manages the EK for the TVC cache disk subsystem, with the external key management disk encryption feature installed. This removes the responsibility of managing the key away from the 3957-Vxx and from the disk subsystem controllers.

Note: The settings for Encryption Server are shared for both tape and external disk encryption.

The following settings are used to configure the TS7740 or TS7720T connection to an EKS (Figure 9-172):

- ▶ **Primary key server address:** The key server name or IP address that is primarily used to access the EKS. This address can be a fully qualified host name or an IP address in IPv4 or IPv6 format. This field is not required if you do not want to connect to an EKS.

A valid IPv4 address is 32 bits and consists of four decimal numbers, each ranging 0 - 255, separated by periods, for example:

98.104.120.12

A valid IPv6 address is a 128-bit hexadecimal value separated into 16-bit fields by colons, for example:

3afa:1910:2535:3:110:e8ef:ef41:91cf

Leading zeros can be omitted in each field so that `:0003:` can be written as `:3:`. A double colon (`::`) can be used once per address to replace multiple fields of zeros. For example, this address:

```
3afa:0:0:0:200:2535:e8ef:91cf
```

can be written as:

```
3afa::200:2535:e8ef:91cf
```

A *fully qualified host name* is a domain name that uniquely and absolutely names a computer. It consists of the host name and the domain name. The domain name is one or more domain labels that place the computer in the domain name server (DNS) naming hierarchy. The host name and the domain name labels are separated by periods and the total length of the host name cannot exceed 255 characters.

- ▶ **Primary key server port:** The port number of the primary key server. Valid values are any whole number 0 - 65535; the default value is 3801. This field is only required if a primary key address is used.

- ▶ **Secondary key server address:** The key server name or IP address that is used to access the EKS when the primary key server is unavailable.

This address can be a fully qualified host name or an IP address in IPv4 or IPv6 format. This field is not required if you do not want to connect to an EKS.

See the primary key server address description for IPv4, IPv6, and fully qualified host name value parameters.

- ▶ **Secondary key manager port:** The port number of the secondary key server. Valid values are any whole number 0 - 65535; the default value is 3801. This field is only required if a secondary key address is used.

- ▶ **Using the Ping Test:** Use the **Ping Test** buttons to check cluster network connection to a key server after changing a cluster's address or port. If you change a key server address or port and do not submit the change before using the **Ping Test** button, you receive the following warning:

To perform a ping test you must first submit your address and/or port changes.

After the ping test has been started, you will receive one of two following messages:

- The ping test against the address "`<address>`" on port "`<port>`" was successful.
- The ping test against the address "`<address>`" on port "`<port>`" from "`<cluster>`" has failed. The error returned *weatherworn text*.

Click **Submit Changes** to save changes to any of these settings.

Consideration: The two EKSs must be set up on separate systems to provide redundancy. Connection to a key manager is required to read encrypted data.

9.3.3 TS7700 definitions

This section describes the basic TS7700 definitions.

Inserting virtual volumes

Use the Insert Virtual Volumes window (Figure 9-173) to insert a range of logical volumes in the TS7700 grid. Logical volumes that are inserted into an individual cluster are available to all clusters within a grid configuration.

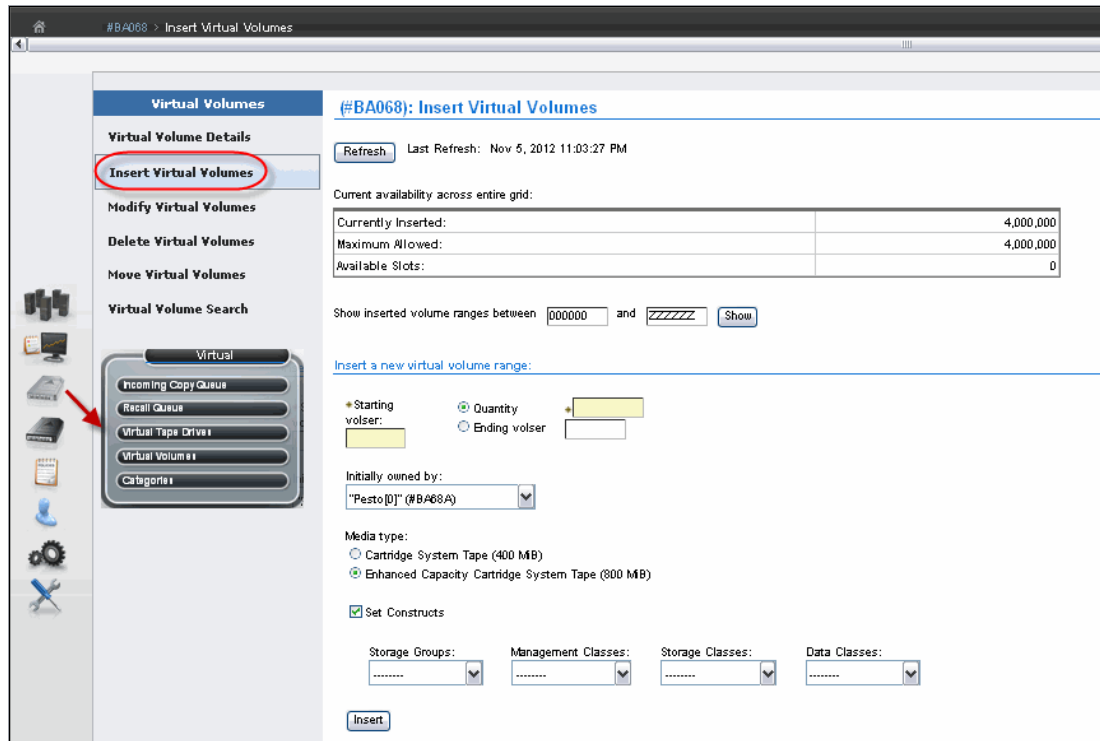


Figure 9-173 TS7700 MI Insert Virtual Volumes window

During logical volume entry processing on z/OS, even if the library is online and operational for a specific host, at least one device must be online (or been online) for that host for the library to send the volume entry attention interrupt to that host. If only the library is online and operational, but there are no online devices to a specific host, that host does not receive the attention interrupt from the library unless a device previously was varied online.

To work around this limitation, ensure that at least one device is online (or 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 only one host is attached to the library that owns the volumes being entered. In general, after the volumes are entered into the library, CBR36xxI cartridge entry messages are expected. The **LIBRARY RESET, CBRUXENT** command from z/OS can be used to reinitiate cartridge entry processing, if necessary. This command causes the host to ask for any volumes in the insert category.

Up to now, as soon as OAM starts for the first time, and being the volumes in the Insert category, the entry processing starts, not allowing for operator interruptions. The **LI DISABLE, CBRUXENT** command can be used before starting the OAM address space. This approach allows for the entry processing to be interrupted before the OAM address space initially starts.

The table at the top of Figure 9-173 on page 546 shows the current information about the number of logical volumes in the TS7700:

- ▶ **Currently Inserted:** The total number of logical volumes that are inserted into the TS7700.
- ▶ **Maximum Allowed:** The total maximum number of logical volumes that can be inserted.
- ▶ **Available Slots:** The available slots that are 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 Grid, enter a logical volume range and click **Show**.

To insert a new logical volume range action, use the following fields:

- ▶ **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 number of logical volumes, beginning with the Starting VOLSER. Enter the quantity of logical volumes to be inserted in the adjacent field. Up to 10,000 logical volumes can be inserted at one time.
- ▶ **Ending VOLSER:** Select this option to insert a range of logical volumes. Enter the ending VOLSER number in the adjacent field.
- ▶ **Initially owned by:** Indicates the name of the cluster that owns the new logical volumes. Select a cluster from the menu.
- ▶ **Media type:** Indicates the media type of the logical volume (volumes). The following values are valid:
 - 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 menu under each construct to select a predefined construct name. The following constructs can be specified:
 - **Storage Group**
 - **Storage Class**
 - **Data Class**
 - **Management Class**

Important: 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 and written from the load point.

To insert a range of logical volumes, complete the following steps:

1. Complete the fields that are listed and click **Insert**. There is a prompt to confirm the decision to insert logical volumes.
2. To continue with the insert operation, click **Yes**. To abandon the insert operation without inserting any new logical volumes, click **No**.

Note: Up to 10,000 logical volumes can be inserted at one time. This applies to both inserting a range of logical volumes and inserting a quantity of logical volumes.

Defining scratch categories

You can use the TS7700 MI to add, delete, or modify a scratch category of virtual volumes. All scratch categories that are defined by using the TS7700 MI inherit the Fast Ready attribute.

Note: The Fast Ready attribute provides a definition of a category to supply scratch mounts. For z/OS, it depends on the definitions. The TS7700 MI provides a way to define one or more scratch (Fast Ready) categories. Figure 9-174 shows the Categories window. A scratch category can be added by using the **Add Scratch Category** menu.

The **MOUNT FROM CATEGORY** command is not exclusively used for scratch mounts. Therefore, the TS7700 cannot assume that any **MOUNT FROM CATEGORY** is for a scratch volume.

When defining a scratch category, an expiration time can be set up, and further define it as an Expire Hold time.

The 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 the IEASYSxx member to be activated.

Tip: Do not add a scratch category by using MI that was previously designated as a *private volume* category at the host. Categories should correspond to the defined categories in the DEVSUPxx from the attached hosts.

Categories	Owning Cluster	Counts	Scratch Expired
Scratch		97390	97170
0062 (No expiration)		220	0
0162 (No expiration)		0	0
Virtual		0	0
Cache Partitions		33	33
Incoming Copy Queue		6808	6808
Recall Queue		0	0
Virtual Tape Drives		0	0
Virtual Volumes		90329	90329
Categories		0	0
0802 (No expiration)		0	0
0902 (1 Hours No Hold)		0	0
Private		9543	
000F		8812	
010F		1	
002F		356	
000E		312	
...		~	

Figure 9-174 Categories

To add, modify, or delete a scratch category of virtual volumes, use the window in Figure 9-174. This window can also be used to view total volumes that are defined by custom, inserted, and damaged categories. The Categories table uses the following values and descriptions:

► **Categories:**

– **Scratch**

Categories within the user-defined private range 0x0001 through 0xEFFF that are defined as scratch (Fast Ready).

– **Private**

Custom categories that are established by a user, within the range of 0x0001 through 0xEFFF.

– **Damaged**

A system category that is identified by the number 0xFF20. Virtual volumes in this category are considered damaged.

– **Insert**

A system category that is identified by the number 0xFF00. Inserted virtual volumes are held in this category until moved by the host into a scratch category.

▶ **Owning Cluster**

Names of all clusters in the grid.

▶ **Counts**

The total number of virtual volumes according to category type, category, or owning cluster.

▶ **Scratch Expired**

The total number of scratch volumes per owning cluster that are expired. The total of all scratch expired volumes is the number of ready scratch volumes.

Number of virtual volumes: The addition of all volumes counts that are shown in the Counts column do not always result in the total number of virtual volumes due to some rare, internal categories not being displayed on the Categories table. Additionally, movement of virtual volumes between scratch and private categories can occur multiple times per second and any snapshot of volumes on all clusters in a grid is obsolete by the time a total count completes.

The Categories table can be used to add, modify, and delete a scratch category, and to change the way that information is displayed.

Figure 9-174 on page 548 shows the Add Category window, which you can open by selecting **Add Scratch Categories** as shown in Figure 9-175.

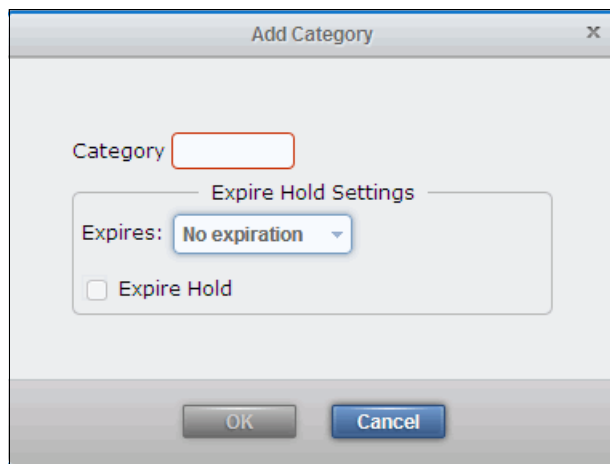


Figure 9-175 Scratch Categories - Add Category

The Add Category window shows these fields:

► **Category**

A four-digit hexadecimal number that identifies the category. The following characters are valid characters for this field:

A-F, 0-9

Important: Do 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**

The amount of time after a virtual volume is returned to the scratch category before its data content is automatically delete-expired.

A volume becomes a candidate for delete-expire after all the following conditions are met:

- The amount of time since the volume entered the scratch category is equal to or greater than the Expire Time.
- The amount of time since the volume's record data was created or last modified is greater than 12 hours.
- At least 12 hours has passed since the volume was migrated out of or recalled back into disk cache.

Note: If No Expiration is selected, volume data never automatically delete-expires.

► **Set Expire Hold**

Select this box to prevent the virtual volume from being mounted or having its category and attributes changed before the expire time elapses.

Selecting this field activates the hold state for any volumes currently in the scratch category and for which the expire time has not yet elapsed. Clearing this field removes the access restrictions on all volumes currently in the hold state within this scratch category.

Note: Trying to mount a non-expired volume that belongs to a scratch category with Expire Hold on results in an error.

If **Expire Hold** is set, the virtual volume cannot be mounted during the expire time duration and is excluded from any scratch counts surfaced to the System z host. The volume category can be changed, but only to a private category, allowing accidental scratch occurrences to be recovered to private.

If **Expire Hold** is not set, then the virtual volume can be mounted or have its category and attributes changed within the expire time duration. The volume is also included in scratch counts surfaced to the System z hosts.

Tip: Add a comment to DEVSUPnn to ensure that the scratch categories are updated when the category values in DEVSUPnn are changed. They always need to be in sync.

Defining the logical volume expiration time

The expiration time is defined from the MI window that is shown in Figure 9-175 on page 549. If the Delete Expired Volume Data setting is not used, logical volumes that have been returned to scratch are still considered active data, allocating physical space in tape cartridges on the tape attach TS7700. In that case, rewriting only this logical volume expires the old data, enabling physical space that is occupied by old data to be reclaimed later.

With the Delete Expired Volume Data setting, the data that is associated with volumes that have been returned to scratch are expired after a specified time period and their physical space in tape can be reclaimed.

The parameter **Expire Time** specifies the amount of time in hours, days, or weeks. The data continues to be managed by the TS7700 after a logical volume is returned to scratch before the data that is associated with the logical volume is deleted. A minimum of 1 hour and a maximum of 32,767 hours (approximately 195 weeks) can be specified.

Remember:

- ▶ Scratch categories are global settings within a multi-cluster grid. Therefore, each defined scratch category and the associated Delete Expire settings are valid on each cluster of the grid.
- ▶ The Delete Expired Volume Data setting applies also to disk only clusters. If it is not used, logical volumes that have been returned to scratch are still considered active data, allocating physical space in the TVC. Therefore, setting an expiration time on a disk only TS7700 is important to maintain an effective cache usage by deleting expired data.

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 that is defined as scratch.

Defining TS7700 constructs

To use the Outboard Policy Management functions, four constructs must be defined:

- ▶ 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 that are defined for each construct are performed by the TS7700. For non-z/OS hosts, the constructs can be manually assigned to logical volume ranges.

Storage Groups

On the z/OS host, the SG construct determines into which tape library a logical volume is written. Within the TS7700T, the SG construct defines the storage pool to which the logical volume is placed.

Even before the first SG is defined, there is always at least one SG present. This is the default SG, which is identified by eight dashes (-----). This SG cannot be deleted, but it can be modified to point to another storage pool. Up to 256 SGs, including the default, can be defined.

Use the window that is shown in Figure 9-176 to add, modify, and delete an SG used to define a primary pool for logical volume premigration.




Figure 9-176 Storage Groups

The SGs table displays all existing SGs available for a selected cluster.

The SGs table can be used to create an SG, modify an existing SG, and delete an SG. The following status information is listed in the SGs table:

- ▶ **Name:** The name of the SG
 - Each SG within a cluster must have a unique name. The following characters are valid for this field:
 - **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 can be selected. If the cluster does not possess a physical library, this column is not visible, and the MI categorizes newly created SGs by using pool 1.
- ▶ **Description:** A description of the SG

Use the menu in the SGs table to add an SG, or to modify or delete an existing SG. To add an SG, complete the following steps:

1. Select **Add** from the menu.
2. Complete the fields for the information that is displayed in the SGs table.

Note: If the cluster is not attached to a physical library, the Primary Pool field is not available in the **Add** or **Modify** menu options.

To modify an existing SG, complete the following steps:

1. Click the radio button from the Select column that appears next to the name of the SG to be modified.
2. Select **Modify** from the menu.
3. Complete the fields for information that shows in the SGs table.

To delete an existing SG, complete the following steps:

1. Select the button in the Select column next to the name of the SG to be deleted.
2. Click **Delete** from the menu.
3. There is a prompt to confirm the decision to delete an SG. If **Yes** is selected, the SG is deleted. **No** discards the request.

Important: Do not delete any existing SG if there are still logical volumes assigned to that SG.

Management Classes

Dual copy for a logical volume within the same TS7700T can be defined in the Management Classes window. In a grid configuration, a typical choice is to copy logical volumes over to the other cluster rather than creating a second copy in the same TS7700T.

However, in a stand-alone configuration, the dual copy capability can be used to protect against media failures. The second copy of a volume can be in a pool that is designated as a Copy Export pool. For more information, see 2.3.32, “Copy Export” on page 90.

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 MC is always available. It is identified by eight dashes (-----) and cannot be deleted. You can define up to 256 MCs, including the default. Use the window that is shown in Figure 9-177 on page 554 to define, modify, or delete the MC that defines the TS7700 copy policy for volume redundancy.

The Current Copy Policy table displays the copy policy in force for each component of the grid. If no MC is selected, this table is not visible. You must select an MC from the MCs table to view copy policy details.

Figure 9-177 shows the MCs table.

Select	Name	Secondary	Description	Scratch mo...	Retain copy ...	"Pesto[0]" (...)	"Squint[1]" (...)	"Celeste[2]"...	"Tom[3]" (#...	"Spike[5]" (#BA68F)
<input type="checkbox"/>	-----	0	The default Management Class	C0, C1, C2, C3, C5	No	Rewind Unload (RUN)	No Copy	No Copy	No Copy	No Copy
<input type="checkbox"/>	DDDI	0	DDDI	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	Deferred	Rewind Unload (RUN)
<input type="checkbox"/>	DDI	0	DDI	C0, C1, C2, C5	No	Deferred	Deferred	Deferred	Deferred	No Copy
<input type="checkbox"/>	DDNI	0	DDNI	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	Deferred	Rewind Unload (RUN)
<input type="checkbox"/>	DDIN	0	DDIN	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	Rewind Unload (RUN)	No Copy
<input type="checkbox"/>	DDIINN	0	DDIINN	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	Rewind Unload (RUN)	No Copy
<input type="checkbox"/>	DDIINN	0	DDIINN	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	Rewind Unload (RUN)	No Copy
<input type="checkbox"/>	DDNINN	4	DDNINN	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	No Copy	No Copy
<input type="checkbox"/>	DDNNIN	4	DDNNIN	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	No Copy	Rewind Unload (RUN)
<input type="checkbox"/>	DDNNIN	0	DDNNIN	C0, C1, C2, C3, C5	No	Deferred	Deferred	Deferred	No Copy	No Copy

Figure 9-177 Management Classes

The MCs table (Figure 9-177) displays defined MC copy policies that can be applied to a cluster. You can use the MCs table to create a new MC, modify an existing MC, and delete one or more existing MCs. The default MC can be modified, but cannot be deleted. The default name of the MC uses eight dashes (-----).

The following status information is displayed in the MCs table:

- ▶ **Name:** The name of the MC
Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field cannot be a number. This is the only field that cannot 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 is not visible, and the MI categorizes the newly created SGs by using pool 0.
- ▶ **Description:** A description of the MC definition
The value in this field must be 1 - 70 characters in length.
- ▶ **Scratch Mount Candidate**
The cluster or clusters that are the candidates for scratch mounts. Clusters that are displayed in this field are selected first for scratch mounts of the volumes that are associated with the MC. If no clusters are displayed, the scratch mount process remains a random selection routine that includes all available clusters. For more information, see "Defining scratch mount candidates" on page 574.
- ▶ **Retain Copy Mode (Yes or No)**
Retain Copy mode accepts the original Copy Consistency Policy that is in place in the cluster where the volume was created. This mode prevents unwanted copies from being created throughout the grid. For more information, see Chapter 2, "Architecture, components, and functional characteristics" on page 15.

The Cluster Copy Policy enables you to define where and when copies are made.

Use the menu in the MCs table to add, modify, or delete MCs.

To add an MC, select **Add** from the menu and click **Go**. Complete the fields for information that you want displayed in the MCs table. You can create up to 256 MCs per TS7700 Grid.

Tip: If cluster is not attached to a physical library, the Secondary Pool field is not available in the **Add** option.

The **Copy Action** menu is next to each cluster in the TS7700 Grid. Use the **Copy Action** menu to select, for each component, the copy mode to use in volume duplication. The following actions are available from this menu:

- ▶ **No Copy:** No volume duplication occurs if this action is selected.
- ▶ **RUN:** Volume duplication occurs when the **Rewind Unload** command is received. The command returns only after the volume duplication completes successfully.
- ▶ **Deferred:** Volume duplication occurs later based on the internal schedule of the copy engine.
- ▶ **Synchronous Copy:** Provides tape copy capabilities up to synchronous-level granularity across two clusters within a multi-cluster grid configuration. For more information about Synchronous mode copy settings and considerations, see “Synchronous mode copy” on page 81.
- ▶ **Time Delayed:** Volume duplication will occur only after the delay time that is specified by the user elapses. This option is only available if all clusters in the grid are running R3.1 or higher level of code.

See “Management Classes window” on page 439 in this chapter for more information about this topic.

Storage Classes

By using the SC construct, you can influence when a logical volume is removed from cache, and assign Cache Partition Residency for logical volumes in a TS7700T cluster.

A default SC is always available. It is identified by eight dashes (-----) and cannot be deleted. Use the window that is shown in Figure 9-178 to define, modify, or delete an SC that is used by the TS7700 to automate storage management through the classification of data sets and objects.

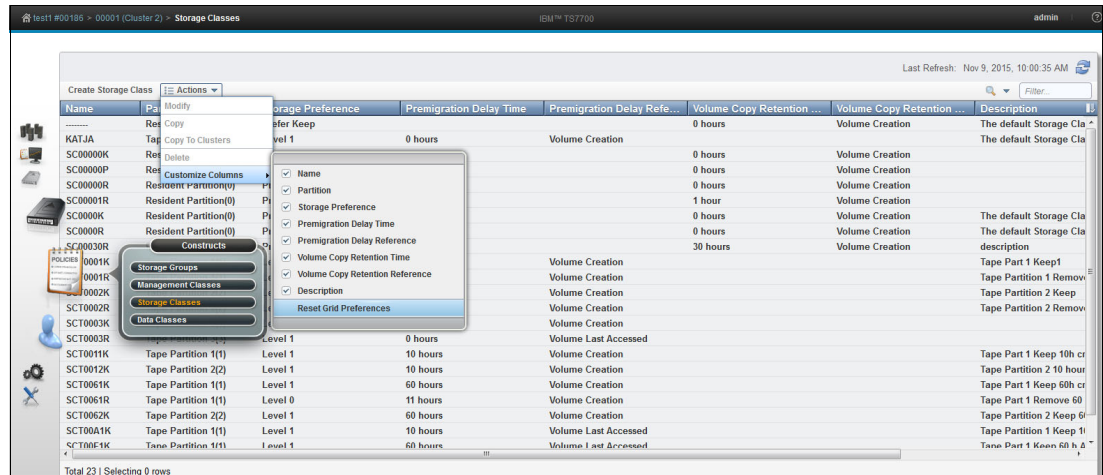


Figure 9-178 Storage Classes window on a TS7700

The SCs table displays defined SCs available to CDSs and objects within a cluster. Although SCs are visible from all TS7700 clusters, only those clusters that are attached to a physical library can alter TVC preferences. A stand-alone TS7700 cluster that does not possess a physical library does not remove logical volumes from the tape cache, so the TVC preference for the disk-only clusters is always Preference Level 1.

Use the SCs table to create an SC, or modify or delete an existing SC. The default SC can be modified, but cannot be deleted. The default SC uses eight dashes as the name (-----).

The following status information is displayed in the SCs table:

- ▶ **Name:** The name of the SC.
Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field might not be a number. The value in this field must be 1 - 8 characters.
- ▶ **Partition:** The name of the partition that is associated with the SC. A partition must be active before it can be selected as a value for this field. This field is displayed only if the cluster is a TS7720 that is attached to a physical library.
- ▶ **Storage Preference:** The preference level for the SC.

Note: A dash (-) indicates that the SC contains a partition that was deleted. Any volumes that are assigned to go to the deleted partition are redirected to the primary partition.

This setting determines how soon volumes are removed from cache after they are copied to tape. This information can be modified only if the selected cluster possesses a physical library. If the selected cluster is a disk-only TS7720, volumes in that cluster's cache display a Level 1 preference. The following values are valid:

– **Use IART**

Volumes are removed according to the IART of the TS7700.

– **Level 0**

Volumes are removed from the TVC when they are copied to tape.

– **Level 1**

Copied volumes remain in the TVC until more space is required. Then, they are the first volumes that are removed to free space in the cache. This is the default preference level that is assigned to new preference groups.

- ▶ **Volume Copy Retention Group:** The name of the group that defines the preferred Auto Removal policy that is applicable to the logical volume.

The Volume Copy Retention Group provides more options to remove data from a TS7720 (disk-only) and for data in Cache Partition 0 (CP0) in a TS7720T, as the active data reaches full capacity. Volumes become candidates for removal if an appropriate number of copies exist on peer clusters *and* the volume copy retention time has elapsed since the volume was last accessed.

Volumes in each group are removed in order based on their least recently used (LRU) access times. The volume copy retention time describes the number of hours that a volume remains in cache before becoming a candidate for removal.

This field is displayed only if the cluster is a disk-only cluster that is part of a grid. A hybrid grid combines TS7700 clusters that both attach (TS7740 or TS7720T) and do not attach (TS7720) to a physical library. If the logical volume is in a scratch (Fast Ready) category and is on a disk-only cluster, retention settings no longer apply to the volume, which becomes a top priority candidate for removal. In this instance, the value that is displayed for the Volume Copy Retention Group is accompanied by a warning icon.

The following list describes the group types:

– **Prefer Remove**

Removal candidates in this group are removed after scratch-volume candidates have been exhausted.

– **Prefer Keep**

Removal candidates in this group are removed after removal of candidates in the Prefer Remove group.

– **Pinned**

Copies of volumes in this group are never removed from the accessing cluster. The volume copy retention time does not apply to volumes in this group. Therefore, volumes in this group that are moved to scratch become priority candidates for removal.

Important: Care must be taken when assigning volumes to this group to avoid cache overruns.

- ▶ **Volume Copy Retention Time:** The minimum amount of time (in hours) after a logical volume copy was last accessed that the copy can be removed from cache.

The copy is said to be expired after this time has passed. The copy then becomes a candidate for removal. Possible values include any values 0 - 65,536. The default is 0.

If the Volume Copy Retention Group has a value of Pinned, this field is disabled.

- ▶ **Premigration Delay Time**

The number of hours until premigration can begin for volumes in the SC, based on the volume time stamp that is designated by Premigration Delay Reference. Possible values are 0 - 65535. If 0 is selected, premigration delay is disabled. This field is visible only for TS7700 Clusters that are attached to a physical library.

- ▶ **Premigration Delay Reference**

The volume operation that establishes the time stamp from which Premigration Delay Time is calculated. This field is visible only for TS7700 Clusters that are attached to a physical library. Possible values include:

- **Volume Creation:** The time at which the volume was created by a scratch mount or write operation from beginning of tape.
- **Volume Last Accessed:** The time at which the volume was last accessed.

- ▶ **Volume Copy Retention Reference**

The volume operation that establishes the time stamp from which Volume Copy Retention Time is calculated. Possible values include:

- **Volume Creation:** The time at which the volume was created by a scratch mount or write operation from beginning of tape.
- **Volume Last Accessed:** The time at which the volume was last accessed.

If the Volume Copy Retention Group displays a value of Pinned, this field is disabled.

- ▶ **Description:** A description of the SC definition. The value in this field must be 0 - 70 characters.

To add an SC, or modify or delete an existing SC, use the menu in the SCs table.

To add an SC, select **Add** from the menu. Complete the fields for the information that is displayed in the SCs table. You can create up to 256 SCs per TS7700 Grid.

To modify an existing SC, click the radio button from the Select column that appears in the same row as the SC that you want to modify. Select **Modify** from the menu. Of the fields that are listed in the SCs table, you can change all of them except for the SC name.

To delete an existing SC, click the radio button from the Select column that appears in the same row as the SC that you want to delete. Select **Delete** from the menu. A dialog box opens where you confirm the SC deletion. Select **Yes** to delete the SC, or select **No** to cancel the delete request.

Important: Do not delete any existing SC if there are still logical volumes assigned to this SC.

See “Storage Classes window” on page 442 in this chapter for more details about this topic.

Data Classes

From a z/OS perspective (SMS-managed tape), the DFSMS DC defines the following information:

- ▶ Media type parameters
- ▶ Recording technology parameters
- ▶ Compaction parameters

For the TS7700, only the Media type, Recording technology, and Compaction parameters are used. The use of larger logical volume sizes is controlled through DC. A default DC is always available. It is identified by eight dashes (-----) and cannot be deleted.

Use the window that is shown in Figure 9-179 to define, modify, or delete a TS7700 DC. The DC is used to automate storage management through the classification of data sets.

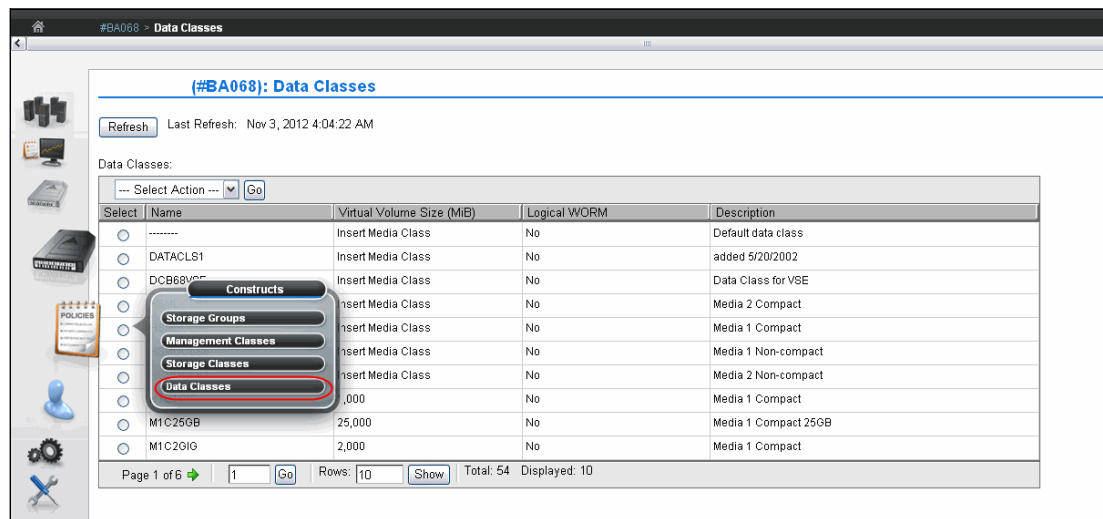


Figure 9-179 Data Classes window

The DC table (Figure 9-179 on page 558) displays the list of DCs defined for each cluster of the grid.

You can use the DCs table to create a DC, or modify or delete an existing DC. The default DC can be modified, but cannot be deleted. The default DC shows the name as eight dashes (-----).

The following status information is displayed in the DCs table:

► **Name:** The name of the DC

Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field cannot be a number. The value in this field must be 1 - 8 characters in length.

► **Virtual Volume Size (mebibytes, MiB):** The logical volume size of the DC

This setting determines the maximum number of MiB for each logical volume in a defined class. The following values are valid:

– **Insert Media Class**

The logical volume size is not defined, so the DC is not defined by a maximum logical volume size. You can use 1,000 MiB, 2,000 MiB, 4,000 MiB, 6,000 MiB, or 25,000 MiB.

Rules: Support for 25,000 MiB logical volumes is allowed without any restriction if all TS7700 clusters in the grid operate at Release 3.3 or higher.

25000 MiB is not supported in mixed code-level grids (with a member earlier than Release 3.2) with one or more TS7740 clusters present in the grid.

For disk-only grids (no tape-attached member), Feature Code 0001 is required in each TS7720 operating at levels earlier than Release 3.2.

► **LWORM**

It specifies whether LWORM is set for the DC. LWORM is the virtual equivalent of WORM tape media, achieved through software emulation.

The following values are valid for this field:

– **Yes**

LWORM is set for the DC. Volumes belonging to the DC are defined as LWORM.

– **No**

LWORM is not set. Volumes belonging to the DC are not defined as LWORM. This is the default value for a new DC.

► **Description:** A description of the DC definition

The value in this field must be 0 - 70 characters in length.

Use the menu on the DCs table to add a DC, or modify or delete an existing DC. Figure 9-180 shows the Add Data Class window.

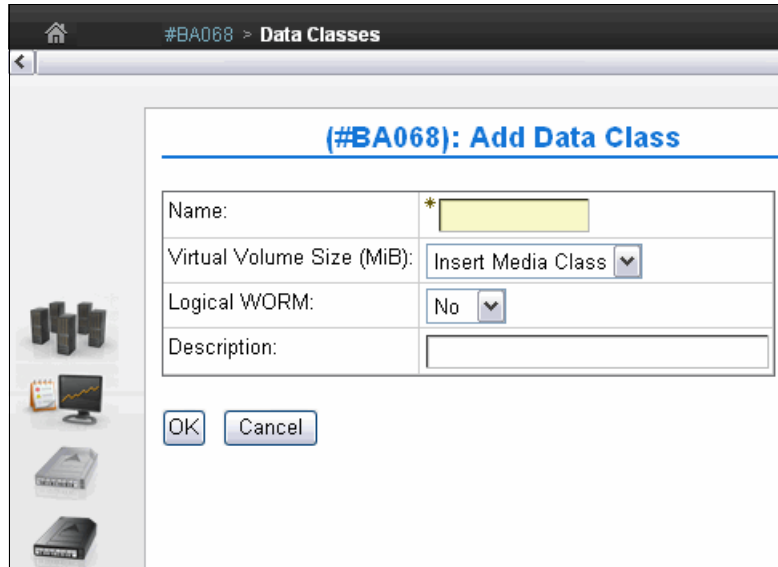


Figure 9-180 Add Data Class window

To add a DC, complete the following steps:

1. Select **Add** from the menu, and click **Go**.
2. Complete the fields for the information that is displayed in the DCs table.

Tip: You can create up to 256 DCs per TS7700 Grid.

To modify an existing DC, complete the following steps:

1. Select the check box in the Select column that appears in the same row as the DC that you want to modify.
2. Select **Modify** from the menu and click **Go**. Of the fields that are listed in the DCs table, you can change all of them except the default DC name.

To delete an existing DC, complete the following steps:

1. Click the radio button from the Select column that appears in the same row as the DC that you want to delete.
2. Select **Delete** from the menu and click **Go**. A dialog box opens where you can confirm the DC deletion. Select **Yes** to delete the DC, or select **No** to cancel the delete request.

Important: Do not delete a DC if there are still logical volumes that are assigned to it.

Activating a TS7700 license key for a new Feature Code

This section describes how to view information about, activate, or remove the following feature licenses from the TS7700 cluster:

- ▶ Peak data throughput increments
- ▶ Logical volume increments
- ▶ Cache enablement
- ▶ Grid enablement
- ▶ SDAC enablement

- ▶ Encryption configuration enablement
- ▶ Dual port grid connection enablement
- ▶ Specific RPQ enablement
- ▶ Maximum amount of queued premigration content

Clarification: *Cache enablement* license key (FC5267) applies only to a TS7740 cluster, where *Enable 1 Tb Pending Tape Capacity* (FC5274) applies only to TS7700T tape attach configuration.

The amount of disk cache capacity and performance capability is enabled by using feature license keys. You receive feature license keys for the features that you have ordered. Each feature increment enables you to tailor the subsystem to meet your disk cache and performance needs.

Use the Feature Licenses window (Figure 9-181) to activate feature licenses in the TS7700. To open the window, complete the following steps:

1. Select **Activate New Feature License** from the list and click **Go**.
2. Enter the license key into the fields that are provided and select **Activate**.

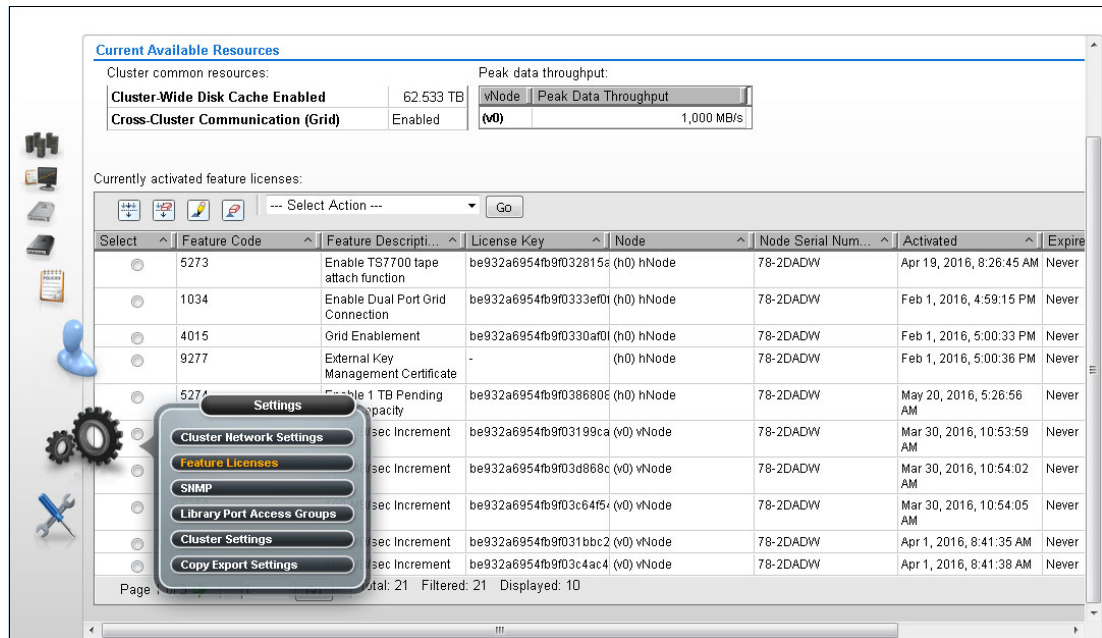


Figure 9-181 Feature Licenses window

To remove a license key, complete the following steps:

1. Select the feature license to be removed.
2. Select **Remove Selected Feature License** from the list, and click **Go**.

Important: Do not remove any installed peak data throughput features because removal can affect host jobs.

Some feature codes are not removable after being installed.

When you select **Activate New Feature License**, the Feature License entry window opens, as shown in Figure 9-182. When you enter a valid feature license key and click **Activate**, the feature is activated.

Tip: Performance Increments become active immediately. Others, such as Cache Increments or FICON dual port enablement, take 10 - 30 minutes to become active.

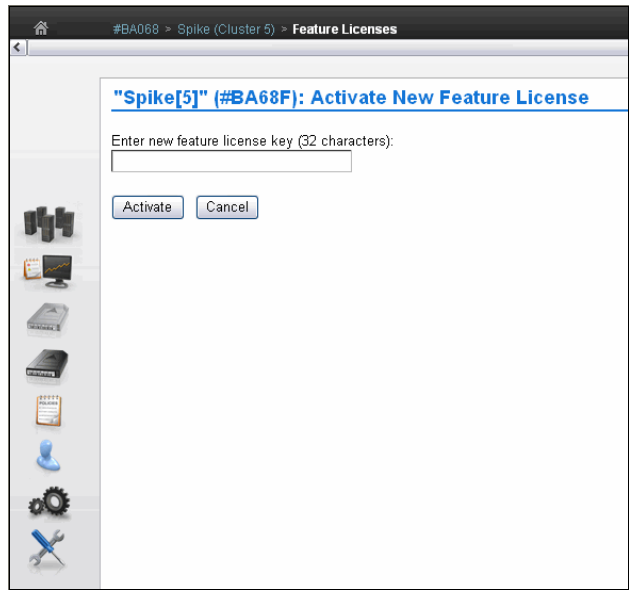


Figure 9-182 Activate New Feature Licenses window

Defining Simple Network Management Protocol

Use the window that is shown in Figure 9-183 to view or modify the SNMP configured on a TS7700 Cluster.

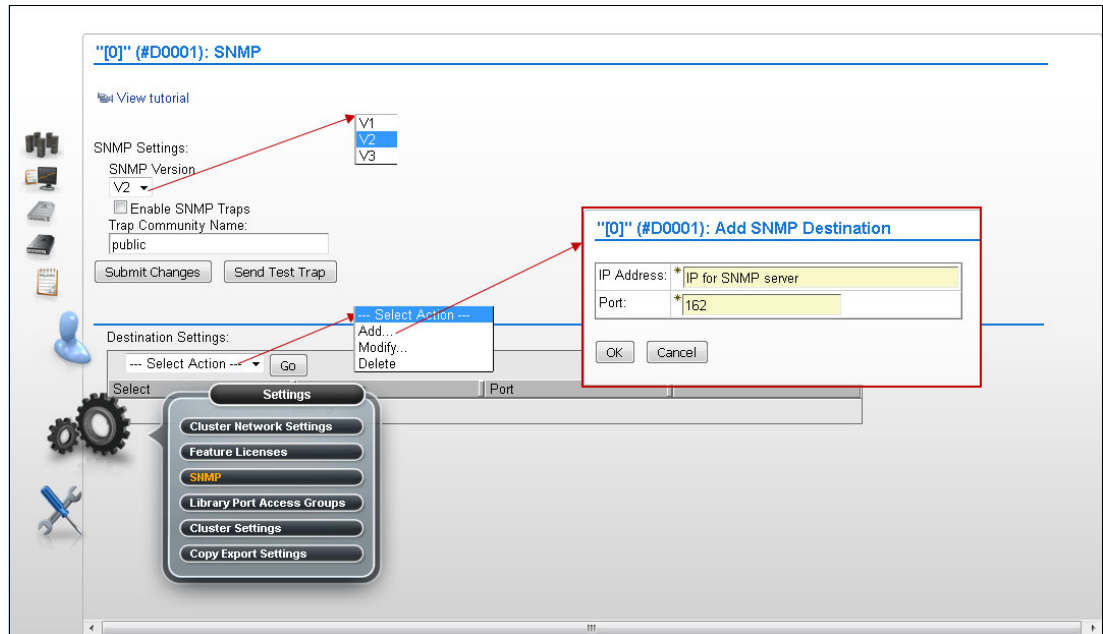


Figure 9-183 SNMP settings

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), shutdown, and code updates. SNMP is a networking protocol that enables a TS7700 to automatically gather and transmit information about alerts and status to other entities in the network.

SNMP Settings section

This section provides information about configuring global settings that apply to SNMP traps on an entire cluster. You can configure the following settings:

- ▶ **SNMP Version:** The SNMP version defines the protocol that is used in sending SNMP requests and is determined by the tool you are using to monitor SNMP traps. Different versions of SNMP traps work with different management applications. The only possible value on TS7700 is V1. No alternative version is supported.
- ▶ **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 does not process an SNMP trap unless it is associated with the correct community. This value must be 1 - 15 characters in length and consists of Unicode characters.
- ▶ **Send Test Trap:** This button sends a test SNMP trap to all destinations listed in the Destination Settings table by using the current SNMP trap values. The **Enable SNMP Traps** check box does not need to be checked to send a test trap. If the SNMP test trap is received successfully and the information is correct, select **Submit Changes**.
- ▶ **Submit Changes:** Select this button to submit changes to any of the global settings, including the SNMP Version, Enable SNMP Traps, and Trap Community Name fields.

Destination Settings section

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. You can configure the following settings:

- ▶ **IP address:** The IP address of the SNMP server. This value can take any of the following formats: IPv4, IPv6, a host name that is resolved by the system (such as `localhost`), or a fully qualified domain name (FQDN) if a domain name server (DNS) is provided. A value in this field is required.

A valid IPv4 address is 32 bits, consists of four decimal numbers, each 0 - 255, separated by periods, for example:

98.104.120.12

A valid IPv6 address is a 128-bit hexadecimal value separated into 16-bit fields by colons, for example:

3afa:1910:2535:3:110:e8ef:ef41:91cf

Leading zeros can be omitted in each field so that `:0003:` can be written as `:3:.` A double colon (`::`) can be used once per address to replace multiple fields of zeros, for example:

3afa:0:0:0:200:2535:e8ef:91cf

can be written this way:

3afa::200:2535:e8ef:91cf

A fully qualified host name is a domain name that uniquely and absolutely names a computer. It consists of the host name and the domain name. The domain name is one or more domain labels that place the computer in the DNS naming hierarchy. The host name and the domain name labels are separated by periods, and the total length of the host name cannot exceed 255 characters.

- ▶ **Port:** This port is where the SNMP trap logs are sent. This value must be a number 0 - 65535. A value in this field is required.

Consideration: A user with read-only permissions cannot modify the contents of the Destination Settings table.

Use the **Select Action** menu in the Destination Settings table to add, modify, or delete an SNMP trap destination. Destinations are changed in the vital product data (VPD) as soon as they are added, modified, or deleted. These updates do not depend on your selecting Submit Changes on the window:

- ▶ **Add SNMP destination:** Select this menu item to add an SNMP trap destination for use in the TS7700 Grid.
- ▶ **Modify SNMP destination:** Select this menu item to modify an SNMP trap destination that is used in the TS7700 Grid.
- ▶ **Confirm delete SNMP destination:** Select this menu item to delete an SNMP trap destination that is used in the TS7700 Grid.

Enabling IPv6

IPv6 and Internet Protocol Security (IPSec) are supported beginning with Release 3.0 of Licensed Internal Code by the 3957-V07 and 3957-VEB configurations of the TS7700.

Tip: The client network must use whether IPv4 or IPv6 for all functions, such as MI, key manager server, SNMP, Lightweight Directory Access Protocol (LDAP), and NTP. Mixing IPv4 and IPv6 is not currently supported.

Figure 9-184 shows how to enable IPv6 in a TS7700.

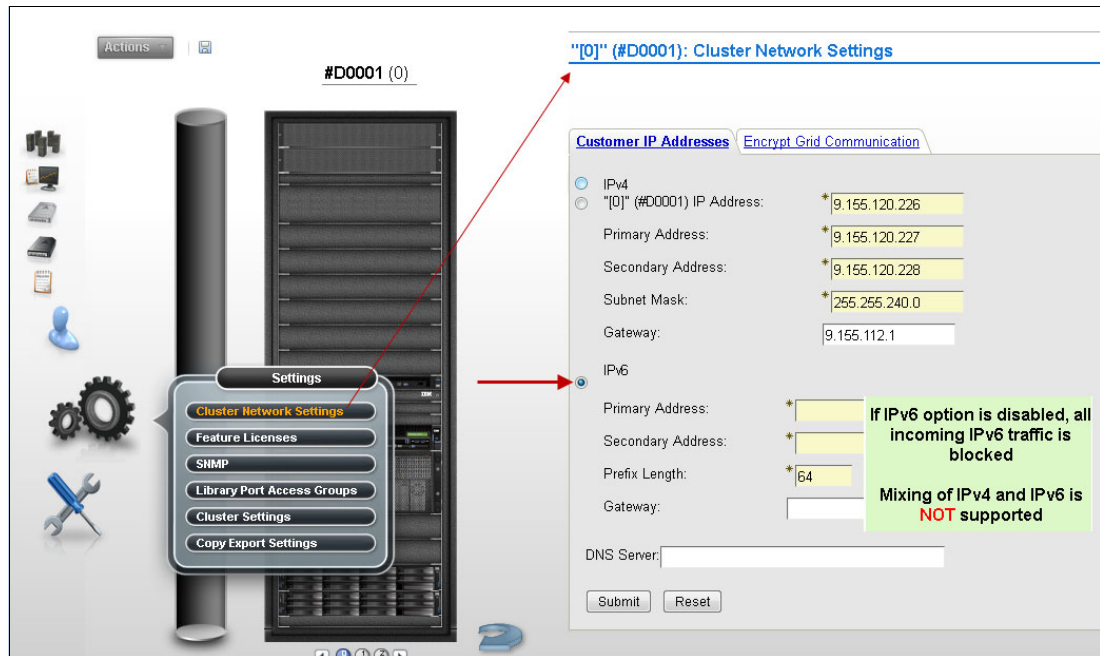


Figure 9-184 Configuring IPv6

For more information about IPv6, see "IPv6 support" on page 141 and Figure 9-119 on page 468.

Enabling IPsec

Beginning with Release 3.0 of Licensed Internal Code, the 3957-V07 and 3957-VEB configurations of the TS7700 support IPsec over the grid links.

Caution: Enabling grid encryption significantly affects the performance of the TS7700.

Figure 9-185 shows how to enable the IPsec for the TS7700 cluster.

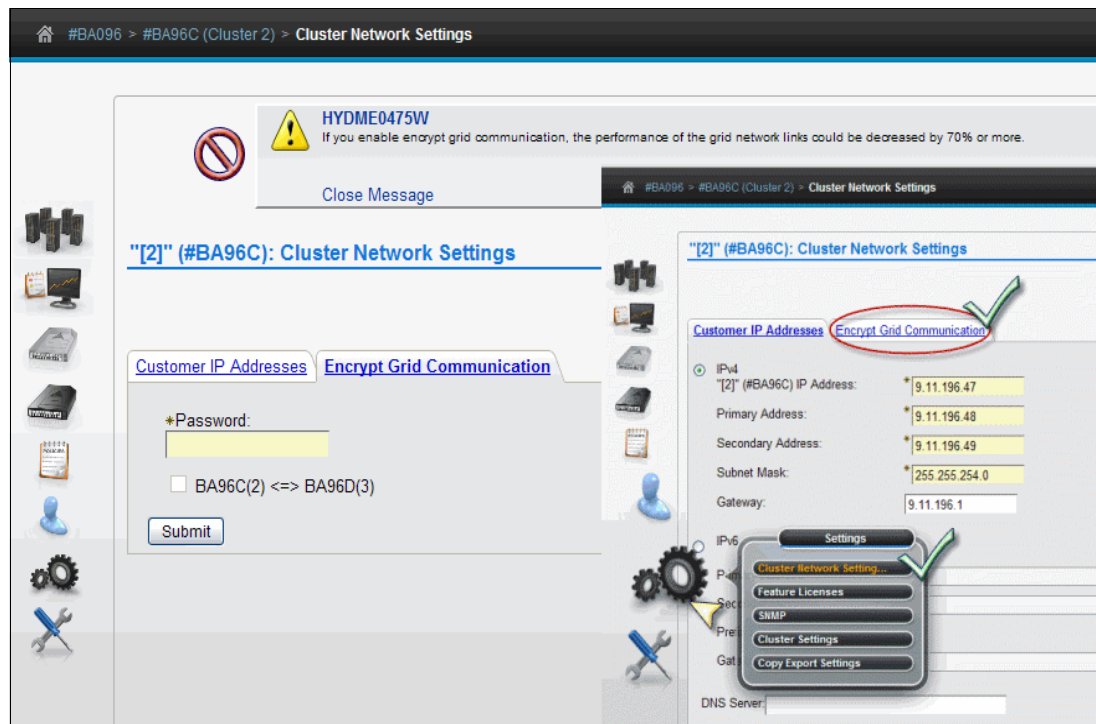


Figure 9-185 Enabling IPsec in the grid links

In a multi-cluster grid, the user can choose which link is encrypted by selecting the boxes in front of the beginning and ending clusters for the selected link. Figure 9-185 shows a two-cluster grid, which is the reason why there is only one option to select.

For more information about IPsec, see “IPsec support for the grid links” on page 142. Also, see the IBM TS7700 IBM Knowledge Center at:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_infrastructure_requirements_network_switches_tcp_ip_ports.html?lang=en

Defining security settings

Use this section to set up and check the security settings for the TS7700 grid. From this window in the MI, you can perform these functions:

- ▶ Add a policy
- ▶ Modify an existing policy
- ▶ Assign an authentication policy
- ▶ Test the security setting before running the application
- ▶ Delete an existing policy

Each cluster in your configuration can have a different security policy assigned to it. However, only one policy can be in effect on a cluster at a time.

Figure 9-186 shows the Security Settings window.

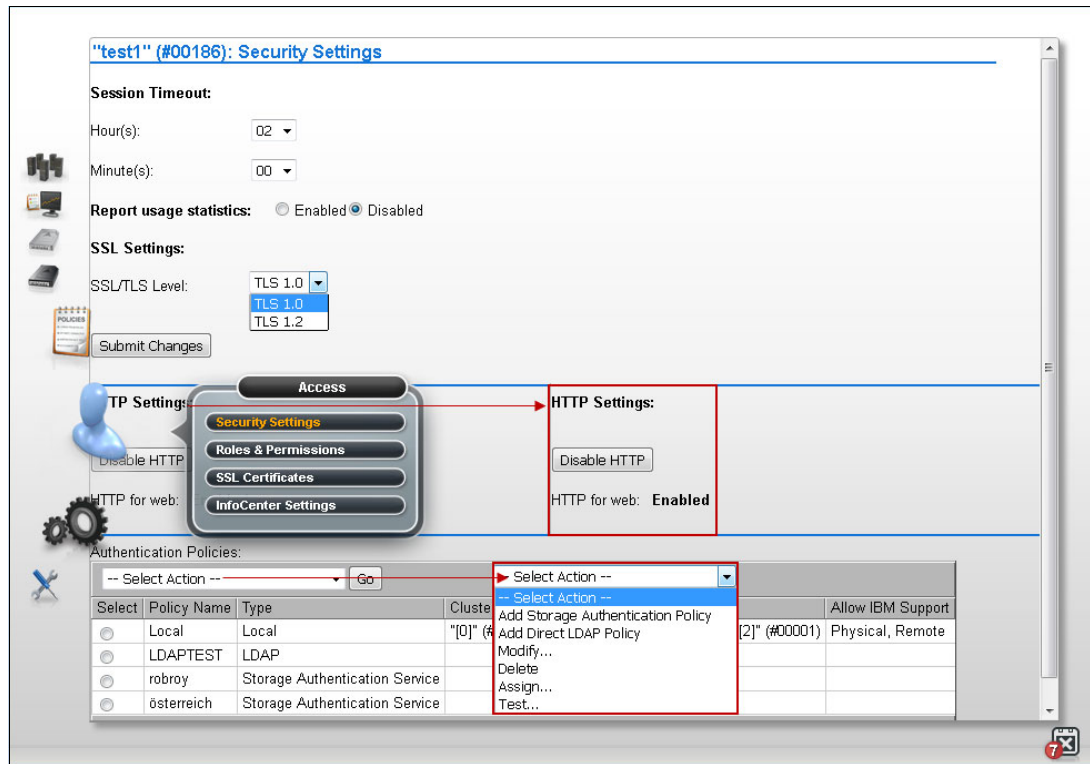


Figure 9-186 Security settings

For Session Timeout, you specify the number of hours and minutes that the MI can be idle before the current session expires and the user is redirected to the login window.

The Authentication Policies table shows the defined policies in the TS7700 Grid. You can set these policies:

- ▶ **Local:** This means that users and their assigned roles are replicated throughout the grid.
- ▶ **External:** This policy stores user and group data on a separate server, verifying the relationship between users, groups, and authorization roles whenever a user logs in to a cluster.

Direct LDAP and Storage Authentication Service policies are included in the external policies.

Important: When a Storage Authentication Service policy is enabled for a cluster, service personnel are **required** to log in with the setup user or group. Be sure that an account has been created for the service personnel **before** enabling storage authentication.

Storage Authentication Service Policy

The Storage Authentication Service Policy uses a centrally managed RBAC that authenticates and authorizes users by using the System Storage Productivity Center to authenticate users to an LDAP server.

Figure 9-187 shows the Add Storage Authentication Service Policy window.

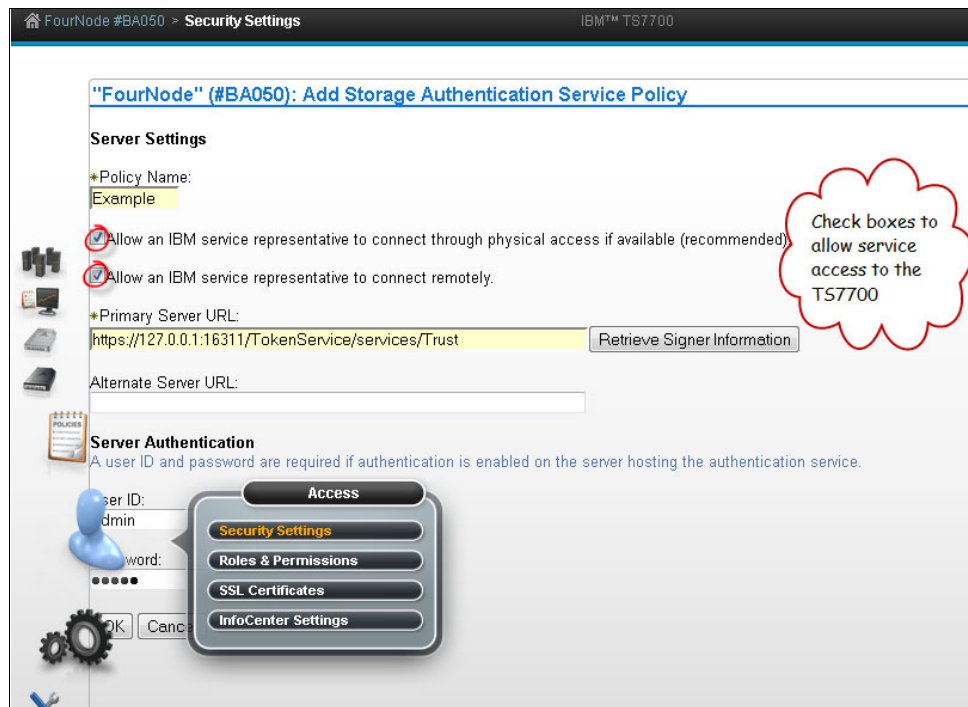


Figure 9-187 Add Storage Authentication Service Policy

Direct LDAP Policy

Figure 9-188 shows the Add Direct LDAP Policy menu. Use this menu to add an RBAC policy that authenticates and authorizes users through direct communication with an LDAP server.

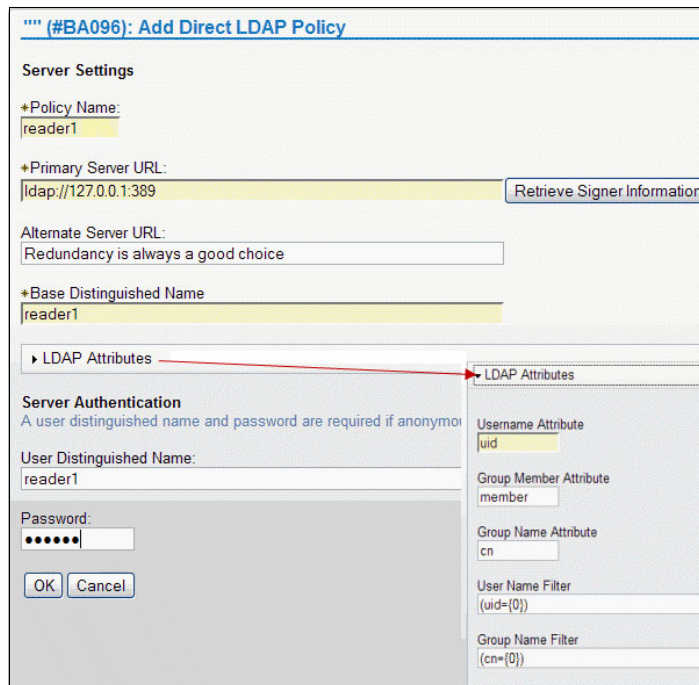


Figure 9-188 Add Direct LDAP Policy

Important: When a Storage Authentication Service policy is enabled for a cluster, service personnel are *required* to log in with the setup user or group. Be sure that an account has been created for the service personnel *before* enabling storage authentication.

The fields in both Figure 9-188 on page 568 and Figure 9-187 on page 568 are defined in the following list:

- ▶ **Policy Name:** The name of the policy that defines the authentication settings. The policy name is a unique value that consists of one to 50 Unicode characters. Heading and trailing blank spaces are deleted, but internal blank spaces are retained. The name of the Local policy is Local. Authentication policy names, either Local or user-created, cannot be modified after they are created.
- ▶ **Primary Server URL:** The primary URL for the Storage Authentication Service. The value in this field consists of 1 - 254 Unicode characters and takes one of the following formats:
 - `https://<server_address>:secure_port/TokenService/services/Trust`
 - `ldaps://<server_address>:secure_port`
 - `ldap://<server_address>:port`

If a domain name server (DNS) address needs to be used here, a DNS must be activated and configured on the Cluster Network settings window. See 9.2.10, “The Settings icon” on page 469.

- ▶ **Alternative Server URL:** The alternative URL for the Storage Authentication Service if the primary URL cannot be accessed. The value is the same as the value described in the previous item.¹
- ▶ **Server Authentication:** Values are required in the *user ID* and *password* fields if IBM WebSphere Application Server security is enabled on the WebSphere Application Server hosting the Authentication Service, or if anonymous access is disabled on the LDAP server:
 - **User ID:** The user name that is used with HTTP basic authentication for authenticating to the Storage Authentication Service. Maximum length of 254 Unicode characters.
 - **Password:** The password that is used with HTTP basic authentication for authenticating to the Storage Authentication Service. Maximum length of 254 Unicode characters.
- ▶ **Direct LDAP:** Values in the following fields are required if secure authentication is used or anonymous connections are disabled in the LDAP server:
 - **User Distinguished Name:** Used to authenticate to the LDAP authentication service. Maximum length of 254 Unicode characters, for example:
`CN=Administrator,CN=users,DC=mycompany,DC=com`
 - **Password:** The password to authenticate to the LDAP authentication service. Maximum length of 254 Unicode characters.
 - **Base Distinguished Name:** The distinguished name (DN) uniquely identifies a set of entries in a domain. Maximum length of 254 Unicode characters.
 - **Username Attribute:** The attribute name that is used for the user name during authentication. This field is required and contains the value *uid*, by default. Maximum length of 61 Unicode characters.

¹ The server address value in the Primary or alternative Server URL can be an IP address or DNS address. Valid IP formats include the following formats:

- IPv4: 32 bits. Four decimal numbers, 0 - 255, separated by periods, for example, 12.345.678.
- IPv6: 128-bit hexadecimal values that are enclosed by brackets, which are separated into 16-bit field by colons, for example, [1234:9abc:0::1:cdef:8]. Leading zeros (0) can be omitted. A double colon (::) means a field of 0s (:0000:).

- **Password Attribute:** The attribute name that is used for the password during authentication. This field is required and contains the value `userPassword`, by default. Maximum length of 61 Unicode characters.
- **Group Member Attribute:** The attribute name that is used to identify the group during authorization. This field is optional and contains the value `member`, by default. Maximum length of 61 Unicode characters.
- **Group Name Attribute:** The attribute name that is used to identify the group during authorization. This field is optional and contains the value `cn`, by default. Maximum length of 61 Unicode characters.
- **User Name Filter:** Used to filter and validate an entered user name. This field is optional and contains the value `(uid={0})`, by default. Maximum length of 254 Unicode characters.
- **Group Name Filter:** Used to filter and validate an entered group name. This field is optional and contains the value `(cn={0})`, by default. Maximum length of 254 Unicode characters.

RACF based LDAP Policy

RACF based LDAP is a particular case of Direct LDAP policy. The procedure is similar to “Direct LDAP Policy” on page 568. However, there are some configurations on the host side regarding the RACF, SDBM, and IBM Security Directory Server (formerly Tivoli LDAP server) that must be in place before defining the RACF based LDAP policy in the MI. For more information, see “Creating a RACF based LDAP Policy” on page 463.

Local policy

The *Local policy* is the default authentication policy. When enabled, it is in effect for all clusters on the grid. It is mutually exclusive with the Storage Authentication Service. Local policy can be modified to add, change, or delete individual accounts, but the policy itself cannot be deleted.

Figure 9-189 shows the Modify Local Policy window.

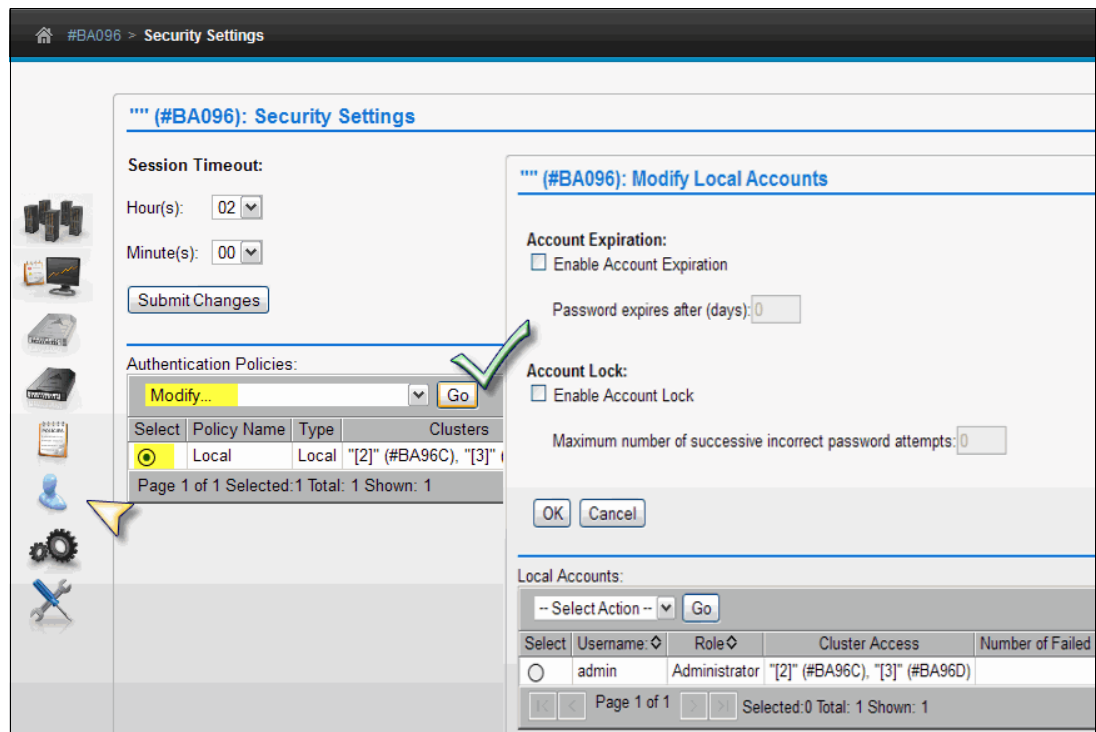


Figure 9-189 Modify Local Accounts

Use this window to modify the Local policy settings for the TS7700 Grid. The following information can be defined:

- ▶ Whether accounts that are defined by a policy can expire, and if so, the number of days that a password can be used before it expires. Possible values are 1 - 999.
- ▶ Whether accounts that are defined by a policy can be locked after several successive incorrect password retries (1 - 9).

9.3.4 TS7700 multi-cluster definitions

The following sections describe TS7700 multi-cluster definitions.

Defining grid copy mode control

When upgrading a stand-alone cluster to a grid, FC4015, Grid Enablement, must be installed on all clusters in the grid. Also, the Copy Consistency Points in the MC definitions on all clusters in the new grid should be set.

The data consistency point is defined in the MCs construct definition through the MI. This task can be performed for an existing grid system. In a stand-alone cluster configuration, but this cluster will show in the Modify MC definition.

To open the MCs window, complete the following steps:

1. Click **Constructs** → **Management Classes** under the **Welcome Admin** menu.
2. Select the MC name and select **Modify** from the **Select Action** menu.

- Figure 9-190 shows how to modify the copy consistency by using the Copy Action table, and then clicking **OK**. In the figure, the TS7700 is part of a three-cluster grid configuration. This additional menu is displayed only if a TS7700 is part of a grid environment (options are not available in a stand-alone cluster).

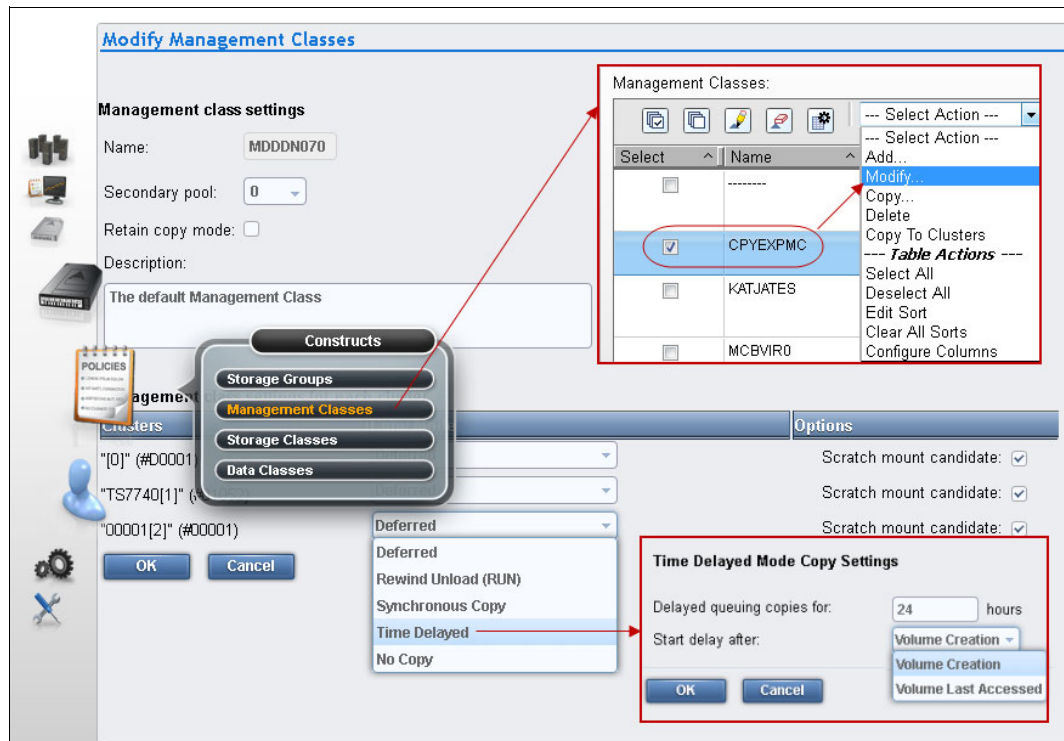


Figure 9-190 Modify Management Classes

As shown in Figure 9-190, the options of consistency points per cluster are:

- ▶ **No Copy:** No copy (NC) is made to this cluster.
- ▶ **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 was unloaded (DEF).
- ▶ **Synchronous Copy:** Provides tape copy capabilities up to synchronous-level granularity across two clusters within a multi-cluster grid configuration.
- ▶ **Time Delayed:** Volume copy occurs after the specified delay time period passes. Options are:
 - **Delay queuing copies for [X] hours**, where X is a number 1 - 65,353.
 - **Start delay after** one of these triggers:
 - The time of the creation of the volume
 - The time where the last access to that volume occurred
- ▶ **No Copy:** No copy is necessary for that cluster.

Tip: A stand-alone TS7700 always uses RUN as the Data Consistency Point.

For more information about this subject, see the following resources:

- ▶ **IBM TS7700 Series Best Practices - TS7700 Hybrid Grid Usage:**
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>
- ▶ **IBM TS7700 Series Best Practices - Copy Consistency Points:**
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101230>
- ▶ **IBM TS7700 Series Best Practices - Synchronous Mode Copy:**
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102098>

Defining Copy Policy Override settings

With the TS7700, the optional override settings that influence the selection of the I/O TVC and replication responses can be defined and set.

Reminder: The items on this window can modify the cluster behavior regarding local copy and certain I/O operations. Some **LI REQ** commands also can do it.

The settings are specific to a cluster in a multi-cluster grid configuration, which means that each cluster can have separate settings, if needed. 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. The following settings can be defined and set:

- ▶ Prefer local cache for scratch mount requests
- ▶ Prefer local cache for private mount requests
- ▶ Force volumes that are mounted on this cluster to be copied to the local cache
- ▶ Enable fewer RUN consistent copies before reporting **RUN** command complete
- ▶ Ignore cache preference groups for copy priority

These settings can be viewed and modified from the TS7700 MI by clicking **Settings** → **Cluster Setting** → **Copy Policy Override**, as shown in Figure 9-191.

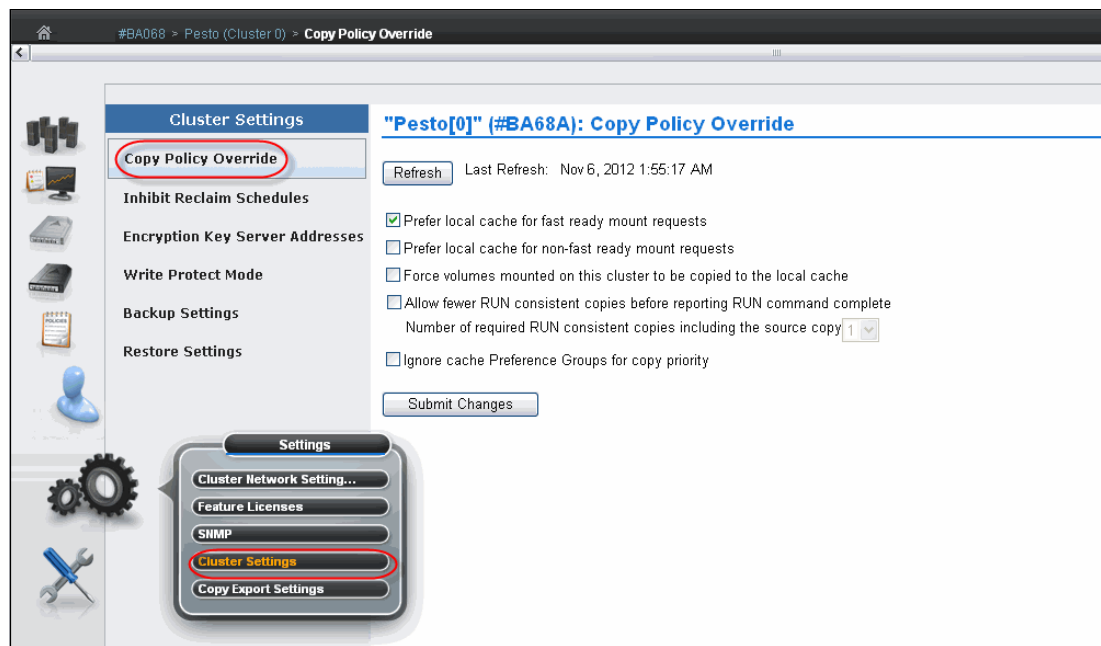


Figure 9-191 Cluster Settings

The following settings can be selected in the MI window:

► **Prefer local cache for fast ready mount requests**

A scratch (Fast Ready) mount selects a local copy if a cluster Copy Consistency Point is not specified as No Copy in the MC for the mount. The cluster is not required to have a valid copy of the data.

► **Prefer local cache for private (non-fast ready) mount requests**

This override causes the local cluster to satisfy the mount request if 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, the default cluster selection criteria applies.

► **Force volumes that are mounted on this cluster to be copied to the local cache**

For a private (non-Fast Ready) mount, this override causes a copy to be run on the local cluster as part of mount processing. For a scratch (Fast Ready) mount, this setting overrides the specified MC with a Copy Consistency Point of Rewind-Unload for the cluster. This does not change the definition of the MC, but influences the Replication policy.

► **Allow fewer RUN consistent copies before reporting RUN command complete**

If selected, the value that is entered for **Number of required RUN consistent copies including the source copy** is used to determine the number of copies to override before the RUN operation reports as complete. If this option is not selected, the MC definitions are used explicitly. Therefore, the number of RUN copies can be from one to the number of clusters in the grid.

► **Ignore cache preference groups for copy priority**

If this option is selected, copy operations ignore the cache preference group when determining the priority of volumes copied to other clusters.

Consideration: 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 cluster is preferred for all I/O operations.

If the TS7700 cluster of the GDPS primary site fails, you must complete the following recovery actions:

1. Vary online virtual devices from a remote TS7700 cluster from the primary site of the GDPS host.
2. Manually start, through the TS7700 MI, a read/write Ownership Takeover, unless AOTM already has transferred ownership.

Defining scratch mount candidates

Scratch allocation assistance (SAA) is an extension of the device allocation assistance (DAA) function for scratch mount requests. SAA filters the list of clusters in a grid to return to the host a smaller list of candidate clusters that are designated as scratch mount candidates.

Scratch mount candidates can be defined in a grid environment with two or more clusters. For example, in a hybrid configuration, the SAA function can be used to direct certain scratch allocations (workloads) to one or more TS7720 tape drives for fast access. Other workloads can be directed to a TS7740 or TS7720T for archival purposes.

Clusters not included in the list of scratch mount candidates are not used for scratch mounts at the associated MC unless those clusters are the only clusters that are known to be available and configured to the host.

See Chapter 10, “Host Console operations” on page 601 for information about software levels that are required by SAA and DAA to function properly, in addition to the **LI REQ** commands that are related to the SAA and DAA operation.

As shown in Figure 9-192, by default all clusters are chosen as scratch mount candidates. Select which clusters are candidates by MC. If no clusters are checked, the TS7700 defaults to all clusters being candidates.

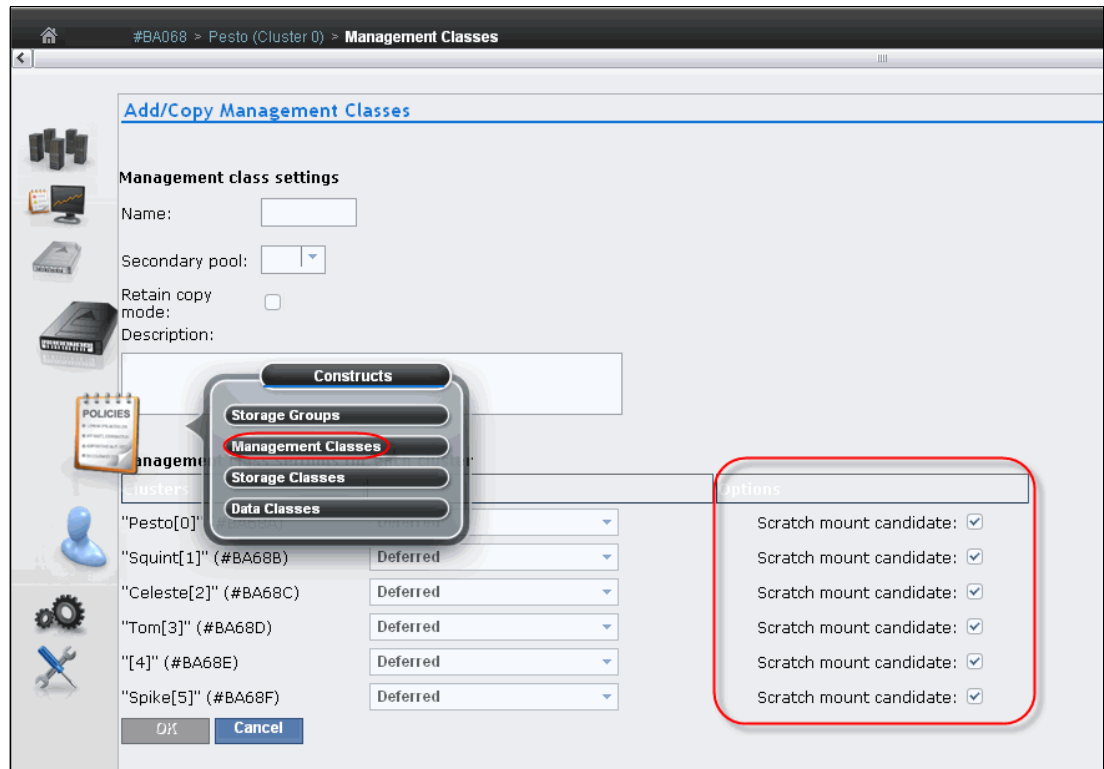


Figure 9-192 Scratch mount candidate list in Add Management Classes window

Each cluster in a grid can provide a unique list of candidate clusters. Clusters with an ‘N’ copy mode, such as cross-cluster mounts, can still be candidates. When defining the scratch mount candidates in an MC, normally you want each cluster in the grid to provide the same list of candidates for load balancing.

Note: Scratch mount candidate list as defined in MI (Figure 9-192) is only accepted upon being enabled by using the **LI REQ** setting.

Retain Copy mode

Retain Copy mode is an optional setting where existing Copy Consistency Points for a volume are accepted rather than 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 wanted. 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 in the Management Classes window for each MC when adding an MC. Figure 9-193 shows the Management Classes window and the Retain Copy mode check box.

Note: The Retain Copy mode option is effective only on private (non-Fast Ready) virtual volume mounts.

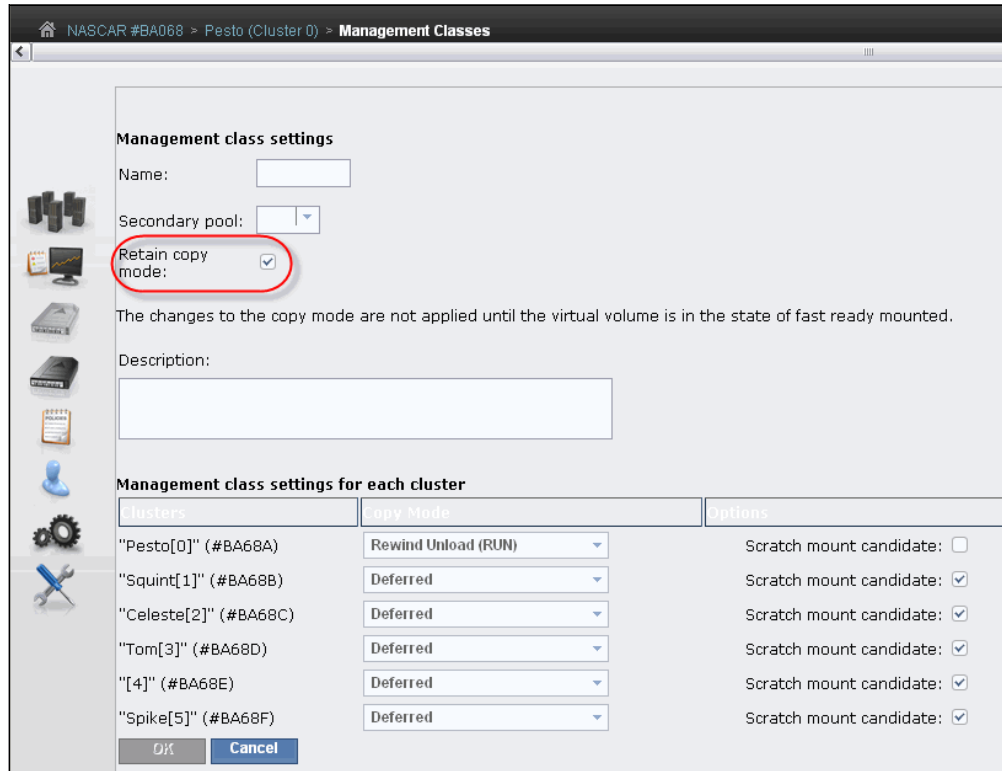


Figure 9-193 Retain Copy mode selection in the Management Classes window

Defining cluster families

Cluster families can be defined in a grid with three or more clusters.

This function introduces a concept of grouping clusters together into families. Using cluster families, a common purpose or role can be assigned to a subset of clusters within a grid configuration. The role that is assigned, for example, production or archive, is used by the TS7700 Licensed Internal Code to make improved decisions for tasks, such as replication and TVC selection. For example, clusters in a common family are favored for TVC selection, or replication can source volumes from other clusters within its family before using clusters outside of its family.

Use the **Cluster Families** option on the **Actions** menu of the Grid Summary window to add, modify, or delete a cluster family. Figure 9-194 shows the menu for the cluster families.

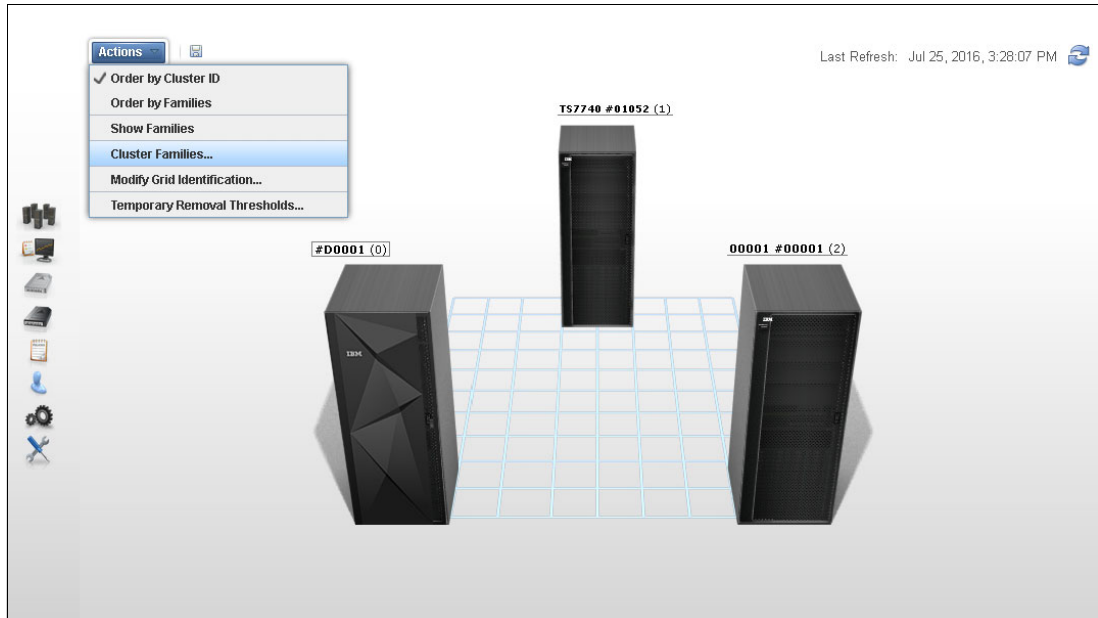


Figure 9-194 Cluster Families menu option

To view or modify cluster family settings, complete the following steps:

1. First, check whether the assigned MI user role is authorized to alter this window.
2. Then, select **Cluster Families** from the **Actions** menu to run the following actions:
 - Adding a family
 - Moving a cluster
 - Deleting a family
 - Saving changes

Adding a family

To add a family, complete the following steps:

1. Click **Add** to create a new cluster family. A new cluster family placeholder is created to the right of any existing cluster families.
2. Enter the name of the new cluster family in the active Name text box. Cluster family names must be 1 - 8 characters in length and consist of Unicode characters. Each family name must be unique.
3. To add a cluster to the new cluster family, move a cluster from the Unassigned Clusters area by following instructions in “Moving a cluster” on page 578.

Consideration: A maximum of eight cluster families can be created.

Moving a cluster

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:

1. Select a cluster: A selected cluster is identified by a highlighted border. Select a cluster from its resident cluster family or the Unassigned Clusters area with any of these methods:
 - Clicking 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
2. Move the selected cluster or clusters by using one these methods:
 - Clicking 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

Deleting a family

To delete an existing cluster family, complete the following steps:

1. Click the **X** in the upper-right corner of the cluster family to delete. If this cluster family contains any clusters, a warning message displays.
2. 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.

Saving changes

Click **Save** to save any changes that are made to the Cluster families window and return it to read-only mode.

Consideration: 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 window remains in edit mode.

Cluster family configuration

Figure 9-195 illustrates the actions to create a family.

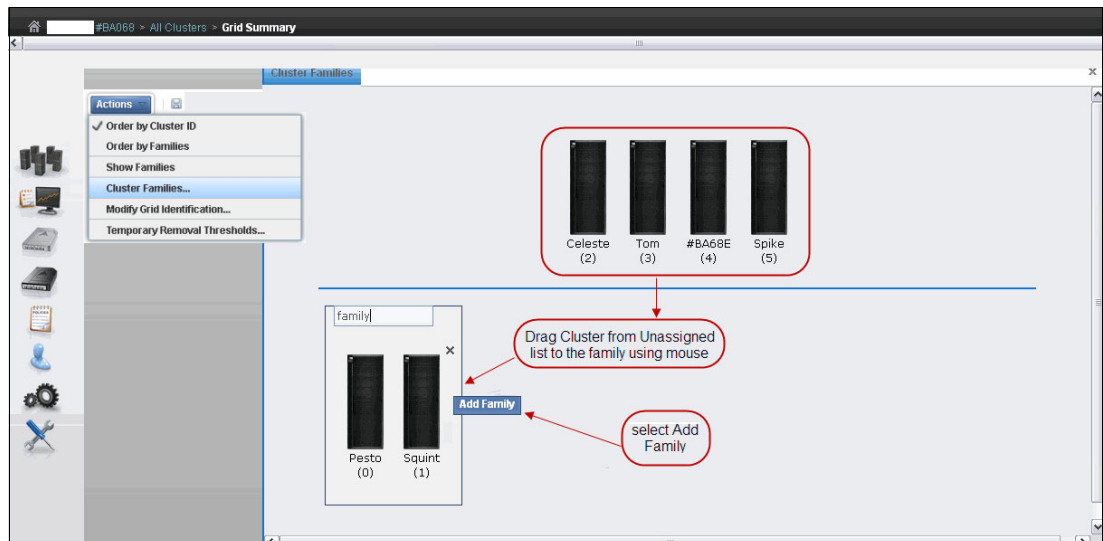


Figure 9-195 Create a cluster family

Figure 9-196 shows an example of a cluster family configuration.

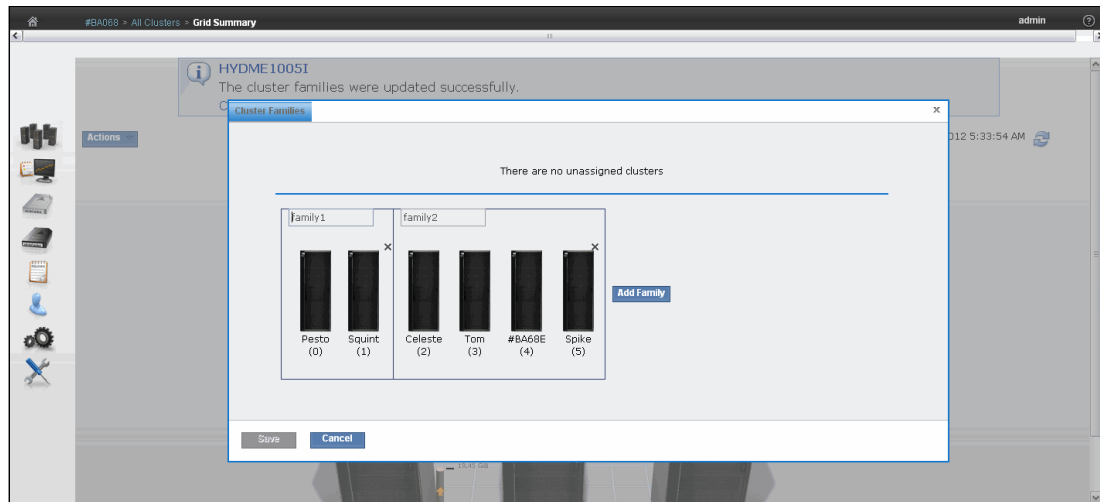


Figure 9-196 Cluster families

Important: Each cluster family needs to contain at least one cluster.

TS7700 cache thresholds and removal policies

This topic describes the boundaries (thresholds) of free cache space in a disk-only TS7700 or TS7700T CP0 partition cluster and the policies that can be used to manage available (active) cache capacity in a grid configuration.

Cache thresholds for a disk only TS7700 or TS7700T resident partition (CP0)

A disk-only TS7700 and the resident partition (CP0) of a TS7700T (tape attach) configuration does not attach to a physical library. All virtual volumes are stored in the cache. Three thresholds define the active cache capacity in a TS7700 and determine the state of the cache as it relates to remaining free space. In ascending order of occurrence, these are the three thresholds:

► **Automatic Removal**

The policy removes the oldest logical volumes from the disk only TS7700 cache if a consistent copy exists elsewhere in the grid. This state occurs when the cache is 4 TB below the out-of-cache-resources threshold. In the automatic removal state, the TS7700 automatically removes volumes from the disk-only cache to prevent the cache from reaching its maximum capacity. This state is identical to the limited-free-cache-space-warning state unless the Temporary Removal Threshold is enabled.

Automatic removal can be disabled within any specific TS7700 cluster by using the following library request command:

```
LIBRARY REQUEST,CACHE,REMOVE,{ENABLE|DISABLE}
```

So that a disaster recovery test can access all production host-written volumes, automatic removal is temporarily disabled while disaster recovery write protect is enabled on a disk-only cluster. When the write protect state is lifted, automatic removal returns to normal operation.

► **Limited free cache space warning**

This state occurs when there is less than 3 TB of free space remaining in the cache. After the cache passes this threshold and enters the limited-free-cache-space-warning state, write operations can use only an extra 2 TB before the out-of-cache-resources state is encountered. When a TS7700 cluster enters the limited-free-cache-space-warning state, it remains in this state until the amount of free space in the cache exceeds 3.5 TB. The following messages can be displayed on the MI during the limited-free-cache-space-warning state:

- HYDME0996W
- HYDME1200W

For more information, see the related information section in the TS7700 IBM Knowledge Center about each of these messages:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7740_removal_policies.html?lang=en

Clarification: Host writes to the disk only TS7700 cluster and inbound copies continue during this state.

► **Out of cache resources**

This state occurs when there is less than 1 TB of free space remaining in the cache. After the cache passes this threshold and enters the out-of-cache-resources state, it remains in this state until the amount of free space in the cache exceeds 3.5 TB. When a TS7720 cluster is in the out-of-cache-resources state, volumes on that cluster become read-only and one or more out-of-cache-resources messages are displayed on the MI. The following messages can display:

- HYDME0997W
- HYDME1133W
- HYDME1201W

For more information, see the related information section in the TS7700 IBM Knowledge Center about each of these messages:

http://www.ibm.com/support/knowledgecenter/STFS69_4,0.0/ts7740_removal_policies.html?lang=en

Clarification: New host allocations do not choose a disk only cluster in this state as a valid TVC candidate. New host allocations that are sent to a TS7700 cluster in this state choose a remote TVC instead. If all valid clusters are in this state or cannot accept mounts, the host allocations fail. Read mounts can choose the disk only TS7700 cluster in this state, but modify and write operations fail. Copies inbound to this cluster are queued as Deferred until the disk only cluster exits this state.

Table 9-15 displays the start and stop thresholds for each of the active cache capacity states defined.

Table 9-15 Active cache capacity state thresholds

State	Enter state (free space available)	Exit state (free space available)	Host message displayed
Automatic removal	< 4 TB	> 4.5 TB	CBR3750I when automatic removal begins
Limited free cache space warning (CP0 for a TS7700 tape attach)	<= 3 TB or <=15% of the size of cache partition 0, whichever is less	>3.5 TB or >17.5% of the size of cache partition 0, whichever is less	CBR3792E upon entering state CBR3793I upon exiting state
Out of cache resources (CP0 for a TS7700 tape attach)	< 1 TB or <= 5% of the size of cache partition 0, whichever is less	> 3.5 TB or >17.5% of the size of cache partition 0, whichever is less	CBR3794A upon entering state CBR3795I upon exiting state
Temporary removal ^a	<(X + 1 TB) ^b	>(X + 1.5 TB) ^b	Console message

a. When enabled.

b. Where X is the value set by the TVC window on the specific cluster.

The Removal policy is set by using the SC window on the TS7700 MI. Figure 9-197 shows several definitions in place.

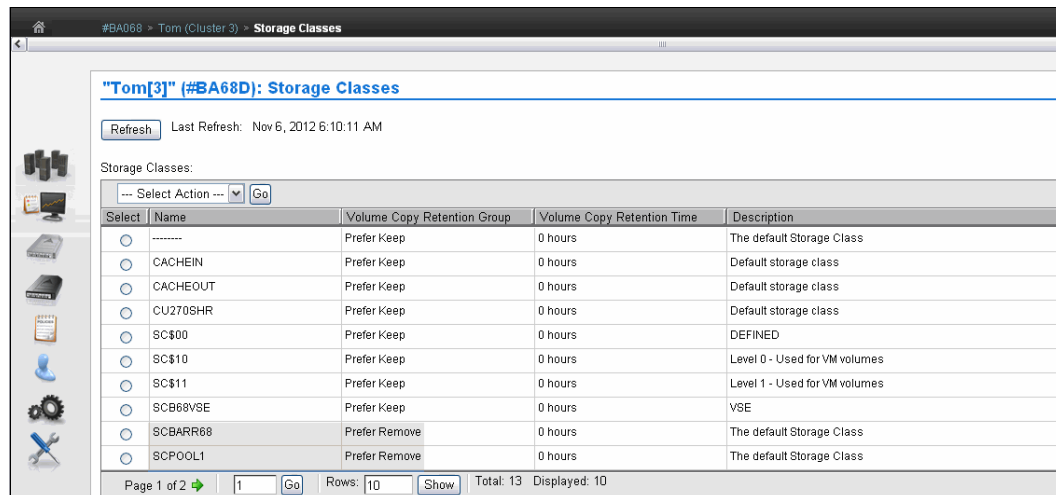


Figure 9-197 Storage Classes in TS7700 with removal policies

To add or change an existing SC, select the appropriate action in the menu, and click **Go**. See Figure 9-198.

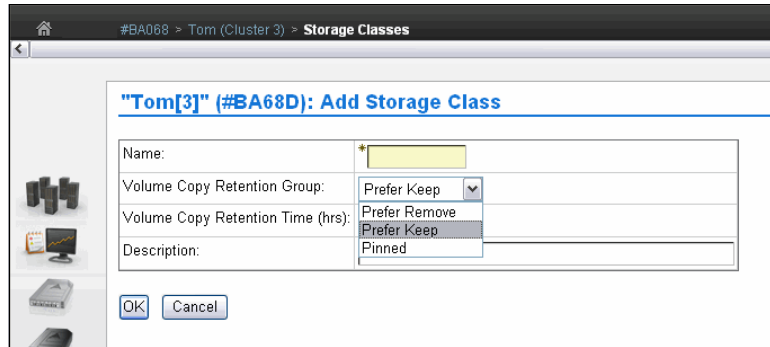


Figure 9-198 Define a new Storage Class with TS7700

Removal Threshold

The Removal Threshold is used to prevent a cache overrun condition in a disk only TS7700 cluster that is configured as part of a grid. By default, it is a 4 TB value (3 TB fixed, plus 1 TB) that, when taken with the amount of used cache, defines the upper limit of a TS7700 cache size. Above this threshold, logical volumes begin to be removed from a disk only TS7700 cache.

Note: Logical volumes are only removed if there is another consistent copy within the grid.

Logical volumes are removed from a disk only TS7700 cache in this order:

1. Volumes in scratch (Fast Ready) categories
2. Private volumes least recently used, by using the enhanced Removal policy definitions

After removal begins, the TS7700 continues to remove logical volumes until the Stop Threshold is met. The Stop Threshold is the Removal Threshold minus 500 GB. A particular logical volume cannot be removed from a disk only TS7700 cache until the TS7700 verifies that a consistent copy exists on a peer cluster. If a peer cluster is not available, or a volume copy has not yet completed, the logical volume is not a candidate for removal until the appropriate number of copies can be verified later.

Tip: This field is only visible if the selected cluster is a disk only TS7700 in a grid configuration.

Temporary Removal Threshold

The Temporary Removal Threshold lowers the default Removal Threshold to a value lower than the Stop Threshold. This resource might be useful in preparation for a disaster recovery testing with FlashCopy, or yet in anticipation of a service activity in a member of the grid. Logical volumes might need to be removed to create extra room in cache for FlashCopy volumes that will be present during a DR rehearsal, or before one or more clusters go into service mode. When a cluster in the grid enters service mode, the remaining clusters can have their ability to make or validate copies and perform auto removal of logical volumes affected. For an extended period, this situation might result in a disk-only cluster getting out of cache resources, considering the worst possible scenario. The *Temporary Removal Threshold* resource is instrumental to help preventing this possibility.

Note: The Temporary Removal Threshold is not supported on the TS7740.

The lower threshold creates extra free cache space, which enables the disk-only TS7700 to accept any host requests or copies during the DR testing or service outage without reaching its maximum cache capacity. The Temporary Removal Threshold value must be equal to or greater than the expected amount of compressed host workload written, copied, or both to the disk-only cluster or CP0 partition during the service outage.

The default Temporary Removal Threshold is 4 TB, which provides 5 TB (4 TB plus 1 TB) of existing free space. The threshold can be set to any value between 2 TB and full capacity minus 2 TB.

All disk-only TS7700 cluster or CP0 partition in the grid that remain available automatically lower their Removal Thresholds to the Temporary Removal Threshold value defined for each. Each cluster can use a different Temporary Removal Threshold. The default Temporary Removal Threshold value is 4 TB (an extra 1 TB more data than the default removal threshold of 3 TB).

Each disk-only TS7700 cluster or CP0 partition uses its defined value until the cluster within the grid in which the removal process has been started enters service mode or the temporary removal process is canceled. The cluster that is initiating the temporary removal process (either a cluster within the grid that is not part of the DR testing, or the one scheduled to go into Service) does not lower its own removal threshold during this process.

Note: The cluster that is elected to initiate Temporary Removal process is not selectable on the list of target cluster for the removal action.

Removal policy settings can be configured by using the **Temporary Removal Threshold** option on the **Actions** menu, which is available on the Grid Summary window of the TS7700 MI. Figure 9-199 shows the Temporary Removal Threshold mode window.

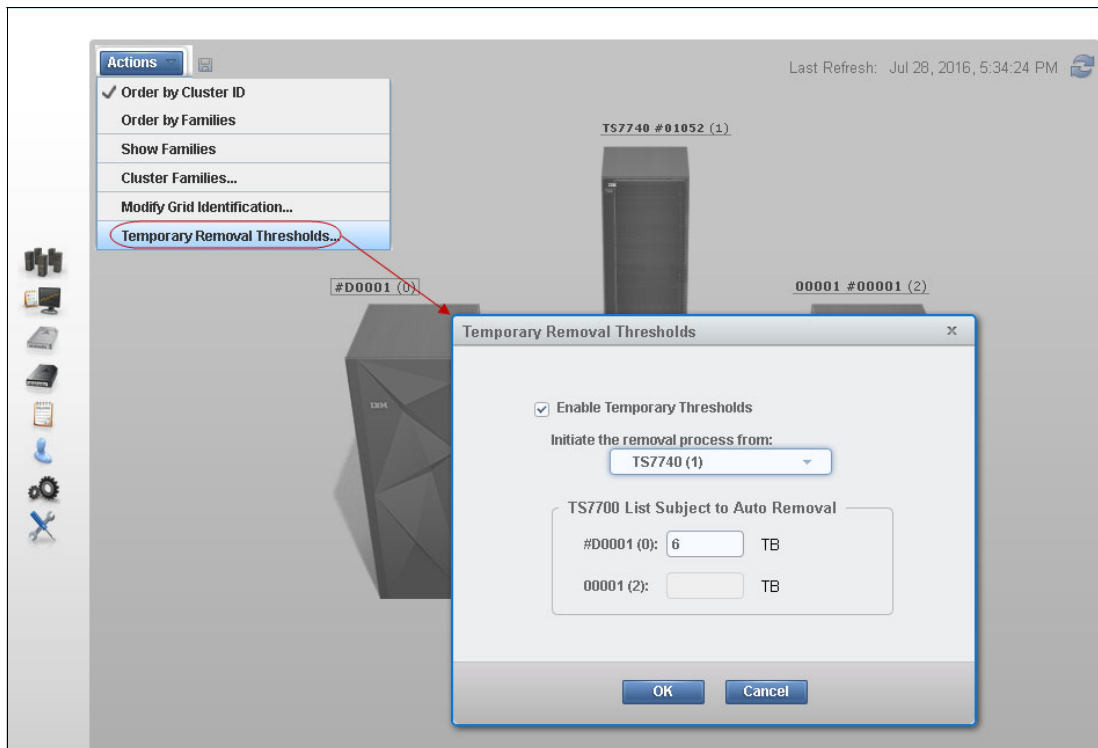


Figure 9-199 Selecting cluster to start removal process and temporary removal threshold levels

The Temporary Removal Threshold mode window includes these options:

- ▶ **Enable Temporary Thresholds**

Check this box and click **OK** to start the pre-removal process. Clear this box and click **OK** to abandon a current pre-removal process.

- ▶ **Initiate the removal process from (cluster to be serviced):**

Select from this menu the cluster that will be put into service mode. The pre-removal process is started from this cluster.

Note: This process does not initiate Service Prep mode.

Even when the temporary removal action is started from a disk-only cluster, this cluster will still be not selectable on the drop-down menu of the TS7700 List Subject to Auto Removal, since the removal action will not affect this cluster.

This area of the window contains each disk-only TS7700 cluster or CP0 partition in the grid and a field to set the temporary removal threshold for that cluster.

Note: The *Temporary Removal Threshold* task ends when the originator cluster enters in Service mode, or the task is canceled on the *Tasks* page in MI.

The Temporary Removal Threshold is not supported on the TS7740 cluster.

9.4 Basic operations

This section explains the tasks that might be needed during the operation of a TS7700.

9.4.1 Clock and time setting

The TS7700 time can be set from an NTP server or by the IBM SSR. It is set to Coordinated Universal Time. See “Date and Time coordination” on page 60 for more details about time coordination.

Note: Use Coordinated Universal Time in all TS7700 clusters whenever possible.

The TS4500 tape library time can be set from management Interface, as shown in Figure 9-200. Notice that TS4500 can be synchronized with NTP server, when available.

More information about TS4500 tape library can be found locally at TS4500 GUI clicking the question mark icon, or online at:

https://www.ibm.com/support/knowledgecenter/STQRQ9/com.ibm.storage.ts4500.doc/ts4500_ichome.html

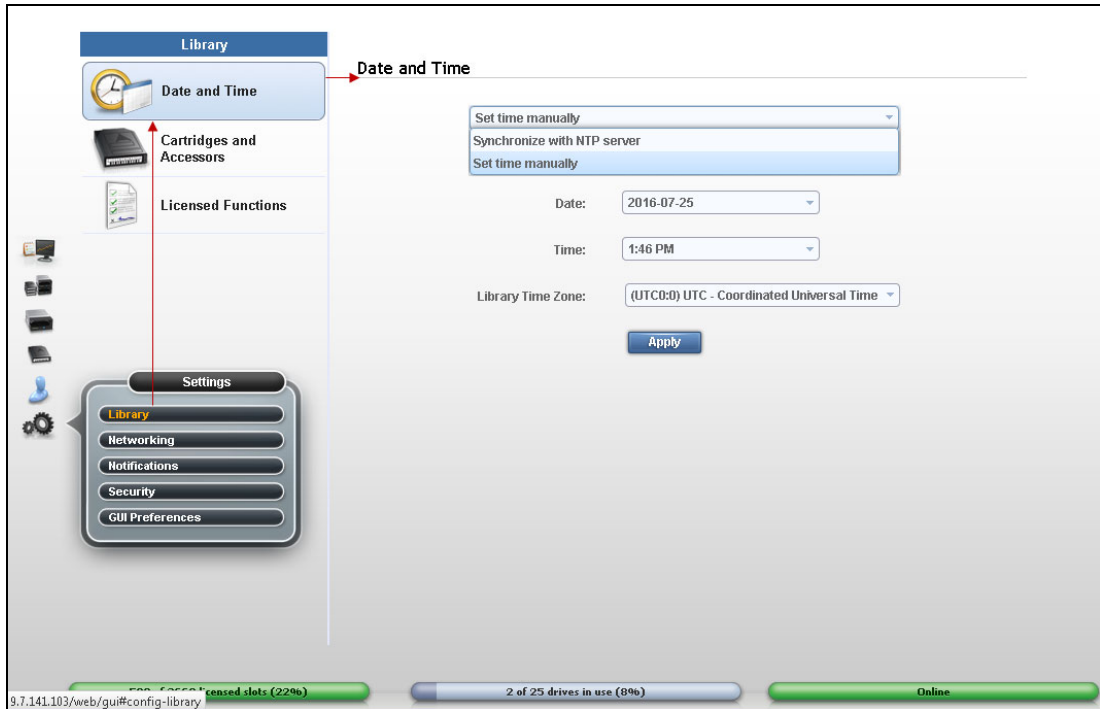


Figure 9-200 Adjusting date and time at TS4500 GUI

On the TS3500 tape library time can be set from IBM Ultra Scalable Specialist work items by clicking **Library** → **Date and Time** as shown in Figure 9-201.

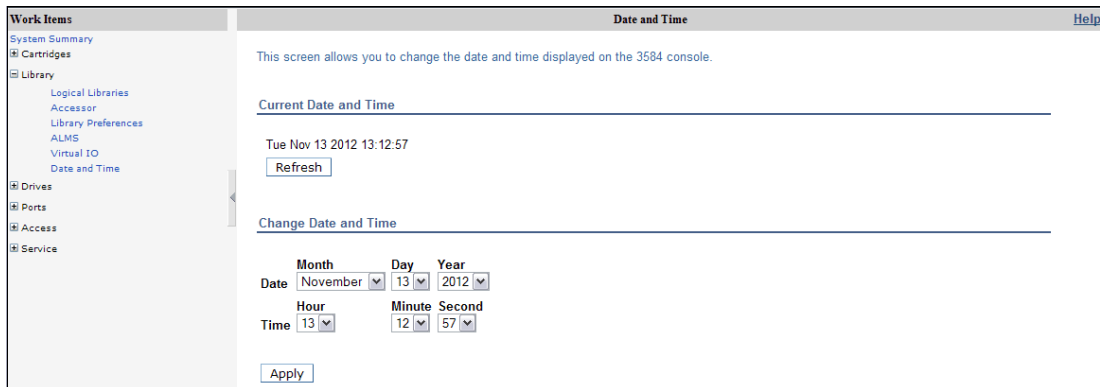


Figure 9-201 TS3500 tape library GUI Date and Time

9.4.2 Library in Pause mode

During the operation, the tape library can be paused, and this might affect the related tape attach cluster, regardless of being in a grid or not. The reasons for the pause can include an enclosure door that is being opened to clear a device after a load/unload failure or to remove 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 TS7740 or TS7720T 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 run, but not physical mounts, no more data can be moved out of the cache after the currently mounted stacked volumes are filled.

The cache is filling up with data that has not been copied to stacked volumes. This results in significant throttling and stopping of any mount activity in the TS7740 cluster or in the tape partitions in the TS7700T cluster. For this reason, it is important to minimize the amount of time that is spent with the library in Pause mode.

9.4.3 Preparing a TS7700 for service

When an operational TS7700 must be taken offline for service, the TS7700 Grid first must be prepared for the loss of the resources that are involved to provide continued access to data. The controls to prepare a TS7700 for service (Service Prep) are provided through the MI. This menu is described in “Service mode window” on page 346.

Here is the message posted to all hosts when the TS7700 Grid is in this state:

```
CBR3788E Service preparation occurring in library library-name.
```

More details about service preparation are in Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Tip: Before starting service preparation at the TS7700, all virtual devices on this cluster must be in the offline state regarding the accessing hosts. Pending offline devices (logical volumes that are mounted to local or remote TVC) with active tasks should be allowed to finish execution and volumes to unload, completing transition to offline state.

Virtual devices in other clusters should be made online to provide mount point to new jobs, shifting workload to other clusters in the grid before start service preparation. After scheduled maintenance finishes and TS7700 can be taken out of service, then virtual devices can be varied back online for accessing hosts.

Preparing the tape library for service

If the physical tape library in a TS7700 Grid must be serviced, the effect on the associated cluster must be evaluated, and the decision about whether to bring the associated cluster (TS7740 or TS7700T) into service should be made. It depends on the duration of the planned outage for the physical tape library, the role played by the tape attached cluster or partition in this particular grid architecture, the policies in force within this grid, and so on.

There might be cases where the best option is to prepare the cluster (TS7740 or TS7700T) for service before servicing the TS3500 tape library. In addition, there might be other scenarios where the preferred option is to service the tape library without bringing the associated cluster in service.

Work with the IBM SSR to identify which is the preferred option in a specific case.

For information about how to set the TS7700 to service preparation mode, see “Cluster Actions menu” on page 343.

9.4.4 The Tape Library inventory

The inventory on the TS4500 tape library can be started from the Management Interface as shown at Figure 9-202. A full tape library or frame inventory can be select.

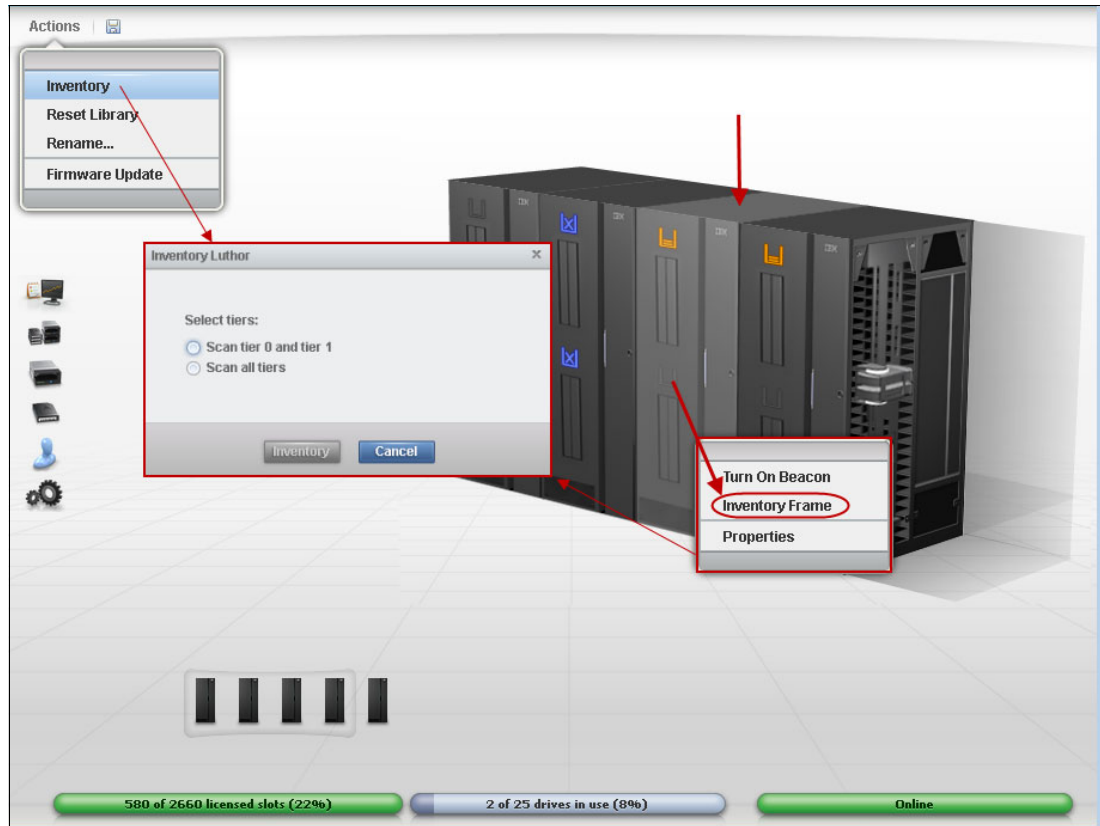


Figure 9-202 TS4500 inventory options

A partial inventory can be performed in any specific frame of the tape library: Left-click the desired frame on the tape library image to select it (it changes colors), and right-click to display the options; then, select **Inventory Frame** from the list.

A complete tape library inventory can be started from the **Actions** button on the top of the page. Both options will pop up a dialog box, asking whether to scan tiers 0 and 1, or all tiers.

The **Scan tier 0 and tier 1** option will check cartridges on the doors and the external layer of the cartridges on the walls of the library. This option will only scan other tiers if a discrepancy is found. This is the preferred option for normal tape library operations, and it can be performed concurrently.

The option **Scan all tiers** will perform full library inventory, shuffling and scanning all cartridges in all tiers. This option is not concurrent (even when selected for a specific frame) and can take a long time to complete, depending on the number of cartridges in the library. Use *scan all tiers* only when a full inventory of the tape library is required.

With TS3500 tape library, use the Tape Library Specialist page that is shown in Figure 9-203 to run Inventory/Audit. Select **All Frames** or a specific frame from the menu.

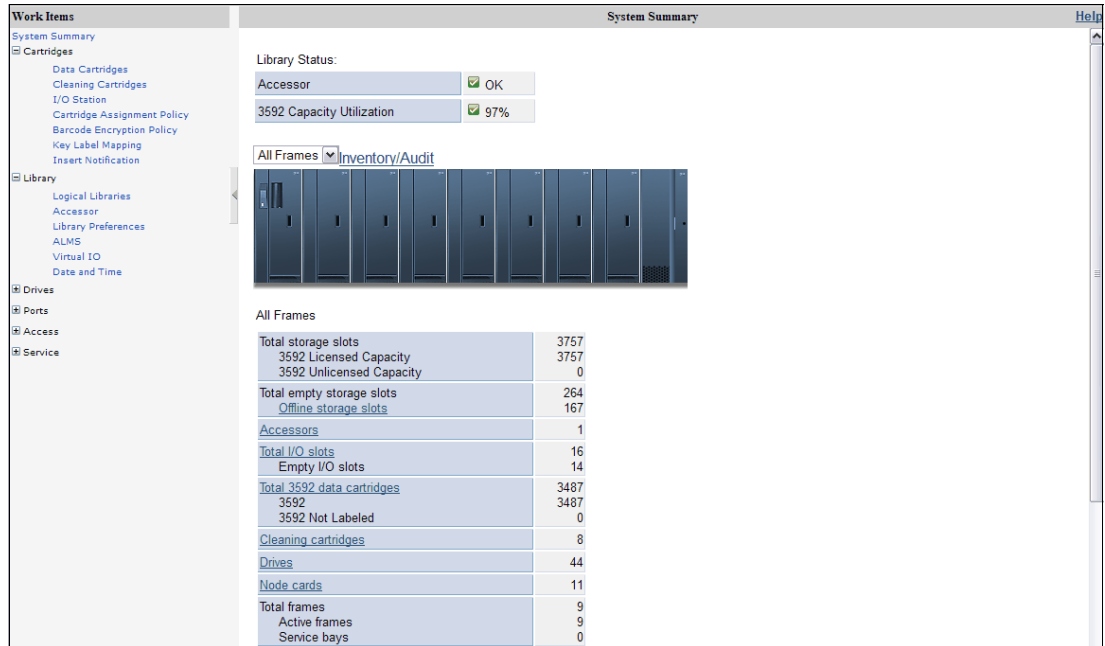


Figure 9-203 TS3500 Tape Library inventory

After you click the **Inventory/Audit** tab, a message is displayed, as shown in Figure 9-204.

Note: Perform inventory if there is no high-density frame installed on the tape library. Perform Inventory inventories high-density frame cells only for the first cartridge unless the first cartridge differs from the stored library inventory. Perform Inventory with Audit inventories all cells in a high-density frame.

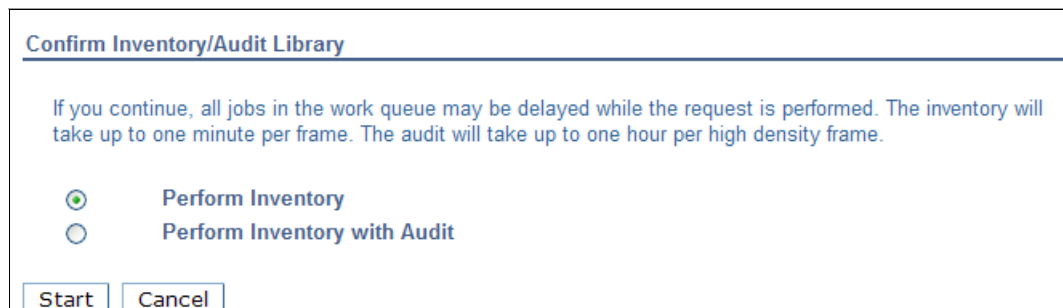


Figure 9-204 TS3500 Tape Library inventory message

Important: As stated in the confirmation window (Figure 9-204), all jobs in the work queue might be delayed while the request runs. The inventory takes up to 1 minute per frame. The audit takes about 45 minutes per high-density frame.

9.4.5 Inventory upload

For information about an inventory upload, see the “Physical Volume Ranges window” on page 422 and Figure 9-86 on page 423.

Click **Inventory Upload** to synchronize the physical cartridge inventory from the attached tape library with the TS7700T database.

Note: Perform the Inventory Upload from the TS3500 tape library to all TS7700T tape drives that are attached to that tape library whenever a library door is closed, manual inventory or Inventory with Audit is run, or a TS7700 cluster is varied online from an offline state.

9.5 Tape cartridge management

Most of the tape management operations are described in 9.1, “User interfaces” on page 320. This section provides information about tape cartridges and labels, inserting and ejecting stacked volumes.

9.5.1 3592 tape cartridges and labels

The data tape cartridge that is used in a 3592 contains the following items (numbers correspond to Figure 9-205):

- ▶ A single reel of magnetic tape
- ▶ Leader pin (1)
- ▶ Clutch mechanism (2)
- ▶ Cartridge write-protect mechanism (3)
- ▶ Internal cartridge memory (CM)

Figure 9-205 shows a J-type data cartridge.




Figure 9-205 Tape cartridge

See Table 4-8 and Table 4-9 on page 132 for a complete list of drives, models, and compatible cartridge types.

Labels

The cartridges use a media label to describe the cartridge type, as shown in Figure 9-206 (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 1 - 6 characters, which are left-aligned on the label. If fewer than 6 characters are used, spaces are added. The media type is indicated by the seventh and eighth characters.




Figure 9-206 Cartridge label

For more information about this topic, see IBM Knowledge Center:

<http://www.ibm.com/support/knowledgecenter>

9.5.2 Manual insertion of stacked cartridges

There are two methods for physically inserting a stacked volume into the 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 pauses the tape library. Therefore, use it only to add or remove large quantities of tape cartridges.

The VOLSER RANGE or CAP defines which volumes are assigned to which logical library partition in the TS4500 or TS3500 tape library. If the VOLSER that was inserted belongs to a range defined in the VOLSER RANGE or CAP, it is assigned to the associated logical library partition as defined after the library or CIO finishes inventory.

Note: Inserted volumes also should belong to a Physical Volume Range that is defined in the corresponding TS7700T cluster. Otherwise, they show as Unassigned in the TS7700 MI.

After closing the doors on the tape library and the physical inventory completes, the tape attach TS7700 processes the upload of the inventory from the tape library to the cluster, before changing the tape library status to AUTO mode in the cluster. The TS7700 updates its database with the information that is uploaded from the tape library. Pre-existing cartridges that were not physically inventoried by the tape library are set into FFFA category (manually removed) by the TS7700 at this time.

Tips:

- ▶ The inventory is performed only on the frame where the door is opened and not on the frames to either side. When inserting cartridges into a frame next to the frame with the front door opened, a manual inventory of the adjacent frame or for the whole TS3500 tape library should be selected through the GUI.
- ▶ For a tape attach TS7700, 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 this criteria, the cartridge is rejected.

Inserting cartridges by using the I/O station

The tape library detects volumes in the I/O station and scans them. With VIO enabled, the tape library moves cartridges into VIO slots for the selected logical library, as defined by CAP or volser range. If any VOLSER was not within any range that is defined by volser range or CAP, cartridges are moved into VIO slots and made available to the existing logical libraries.

With VIO disabled, tape volumes are moved from the I/O station into the tape library only upon a host command (in this case, by a TS7700 cluster) or by the tape library GUI. Within the tape library, VOLSER RANGE and CAP define which volumes are assigned to which logical library. If the VOLSER is included in the defined range, it is assigned to the proper logical library partition. If any VOLSER is not in a range that is defined by the CAP or volser range, the operator is notified by an Insert Notification message on the operator panel (for the TS3500), and prompted to assign that cartridge to a logical library.

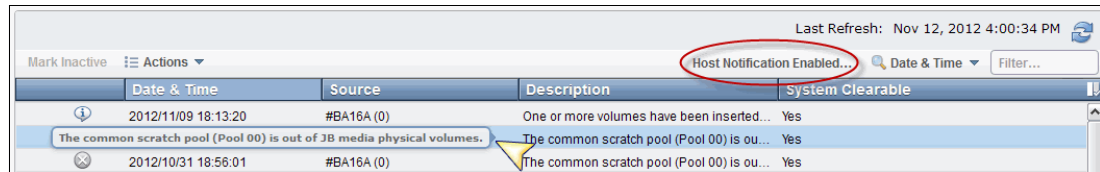
Note: The TS3500 tape library disables Insert Notification for High-Density frame configurations.

Every TS7700 tape-attached cluster must have an exclusive, dedicated logical library in the tape logical library, with its own set of cartridges. Therefore, in a library with more than one logical library, be sure that the VOLSER RANGE or CAP information is accurate and updated, avoiding cartridge assignment errors. This minimizes conflicts by ensuring that the cartridge is accessible only by the intended partition.

Consideration: Unassigned cartridges can exist in the tape library, but unassigned cartridges can have different meanings and need different actions. For more information, see *IBM TS3500 tape library with ALMS Operator Guide, GA32-0594*.

9.6 Cluster intervention scenarios

This section describes some operator intervention scenarios that might happen. Most of the errors requiring operator attention are reported on the MI or through a Host Notification, which is enabled from the Events window of the MI. For a sample of one Event message that needs an operator intervention, see Figure 9-207.



Date & Time	Source	Description	System Clearable
2012/11/09 18:13:20	#BA16A (0)	One or more volumes have been inserted...	Yes
2012/10/31 18:56:01	#BA16A (0)	The common scratch pool (Pool 00) is out of JB media physical volumes.	Yes
		The common scratch pool (Pool 00) is ou...	Yes

Figure 9-207 Example of an operator intervention

9.6.1 Hardware conditions

Some potential hardware attention scenarios are described in the following sections. The main source that is available for reference about the operational or recovery procedures is the IBM TS7700 R4.0 IBM Knowledge Center. The TS7700 IBM Knowledge Center is available directly from the TS7700 MI by clicking the question mark (?) symbol in the upper-right corner of the top bar of the MI or online at the following website:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/hydra_c_ichome.html

Most of the unusual conditions are reported to the host through Host Notification (which is enabled on the Events MI window). In a z/OS, the information messages generate the host message CBR3750I Message from library library-name: message-text, which identifies the source of the message and shows the information about the failure, intervention, or some specific operation that the TS7700 library is bringing to your attention.

The meaningful information that provided by the tape library (the TS7700 in this book) is contained in the message-text field, which can have 240 characters. This field includes a five-character message ID that might be examined by the message automation software to filter the events that should get operator attention. The message ID classifies the event being reported by its potential impact to the operations. The categories are critical, serious, impact, warning, and information. For more information, see the IBM TS7700 4.0 IBM Knowledge Center.

For the list of informational messages, see *The IBM TS7700 Series Operator Informational Messages* white paper, available at:

<https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101689>

IBM 3592 tape drive failure (TS7740 or TS7700T)

When the TS7700 determines that one of its tape drives is not operating correctly and requires service (due to read/write errors, fiber interface, or another hardware-related reason), the drive is marked offline and an IBM 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 continues with a reduced number of drives until the repair action on the drive is complete. To recover, the IBM SSR repairs the failed tape drive and makes it available for the TS7700 to use it again.

Physical volume in read-only status

The message 0P0123 Physical volume in read-only status due to successive media errors reports that a specific cartridge belonging to the TS7740 or TS7700T has exceeded the media error threshold, or encountered a permanent error during write or read operations, or is damaged. The faulty condition is reported by the tape drive to the cluster, and the cartridge is flagged Read-Only by the running cluster code. A read-only status means that new data is not written to that suboptimal media.

By default, this cartridge is corrected by an internal function of the TS7700 named Automated ROR. Make sure that the IBM SSR has enabled Automated ROR in the cluster. Automated ROR is the process by which hierarchical storage management (HSM) recalls all active data from a particular physical volume that has exceeded its error thresholds, encountered a permanent error, or is damaged.

This process extracts all active data (in the active logical volumes) that is contained in that read-only cartridge. When all active logical volumes are successfully retrieved from that cartridge, the Automated ROR process ejects the suboptimal physical cartridge from the tape library, ending the recovery process with success. Messages 0P0100 A read-only status physical volume xxxxxx has been ejected or 0P0099 Volser XXXXXX was ejected during recovery processing reports that the volume was ejected successfully.

After the ejection is complete, the cartridge VOLID is removed from the TS7700 physical cartridges inventory.

Note: By design, the tape attach TS7700 never ejects a cartridge that contains any active data. The requirement to eject a physical cartridge is to move off all active data to another physical cartridge in the same pool, and only then the cartridge can be ejected.

If Automated ROR process successfully ejects a read-only cartridge, there is no further actions to be taken, except inserting a new cartridge to replace the ejected one.

The ROR ejection task runs at a low priority to avoid causing any impact on the production environment. The complete process, from cartridge being flagged Read-Only to the 0P0100 A read-only status physical volume xxxxxx has been ejected message, signaling the end of the process, can take several hours (typically one day to complete).

If the process fails to retrieve the active logical volumes from that cartridge due to a damaged media or unrecoverable read error, the next actions depend on the current configuration that is implemented in this cluster, whether stand-alone or part of a multi-cluster grid.

In a grid environment, the ROR reaches into other peer clusters to find a valid instance of the missing logical volume, and automatically copies it back into this cluster, completing the active data recovery.

If recovery fails because there is no other consistent copy that is available within the grid, or this is a stand-alone cluster, the media is not ejected and message 0P0115 The cluster attempted unsuccessfully to eject a damaged physical volume xxxxxx is reported, along with 0P0107 Virtual volume xxxxxx was not fully recovered from damaged physical volume yyyyyy for each logical volume that failed to be retrieved.

In this situation, the physical cartridge is not ejected. A decision must be made regarding the missing logical volumes that are reported by 0P107 messages. Also, the defective cartridge contents can be verified through the MI Physical Volume Details window by clicking the **Download List of Virtual Volumes** for that damaged physical volume. Check the list of the logical volumes that are contained in the cartridge, and work with the IBM SSR if data recovery from that damaged tape should be attempted.

If those logical volumes are not needed anymore, they should be made into scratch volumes by using the TMS on the z Systems host. After this task is done, the IBM SSR can redo the ROR process for that defective cartridge (which is done from the TS7700 internal maintenance screen, and through an MI function). This time because these logical volumes that are not retrieved do not contain active data, the Automated ROR completes successfully, and the cartridge is ejected from the library.

Note: Subroutines of the same Automated ROR process are started to reclaim space in the physical volumes and to perform some MI functions, such as eject or move physical volumes or ranges from the MI. Those cartridges are made read-only momentarily during the running of the function, returning to normal status at the end of the process.

Power failure

User data is protected during a power failure because it is stored on the TVC. Any host jobs reading or writing to virtual tapes fail as they fail with a real IBM 3490E, and they must be restarted after the TS7700 is available again. When power is restored and stable, the TS7700 must be started manually. The TS7700 recovers access to the TVC by using information that is available from the TS7700 database and logs.

TS7700 Tape Volume Cache errors

Eventually, one DDM or another component might fail in the TS7700 TVC. In this situation, the host is notified by the TS7700, and the operator sees the HYDIN0571E Disk operation in the cache is degraded message. Also, the MI shows the Health Status bar (lower-right corner in Figure 9-208 on page 598) in yellow that warns you about a degraded resource in the subsystem. A degraded TVC needs an IBM SSR engagement. The TS7700 continues to operate normally during the intervention.

The MI has improved the accuracy and comprehensiveness of Health Alert messages and Health Status messages. For example, new alert messages report that a DDM failed in a specific cache drawer, which is compared to a generic message of degradation in previous levels. Also, the MI shows enriched information in graphical format, such as in Figure 9-28 on page 353.

Accessor failure and manual mode (TS7740 or TS7700T)

If the physical tape library does not have the dual accessors installed, a failure of the accessor results in the library being unable to mount automatically physical volumes. If the high availability dual accessors are installed in the tape library, the second accessor takes over. Then, the IBM SSR should be notified about repairing the failed accessor.

Gripper failure (TS7700T)

The TS3500 and TS4500 tape library have dual grippers. If a gripper fails, library operations continue with the remaining gripper. While the gripper is being repaired, the accessor is not available. If the dual accessors are installed, the second accessor is used until the gripper is repaired. For more information about operating the tape library, see the documentation for TS3500 or TS4500.

Out of stacked volumes (TS7700T)

If the tape library runs out of stacked volumes, copying to the 3592 tape drives fail, and an intervention-required message is sent to the host and the TS7700 MI. All further logical mount requests are delayed by the Library Manager until more stacked volumes are added to the tape library that is connected to the TS7700T. To recover, insert more stacked volumes. Copy processing can then continue.

Important: In a TS7700T cluster, only the tape attached partitions are affected.

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 onto the take-up spool inside the drive. If this pin gets dislodged or damaged, follow the instructions in *IBM Enterprise Tape System 3592 Operators Guide*, GA32-0465, to correct the problem.

Important: Repairing a 3592 tape must be done only for data recovery. After the data is moved to a new volume, eject the repaired cartridge from the TS7700 library.

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 or TS7700T cannot recover the contents that are configured as a stand-alone cluster. If this TS7700 cluster is part of a grid, the damaged tape contents (active logical volumes) are retrieved from other clusters, and the TS7700 has those logical volumes brought in automatically (given that those logical volumes had another valid copy within the grid).

Otherwise, this is the same for any tape drive media cartridges. Check the list of the logical volumes that are contained in cartridge, and work with the IBM SSR if data recovery from that broken tape should be attempted.

Logical mount failure

When a mount request is received for a logical volume, the TS7700 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, the way that a mount error condition is handled is different than with the prior generations of VTS.

With the prior generation of VTS, the VTS always indicated to the host that the mount completed even if a problem had occurred. When the first I/O command is sent, the VTS fails that I/O because of the error. This results in a failure of the job without the opportunity to try to correct the problem and try the mount again.

With the TS7700 subsystem, if an error condition is encountered during the execution of the mount, rather than indicating that the mount was successful, the TS7700 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 9-2.

Example 9-2 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:

- ▶ For example, look at CBR4171I. Reason codes are documented in IBM Knowledge Center. As an exercise, assume RSN=32. In IBM Knowledge Center, the reason code is as follows:

Reason code x'32': Local cluster recall failed; the stacked volume is unavailable.

- ▶ CBR4196D: Error code shows in the format 14xxIT:
 - 14 is the permanent error return code.
 - xx is 01 if the function was a mount request or 03 if the function was a wait request.
 - IT is the permanent error reason code. The recovery action to be taken for each CODE.
 - In this example, it is possible to have a value of 140194 for the error code, which means xx=01: Mount request failed.
- ▶ IT=94: Logical volume mount failed. An error was encountered during the execution of the mount request for the logical volume. The reason code that is associated with the failure is documented in CBR4171I. The first book title includes the acronyms for message IDs, but the acronyms are not defined in the book.

For CBR messages, see *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, SA38-0671, for an explanation of the reason code and for specific actions that should be taken to correct the failure. See *z/OS DFSMSdfp Diagnosis*, SC23-6863, for OAM return and reason codes. Take the necessary corrective action and reply 'R' to try again. Otherwise, reply 'C' to cancel.

Tip: Always see the appropriate documentation (TS7700 IBM Knowledge Center and MVS System Messages) for the meaning of the messages and the applicable recovery actions.

Orphaned logical volume

This situation occurs when the TS7700 database has a reference to a logical volume but no reference to its physical location. This can result from hardware or internal processing errors. For more information about orphaned logical volume messages, contact your IBM 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 TS7740 or TS7700T MI and sent to the host console (Example 9-3).

Example 9-3 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 TS7740 or TS7700T Licensed Internal Code. No user action is needed because 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. The stacked physical volume goes through the ROR process and is ejected after all active data is recalled. This process is handled automatically by the TS7700.

9.6.2 TS7700 LIC processing failure

If a problem develops with the TS7700 Licensed Internal Code (LIC), the TS7700 sends an intervention-required message to the TS7700 MI and host console, and attempts to recover. In the worst case, this situation involves a restart of the TS7700 itself. If the problem persists, your IBM SSR should be contacted. The intervention-required message that is shown in Example 9-4 is sent to the host console.

Example 9-4 VTS software failure

```
CBR3750I MESSAGE FROM LIBRARY lib: Virtual Tape z Systems has a CHECK-1 (xxxx)
failure
```

The TS7700 internal recovery procedures handle this situation and restart the TS7700. For more information, see Chapter 13, “Disaster Recovery Testing” on page 787.

9.7 TS7700 Management Interface considerations

In the TS7700 MI, some operations or functions might be unavailable or disabled, depending on the cluster configuration (disk only or tape attach TS7700 clusters). Functions or operations that apply only to a disk-only configuration might be unavailable or disabled on a tape attached configuration (TS7740 or a tape partition in a TS7700T) and vice versa. Also, some operations or functions might be available or disabled depending on whether the cluster is part of a grid configuration or not.

The following figures show the Cluster Summary window for disk-only and tape-attached configurations, highlighting some particularities.

Monitoring the health with TS7720

Figure 9-208 shows the Health Status bar as displayed in a TS7720 configuration. The TS7720 on the left is configured with an expansion frame. Compare it with the illustration of a TS7720 with only the base frame on the right.

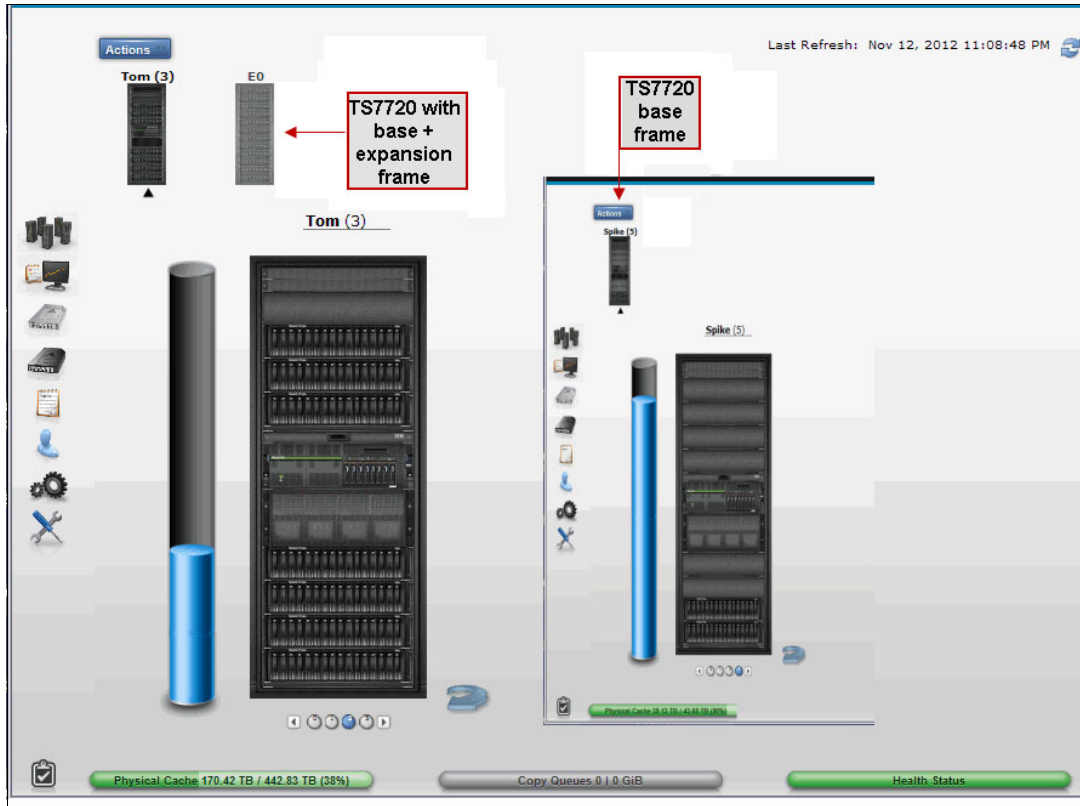


Figure 9-208 TS7720 MI health and monitoring options

More details about the health state of each component of the TS7700 can be obtained by hovering your cursor over the picture of the cluster. Also, components in the back of the frame can be viewed by clicking the blue circular arrow near the lower-right corner of the frame. This arrow flips the picture, showing the back side. Again, hover the mouse over the components for the health details of those components.

Compare Figure 9-208 on page 598 with Figure 9-209, and notice that a TS3500 tape library icon is displayed for the TS7740 or the TS7720T. Hover your cursor over the TS3500 icon to see information about health of the tape library and tape drives.

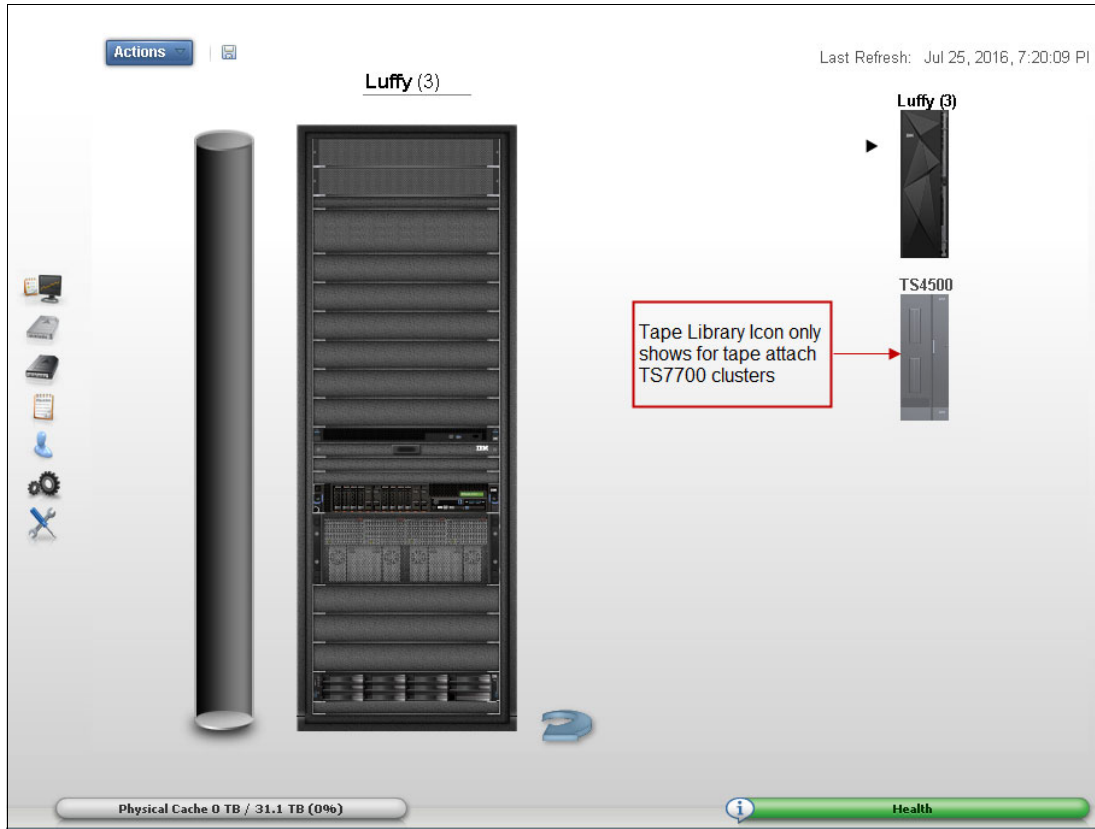


Figure 9-209 Tape Attached MI

For the remaining possible options, see 9.2, "TS7700 Management Interface" on page 326.



Host Console operations

This chapter provides information about how to operate the IBM TS7700, with emphasis on commands and procedures that are initiated from the host operating system.

This chapter includes the following sections:

- ▶ System-managed tape
- ▶ Messages from the library
- ▶ EXPIRE HOLD and scratch processing considerations
- ▶ Scratch count mismatch
- ▶ Effects of changing categories
- ▶ Library messages and automation
- ▶ Return-to-Scratch Enhancement
- ▶ Deleting Virtual Volumes

10.1 System-managed tape

This section describes the commands that are used to operate a tape library in an IBM 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 some of the more useful DFSMS and MVS commands.

10.1.1 DFSMS operator commands

Some of the commands contain *libname* as a variable. In this case, the storage management subsystem (SMS)-defined library name is required. The output for some of these commands differs slightly depending on whether you reference a TS7700 composite library or distributed library. For more information about DFSMS commands, see *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

Information from the IBM TS4500/TS3500 Tape Library is contained in some of the outputs. However, you cannot switch the operational mode of the TS4500/TS3500 Tape Library with z/OS commands.

Consideration: DFSMS and MVS commands apply only to SMS-defined libraries. The library name that is defined during the definition of a library in Interactive Storage Management Facility (ISMF) is required for *libname* in the DFSMS commands. The activation of a source control data set (SCDS) with this *libname* must have already been performed for SMS to recognize the library.

The following DFSMS operator commands support the tape library:

► **DISPLAY SMS, LIBRARY(*libname*|ALL), STATUS**

This is an SMS Configuration level view, which indicates whether the SMS-defined libraries are online, offline, or pending offline, on each of the systems in the configuration.

STATUS is the default parameter.

► **DISPLAY SMS, LIBRARY(ALL), DETAIL**

The DETAIL display, although a single-system view, gives slightly more information. This is the suggested method to display a high-level overview of each library that is defined to SMS in the configuration. See Example 10-1.

Example 10-1 D SMS,LIB(ALL),DETAIL

```
D SMS,LIB(ALL),DETAIL
CBR1110I OAM library status: 738
TAPE      LIB  DEVICE  TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON OP
LIBRARY   TYP  TYPE    DRV  DRV  DRV  SLOTS  SLOTS  VOLS
CLIB00    VCL  GRID    512  0    0    0      0      0    N  Y
DTS7720   VDL  3957-VEB  0    0    0    559    516    0    Y  Y
D0001     VDL  3957-V07  0    0    0    1000   960    0    Y  N
D0002     VDL  3957-V07  0    0    0    1000   880    0    Y  N
E0001     VDL  3957-V07  0    0    0    1000   883    0    Y  N
E0002     VDL  3957-V07  0    0    0    1000   880    0    Y  N
HYDRAE    VDL  3957-V07  0    0    0    185    129    0    Y  Y
HYDRAG    VCL  GRID    512  2    2    0      0    45547 Y  Y
```

► **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). The status lines at the bottom of the output are surfaced based on information that is obtained directly from the library. See Example 10-2.

Example 10-2 D SMS,LIB(libname)DETAIL

```
D SMS,LIB(HYDRAG),DETAIL
CBR1110I OAM library status: 754
TAPE      LIB  DEVICE  TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON OP
LIBRARY   TYP  TYPE    DRV  DRV  DRV  SLOTS  SLOTS  VOLS
HYDRAG    VCL  GRID    512  2   2    0      0   45547  Y  Y
-----
MEDIA          SCRATCH      SCRATCH      SCRATCH
TYPE           COUNT      THRESHOLD    CATEGORY
MEDIA1         10          0            0021
MEDIA2        45537      0            0022
-----
DISTRIBUTED LIBRARIES:  HYDRAE  DTS7720
-----
LIBRARY ID: 00186
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:          1
CORRUPTED TOKEN VOLUME COUNT:         24
-----
Library supports import/export.
Library supports outboard policy management.
Library supports logical WORM.
Library enabled for scratch allocation assistance.
```

► **DISPLAY SMS,VOLUME(volser)**

This command displays all of the information that is stored about a volume in the tape configuration database (TCDB), also known as the VOLCAT, as well as information obtained directly from the library, such as the LIBRARY CATEGORY, LM constructs (SMS constructs stored in the library), and LM CATEGORY. See Example 10-3.

Example 10-3 D SMS,VOL

```
D SMS,VOL(B00941)
RESPONSE=MZPEVS2
CBR1180I OAM tape volume status: 195
VOLUME  MEDIA  STORAGE  LIBRARY  USE  W  C  SOFTWARE  LIBRARY
        TYPE   GROUP    NAME     ATR  P  P  ERR STAT  CATEGORY
B00941  MEDIA2  SGG00001  HYDRAG  P   N  N  NOERROR  PRIVATE
-----
RECORDING TECH:      36 TRACK          COMPACTION:          YES
SPECIAL ATTRIBUTE:  NONE          ENTER/EJECT DATE:   2011-02-14
CREATION DATE:      2011-02-14    EXPIRATION DATE:    2014-11-15
LAST MOUNTED DATE:  2014-11-10  LAST WRITTEN DATE:  2014-11-10
SHELF LOCATION:
OWNER: DENEKA
LM SG: SGG00001  LM SC: SC00030R  LM MC: MNDNN020  LM DC: D000N004
LM CATEGORY: 002F
```

► **DISPLAY SMS,OAM**

This command, which is shown in Example 10-4, is primarily useful for checking the status of the object access method (OAM) user exits.

Example 10-4 D SMS,OAM

```
D SMS,OAM
RESPONSE=MZPEVS2
CBR1100I OAM status: 744
TAPE TOT  ONL  TOT  TOT  TOT  TOT  TOT  ONL  AVL  TOTAL
      LIB  LIB  AL  VL  VCL  ML  DRV  DRV  DRV  SCRTCH
        3   1   0   0   3   0  1280   2   2   45547
There are also 7 VTS distributed libraries defined.
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.
```

► **VARY SMS,LIBRARY(1ibname),ONLINE/OFFLINE**

From the host standpoint, the vary online and vary offline commands for a TS7700 library always use the library name as defined through ISMF.

This command acts on the SMS library, which is referred to as *libname*. Using this command with the OFFLINE parameter stops tape library actions and gradually makes all of the tape units within this logical library unavailable. This simple form is a single-system form. The status of the library remains unaffected in other MVS systems.

Note: A composite and distributed IBM Virtual Tape Server (VTS) library can be varied online and offline like any VTS library, though varying a distributed library offline from the host really has no meaning (does not prevent outboard usage of the library). Message CBR3016I warns the user when a distributed library is offline during OAM initialization or varied offline.

Using this command with the ONLINE parameter 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 a 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. However, this is only the MVS part of error recovery. You must clear the hardware, software, or operational error within the physical library and TS7700 before you bring the library online to MVS.

► **VARY SMS,LIBRARY(1ibname,sysname,...),ON/OFF and VARY SMS,LIBRARY(1ibname,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 runs the VARY action on all systems within the SMSplex.

The **VARY SMS** command enables the short forms ON as an abbreviation for ONLINE and OFF as an abbreviation for OFFLINE.

► **LIBRARY EJECT,volser{,PURGE|KEEP|LOCATION}{,BULK}**

This command is used to request the ejection of a volume from a tape library. The following options are available for this command:

- Eject to the convenience I/O station for physical volumes. Delete from the tape library for logical volumes that are considered scratch.
- Eject to the bulk output station (BULK or B) for physical volumes. Delete from the tape library for logical volumes that are considered scratch.
- 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 write to operator with reply (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 information in this field (LOCATION or L). The operator must enter the new information as a reply to WTOR.

If none of the variations (PURGE, KEEP, or LOCATION) is 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 performed through program interfaces, such as ISMF, a tape management system (TMS), or a batch job.

10.1.2 MVS system commands

The following commands are described in detail in *z/OS MVS System Commands*, SA22-7627:

► **VARY *unit*,ONLINE/OFFLINE**

The **VARY *unit*** command is no different from what it was before. 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 only for library reasons cannot be varied online by running **VARY *unit*,ONLINE**. Only **VARY SMS,LIBRARY(...),ONLINE** can do so.

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.

► **LIBRARY DISPDRV,*library_name***

The **LIBRARY DISPDRV (LI DD)** command indicates whether a device is online or offline, and the reason if it is offline. With OAM APAR OA47487 and Release 3.3 installed, a new keyword **MOUNTED** is added to this command. This keyword specifies that status information should be displayed for volumes that are mounted in the TS7700 for the specified library (composite or distributed). Information pertaining to the distributed library that owns the device for the mount, and distributed library information that is associated with the primary and the secondary Tape Volume Cache (TVC), is displayed.

There is an **ALL** parameter that can be passed with **MOUNTED**. **ALL** specifies that additional drives can be displayed that are not owned by the distributed library that is targeted in the command. The additional drives are displayed if the distributed library that is specified is the primary or secondary TVC for the mounted volume.

As with the existing **LIBRARY DISPDRV** command, the host must be able to communicate with the device and the device can be online or offline. Unlike the existing **LIBRARY DISPDRV** command, the OAM address space must have been started on the host, but does not have to be currently active.

The intent of the addition of **MOUNTED** to the **LI DD** command is to provide a way to tell which devices and volumes are mounted where in the grid without having to query individual volumes. This parameter should aid in the process of placing a cluster into service mode by simplifying the process of identifying which devices must be varied offline.

Here is the syntax for the **LIBRARY DISPDRV MOUNTED,ALL** command:

```
LIBRARY DISPDRV,library_name,MOUNTED and
LIBRARY DISPDRV,library_name,MOUNTED,ALL
```

or

```
LI DD,library_name,M
LI DD,library-name,M,A
```

library_name can be a composite or distributed library. The following examples illustrate the differences in the display output. Here is the configuration for the clusters in the example COMPLIB1 grid:

- DISTLIB1 – (1C00 – 1CFF): Devices 1C05, 1C10, 1C25, 1C30, 1C45, and 1C48 are mounted. The mounts are satisfied by the primary TVC being DISTLIB1. Synchronous mode copy is not used for allocations that are directed to this cluster.
- DISTLIB2 – (1D00 - 1DFF): Devices 1D03, 1D1C, 1D22, 1D35, and 1D42 are mounted. The primary TVC is DISTLIB1 for some of the volumes and DISTLIB2 for others. The mounted volumes are in synchronous mode and copied to DISTLIB3.
- DISTLIB3 – (1E00 - 1EFF): Devices 1E1F, 1E21, 1E30, 1E56, and 1E68 are mounted. The primary TVC is a combination of all three clusters. Synchronous mode copy is not used for allocations that are directed to this cluster.

In Example 10-5, all volumes that are mounted in the grid are displayed along with the distributed library on which they are mounted.

Example 10-5 LIBRARY DISPDRV MOUNTED command against a composite library

```
LIBRARY DISPDRV,COMPLIB1,MOUNTED
CBR1230I Mounted status:
DRIVE  COMPLIB  ON  MOUNT  DISTLIB  PRI-TVC  SEC-TVC
NUM    NAME      VOLUME Name    DISTLIB  DISTLIB
1C05   COMPLIB1  Y  A00100  DISTLIB1  DISTLIB1
1C10   COMPLIB1  Y  A00108  DISTLIB1  DISTLIB1
1C25   COMPLIB1  Y  A00115  DISTLIB1  DISTLIB1
1C30   COMPLIB1  Y  A00050  DISTLIB1  DISTLIB1
1C45   COMPLIB1  Y  A00142  DISTLIB1  DISTLIB1
1C48   COMPLIB1  Y  A01001  DISTLIB1  DISTLIB1
1D03   COMPLIB1  Y  A00118  DISTLIB2  DISTLIB1  DISTLIB3
1D1C   COMPLIB1  Y  A00124  DISTLIB2  DISTLIB2  DISTLIB3
1D22   COMPLIB1  Y  A00999  DISTLIB2  DISTLIB1  DISTLIB3
1D35   COMPLIB1  Y  A00008  DISTLIB2  DISTLIB2  DISTLIB3
1D42   COMPLIB1  Y  A00175  DISTLIB2  DISTLIB1  DISTLIB3
1E1F   COMPLIB1  Y  A00117  DISTLIB3  DISTLIB3
```

1E21	COMPLIB1	Y	A02075	DISTLIB3	DISTLIB1
1E30	COMPLIB1	Y	A01070	DISTLIB3	DISTLIB1
1E56	COMPLIB1	Y	A00004	DISTLIB3	DISTLIB2
1E68	COMPLIB1	Y	A00576	DISTLIB3	DISTLIB3

In Example 10-6, the command is directed to a specific distributed library in the grid. You see the mounts for devices only in that specific distributed library.

Example 10-6 LIBRARY DISPDRV MOUNTED command against a distributed library

```
LIBRARY DISPDRV,DISTLIB3,MOUNTED
CBR1230I Mounted status:
DRIVE  COMPLIB  ON  MOUNT  DISTLIB  PRI-TVC  SEC-TVC
NUM    NAME          VOLUME Name      DISTLIB  DISTLIB
1E1F   COMPLIB1  Y   A00117 DISTLIB3 DISTLIB3
1E21   COMPLIB1  Y   A02075 DISTLIB3 DISTLIB1
1E30   COMPLIB1  Y   A01070 DISTLIB3 DISTLIB1
1E56   COMPLIB1  Y   A00004 DISTLIB3 DISTLIB2
1E68   COMPLIB1  Y   A00576 DISTLIB3 DISTLIB3
```

In Example 10-7, you add the **ALL** keyword to the command. It now includes all mounts where the DISTLIB, PRI-TVC, or SEC-TVC is the distributed library that is specified on the command.

Example 10-7 LIBRARY DISPDRV MOUNTED ALL command against a distributed library

```
LIBRARY DISPDRV,DISTLIB3,MOUNTED,ALL
CBR1230I Mounted status:
DRIVE  COMPLIB  ON  MOUNT  DISTLIB  PRI-TVC  SEC-TVC
NUM    NAME          VOLUME Name      DISTLIB  DISTLIB
1D03   COMPLIB1  Y   A00118 DISTLIB2 DISTLIB1 DISTLIB3
1D1C   COMPLIB1  Y   A00124 DISTLIB2 DISTLIB2 DISTLIB3
1D22   COMPLIB1  Y   A00999 DISTLIB2 DISTLIB2 DISTLIB3
1D35   COMPLIB1  Y   A00008 DISTLIB2 DISTLIB2 DISTLIB3
1D42   COMPLIB1  Y   A00075 DISTLIB2 DISTLIB1 DISTLIB3
1E1F   COMPLIB1  Y   A00117 DISTLIB3 DISTLIB3
1E21   COMPLIB1  Y   A02075 DISTLIB3 DISTLIB1
1E30   COMPLIB1  Y   A01070 DISTLIB3 DISTLIB1
1E56   COMPLIB1  Y   A00004 DISTLIB3 DISTLIB2
1E68   COMPLIB1  Y   A00576 DISTLIB3 DISTLIB3
```

For a complete description of this command and its output, see APAR OA47487.

► **DISPLAY M=DEV(xxxx)**

The **D M=DEV** command is useful for checking the operational status of the paths to the device. See Example 10-8.

Example 10-8 D M=DEV

```
D M=DEV(2500)
IEE174I 04.29.15 DISPLAY M 626
DEVICE 02500 STATUS=OFFLINE
CHP          B2  B3  B8  B9
ENTRY LINK ADDRESS  20  21  22  23
DEST LINK ADDRESS  D4  D5  D6  D7
PATH ONLINE          Y  Y  N  N
CHP PHYSICALLY ONLINE Y  Y  Y  Y
PATH OPERATIONAL    Y  Y  Y  Y
```

```

MANAGED             N   N   N   N
CU NUMBER           2500 2500 2500 2500
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND           = NOT AVAILABLE
SCP TOKEN NED       = 003490.C2A.IBM.78.0000000H6395.0000
SCP DEVICE NED      = 003490.C2A.IBM.78.0000000H6395.0000

```

► **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 is shown as 359L.

► **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 and is in private status in the TCDB.

► **UNLOAD *devnum***

The **UNLOAD** command enables you to unload a drive, if the Rewind Unload (RUN) process was not successful initially.

10.1.3 Host Console Request function

The **LIBRARY REQUEST** host console command (**LI REQ**) provides a simple way for an operator to determine the status of the TS7700 to obtain information about the resources of the TS7700, and to run an operation in the TS7700. It can also be used with automation software to obtain and analyze operational information that can then be used to alert a storage administrator that something must be examined further.

With the 3.2 code release, the TS7700 Management Interface (MI) enables an operator to issue a Library Request host console command as through it was issued from the z/OS host. The result of the command is displayed on the MI window.

The command accepts the following parameters:

- A library name, which can be a composite or a distributed library.
- It also enables 1 - 4 keywords, with each keyword being a maximum of 8 characters.

The specified keywords are passed to the TS7700 identified by the library name to instruct it about what type of information is being requested or which operation is to be run. Based on the operation that is requested through the command, the TS7700 then returns information to the host that is displayed as a multiline write to operator (WTO) message.

Note: The information that is presented in the WTO message comes directly from the hardware as a response to the LI REQ command. If you have a question about the information that is presented to the host in the WTO messages that are generated, contact hardware support.

This section describes some of the more useful and common **LI REQ** commands that a client uses. 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:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

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 10-9 on page 609.

Example 10-9 Host Console Request function syntax

```

>> _ LIBRARY _ REQUEST _ , _ library_name _____ >
      | _LI _____ | | _REQ _____ |

> _ _ , keyword1 _____ ><
      | _ , keyword2 _____ | | _ , L = _ a _____ |
      | _ , keyword3 _____ | | _ name _____ |
      | _ , keyword4 _____ | | _ name - a _____ |
  
```

The following parameters are required:

REQUEST REQ	Specifies a request to obtain information from the TS7700, or to run an outboard operation.
library_name	Specifies the library name that is associated with the TS7700 to which the request needs to be directed. The library name that is specified can 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 run on the TS7700.

The following parameters are optional. The optional parameters depend on the first keyword specified. Based on the first keyword that is specified, zero or more of the additional keywords might be appropriate:

keyword2	Specifies additional information in support of the operation that is specified with the first keyword.
keyword3	Specifies additional information in support of the operation that is specified with the first keyword.
keyword4	Specifies additional information in support of the operation that is specified with the first keyword.
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, the composite library name must be specified.
- ▶ If the request is specific to a distributed library, 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.

Overview of the Host Console Request functions

Table 10-1 lists some of the LI REQ commands, and gives a short description of each of them.

Table 10-1 Overview of Host Console Request functions

Keyword1	Keyword2	Keyword3	Keyword4	Description	Comp Lib	Dist Lib	TS7720 TS7720T TS7740
CACHE				Requests information about the current state of the cache and the data that is managed within it.	N/A	Y	ALL
COPY	ACTIVITY	See the User's Guide		Requests information about Active Copy jobs.	N/A	Y	ALL
COPY	SUMMARY			Requests information about all the copy jobs.	N/A	Y	ALL
LVOL	volser	FLASH		Requests information about a specific logical volume.	Y	N/A	ALL
LVOL	volser	PREFER MIGRATE REMOVE REMOVE REMOVE	PROMOTE INFO	Requests a change in the cache management for a logical volume.	N/A N/A N/A N/A N/A	Y Y Y Y Y	ALL TS7740 TS7720 TS7720 TS7720
LVOL	volser	COPY	KICK FORCE	KICK requests to move a logical volume to the front of the copy queue. FORCE puts a copy job against a removed logical volume and promotes it to the front of the copy queue. It is useful when it is required to get a removed volume back into a TS7720 by copying it from another consistent cluster.	N/A	Y	ALL
PDRIVE				Requests information about the physical drives and their current usage associated with the specified distributed library.	N/A	Y	TS7740 and TS7720T
POOLCNT	00-32			Requests information about the media types and counts, which are associated with a specified distributed library, for volume pools beginning with the value in keyword2.	N/A	Y	TS7740 and TS7720T

Keyword1	Keyword2	Keyword3	Keyword4	Description	Comp Lib	Dist Lib	TS7720 TS7720T TS7740
PVOL	volser			Requests information about a specific physical volume.	N/A	Y	TS7740 and TS7720T
PVOL	volser	DELETE		Requests the specified physical volume record to be deleted from the TS7700 database. The specified physical volume must be empty and not physically in the library.	N/A	Y	TS7740 and TS7720T
RECALLQ	volser			Requests the content of the recall queue, starting with the specified logical volume. Keyword2 can be blank.	N/A	Y	TS7740 and TS7720T
RECALLQ	volser	PROMOTE		Requests that the specified logical volume be promoted to the top of the recall queue.	N/A	Y	TS7740 and TS7720T
RRCLSUN	ENABLE DISABLE STATUS			In response to the RRCLSUN request, the cluster that is associated with the distributed library enables, disables, or displays the status of the force residency on recall feature.	N/A	Y	TS7740 and TS7720T
SETTING	ALERT, CACHE, THROTTLE DEVALLOC RECLAIM CPYCNT COPY LINK DELEXP EXISTDEL	See Settings descriptions after this table	See Settings descriptions after this table	Settings to control functions in the grid.	N/A	Y	ALL
SETTINGS2	SCRATCH CACHE	PFRLOCO MAXLGMC	ENABLE DISABLE	Additional settings to control functions in the grid.	N/A	N/A	ALL
STATUS	GRID			Requests information about the copy, reconcile, and ownership takeover status of the libraries in a grid configuration.	Y	N/A	ALL

Keyword1	Keyword2	Keyword3	Keyword4	Description	Comp Lib	Dist Lib	TS7720 TS7720T TS7740
STATUS	GRIDLINK			Requests information about the status and performance of the links between the TS7700 tape drives in the grid configuration.	N/A	Y	ALL
COPYRFSH	volser	See the User's Guide	See the User's Guide	Refreshes copy policy and queue a copy job on the copy target clusters without mounting or dismounting a volume.	N/A	Y	ALL
DRSETUP	CCCCCCC (DR family name)	ADD REMOVE	0-7 (cluster ID)	Adds/removes a cluster to/from the disaster recovery (DR) family.	Y	N/A	ALL
		WP	ENABLE/ DISABLE	Enables/disables write protect mode within the DR family.	Y	N/A	ALL
		FLASH	ENABLE/ DISABLE	Enables/disables Flash Copy within the DR family.	Y	N/A	ALL
		DOALL	ENABLE/ DISABLE	Enables/disables write protect mode and FlashCopy with a single command.	Y	N/A	ALL
		LIVECOPY	FAMILY/ NONE	Enables or disallows the use of a live copy within the DR family.	Y	N/A	ALL
		SELFLIVE	ENABLE/ DISABLE	Enables or disables accessing live copy created after time zero.	N	Y	ALL
	SHOW	CCCCCCC (DR family name)		Views information about the DR family.	Y	N/A	ALL
PARTRFSH				Changes cache partition assignment.			TS7720T

Overview of the Host Console SETTING request

The SETTING request provides information about many of the current workflow and management settings of the cluster that is specified in the request and the ability to modify the settings. It also enables alerts to be set for many of the resources that are managed by the cluster.

In response to the SETTING request, the cluster that is associated with the distributed library names that are specified on the request modifies its settings based on the additional keywords specified. If no additional keywords are specified, the request returns the current settings. When a value is specified, lead blanks or zeros are ignored.

ALERT settings

Thresholds can be set for many of the resources that are 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 effect to the operations of the attached hosts. A second provides a warning that the resource has exceeded a value that might result in an effect to the operations of the attached hosts. When the second warning is reached, the warning message is repeated every 15 minutes.

These threshold settings are described in detail in Chapter 11, “Performance and monitoring” on page 635.

In the ALERT settings, you can also specify how messages are treated, and influence other aspects of TS7700 behavior.

Figure 10-1 shows the alert thresholds available for various resources that are managed by the cluster.

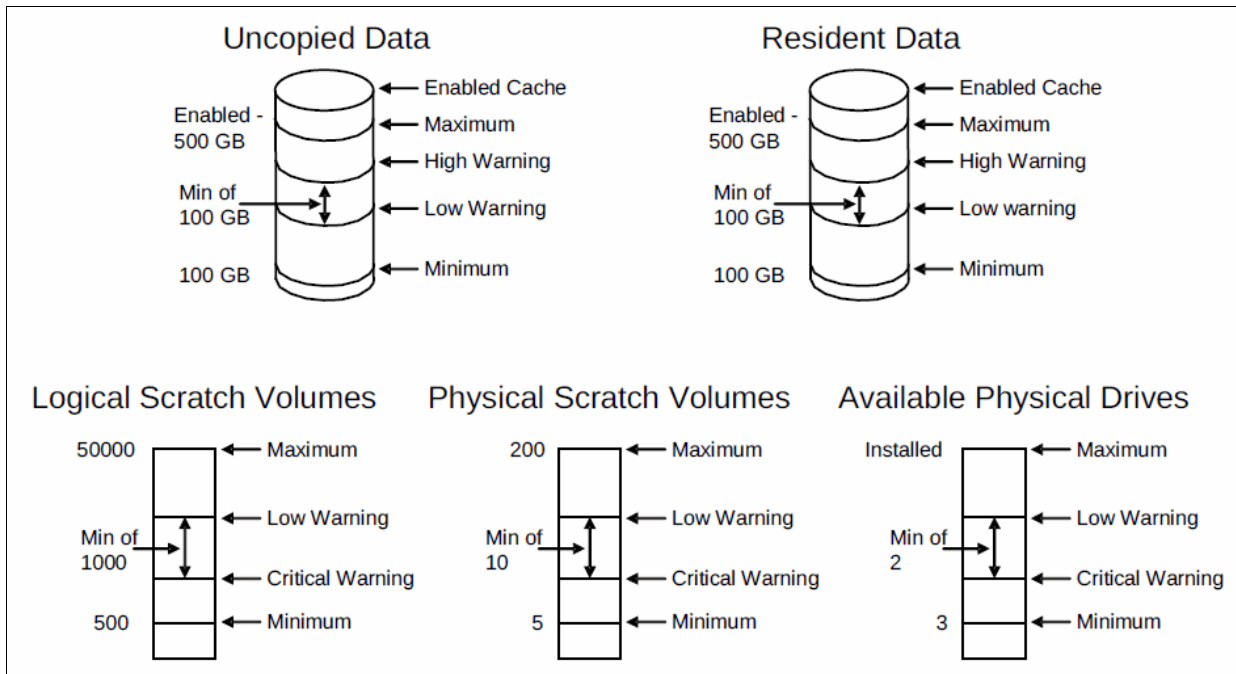


Figure 10-1 Alert setting diagram

Remember: All settings are persistent across system restarts, service actions, or code updates. The settings are not carried forward as part of disaster recovery (DR) from copy-exported tapes or the recovery of a system.

Table 10-2 shows the ALERT thresholds that are supported.

Table 10-2 ALERT thresholds

Keyword3	Keyword4	Description
COPYHIGH	value	<p>Uncopied Data High Warning Limit This is the threshold, in gigabytes of data in cache, that needs to be copied to other TS7700 tape drives in a grid configuration, at which point the TS7700 generates a message indicating that the amount of uncopied data has exceeded a high warning limit. The following message text is shown:</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL5000 Uncopied data of xxxxxxxx GB above high warning limit of yyyyyyyy GB. ▶ When below the threshold: AL5001 No longer above uncopied data high warning limit of yyyyyyyy GB.
COPYLOW	value	<p>Uncopied Data Low Warning Limit This is the threshold, in gigabytes of data, in cache that must to be copied to other TS7700 tape drives in a grid configuration, at which the TS7700 generates a message indicating that the amount of uncopied data has exceeded a low warning limit. The following message text is shown:</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL0000 Uncopied data of xxxxxxxx GB above low warning limit of yyyyyyyy GB. ▶ When below the threshold: AL0001 No longer above uncopied data low warning limit of yyyyyyyy GB.
PDRVCRIT	value	<p>Available Physical Drive Critical Warning Limit This is the threshold, in number of physical drives, at which the TS7700 generates a message indicating that the number of available physical drives has fallen below the critical warning limit. The following message text is shown:</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.
PDRVLOW	value	<p>Available Physical Drive Low Warning Limit This is the threshold, in number of physical drives, at which the TS7700 generates a message indicating that the number of available physical drives has fallen below the low warning limit. The following message text is shown:</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.
PSCRCRIT	value	<p>Physical Scratch Volume Critical Warning Limit This is the threshold, in number of scratch physical volumes, at which the TS7700 generates a message indicating that the number of available scratch physical volumes has fallen below the critical warning limit. The following message text is shown:</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>Tip: The TS7700 enters panic reclaim if the number of scratch volumes available to a defined pool is less than two, including ones that it can borrow from pool 0.</p>

Keyword3	Keyword4	Description
PSCRLOW	value	<p>Physical Scratch Volume Low Warning Limit</p> <p>This is the threshold, in number of scratch physical volumes, at which the TS7700 generates a message indicating that the number of available scratch physical volumes has fallen below the low warning limit. The following text is shown:</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>Tip: The TS7700 enters panic reclaim if the number of scratch volumes available to a defined pool is less than two, including ones that it can borrow from pool 0.</p>
RESDHIGH	value	<p>Resident Data High Warning Limit</p> <p>This is the threshold, in gigabytes of resident data, at which the TS7700 generates a message indicating that the amount of resident data has exceeded a high warning limit. The following message text is shown:</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL5008 Resident data of xxxxxxxx GB above high warning limit of yyyyyyyy GB. ▶ When below the threshold: AL5009 No longer above resident data high warning limit of yyyyyyyy GB.
RESDTHIGH	value	<p>This is the same threshold with RESDHIGH, but only applicable to the total TS7720 TA tape attached cache partitions (CPx). The following message test is shown:</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL5015 Sum of resident data in tape partitions of xxxxxxxx GB above high warning limit of yyyyyyyy GB. ▶ When below the threshold: AL5016 Sum of resident data in tape partitions no longer above resident data high warning limit of yyyyyyyy GB.
RESDLOW	Value	<p>Resident Data Low Warning Limit</p> <p>This is the threshold, in gigabytes of resident data, at which the TS7700 generates a message indicating that the amount of resident data has exceeded a low warning limit. The following message text is shown:</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.
RSDTLOW	Value	<p>This is the same threshold with RESDLOW, but only applicable to the total TS7720 TA tape attached cache partitions (CPx). The following message text is shown:</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL0012 Sum of resident data in tape partitions of xxxxxxxx GB above low warning limit of yyyyyyyy GB. ▶ When below the threshold: AL0013 Sum of resident data in tape partition no longer above low warning limit of yyyyyyyy GB.
PCPYLOW	Value	<p>Pending Copy Low</p> <p>This is the threshold in gigabytes of volumes in the copy queue. The following message text is presented when the level falls below or rises above the threshold:</p> <ul style="list-style-type: none"> ▶ AL0011 Distributed Library xx has successfully fallen below the inbound copy backlog low warning limit of zzzz GB. ▶ AL0010 Distributed Library xx has a total pending inbound copy backlog of yyyy GB, which is above the low warning limit of zzzz GB.

Keyword3	Keyword4	Description
PCPYCRIT	Value	Pending Copy Critical This is the upper limit in gigabytes for volumes in the copy queue. The same messages, AL0011 and AL0010, are presented (see the PCPYLOW keyword section).
DEFDEG	ENABLE DISABLE	<i>Synchronous deferred</i> or the <i>immediate deferred</i> condition occurs: <ul style="list-style-type: none"> ▶ When the ENABLE keyword is specified, the degraded state is reported to the host through the operational state change attention. ▶ When the DISABLE keyword is specified, the degraded state is not reported to the host through the operational state change attention. The following messages are shown: <ul style="list-style-type: none"> ▶ G0005 Distributed Library xx has entered the immediate deferred state. ▶ G0032 Distributed Library xx has entered the synchronous deferred state due to volser yyyyyy are generated regardless of the DEFDEG setting.
LINKDEG	ENABLE DISABLE	Prevents a composite library from entering the link degraded state when Grid link degradation occurs. When the ENABLE keyword is specified, the degraded state is reported to the host through the operational state change attention. When the DISABLE keyword is specified, the degraded state is NOT reported to the host through the operational state change attention. The following operator messages are still generated regardless of the LINKDEG setting: <ul style="list-style-type: none"> ▶ G0030 Library XXXXX, PPP, AAA Grid Link are degraded. ▶ G0031All grid links for this cluster have left the degraded state.
REMOVMSG	ENABLE DISABLE	Prevents a distributed library from reporting Automatic Removal start and stop events to the MI and operator messages to the host when Auto Removal occurs.

CACHE settings

If the second keyword of **CACHE** is specified, the cluster modifies how it controls the workflow and content of the TVC. The supported **CACHE** settings are shown in Table 10-3.

Table 10-3 *CACHE settings*

Keyword3	Keyword4	Description
COPYFSC	ENABLE, DISABLE	Copies To Follow Storage Class Preference When the ENABLE keyword is specified, the logical volumes that are copied into the cache from a peer TS7700 are managed by using the actions that are defined for the Storage Class (SC) construct that is associated with the volume, as defined at the TS7700 receiving the copy. When the DISABLE keyword is specified, the logical volumes that are copied into the cache from a peer TS7700 are managed as PGO (prefer to be removed from cache). The default is disabled.
PMPRIOR	value	Premigration Priority Threshold This is the threshold, in gigabytes of unpremigrated data, at which the TS7700 begins increasing the number of premigration tasks that are allowed to compete with host I/O for cache and processor cycles. Tip: Do not change this setting from the default unless you understand the effect that the change will have on the operation of the TS7700. 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.

Keyword3	Keyword4	Description
PMTHLVL	value	<p>Premigration Throttling Threshold This is the threshold, in gigabytes of unpremigrated data, at which the TS7700 begins introducing a delay in responding to host write operations on all virtual tape device addresses of the TS7700.</p> <p>Tip: Do not change this setting from the default unless you understand the effect that the change will have on the operation of the TS7700. 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 SC associated with the recalled volume. When the DISABLE keyword is specified, logical volumes that are recalled into cache are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7700. The default is disabled.</p>
REMOVE	ENABLE, DISABLE	<p>Automatic removal starts when cache usage size crosses the removal threshold. When the ENABLE keyword is specified, automatic removal is enabled on this disk-only cluster. When the DISABLE keyword is specified, automatic removal is disabled on this disk-only cluster. The default value is enabled.</p>
REMVTHR	Value	<p>Automatic removal starts when the cache usage size crosses the removal threshold. When automatic removal is enabled on this disk-only cluster, logical volume removal starts when the free cache space is below the removal threshold plus 1000 GB (1000 GB is the out-of-cache warning threshold). The default value is 3000 (in GB). A value of 2000 - 10000 can be set. As an example, a value of 2000 means that the TS7700 starts the automatic removal when there is less than 3000 GB (2000 GB + 1000 GB) of free cache space. If a value less than 2000 is given, it is set to 2000. If a value more than 10000 is given, it is set to 10000.</p>
CPYPRIOR	ENABLE, DISABLE	<p>Limit premigration resources under Grid copy activity When the ENABLE keyword is specified, the number of premigration tasks starts decreasing when Grid copy activity crosses the threshold that is defined by the CPYPRITH command. When the DISABLE keyword is specified, the number of premigration tasks do not decrease. The default value is DISABLED.</p>
CPYPRITH	Value	<p>Premigration resources start being limited when Grid copy activity crosses the threshold When CPYPRIOR is enabled, the number of premigration tasks starts decreasing when the total Grid copy activity (the data transfer rate to read/write the data from/into the local cluster's cache) is above the threshold.</p>
RBPRIOR	Value	Cache Rebuild Priority
RBTHLVL	Value	Cache Rebuild Priority Threshold

THROTTLE settings

If a second keyword of **THROTTLE** is specified, the cluster modifies how it controls the data flow rates into and out to the cluster. Supported THROTTLE settings are shown in Table 10-4.

Table 10-4 THROTTLE settings

Keyword3	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 tape drives is enabled. When the DISABLE keyword is specified, throttling when the cache is full of uncopied data to other TS7700 tape drives is disabled. The default is enabled. Tip: Full Cache Copy Throttling is also disabled for 24 hours after one of the other TS7700 tape drives has been in service mode. This is to prevent immediate host throttling when the TS7700 being serviced is returned to use.
DCOPYT	value	Deferred Copy Throttle The default value is 125 milliseconds.
ICOPYT	ENABLE, DISABLE	Immediate Copy Throttling The default is enabled.
DCTAVGTD	value	Deferred Copy Throttling Average Threshold A value of 0 sets the threshold to the default. A value of 1 - 500 can be set.

DEVALLOC (Device Allocation) settings

If a second keyword of **DEVALLOC** is specified, the cluster modifies how it runs scratch allocation assistance (SAA) for scratch tapes or device allocation assistance (DAA) for private tapes. For details about SAA and DAA, see Chapter 11, "Performance and monitoring" on page 635. The DEVALLOC settings that are shown in Table 10-5 are supported.

Table 10-5 DEVALLOC settings

Keyword3	Keyword4	Description
SCRATCH	ENABLE, DISABLE	Device Allocation Assist for Scratch Volumes The default is disabled.
PRIVATE	ENABLE, DISABLE	Device Allocation Assist for Private Volumes The default is enabled.

Reclaim settings

If a second keyword of **RECLAIM** is specified, the cluster modifies how the reclaim background tasks controls the workflow and content of the TVC.

Note: Also, if a valid RECLAIM request is received while reclaims are inhibited, that request takes effect as soon as reclaims are no longer inhibited by the Inhibit Reclaim schedule.

The RECLAIM settings that are shown in Table 10-6 are supported.

Table 10-6 RECLAIM settings

Keyword3	Keyword4	Description
RCLMMAX	value	Reclaim Maximum Tasks Limit

CPYCNT (Copy Thread Count) settings

If a second keyword of **CPYCNT** is specified, the domain modifies how many concurrent threads are allowed to process either RUN or Deferred copies over the grid.

The CPYCNT settings that are shown in Table 10-7 are supported.

Table 10-7 CPYCNT settings

Keyword3	Keyword4	Description
RUN	Number of Concurrent RUN Copy Threads	The number of concurrent copy threads for processing RUN copies The allowed values for copy thread counts are 5 - 128. The default value is 20 for clusters with two 1-gigabit Ethernet (GbE) links, and 40 for clusters with four 1 Gb Ethernet links or two 10 Gb Ethernet links.
DEF	Number of Concurrent Deferred Copy Threads	The number of concurrent copy threads for processing Deferred copies The allowed values for copy thread counts are 5 - 128. The default value is 20 for clusters with two 1 Gb Ethernet links, and 40 for clusters with four 1 Gb Ethernet links or two 10 Gb Ethernet links.

Table 10-8 lists the supported settings for COPY.

Table 10-8 Copy settings

Keyword3	Keyword4	Description
IMMSNS	All UNEXP NONE	Immediate-Deferred State Reporting Method This is the control method to report the immediate-deferred state in the CCW (RUN) ERA35 sense data. With Release 1.6 code, TS7700 reports all the immediate-deferred state in the CCW (RUN) ERA35 sense data. Since Release 1.7, subsequent modification level 5 (8.7.0.155), or 2.0 release level the reporting method can be modified: <ul style="list-style-type: none"> ▶ If keyword4 of ALL is specified, all the immediate-deferred state is reported in the ERA35 sense data the same as Release 1.6. ▶ If keyword4 of UNEXP is specified, only the immediate-deferred state induced unexpectedly is reported in the ERA35 sense data. ▶ If keyword4 of NONE is specified, no immediate-deferred state is reported in the ERA35 sense data except the case where no valid source to copy is available. ▶ The default value is NONE. For more information, see <i>IBM Virtualization Engine TS7700 Series Best Practice Understanding, Monitoring, and Tuning the TS7700 Performance</i> , WP101465, found at: http://www.ibm.com/support/techdocs
SCRATCH	ALWAYS/ NEVER/ NONTLDY	Control the replication of logical volumes in the scratch category: <ul style="list-style-type: none"> ▶ If a keyword4 of ALWAYS is specified, all logical volumes regardless of the category are replicated in the grid the same as the previous code level. ▶ If a keyword 4 of NEVER is specified, no logical volumes in the scratch category are replicated in the grid. ▶ If a keyword 4 of NONTLDY is specified, only the logical volumes in the scratch category with Time Delayed copy policy aren't replicated in the grid. The logical volumes with other copy modes are still replicated. The request is supported only when all the clusters in the grid have the code level of 8.32.X.X or later. The default value is ALWAYS.
TIMEOUT	value	Volume Copy Timeout Time This is the timeout value in minutes for logical volume copies between clusters to complete. The allowed values for copy timeout are 30 - 999 minutes.

Link failover settings

If a second keyword of **LINK** is specified, the cluster modifies how to react in a link failure during a remote mount.

Table 10-9 shows the supported settings for **LINK**.

Table 10-9 *LINK settings*

Keyword 3	Keyword4	Description
FAILOVER	ENABLE DISABLE	IP Link Failover for Remount Mounts If keyword4 of ENABLE is specified, a cluster at code level 8.21.x.x or greater uses the failover capability in a link failure during a remote mount. Keyword4 of DISABLE removes the failover capability. The default behavior is ENABLE .

Delexp (Delete Expire) count settings

In response to a request where a composite library is specified, the Delete-Expire setting is modified as described in Table 10-10.

Table 10-10 *Delete-Expire setting*

Keyword3	Keyword4	Description
COUNT	Value	Delete Expire Count The Delete Expire Count can be set to any value from the default value of 1000 to the maximum value of 2000.

Existdel settings

In response to this request where a distributed library is specified, the cluster modifies how to handle the data of E (Exist) copy mode volume. See Table 10-11.

Table 10-11 *Existdel Settings*

Keyword3	Keyword4	Description
CRITERIA	STALE/ ALWAYS/ NONE	Delete E (Exist) copy mode volume at mount/demount. ▶ STALE: Delete only a E copy mode volume when it is inconsistent. ▶ ALWAYS: Always delete E copy mode volume if all other non-E copy mode sites are consistent. ▶ NONE: Never delete consistent or inconsistent E copymode volumes. The request is supported only when all clusters in the domain have the code level of 8.31.x.x or later. The default is STALE.
WHEN	ATCLOSE/ AUTO	This is the setting to determine when E copy mode volumes that satisfy the condition set by the "CRITERIA" keyword can be deleted. E copy mode volume is deleted at the timing based on the following settings: ▶ ATCLOSE: E copy volume is deleted at the volume mount/demount. It is the same behavior with the previous release 8.31. ▶ AUTO: In addition to the volume mount/demount timing, TS7700 periodically checks E copy mode volume, then deletes it if it satisfies the condition set by CRITERIA. The check runs once per 24 hours and it deletes up to 100 E copy mode volumes all at the same time. The request is supported only when all clusters in the domain have the code level of 8.32.x.x or later. The default is ATCLOSE.

10.1.4 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 are in a library that supports outboard policy management.

The processing for the **LIBRARY LMPOLICY** command runs the Library Control System (LCS) external services FUNC=CUA function. Any errors that the Change Use Attribute (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 by 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 10-10.

Example 10-10 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 in the TS7700.
- ▶ **Volser**
Volser specifies the volume serial number of a private volume that is in a TS7700.
- ▶ 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
 - **SC={storage class name | *RESET*}**
Specifies a construct name for the SC parameter
 - **MC={Management Class name | *RESET*}**
Specifies a construct name for the MC parameter
 - **DC={Data Class name | *RESET*}**
Specifies a construct name for the DC parameter

If the request is successful, the construct name is changed to the requested name. If you specify the ***RESET*** keyword, you are requesting that OAM set this construct to the default, which is blanks.

The values that you specify for the SG, SC, MC, and DC policy names must meet the storage management subsystem (SMS) naming convention standards:

- ▶ Alphanumeric and national (special) characters only
- ▶ Name must begin with an alphabetical or national (special) character (\$, *, @, #, or %)
- ▶ No leading or embedded blanks
- ▶ Eight characters or less

10.1.5 Useful DEVSERV QUERY commands

Some of the more useful **DEVSERV QUERY** commands are described in this section.

DEVSERV QTAPE command

The **DEVSERV QTAPE** or **DS QT** command allows a query of the basic configuration of the SMS tape library as it has been defined in the input/output definition file (IODF). With the **RDC** operand, it is useful for viewing the Composite Library ID and Lib Port ID.

The following command shows the syntax:

```
DS QT,devnum,1,RDC
```

The following are the values in the command:

DS	Device service
QT	Query tape
devnum	Device address
1	Number of devices to be displayed
RDC	Read device characteristics

Figure 10-2 shows the output of a **DS QT** system command.

```
DS QT,1C01,RDC
IEE459I 15.03.41 DEVSERV QTAPE 570
UNIT DTYPE DSTATUS CUTYPE DEVTYPE CU-SERIAL DEV-SERIAL ACL LIBID
1C01 3490L ON-RDY 3957C2A 3592 * 0178-272BP 0178-272BP I 3484F
  READ DEVICE CHARACTERISTIC
34905434905400E0 1FD88080B61B41E9 00045AC000000000 3957413592410002
03484F0101000000 4281000004000000 0400000000000000 0000000000000000
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      1 DEVICE(S) WITH DEVICE EMULATION ACTIVE
```

01 - Distributed LIBRARY-ID
01 - LIBPORT-ID
3484F - Composite LIBRARY-ID

Figure 10-2 Sample DEVSERV QT command output

Clarification: 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 10-2, the response shows LIBPORT-ID 01 for logical drive 9600. LIBPORT-ID 01 is associated with Cluster 0. The association between distributed libraries and LIBPORT-IDs is discussed in 6.4.1, “Defining devices through HCD” on page 218.

From the **DS QT** command in Figure 10-2, you can derive the LIBRARY-ID for the composite library and the LIBPORT-ID of the logical control unit (LCU) 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.

Tip: You can get the real device type from the Host Console Request function **LI REQ,<distributed library name>,PDRIVE** in the distributed library.

DEVSERV QLIB,CATS command

The command **DS QLIB,CATS** allows you to view and change logical VOLSER categories without need to initial program load (IPL) the system. Example 10-11 shows how to list all of the categories that are used in a system.

Example 10-11 Sample output of DEVSERV QLIB,CATS

```
DS QL,CATS
IEE459I 10.56.27 DEVSERV QLIB 626
5001 5002 5003 5004 5005 5006 5007 5008 5009 500A 500B 500C 500D
500E 500F
```

After you have the actual categories, you can change them. To perform this task, change the first 3 digits of the category. However, the last digit must remain unchanged because it represents the media type.

Example 10-12 shows the command that changes all categories to 111 for the first 3 digits.

Example 10-12 Sample output of DEVSERV QLIB,CATS(111)*

```
DS QL,CATS(111*)
IEE459I 10.57.35 DEVSERV QLIB 899
1111 1112 1113 1114 1115 1116 1117 1118 1119 111A 111B 111C 111D
111E 111F
```

Ensure that this change is also made in the DEVSUPxx PARMLIB member. If it is not, the next IPL reverts categories to what they were in DEVSUPxx. For a further description of changing categories, see 10.5, “Effects of changing categories” on page 631.

DEVSERV QLIB,LIST command

Example 10-13 shows how you list all of the active composite libraries by using the **DS QL,LIST** command. The **QLIB** command uses the LIBRARY-IDs (LIBIDs), not the TAPE LIBRARY NAME that was used in the **D SMS,LIB** command.

Example 10-13 DEVSERV QLIB,LIST

```
DS QL,LIST
IEE459I 09.39.33 DEVSERV QLIB 933
The following are defined in the ACTIVE configuration:
*BA062 *CA045 *BA060 *BA045 *BA003 *BA031 *BA032 *BA002 *BA039 *BA038
*BA010 BA066 BA051 BA004
```

Note: The asterisks in the QLIB displays indicate libraries that are attached to the host.

For a complete description of all of the **DS QLIB** commands, see Appendix D, “DEVSERV QLIB command” on page 867.

10.1.6 Scratch volume recovery for logical volumes

If you determine that a volume was mistakenly returned to scratch, you can sometimes return the volume to private status to recover its contents. If EXPIRE HOLD has been ENABLED, the volume will not be reused before the EXPIRE time has been reached. The method to recover depends on the TMS 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 TMS from returning the volume to scratch during the next few days. Further details follow about how to recover a logical volume that was scratched in error.

Checking a VOLSER to determine whether it has been reused

The first step in determining whether a volume can be recovered is to check that the VOLSERs that you want to recover have not already been reused as SCRATCH mounts. Issue the **D SMS,VOL(vo1ser)** command for each volume that you want to check on. See Example 10-14.

Example 10-14 Check whether VOLSERs are still available

```
D SMS,VOL(A0000P)
CBR1180I OAM tape volume status: 095
VOLUME MEDIA      STORAGE  LIBRARY  USE  W  C   SOFTWARE  LIBRARY
      TYPE        GROUP    NAME     ATR  P  P   ERR STAT  CATEGORY
A0000P MEDIA2     *SCRCH*  HYDRAG   S   N  N   NOERROR   SCRME2
```

USE ATR of S indicates that the volume is still in SCRATCH status and has not yet been reused. Therefore, you have a chance to recover the volume contents if there is a consistent copy in the TS7700. If the display for this command says USE ATR of P, it has been reused and you cannot recover the contents of the volume by using host software procedures.

Checking for a consistent copy of the volume

The second step in determining whether a volume can be recovered is to check that the VOLSERs that you want to recover have a consistent copy that is in the TS7700 somewhere. The best command to use to check field-known consistent copies is the **LI REQ,LVOL** command, as shown in Example 10-15.

Example 10-15 Check for a consistent copy

```
LI REQ,HYDRAG,LVOL,A0000P

LOGICAL VOLUME INFORMATION V4 0.1
LOGICAL VOLUME: A0000P
MEDIA TYPE: ECST
COMPRESSED SIZE (MB): 627
MAXIMUM VOLUME CAPACITY (MB): 4000
CURRENT OWNER: 00001
.
CATEGORY: 0022
LAST MOUNTED (UTC): 2014-10-02 15:29:28
LAST MODIFIED (UTC): 2014-10-02 15:29:04
LAST MODIFIED VNODE: 00
LAST MODIFIED DEVICE: 000B
TOTAL REQUIRED COPIES: 1
KNOWN CONSISTENT COPIES: 1
```

If KNOWN CONSISTENT COPIES is zero, you cannot recover this volume because it has been DELETE EXPIRED already.

Changing STATUS of scratched volumes to MASTER

The next step in recovering the volumes is to change the status from SCRATCH back to MASTER. The next steps will vary depending on your TMS. If your TMS is DFSMS Removable Media Manager (DFSMSrmm), you might use DFSMSrmm to search on a volume string.

Then, put all of the scratch volumes matching that string into a file with a TSO subcommand to change their status back to MASTER, and set an Expiration Date to some future value (to prevent the next run of DFSMSrmm Housekeeping from sending the volume back to SCRATCH), as shown in Example 10-16.

Example 10-16 Change status

```
//STEPA EXEC PGM=IKJEFT01
//SYSTSPRT DD DUMMY
//RMMCLIST DD DSN=DENEKA.RMMCV,
// DISP=(NEW,CATLG,DELETE),
// SPACE=(CYL,(2,2),RLSE),
// UNIT=SYSDA
//SYSTSIN DD *
PROF NOMSGID
RMM SV VOLUME(A*) OWNER(*) LIM(1) HOME(HYDRAG) STATUS(SCRATCH) -
CLIST('RMM CHANGEVOLUME ',' STATUS(MASTER) EXPDT(14355)
/*
```

The output in the RMMCLIST DD is as follows:

```
READY
RMM CHANGEVOLUME A0000P STATUS(MASTER) EXPDT(14355)
```

Use the job control language (JCL) shown in Example 10-17 to run the previously generated CLIST. This can be done in the same job as the RMM SV command if no editing of the generated list was needed to remove volumes without a consistent copy found. (Altering the status of such volumes to MASTER needlessly uses a scratch volser because the volume contents have already been expire-deleted.) This will also change the volume from scratch to private in the TCDB.

Example 10-17 JCL for CLIST

```
//STEPB EXEC PGM=IKJEFT01,DYNAMNBR=60
//SYSTSPRT DD DSN=DENEKA.RMMCV.OUT,
// DISP=(NEW,CATLG,DELETE),
// SPACE=(CYL,(10,2),RLSE),
// UNIT=SYSDA
//SYSTSIN DD DISP=SHR,
// DSN=DENEKA.RMMCV
```

The **D SMS,VOL** command can now be used to verify that the VOLSER was changed from S to P, as shown in Example 10-18.

Example 10-18 Verify the change

```
D SMS,VOL(A0000P)
VOLUME MEDIA STORAGE LIBRARY USE W C SOFTWARE LIBRARY
      TYPE GROUP NAME ATR P P ERR STAT CATEGORY
A0000P MEDIA2 SGG00001 HYDRAG P N N NOERROR PRIVATE
```

10.1.7 Ejecting logical volumes

Logical volumes are not physical entities that can be individually removed from the library. They can also be on stacked volumes with many other logical volumes. An EJECT should only be issued for a logical volume in the TS7700 if the intent is to delete the volume from the TS7700.

Because of the permanent nature of the EJECT, the TS7700 allows you to EJECT only a logical volume that is in either the INSERT or SCRATCH category. If a logical volume is in any other status, the EJECT fails. If you eject a scratch volume, you cannot recover the data on that logical volume.

Tip: Logical volumes that are in the error category (000E) can be moved back to the scratch category by using **ISMF ALTER** to move them from Scratch to Scratch category.

Tapes that are in INSERT status can be ejected by the resetting of the return code through the CBRUXENT exit. This exit is 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 Virtual Volumes window” on page 400.

After the tape is in SCRATCH status, follow the procedure for EJECT processing 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:

- ▶ The z/OS command **LIBRARY EJECT,volser**
- ▶ 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** sent to a logical volume in PRIVATE status fails with this message:

```
CBR3726I Function incompatible error code 6 from library <library-name> for volume
<volser>
```

Clarification: 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 eventually fails the eject, but the tape management system has already marked the volume as ejected.

If your tape management system is DFSMSrmm, you can use the commands that are shown in Example 10-19 to clean up the Removable Media Management (RMM) control data set (CDS) for failed logical volume ejects, and to resynchronize the TCDB and RMM CDS.

Example 10-19 Clean up the 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 whether the volume is still resident in the VTS library and, if so, it corrects the RMM CDS.

Issuing a large number of ejects at one time can cause some resource effect on the host. A good limit for the number of outstanding eject requests is no more than 10,000 per system. More ejects can be initiated when others complete. The following commands can be used on the z Systems hosts to list the outstanding and the active requests:

```
F OAM,QUERY,WAITING  
F OAM,QUERY,ACTIVE
```

10.2 Messages from the library

This section describes TS7700 enhanced message support and relevant messages.

10.2.1 CBR3750I Console Message

When the host receives a message from the library that is either informational or indicates an abnormal condition of some type, the host will surface this message (via OAM) within the CBR3750I message. This message has the following format:

```
CBR3750I Message from library library-name: message.
```

This indicates that a *message* has been sent from library *library-name*. Either the operator, at the library manager console has entered a message that is to be broadcast to the host, or the library itself, has broadcast a message to the host to relay status information or report an error condition. A list of the messages that can be broadcast from the library to the host is contained in the IBM TS7700 Series Operator Informational Messages white paper, which can be accessed at the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101689>

If additional assistance is required regarding the contents of the messages surfaced, engage IBM hardware support. The only role that OAM provides is to surface on the host the library-generated message through the CBR3750I message.

10.2.2 Alert setting messages

The SETTING function provides a new set of messages. These messages are described in 10.1.3, “Host Console Request function” on page 608.

The following example shows the message format:

CBR3750I Message from library *lib-id*: ALxxxx *message description*

10.2.3 TS7700 Host Console messages

Some of the TS7700 specific messages that you might see are listed. For a complete and current list, see the appropriate volume of *z/OS MVS System Messages*.

Incompatibility error message

In 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 entered that requests an operation that is understood by the subsystem microcode, but cannot be run.

The explanation for the *error-code* can be found in the TS7700 Customer IBM Knowledge Center under **Reference** → **Perform library function codes** → **Error recovery action codes** → **Function Incompatible**.

Warning VTS operation degraded messages

When a VTS is operating in a degraded state, the following message is generated:

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

Warning cache use capacity (TS7720)

For the TS7720, 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.

Out of physical volumes

When a distributed library that is associated with a cluster runs out of scratch stacked physical volumes, operations of the TS7760T, TS7740, or TS7720T are affected. As part of normal processing, data is copied from cache to physical volumes in a primary pool that is managed by the TS7700. A copy might also be made to a physical volume in a secondary pool if the dual copy function is specified by using Management Class (MC).

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 you 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 remains on the MVS console until empty physical volumes are 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 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 MI.

If the TS7760T, TS7740, or TS7720T 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.
- ▶ The grid prefers clusters that are not in the out-of-empty-stacked-volume state in choosing a TVC cluster, but the grid can 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 fails. If data is being written to the remote cluster's TVC, the writes are allowed.

However, 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. In this case, all host I/O that uses that cluster's TVC becomes throttled to prevent a cache overrun.

Monitor the number of empty stacked volumes in a library. If the library is close to running out of a physical volume media type, either expedite the reclamation of physical stacked volumes or add more volumes. You can use the Bulk Volume Information Retrieval (BVIR) function to obtain the physical media counts for each library. The information that is obtained includes the empty physical volume counts by media type for the common scratch pool and each defined pool.

If your Pool properties have a Second Media that is defined, and the primary media type is exhausted, the library does not go into degraded status for out of scratch.

Above Threshold Warning state

The TS7760T, TS7740, or TS7720T 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 TS7760T, TS7740, or TS7720T 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:OP0160 Above threshold for uncopied data in cache, throttling possible
```

10.3 EXPIRE HOLD and scratch processing considerations

This topic deals with the interaction of SCRATCH processing and the EXPIRE settings on the TS7700. The topics of EXPIRE time and EXPIRE HOLD were introduced in 2.3.21, “Expired virtual volumes and the Delete Expired function” on page 75, and should be referenced for a basic understanding of these two settings. Consider the following cases:

- ▶ EXPIRE HOLD option is enabled and the TS7760T, TS7740, or TS7720T is low on scratch volumes.
- ▶ EXPIRE HOLD option is enabled and Cache Utilization is beyond the wanted threshold in the TS7760D or TS7720D.

EXPIRE HOLD and low on scratch tapes in the TS7760T, TS7740, or TS7720T

The EXPIRE HOLD option is used to ensure that a logical volume that is sent to the SCRATCH pool cannot be reused or deleted before the grace period that is specified as the EXPIRE time has passed. This can sometimes create a problem where the library might run out of available volumes to satisfy mount requests to write new data. If the EXPIRE time has been set too long, and there are not enough logical volumes being released to keep a healthy level of SCRATCH volumes available, it might be necessary to reduce the EXPIRE time.

However, this change affects only new volumes that are going into the SCRATCH pool. The existing volumes in the pool continue to be held until the original EXPIRE time has passed. However, if EXPIRE HOLD is cleared, these volumes can then be added to the candidate list for SCRATCH mounts. Therefore, clearing the EXPIRE HOLD option immediately helps to alleviate the low on scratch condition, but it no longer protects data that has inadvertently been sent to SCRATCH. The recovery of user data on volumes in the SCRATCH pool might no longer be certain.

EXPIRE HOLD and cache utilization in the TS7760 or TS7720

When the EXPIRE HOLD option is enabled, the cache in a TS7700 is used up in part by holding data from logical volumes that have been sent to the SCRATCH pool. In addition to the risk of running out of logical volumes to mount, there is a risk of running out of cache in a TS7760D or TS7720D when AUTOREMOVAL is not enabled.

Again, EXPIRE time should be considered, and if a TS7760 or TS7720 is consistently running high on cache utilization, this EXPIRE time should be adjusted. In this case, the clearing of the EXPIRE HOLD setting does not immediately alleviate the high cache utilization condition. The effect of disabling the EXPIRE HOLD from the cache perspective is to enable volumes in the SCRATCH pool to begin entering the candidate list for expire delete processing.

Once per hour, a task runs in the library that processes some number of volumes from this list, and reduces cache utilization by deleting the expired volumes from cache. The number of volumes that are deleted per hour is by default 1000. The number of volumes that are moved to this candidate list is customizable (1000 - 2000), and is controlled by using the **LI REQ DELEXP COUNT** command that is documented in 10.1.3, “Host Console Request function” on page 608, and in the following white paper:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

The EXPIRE time is the grace period that enables the recovery of the data in the case of procedural error. Careful consideration needs to be made to ensure that this value is long enough to allow for such errors to be detected, and the data recovered, before the DELETE EXPIRE process removes the logical volume permanently.

10.4 Scratch count mismatch

There is often some discrepancy between the values reported for scratch counts from various sources. One reason for this is that the ISMF panels and the TS7700 MI report the total number of scratch volumes. However, the **D SMS, LIBRARY** command reports only the number of usable scratch volumes.

When EXPIRE HOLD is in effect, the total number of scratch volumes differs from the total number of usable SCRATCH volumes because volumes for which the EXPIRE time has not yet elapsed are not eligible to be mounted. For this reason, the most accurate source of scratch counts for a TS7700 is always the **D SMS, LIBRARY** report.

10.5 Effects of changing categories

As described in 4.3, “Planning for software implementation” on page 160, categories are assigned according to the DEVSUPxx parmlib member on the host that performs cartridge entry processing. These categories can be changed dynamically by using the **DEVSERV QLIB, CATS** command, and modified in the DEVSUPXX member.

Special consideration should be given to the effects that the change will have on the host system. The most common problem is that all of the logical volumes in the scratch pool belong to the initially defined categories, and requests for scratch mounts fail with CBR4196D error code 140169.

There are several ways to resolve such an issue. If the categories were changed because there is a desire to partition the library, then a new scratch pool must be created for this host by adding a range of volumes that are accepted by the TMS. If the old scratch pool was intended to be used by this host, then the category can be updated by using the ISMF panels to ALTER the USE ATTRIBUTE of a volume or range of volumes from S (Scratch) to S (Scratch).

This resets the category of the volume to match the currently defined categories on the host. If a large range of volumes needs to be changed, consider using the CBRSPCLS utility to perform such a change. For more information about how to use this utility, see *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

A modification of the volume entries in the TCDB using IDCAMS does not change categories, and should not be used for this purpose.

10.6 Library messages and automation

There are some messages that can be useful to automate. In particular, there are three that represent issues with the library being unable to call home. All library messages, including these three, are prefaced with CBR3750I:

- ▶ OP0463: A TSSC error has occurred. The TSSC is unable to call home.
- ▶ OP0625: A system reboot interrupted a call home.
- ▶ OP0550: Service Call Home, TSSC was unable to generate a PMR number.

The advised action for these messages is to contact the service center and engage support to determine why the call home was being attempted.

10.7 Return-to-scratch enhancement

When a volume is returned to scratch, two I/Os calls to the library for each volume:

- ▶ The I/O call to move the volume from private to scratch
- ▶ The I/O call to obtain the number of volumes in the scratch category

APAR OA48240 (z/OS V1R13+) provides the ability to eliminate the second I/O call. After the APAR is applied, this can be done with the LIBRARY DISABLE,CATCOUNT command. This can decrease the overall duration of return-to-scratch processing. If re-enabling the second I/O call is wanted, this can be done with the LIBRARY RESET,CATCOUNT command.

Even though the second I/O call might be disabled, OAM is able to stay updated as to the current count of scratch tapes in the library by using a monitoring task in the OAM address space that queries the library for the current scratch count every 10 minutes. OAM continues to update the scratch count when a volume is changed from scratch to private.

10.8 Deleting virtual volumes

When a virtual volume has been moved from the insert category into a scratch or private category, it can no longer be deleted from the TS7700 MI. In order to delete these volumes, you must verify that a TCDB volume entry exists for the volume to be deleted, as well as verify that it is in scratch status. After these have been verified, you can use the EJECT function from a host that is connected to the TS7700. Because these volumes are virtual, rather than being ejected by the host EJECT function, the virtual volumes are deleted.

The following steps describe how to create the volume entry for the volume (if it is not already present in the TCDB), ALTER the volume to SCRATCH status, and EJECT the volume from the host using different methods:

1. Use the following JCL to invoke IDCAMS to create the volume entry in the TCDB:

```
//CREATVOL JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
CREATE VOLUMEENTRY -
(NAME(Vxxxxxx) -
LIBRARYNAME(libname) -
MEDIATYPE(mediatype) -
LOCATION(LIBRARY)
```

2. Use ISMF to ALTER the use attribute from SCRATCH to SCRATCH. This invokes the CBRUXCUA exit to communicate with the TMS. If the TMS indicates that the change can process, the category is changed in the library to the category defined for the corresponding media type in this host's DEVSUPxx parmlib member.
3. Use one of the following methods to EJECT the tape:
 - a. Use the TMS to EJECT the volume. If RMM is used, the following command issues an EJECT for the volume:

```
RMM DV volser FORCE EJECT
```
 - b. Use ISMF to EJECT each volume
 - a. Use CBRXLCs macro to EJECT each volume

These steps can also be followed after a DR test if a DR host is shut down before deleting the virtual volumes created during the DR test. These volumes continue to use space on the TS7700 until they are deleted.



Performance and monitoring

This chapter describes the factors that determine and influence the performance of the IBM TS7700. It also describes what actions to take, when necessary, to improve TS7700 performance.

In addition, this chapter also covers the possible settings, alerts, and messages that should be considered for exception handling or automation.

This chapter includes the following sections:

- ▶ Overview
- ▶ TS7700 performance characteristics
- ▶ Basic performance overview
- ▶ Monitoring TS7700 performance
- ▶ Cache capacity
- ▶ Cache throughput / Cache bandwidth
- ▶ TS7700 throughput: Host I/O increments
- ▶ Grid link and replication performance
- ▶ Considerations for the backend TS7740 / TS7700T
- ▶ Throttling the TS7700
- ▶ Adjusting parameters in the TS7700
- ▶ Monitoring after service or outage
- ▶ Performance evaluation tool: Plotting cache throughput from VEHSTATS
- ▶ Bulk Volume Information Retrieval
- ▶ Alerts and exception and message handling
- ▶ IBM Tape Tools
- ▶ Using Volume Mount Analyzer
- ▶ Using VEHSTATS and VEHGRXCL for monitoring and reporting
- ▶ IBM z/OS commands for monitoring
- ▶ What to look for and where
- ▶ Virtual Device Allocation in z/OS with JES2

11.1 Overview

R4.0 introduces the TS7760. The TS7760 models (TS7760D and TS7760T) replace all previous models, including the TS7740.

With R3.2 the TS7700 Tape-Attach was introduced. In general, the tape attach models are based on the same physical hardware than the disk-only models, and the basic performance numbers are identical. In addition to the TS770D, the TS7700T needs to write data from the cache to the backend drives, needs to process recalls, and needs additional resources for reclaims. It is necessary to understand how these actions affect the overall performance of the TS7700T.

R3.1 introduced the next generation 8-gigabit (Gb) Fibre Channel connection (FICON) adapters. The TS7700 4-port 8 Gb FICON adapter is the same type as the IBM System Storage DS8700 family 8-port 8 Gb FICON adapter. This adapter provides the same cyclic redundancy check (CRC) protection and inline compression as previous TS7700 FICON adapters provided.

R3.1 supports two ports per 8-Gb FICON adapter only. In a fully populated TS7700, with four 8-Gb adapters, there are eight ports available for TS7700 host connections.

Each port on the 8-Gb FICON adapter supports 512 logical paths, which are twice the number of logical paths that are supported by the 4-Gb FICON adapters. When fully configured with 8 8-Gb FICON channels, the TS7700 supports 4096 logical paths.

This means that you have more flexibility when connecting large numbers of logical partitions (LPARs).

This chapter includes the newest overall performance information, especially for the TS7760 models:

- ▶ An overview of the shared tasks that are running in the TS7700 server
- ▶ A description of a TS7700 monitoring and performance evaluation methodology
- ▶ Understanding the speciality for the TS7700T regarding the different throttling impacts of CP0 and CPx
- ▶ Performance monitoring with the TS7700 GUI
- ▶ Additional information about performance alerting and thresholds
- ▶ Information about the handling of Sync-Deferred and Immediate-deferred copy handling
- ▶ A review of bulk volume information retrieval (BVIR) and VEHSTATS reporting
- ▶ VEHAUDIT and BVIRAUDIT usage
- ▶ A detailed description of the device allocation possibilities regarding the TS7700 capabilities

A brief overview of the tasks in the TS7700 is provided so that you can understand the effect that contention for these resources has on the performance of the TS7700.

The monitoring section can help you understand the performance-related data that is recorded in the TS7700. It discusses the performance issues that might arise with the TS7700. This chapter can also help you recognize the symptoms that indicate that the TS7700 configuration is at or near its maximum performance capability. The information that is provided can help you evaluate the options available to improve the throughput and performance of the TS7700.

Information about extra threshold alerting is provided to help you to implement automation-based monitoring in IBM z/OS. Scenarios are described to show the effect of various algorithms on z/OS and the TS7700 device allocation. These scenarios help you to understand how settings and definitions affect device allocation.

11.2 TS7700 performance characteristics

The TS7700 can provide significant benefits to a tape processing environment. In general, performance depends on such factors, such as total system configuration, Tape Volume Cache (TVC) capacity, the number of physical tape drives available to the TS7700, the number of channels, the read/write ratio, and data characteristics, such as blocksize and mount pattern.

You might experience deviations from the presented figures in your environment. The measurements are based on a theoretical workload profile, and cannot be fully compared with a varying workload. The performance factors and numbers for configurations are shown in the following pages.

Based on initial modeling and measurements, and assuming a 2.66:1 compression ratio, Figure 11-1 on page 638 shows the evolution in the write performance with the TS7700 family, which is also described in more detail in *IBM TS7700 R4 (TS7760) Performance*, WP102652. The following charts are for illustrative purposes only. Always use the most recently published performance white papers available on the Techdocs website at the following address:

<http://www.ibm.com/support/techdocs/atmastr.nsf/Web/TechDocs>

Figure 11-1 shows write performance history.




Figure 11-1 VTS and TS7700 maximum host write throughput

Figure 11-1 shows the evolution of performance in the TS7700 IBM family compared with the previous member of the IBM Tape Virtualization family, the IBM Virtual Tape Server (VTS). All runs were made with 128 concurrent jobs, using 32 kibibyte (KiB) blocks, and queued sequential access method (QSAM) BUFNO = 20. Before R 3.2, volume size is 800 mebibytes (MiB), made up of 300 MiB volumes @ 2.66:1 compression. In R 3.2, the volume size is 2659 MiB (1000 MiB volumes @ 2.66:1 compression).

Figure 11-2 shows the read hit performance numbers in a similar fashion.

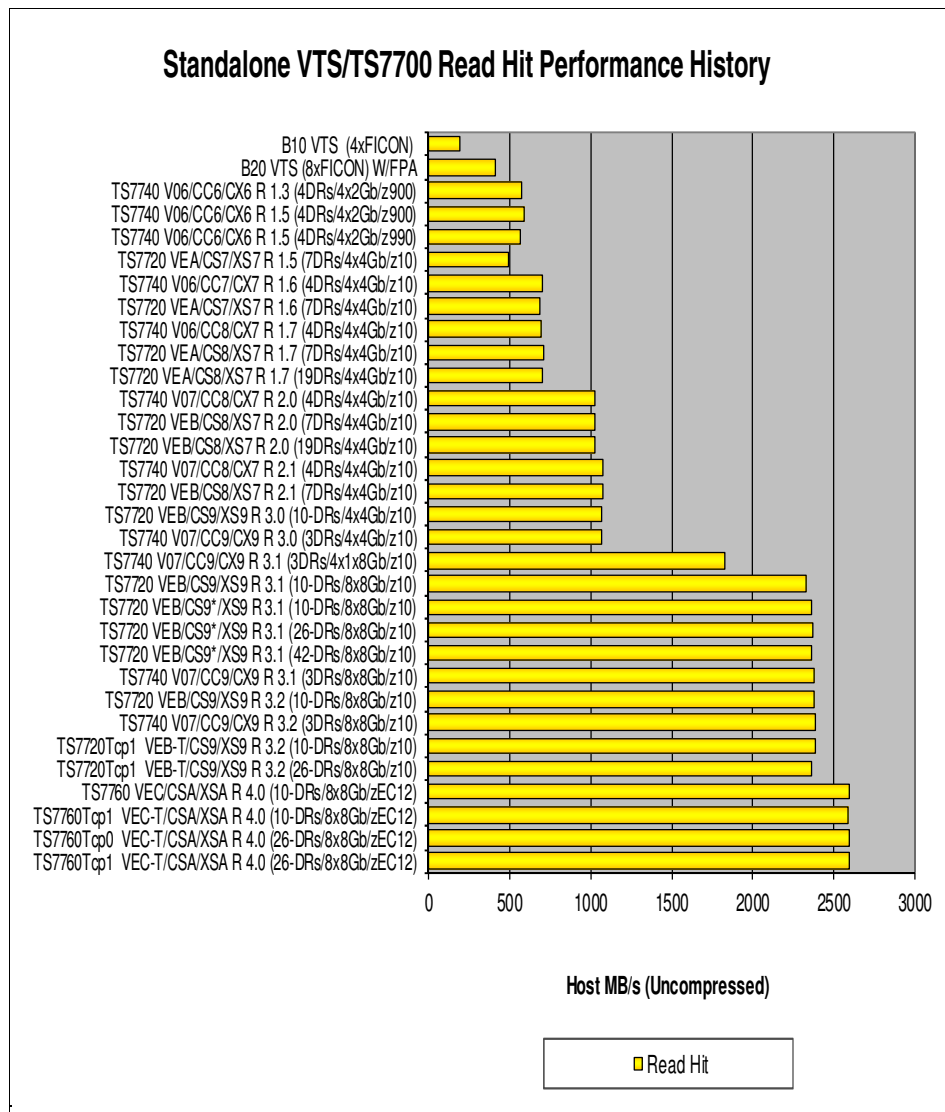


Figure 11-2 VTS and TS7700 maximum host read hit throughput

The numbers that are shown in Figure 11-2 were obtained with 128 concurrent jobs in all runs, each job using 32 KiB blocks, and QSAM BUFNO = 20. Before R 3.2, volume size is 800 MiB (300 MiB volumes @ 2.66:1 compression). Since R 3.2 the volume size is 2659 MiB (1000 MiB volumes @ 2.66:1 compression).

From a performance aspect, the architecture offers these important characteristics:

- ▶ With the selection of IBM DB2 as the central repository, the TS7700 provides a standard Structured Query Language (SQL) interface to the data, and all data is stored and managed in one place. DB2 also allows for more control over performance.
- ▶ The cluster design with virtualization node (vnode) and hierarchical data storage management node (hnode) provides increased configuration flexibility over the monolithic design of the VTS.
- ▶ The use of Transmission Control Protocol/Internet Protocol (TCP/IP) instead of FICON for site-to-site communication eliminates the requirement to use channel extenders.

11.3 Basic performance overview

Performance of TS7700 has several characteristics, and is influenced by many aspects. The following section gives a brief overview of the TS7700 performance aspects, and describes some of the dependencies. For more information about TS7700 performance, see the current version of *TS7700 Understanding, Monitoring, and Tuning Performance* white paper:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101465>

These are the five major aspects that influence the overall performance:

- ▶ TS7700 components and task distribution
- ▶ Replication modes and grid link considerations
- ▶ Workload profile from your hosts
- ▶ Lifecycle Management of your data
- ▶ Parameters and customization of the TS7700

11.3.1 TS7700 components and task distribution

While writing scratch volumes, or premigrating and recalling virtual volumes from physical stacked volumes, hardware components are shared by tasks running on the TS7700. Some of these tasks represent users' work, such as scratch mounts, and other tasks are associated with the internal operations of the TS7700, such as reclamation in a TS7700T and TS7740.

See Figure 11-3 on page 641 for an overview of all of the tasks. The tasks that TS7700 runs, the correlation of the tasks to the components that are involved, and tuning points that can be used to favor certain tasks over others are all described.

The tasks are in general the same for all models of the TS7700. For the TS770D, the backend tasks are not applicable.

Data Movement Inside the TS7700

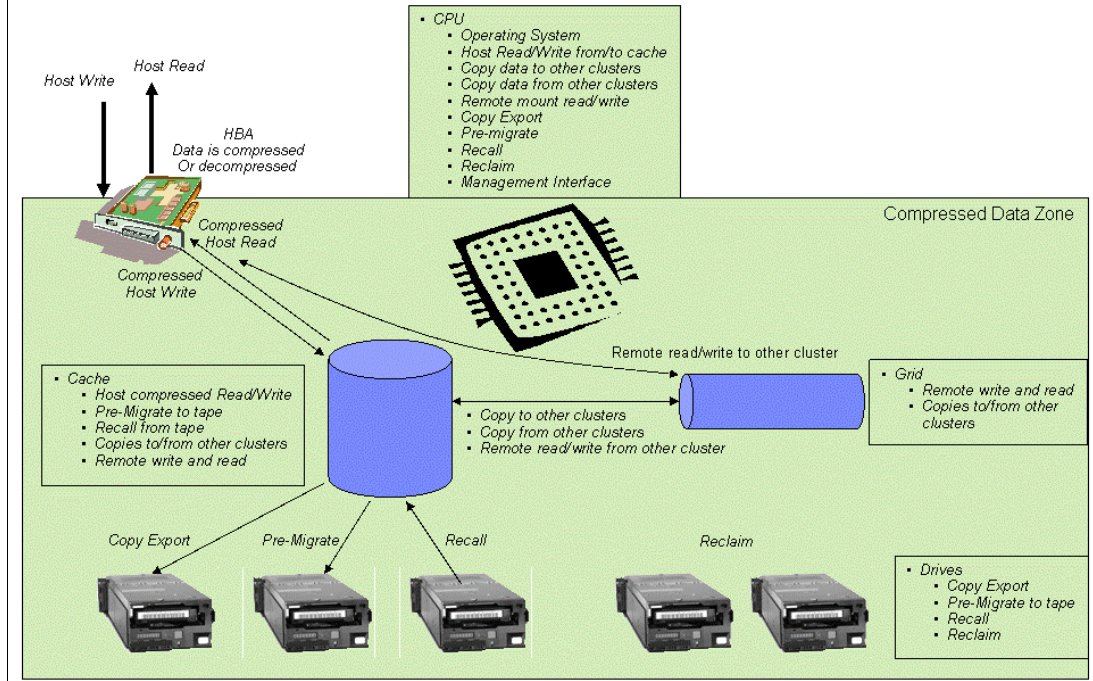


Figure 11-3 Tasks that are performed by the TS7700

All of these tasks share resources, especially the TS7700 Server processor, the TVC, and the physical tape drives attached to a TS7700T or TS7740. Contention might occur for these resources when high workload demands are placed on the TS7700. To manage the use of shared resources, the TS7700 uses various resource management algorithms, which can have a significant impact on the level of performance that is achieved for a specific workload.

In general, the administrative tasks (except premigration) have lower priority than host-related operations. In certain situations, the TS7700 grants higher priority to activities to solve a problem state, including the following scenarios:

- ▶ **Panic reclamation:** The TS7740 or TS7700T detects that 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 uncopied volumes being removed from cache, the TS7700T and TS7740 throttle data coming into the cache. For the TS7700T, this type of throttling occurs only to Host I/O related to the CPx partitions. Data that is written to CP0 is not throttled in this situation.

A complete description of the tasks processing can be found in the *TS7700 Understanding, Monitoring, and Tuning Performance* white paper:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101465>

The TS7700T is not yet reflected in that white paper version.

11.3.2 Grid considerations and replication modes

Data can be copied between the clusters by using the Synchronous, RUN (also known as *Immediate*), Deferred Copy, or Time Delayed Copy policy settings. Every one of these copy modes has specific characteristics, and influences your RPO for specific application workload.

In addition, the chosen copy mode might also have a direct influence on the performance of the TS7700 grid. Although some of the modes enable the copies to be run at a non-peak time (all types of deferred copies), other copy modes are enforced to run before job end (Synchronous, RUN).

This implies that resources from the TS7700 (cache bandwidth, CPU, and grid link bandwidth) need to be available at this point in time.

The following paragraphs provide a quick recap of replication modes:

Synchronous mode means that the data is copied instantly to a second cluster. Synchronous mode copies occur during the writes from the host, and are synchronized when a tape sync event occurs. Due to this behavior, synchronous mode copy has some performance benefits over RUN or deferred copy modes.

Typically, when the Rewind Unload (RUN) is sent, the Synchronous copy has already completed. With an immediate copy (RUN), the second copy is not started until the RUN command is received. 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 before job end. Deferred copy mode allows for a temporarily higher host data rate than Synchronous or RUN copy mode, and can be useful for meeting peak workload demands. Consistently exceeding the capacity of your configuration can result in a copy queue, which cannot be depleted before the next workload peak. Deferred copies are controlled during heavy host I/O with *Deferred Copy Throttling* (DCT). For more information, see “Tuning possibilities for copies: Deferred Copy Throttling” on page 682.

The number of concurrent copy tasks for deferred and RUN copies can be altered by an authorized user by using a host console Library Request command. When altering the copy tasks, consider the Grid network configuration and throughput capabilities to make the best use of the available resources and not over-commit the network or source cluster for copy activity.

The customer can influence the number of copies for concurrent deferred and RUN copies, but not for the number of concurrent Synchronous mode copies. In addition, mechanisms exist to control whether deferred copies are running during peak Host I/O. Be aware, that delaying deferred copies might have an impact on your RPO.

The *Time Delayed Replication mode* is useful to produce copies to multiple TS7700 clusters only after a predefined timeline has expired. This might reduce the number of needed copies, but (if you specify hours after create/access) the copies can be produced in a different timeline. You cannot specify a specific start.

In certain situations, requested Synchronous mode copy or RUN copies cannot be processed. In this case (depending on your customization), the jobs will either not run, or they produce Synchronous-Deferred or Immediate-Deferred copies. These copies are processed later, when the situation is relieved. For more information about the Synchronous deferred and Immediate Deferred copies, see 11.15.2, “Handling Replication Exceptions” on page 715.

The copies are processed in the following order:

1. Synchronous-Deferred PG0
2. Synchronous-Deferred PG1
3. Immediate PG1
4. Immediate-Deferred PG0
5. Immediate-Deferred PG1
6. Deferred PG0
7. Deferred PG1

In the grid, extra tasks can be performed:

- ▶ Remote or Cross-cluster mounts:
 - Using another cluster's cache
 - Another cluster that uses this cluster's cache
- ▶ Cluster coordination traffic:
 - Ownership takeover
 - Volume attribute changes
 - Logical volume insert
 - Configuration

Clarification: Cross-cluster mounts to other clusters do not move data through local cache. Also, reclaim data does not move through the cache.

Cluster families and cooperative replication

Considering a composite library with two or more clusters at a local site and two or more members at a remote site. If more than one cluster needs a copy of a volume at the remote site, cluster families make it possible to send only one copy of the data across a long-distance grid link network. When deciding where to source a volume, a cluster gives higher priority to a cluster in its family over a cluster in another family.

Family members are given higher weight when deciding which cluster to prefer for TVC selection.

Members of a family source their copies within the family when possible. In this manner, data does not have to travel across the long link between sites, optimizing the use of the data link and shortening the copy time. Also, the family members cooperate among themselves, each pulling a copy of separate volume and exchanging them later among family members.

With cooperative replication, a family prefers 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 copy among themselves. Therefore, second copies of volumes within a family are deferred in preference to new volume copies into the family.

When a copy within a family is queued for 12 hours or more, it is given equal priority with copies from other families. This prevents family copies from stagnating in the copy queue.

See the following resources for details about cluster families:

- ▶ *IBM Virtualization Engine TS7700 Series Best Practices -TS7700 Hybrid Grid Usage:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>

11.3.3 Workload profile from your hosts

In general, there are two types of workloads:

- ▶ **Planned workload:** This type of workload is driven by predictable action, mostly batch jobs. These planned actions can be influenced by the operation team regarding the execution time of the jobs.
- ▶ **Unplanned workload:** The other type of workload is the user driven workload, for example, hierarchical storage management (HSM) or object access method (OAM) processing requests, and event-driven workload, for example, database log archiving or System Management Facilities (SMF) processing exists.

Unplanned read workload might have peaks that can affect on the one hand the response times of these actions (read / recall times). However, these actions can also influence the deferred copy times and, in a TS7740, the reclamation execution. Changes in the workload profile might affect the replication time of deferred copies and can lead to throttling situations. Therefore, review the performance charts of the TS7700 to identify workload profile changes, and to take appropriate performance tuning measurements if necessary.

11.3.4 Lifecycle Management of your data

This specific aspect is important for a TS7740 or TS7700T. Depending on your amount of data (and logical volumes) with a short expiration date, the TS7740/TS7700T needs to run more reclamation. This has an impact to your back-end drives and TS7740 and TS7700T processor cycles.

In a hybrid grid, this data can be placed in the TS7700D, and can be replicated with the Time Delayed copy mode. This can reduce the backend activities in the TS7740 and TS7700T. Using the delay premigration queue on a TS7700T can also reduce the back-end activities for migrate and reclaim.

11.3.5 Parameters and customization of the TS7700

The TS7700 offers you a broad variety of tuning possibilities, especially for the cache management, the replication activities, and the backend activities.

The following list describes some examples of TS7700 tuning activities:

- ▶ Preference group of the data (data is preferably in cache or not).
- ▶ Number of the tape cache partitions in a TS7700T.
- ▶ Using the premigration delay feature in the TS7700T.
- ▶ Premigration threshold and control of premigration tasks for TS7700T and TS7740
- ▶ Deferred Copy Throttling (to prioritize the host workload).
- ▶ Number of concurrent copy tasks.
- ▶ Schedule for reclamation.
- ▶ Number of physical volume pools.

Consider that some of these activities have dependencies.

If you change the preconfigured values, review your adjustment with the performance monitoring tools.

For more information, see the *TS7700 Understanding, Monitoring, and Tuning Performance* white paper:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101465>

11.3.6 Terminology of throughput

A TS7700 disk-cache-only cluster has a fairly consistent workload throughout.

Because the TS7740 and TS7700T contain physical tapes to which the cache data is periodically written, recalls from tape to cache occur, and Copy Export and reclaim activities occur, the TS7740 and TS7700T exhibits four basic throughput rates:

- ▶ Peak write
- ▶ Sustained write
- ▶ Read-hit
- ▶ Recall

These four rates are described in the following sections.

Peak and sustained write throughput

For the TS7740 and TS7700T, 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 and TS7700T 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 of non-premigrated data is reached, premigration starts, which can reduce the host write data rate. This threshold is called the *premigration priority threshold*, and has a default value of 1600 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 can 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 2000 GB. These two thresholds can be used with the peak data rate to project the duration of the peak period. Both the priority and throttling thresholds can be increased through a host command-line request, which is described later in this chapter.

Read-hit and recall throughput

Similar to write activity, there are two types of TS7740 and TS7700T read performance:

- ▶ *Read-hit* (also referred to as *peak*) occurs when the data that is requested by the host is 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. Performance depends 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.

11.3.7 Throttling in the TS7700

Throttling is the mechanism adopted to control and balance several tasks that run at the same time within the TS7700, prioritizing certain tasks over others. These mechanisms are called upon only when the system reaches higher levels of utilization, where the components are used almost to their maximum capacity and bottlenecks start to show. The criteria balances the user needs with the vital resources that are needed for the TS7700 to function.

This control is accomplished by delaying the launch of new tasks and prioritizing more important tasks over the other tasks. After the tasks are dispatched and running, control over the execution is accomplished by slowing down a specific functional area by introducing calculated amounts of delay in the operations. This alleviates stress on an overloaded component, or leaves extra central processor unit (CPU) cycles to another needed function, or waits for a slower operation to finish.

The subsystem has a series of self-regulatory mechanisms that try to optimize the shared resources usage. Subsystem resources, such as CPU, cache bandwidth, cache size, host channel bandwidth, grid network bandwidth, and physical drives, are limited, and they must be shared by all tasks moving data throughout the subsystem.

The resources implicitly throttle by themselves when reaching their limits. The TS7700 introduces various explicit throttling methods to give higher priority tasks more of the shared resources:

- ▶ Incoming traffic from the host (host throttling)
- ▶ RUN copy processing from other cluster (copy throttling)
- ▶ Copy processing of deferred copies to other cluster (deferred copy throttling)

TS7700T specific throttling behaviors

From a throttling perspective, the TS7700T is different than a TS7700 disk only model or a TS7740. Resident partition and Tape partitions have two independent throttling measurements, and are treated differently. Reaching a throttling limit on the Tape partitions, for example PMTHLVL, has no impact on the workload directed to the resident partition and vice versa.

The following rules apply:

- ▶ Throttling initiated by reaching the maximum Host Throughput applies to all partitions (Resident and Tape partitions).
- ▶ Throttling initiated by reaching any premigration limit (PMTHLVL or maximum size of the premigration feature queue) impacts only tape partitions, but not the workload directed to CP0.
- ▶ Copy Throttling and Deferred copy throttling have a common measurement regardless of whether the workload is created in CP0 or CP1.

Important: Resident partition (CP0) and Tape Partitions (CP1 - CP7) are monitored and handled separately in a TS7700T.

However, even if the PMTHLVL throttling does not apply to the CP0 of a TS7700T, there is still an indirect influence because of the shared cache modules.

Consider the following items when you configure or monitor a TS7700T resource usage:

- ▶ Workloads that create data (either host I/O, remote writes, or copy processes from other clusters) in the CP0 uses resources of the cache bandwidth (write to cache).
- ▶ After PMTHLVL is crossed for the CPx, the Host I/O creating data in the CPx is throttled, but there will be no throttling to the jobs creating data in CP0.
- ▶ In small configurations (for example, up to four drawers), this can lead to the situation where the jobs running to CP0 use the cache bandwidth resources, and resources for premigration might be limited. This is especially true for a TS7720T due to the used cache.
- ▶ If the unpremigrated amount of data still increases, the throttling of the workload into the CPx increases too. This might reach a number of delays to the jobs where the jobs creating data in CPx are seriously impacted.

TS7740/TS7700T tape drives usage considerations

The physical tape drives are managed by the TS7740/TS7700T internal management software, and cannot be accessed from any other attached host. These drives are used exclusively by the TS7740/TS7700T for the mounts that are required for copying virtual volumes to stacked volumes, recalling virtual volumes into the cache, and reclaiming stacked volume space.

The availability of TS7740/TS7700T physical tape drives for certain functions can significantly affect TS7740/TS7700T performance.

The two major maintenance or “housekeeping” tasks at work are the premigration of data from cache to tape, and deferred copies to and from other clusters. The TS7740 and the TS7700T delay these housekeeping tasks to preference host I/O while no thresholds are reached (PMTHLVL).

The TS7740/TS7700T manages the internal allocation of these drives as required for various functions, but it usually reserves at least one physical drive for recall and one drive for premigration.

TVC management algorithms also influence the allocation of physical tape drives, as described in the following examples:

- ▶ Cache freespace low: The TS7740/TS7700T increases the number of drives available to the premigration function, and reduces the number of drives available for recalls.
- ▶ Premigration threshold crossed: The TS7740/TS7700T 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 is also reduced during reclamation.

The number of drives for premigration can be restricted on a physical pool base. If the number of drives available for premigration is restricted, or the physical drives are already used by other processes, it can lead to limiting the number of virtual volumes in the cache to be premigrated. This might lead to premigration throttling (host I/O is throttled), and later it can lead to free space or copy queue throttling.

If no physical drive is available when a recall is requested, elongated virtual mount times for logical volumes that are being recalled can be the result.

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 by 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 must be constant.

Reclaim policies must be set in the Management Interface (MI) for each volume pool. Reclamation occupies drives and can affect performance. Using multiple physical pools can cause a higher usage of physical drives for premigration and reclaim.

In general, the more pools are used, the more drives are needed. If all drives are busy, and a recall is requested, the reclaim process is interrupted. That can take some seconds to minutes, because the actual moving logical volume needs to be finished, and then the cartridge needs to be exchanged.

The Inhibit Reclaim schedule is also set from the MI, and it 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.

Note: Secure Data Erase does not acknowledge your settings and can run erasure operations if 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/TS7700T.

11.4 Monitoring TS7700 performance

The IBM TS7700 series is part of a line of tape virtualization products that has revolutionized the way that mainframes use their tape resources. As the capability of tape virtualization has grown, so has the need to efficiently manage the large number of logical volumes that the system supports. Internal to the TS7700, a large amount of information is captured and maintained about the state and operational aspects of the resources within the TS7700.

The TS7700 provides a MI based on open standards through which a storage management application can request specific information that the TS7700 maintains. The open standards are not currently supported for applications running under IBM z/OS, so an alternative method is needed to provide the information to mainframe applications.

You can use the following interfaces, tools, and methods to monitor the TS7700:

- ▶ IBM TS4500 and TS3500 tape library Specialist (TS7740/TS7700T only)
- ▶ TS7700 M)
- ▶ Bulk Volume Information Retrieval function (BVIR)
- ▶ IBM Tape Tools: VEHSTATS, VEHAUDIT, and VEHGRXCL
- ▶ Host Console Request Commands

The specialist and MI 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. Finally, the VEHSTATS tools can be used to format the BVIR responses, which are in a binary format, to create usable statistical reports.

With the VEHSTATS data, there are now performance evaluation tools available on Techdocs that quickly create performance-related charts. Performance tools are provided to analyze 24 hours worth of 15-minute data, seven days worth of one-hour interval data, and 90 days worth of daily summary data. For spreadsheets, data collection requirements, and trending evaluation guides, see the following Techdocs website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4717>

All interfaces, tools, and methods to monitor the TS7700 are explained in detail next. An overview of these interfaces, tools, and methods is shown in Figure 11-4.

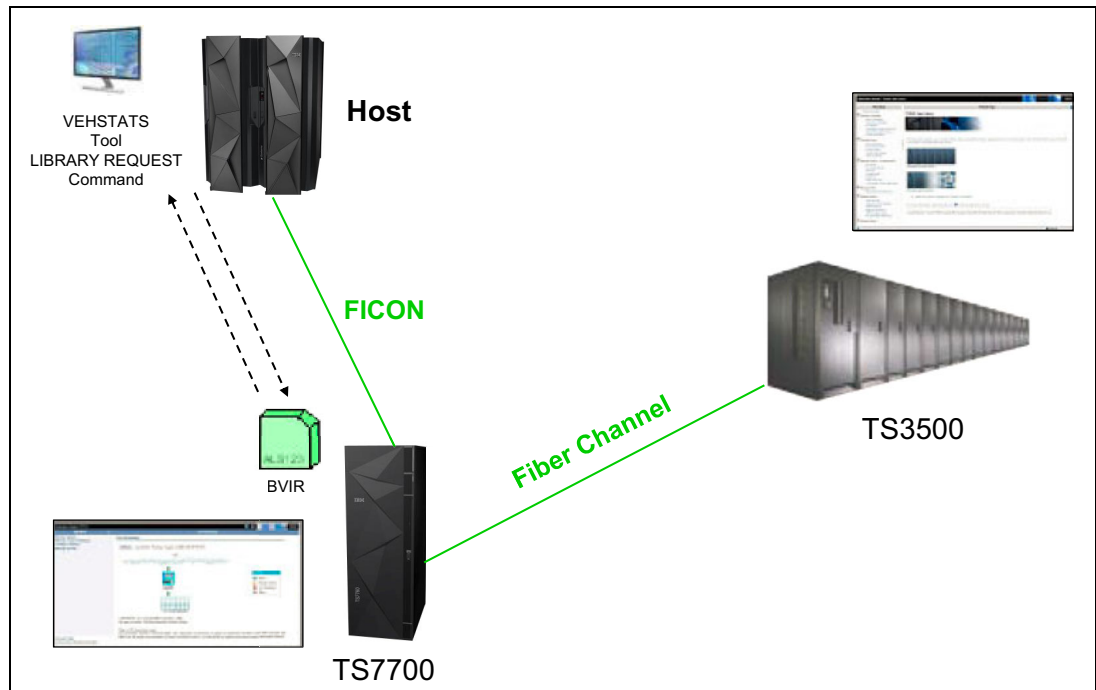


Figure 11-4 Interfaces, tools, and methods to monitor the TS7700

11.4.1 Base information: Types of statistical records

All of the mentioned interfaces and tools process the two types of statistics that are provided by the TS7700:

- ▶ Point-in-time statistics
- ▶ Historical statistics

Point-in-time statistics

These statistics are performance-related. The point-in-time 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 it 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. You can retrieve the point-in-time statistics from the TS7700 at any time by using the BVIR facility. A subset of point-in-time statistics is also available on the TS7700 MI.

The response records are written in binary undefined (U) format of maximum 24,000 bytes.

Tips:

- ▶ If a cluster or node is not available at the time that the point-in-time statistics are recorded, except for the headers, all the data fields for that cluster or node are zeros.
- ▶ The request records are written in fixed-block architecture (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.

Historical statistics

Historical (HIS) 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. 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 are stored for 90 days in a TS7700 database.

The user can also retrieve these records by using BVIR. A subset of the historical statistics is also available on the TS7700 MI. More information is available in 11.5, “Cache capacity” on page 665.

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. The response records are written in a binary undefined (U) format of a maximum of 24,000 bytes.

Tips:

- ▶ If a cluster or node is not available at the time that the historical statistics are recorded, except for the headers, all the data fields for that cluster or node are zeros.
- ▶ The TS7700 retains 90 days worth of historical statistics. If you want to keep statistics for a longer period, be sure that you retain the logical volumes that are used to obtain the statistics.
- ▶ The request records are written in 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.

Both point-in-time statistics and historical statistics are recorded. The point-in-time 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. For more information about how to format and analyze these records, see 11.15, “Alerts and exception and message handling” on page 713.

Each cluster in a grid has its own set of point-in-time and historical statistics for both the vnode and hnode.

For a complete description of the records, see *IBM Virtualization Engine TS7700 Series Statistical Data Format* white paper and *VEHSTATS Decoder*, which are available on the following Techdocs web pages:

<http://www.ibm.com/support/techdocs/atmastr.nsf/Web/Techdocs>

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100829>

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD105477>

11.4.2 Using the TS4500 Management GUI

The TS4500 Management GUI has a new design, and have now the same look and feel as the TS7700 series. To gain an overview of the system, hover with the mouse over the TS4500 icon. Depending on your position, the actual health status is provided, but configuration information is also shown. Figure 11-5 shows an example.



Figure 11-5 TS4500 Overview example

In the TS4500 Logical Library view, you can find the information regarding the number of cartridges, drives, and maximum cartridges. Use the FILTER option for selecting the columns.

Figure 11-6 shows this panel, and the selected properties.

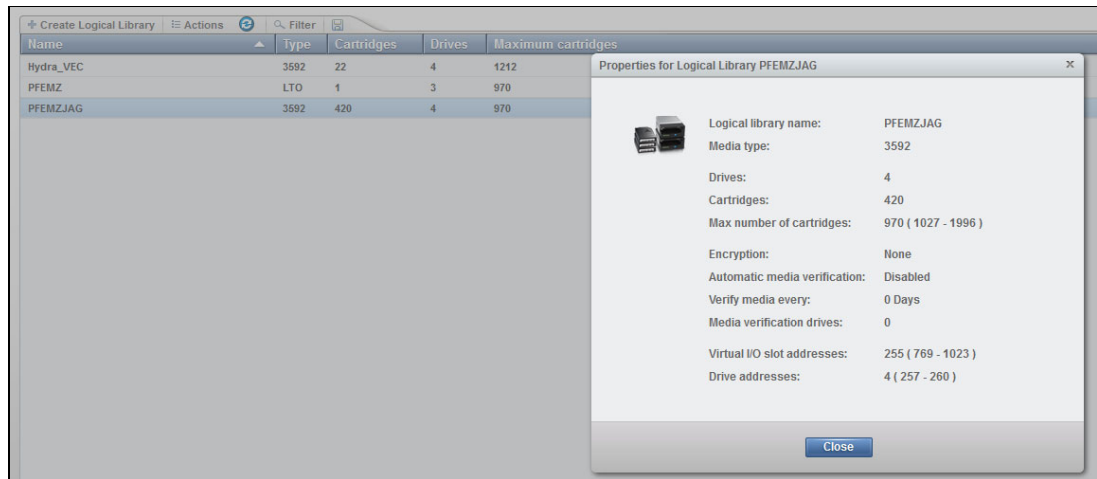


Figure 11-6 TS4500 logical library display with properties

11.4.3 Using the TS3500 Tape Library Specialist for monitoring

The Tape Library Specialist (TS3500 Tape Library Specialist), only available with the TS7740/TS7700T, allows users to manage and monitor items that are related to the TS3500 tape library. Initially, the web user interface to the TS3500 tape library only supported a single user at any 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 11-7 shows the TS3500 tape library System Summary window.

Library Status:

Accessors	OK
3592 Capacity Utilization	63%

View Configuration and Cartridge Counts:

All Frames [Inventory/Audit](#)

All Frames

Total storage slots	2037
3592 Licensed Capacity	2037
3592 Unlicensed Capacity	0
Total empty storage slots	780
Offline storage slots	0
Accessors	2
Total I/O slots	16
Empty I/O slots	15
Total 3592 data cartridges	1280
3592	1280
3592 Not Labeled	0
Cleaning cartridges	10
Drives	36
Node cards	11
Total frames	8
Active frames	6
Service bays	2

Figure 11-7 TS3500 Tape Library Specialist System Summary window

The TS3500 Tape Library Specialist session times out after a default setting of 10 minutes. This is different from the TS7700 MI. You can change the default values through the TS3500 Tape Library Specialist by selecting **Manage Access** → **Operator Panel Security**, which opens the window that is shown in Figure 11-8.




Figure 11-8 TS3500 Tape Library Specialist Operator Panel Security window

Some information that is 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 computer and then used as input for snapshot analysis for the TS3500. This information refers to the TS3500 and its physical drive usage statistics from a TS3500 standpoint only.

For more 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 megabytes (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 lifecycle
 - Number of Write and Read permanent errors of a specific cartridge in the lifecycle
- ▶ Fibre Port statistics: Download only

The Fibre Port statistics include fiber errors, aborts, resets, and recoveries between the TS7700 and the physical tape drives in the TS3500 tape library.

Consideration: This statistic does not provide information from the host to the TS7700 or from the host to the controller.

- ▶ Library statistics, on an hourly basis: Download only
 - Total Mounts
 - Total Ejects
 - Total Inserts
 - Average and Maximum amount of time that a drive was mounted on a drive (residency)
 - Average and Maximum amount of time that was needed to run a single mount
 - Average and Maximum amount of time that was needed to run an eject

These statistics can be downloaded to a workstation for more analysis. These statistics are not included in the BVIR records processed by the TS7700.

11.4.4 Using the TS7700 Management Interface to monitor IBM storage

The TS7700 MI belongs to the family of tools that are used for reporting and monitoring IBM storage products. These tools do not provide reports, but they can be used for online queries about the status of the TS7700, its components, and the distributed libraries. They also provide information about the copies that have not completed yet and the amount of data to be copied.

The TS7700 MI is based on a web server that is installed in each TS7700. You can access this interface with any standard web browser.

The TS7700 MI is a Storage Management Initiative - Specification (SMI-S) - compliant interface that provides you with 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 configuration, two possible web interfaces are available:

- ▶ The TS4500 Management GUI or the TS3500 Tape Library Specialist
- ▶ The TS7700 MI

A link is available to the physical tape library management interface from the TS7700 MI, as shown at the lower left corner of Figure 11-9 on page 656. This link might not be available if not configured during TS7740/TS7700T installation, or for a TS7700D.

The Performance and Statistics windows of the TS7700 MI are described.

Performance and statistics

Information that relates to viewing performance information and statistics for the TS7700 for single and multi-cluster grid configurations is described. The graphical views display snapshots of the processing activities from the last 15 minutes if nothing else is stated when describing the windows. You can access the following selections by going to the Performance & Statistics section in the TS7700 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 11-9.



Figure 11-9 TS7700 MI Performance

Historical Summary

This window (Figure 11-10 on page 657) shows the various performance statistics over a period of 24 hours. Data is retrieved from the Historical Statistic Records. It presents data in averages over 15-minute periods.

Multiple views can be selected, also dependent on the installed cluster type. Such as displaying for the maximum of a whole day:

- ▶ Throughput performance
- ▶ Throttling information
- ▶ Copy Queue

The value that is shown in the performance graph, Figure 11-10 on page 657, can be changed by selecting the different metrics. To do so, press “Select Metric” and choose the requested values. The maximum number of values you can view in one picture is ten. You can also save the graph by pressing the download button (disc).

Figure 11-10 shows the Historical Summary “Throughput” window for a TS7760T




Figure 11-10 TS7700 MI Historical Summary of TS7760T

The precustomized “Throughput” enables you to see all cache bandwidth relevant information

- ▶ Primary cache device write
- ▶ Primary cache device read
- ▶ Remote read
- ▶ Remote write
- ▶ Link copy in
- ▶ Link copy out
- ▶ Write to tape
- ▶ Recall from tape

A pre-customized “Throttling” graph is shown in Figure 11-11 on page 657. It shows which type of throttling happened during the selected interval, and the throttling impact in milliseconds.




Figure 11-11 TS7700 MI Historical Summary Throttling overview

A precustomized “Copy Queue” is shown in Figure 11-12.

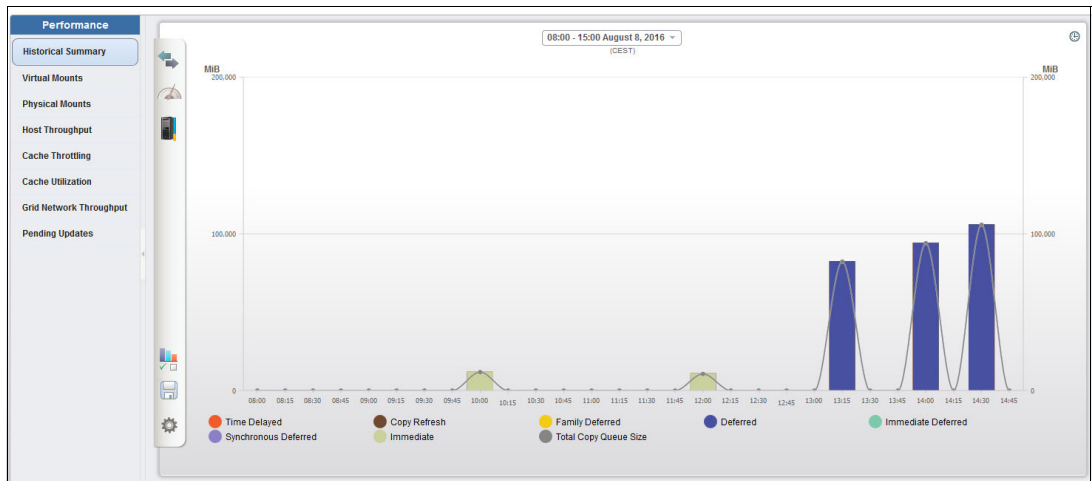


Figure 11-12 TS7700 MI Historical Summary Copy Queue graph

This copy queue view shows how many MiB are sitting in the queue for a specific copy consistency policy.

All of these metrics can be changed by selecting different metrics. Remember, that you can select only 10 items for one graph, as shown in Figure 11-13.




Figure 11-13 Historical summary overview

Although this is a snapshot of one day or less, the performance evaluation tools on TECHDOCS provide you a 24-hour or 90-day overview of these numbers. Review the numbers to help you with these tasks:

- ▶ Identify your workload peaks and possible bottlenecks.
- ▶ See trends to identify increasing workload.
- ▶ Identify schedule times for reclaim.

For more information about using the tool, see 11.6.1, “Interpreting Cache throughput: Performance graph” on page 669. The Performance evaluation tool does not support new content regarding the TS7700T yet.

Virtual Mounts window

Use this window (Figure 11-14) to view virtual mount statistics for the TS7700. The virtual 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. This example is from a TS7700 Cluster that is part of a six-cluster grid configuration.

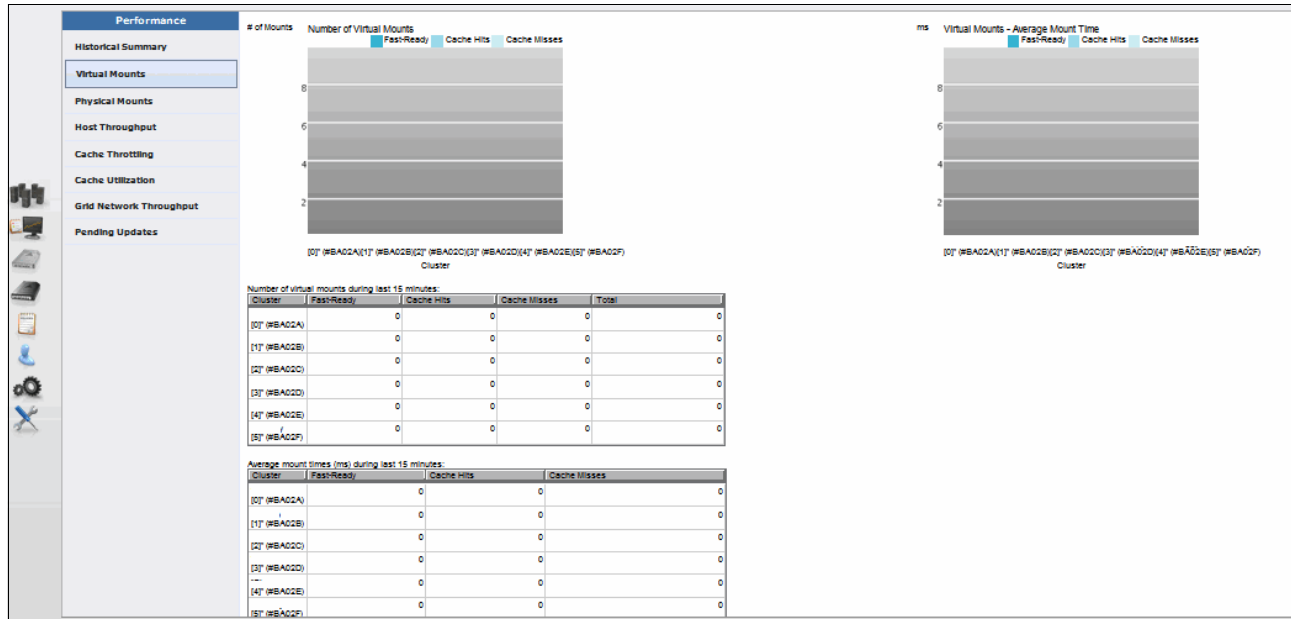


Figure 11-14 TS7700 MI Virtual Mounts window

The “Number of logical mounts during last 15 minutes” table has the following information:

Cluster	The cluster name
Fast Ready	Number of logical mounts that are completed by using the scratch (Fast Ready) method
Cache Hits	Number of logical mounts that are completed from cache
Cache Misses	Number of mount requests that were not fulfilled from cache
Total	Total number of logical mounts

The “Average mount times (ms) during last 15 minutes” table has the following information:

Cluster	The cluster name
Fast Ready	Average mount time for scratch (Fast Ready) logical mounts
Cache Hits	Average mount time for logical mounts that are completed from cache
Cache Misses	Average mount time for requests that are not fulfilled from cache

This view gives you an overview only if you run out of virtual drives in a cluster. Depending on your environment, it does not show you, if in a specific LPAR or sysplex, there might be a shortage of virtual drives. Especially if you define virtual drives in a static way to an LPAR (without an allocation manager), a certain LPAR might not have enough drives. To ensure that a specific LPAR has enough virtual drives, analyze your environment with Tapetools MOUNTMON.

Physical Mounts window

Use this window (Figure 11-15) to view physical mount statistics for the TS7740. 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 11-15 is from a TS7740 cluster that is part of a multi-cluster grid configuration (four-cluster grid).




Figure 11-15 TS7740 MI 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

Review the used numbers of physical drives to help you with the following tasks:

- ▶ Identify upcoming bottlenecks.
- ▶ Determine whether it is appropriate to add or reduce physical pools. Using a larger number of pools requires more physical drives to handle the premigration, recall, and reclaim activity.
- ▶ Determine possible timelines for Copy Export operations.

Host Throughput window

You can use this window (Figure 11-16 on page 661) to view host throughput statistics for the TS7700. 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 11-16 is from a TS7700 cluster that is part of a multi-cluster grid configuration (four-cluster grid).

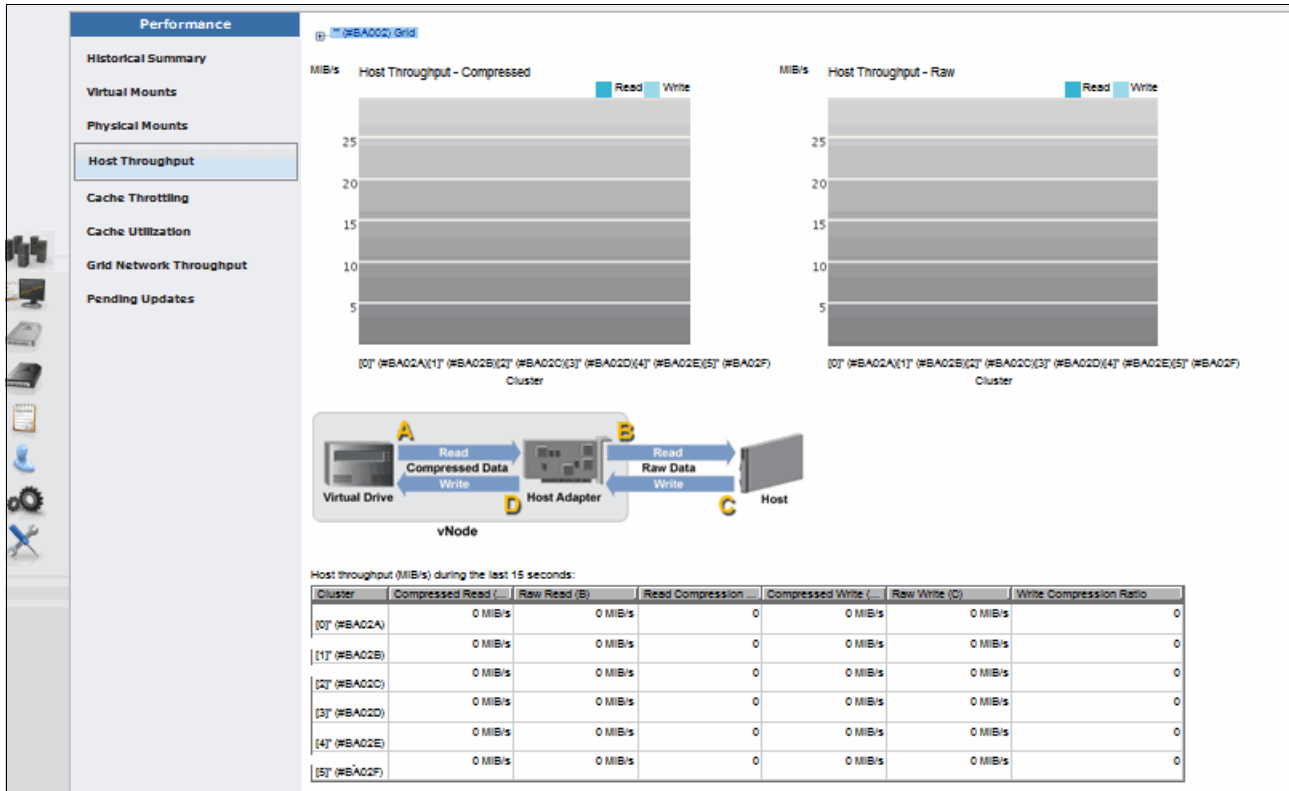


Figure 11-16 TS7700 MI Host Throughput window

The host throughput data is displayed in two bar graphs and one table. The bar graphs are for raw data that is coming from the host to the host bus adapter (HBA), and for compressed data that is 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, and host adapter port. The table cells include the following items:

- Cluster** The cluster or cluster component for which data is being displayed (vnode, host adapter, or host adapter port)
- Compressed Read (A)** Amount of data that is read between the virtual drive and the HBA
- Raw Read (B)** Amount of data that is read between the HBA and host
- Read Compression Ratio** Ratio of compressed read data to raw read data
- Compressed Write (D)** Amount of data that is written from the HBA to the virtual drive
- Raw Write (C)** Amount of data that is written from the host to the HBA
- Write Compression Ratio** Ratio of compressed written data to raw written data

Although this is a snapshot, the performance evaluation tools on Techdocs provide you with a 24-hour, 7-day, or 90-day overview about these numbers.

Review these numbers to help you to identify the following conditions:

- ▶ Identify the compression ratio in your environment for cache and stacked volume planning.
- ▶ Identify any bottlenecks on the host throughput (FC enablement).

Cache Throttling window

You can use this window (Figure 11-17) to view cache throttling statistics for the TS7700. This example is from a TS7700 cluster that is part of a multi-cluster grid configuration (four-cluster grid).

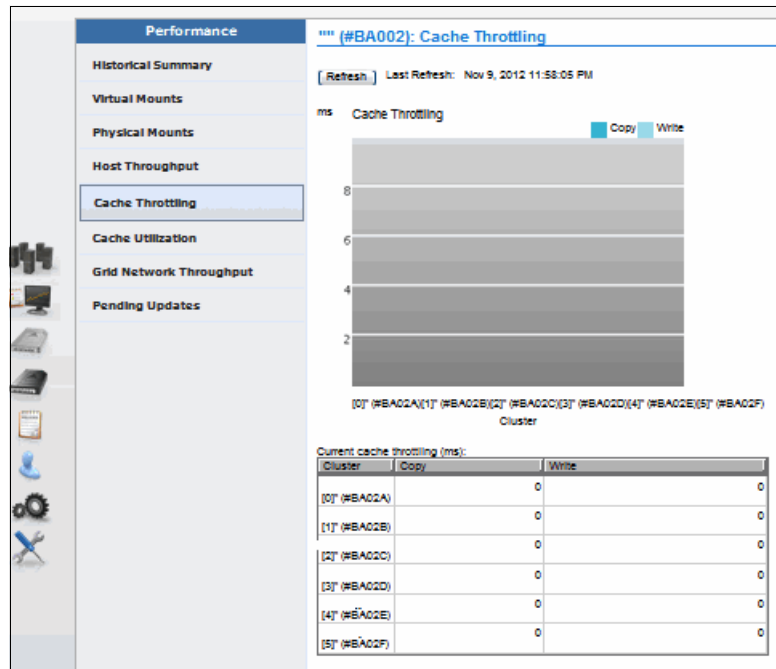


Figure 11-17 TS7700 MI Cache Throttling window

Cache throttling is a time interval that is applied to TS7700 internal functions to improve throughput performance to the host. The cache throttling statistics for each cluster that relate 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 that is inserted between internal copy operations
- Write** The amount of time that is inserted between host write operations

Cache Utilization window

You can use this window (Figure 11-18 on page 663) to view cache utilization statistics for the TS7700D or the TS7740. If a TS7700T is installed this selection is called “Cache Partitions” and is explained in the next section.

This example is from a TS7740 cluster that is part of a multi-cluster grid configuration (four-cluster grid).




Figure 11-18 TS7740 MI Cache Utilization window

The cache utilization statistics can be selected for each cluster. Various aspects of cache performance are displayed for each cluster. Select them from the **Select cache utilizations statistics** menu. The data is displayed in both bar graph and table form, and can be displayed also by preference groups 0 and 1.

The following cache utilization statistics are available:

- ▶ Cache Preference Group possible values:
 - 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 that are assigned to the cache preference group.
- ▶ Median duration that volumes have remained in cache: Rolling average of the cache age of volumes that are migrated out of this cache preference group for the specified amount of time (last 4 hours, 48 hours, and 35 days).
- ▶ Number of logical volumes migrated: Rolling average of the number of volumes that are migrated to this cache preference group (4 hours, 48 hours, and 35 days). Bar graph is not used.

Clarification: Median Duration in Cache and Number of Logical Volumes Migrated statistics have a table column for each of the time periods that are mentioned in parentheses.

Review this data with the performance evaluation tool from Techdocs to identify the following conditions:

- ▶ Cache shortages, especially in your TS7700D
- ▶ Improvement capabilities for your cache usage through the adjustment of copy policies

Cache Partitions

You can use this window (Figure 11-18) to view Tape cache partitions and their utilization. Depending on your filter list, you might see the following output:

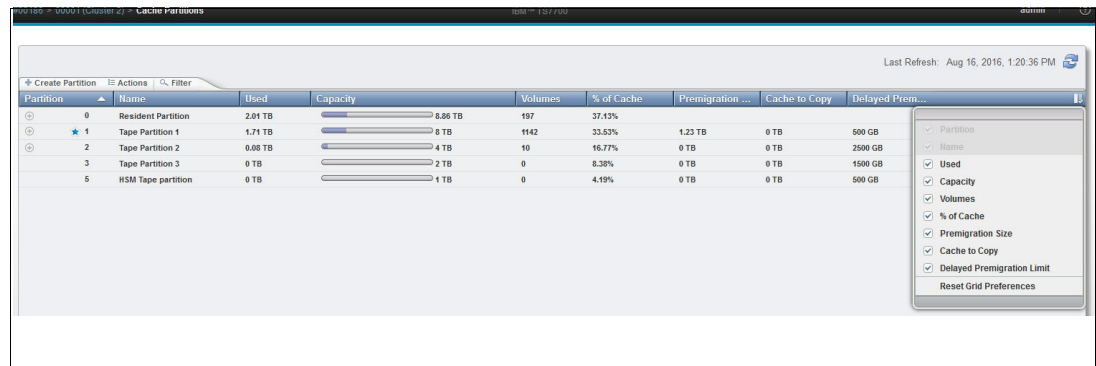


Figure 11-19 Tape Partition view in a TS7700T

For more information refer to “Cache Partition” on page 370.

Grid Network Throughput window

Use this window (Figure 11-20) to view network path utilization (Grid Network Throughput) statistics for the TS7700 Cluster.

Consideration: The Grid Network Throughput option is not available in a stand-alone cluster.

This window presents information about cross-cluster data transfer rates. This selection is present only in a multi-cluster grid configuration. If the TS7700 grid has only one cluster, there is no cross-cluster data transfer through the Ethernet adapters.

The example in Figure 11-20 is from a TS7700 Cluster that is part of a multi-cluster grid configuration (four-cluster grid).

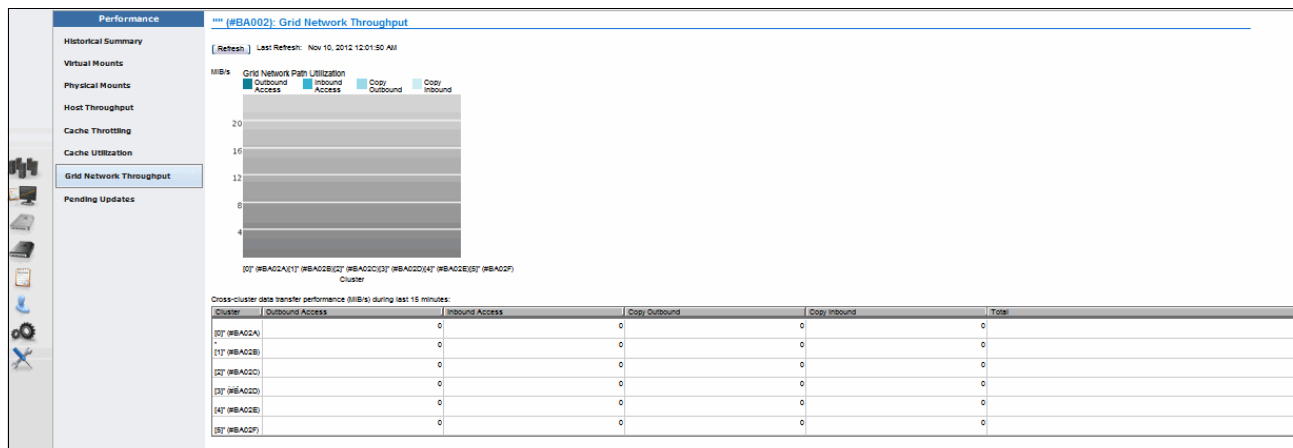


Figure 11-20 TS7700 MI grid network throughput in a six-cluster grid

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

Review this data with the performance evaluation tools on Techdocs to identify the following conditions:

- ▶ Identify upcoming performance problems due to grid link usage.
- ▶ Identify the amount of transferred data to review your settings, such as DAA, SAA, override policies, and Copy Consistency Points.

11.5 Cache capacity

The amount of used cache is depending on the installed configuration and the TS7700 models. In general, the TS7740 and TS7700T are configured with smaller cache sizes, because the main portion of the data will be destaged to the backend environment.

For TS7700D or the TS7700T CP0 the aim is that all data is kept in cache. However it is possible to use the function of autoremoval, to allow data in a TS7700D or TS7700T CP0 to be removed, if otherwise a short on storage condition would occur.

In the TS7740 and TS7700T CPx the storage class with the storage preference determine if a logical volume is kept in cache (PG1) or is migrated as soon as possible.(PG1)

In the TS7700D or TS7700T CP0, you can control if autoremoval is allowed or not. If autoremoval is enabled, it is possible to define, if a single logical volume is not eligible for autoremoval (Pinned), is eligible but should be kept if possible (prefer keep) or should be removed first if storage is needed (prefer remove).

Other than this, additional cache control exists:

- ▶ Tape partitions in the TS7700T to restrict the amount of data that can be used by a specific environment (LPAR/ Application/ HLQ), depending on the ACS routines
- ▶ How much space is protected in the CP0 that cannot be used for overspill
- ▶ Delay premigration in CPx, to keep data in cache for a specific time
- ▶ How copies are treated in the clusters (LI REQ SETTING: COPYFSC)
- ▶ How recalled data is treated in the cluster (LI REQ SETTING: RECLPG)
- ▶ How unwanted copies are treated (LI REQ SETTING: EXISTDEL)

11.5.1 Interpreting Cache Usage: MI

In the MI, you have several possibilities to view the cache usage. In the Cluster summary, you can determine how much space is consumed. In a TS7700T, you can use the Tape Partition screen to get more detailed information. In the TS7740 and TS7700D, you find similar information in the Performance Section “Cache Utilization.”

Remember that this information is only a snapshot view.

11.5.2 Interpreting Cache Usage: VEHSTATS

In the H30TVCx reports, you find information about the cache usage. In a TS7700D and TS7740 the whole cache is reported in the H30TVC.

In a TS7700T, the CP0 usage is reported in the H30TVC1, and CP1 is reported in H30TVC2 and so on.

There is no explicit usage of the cache capacity for each partition reported. The total TVC usage is reported in TOTAL TVC_GB USED. Also, you find information regarding the following data:

- ▶ How many logical volumes are kept in the different preference classes, depending on the models
- ▶ How long logical volumes are kept in the different storage preferences (4 hours, 48 hours and 35 days for TS7740 or TS7700T CPx).
- ▶ How many logical volumes have been removed with autoremoval

11.5.3 Interpreting Cache Usage: LI REQ,distlib,CACHE

The LI REQ command with the subcommand CACHE shows you the actual usage of the cache of the cluster that is defined in the command. Figure 11-21 on page 667 shows an example output.

In the example that is shown in Figure 11-21 on page 667, multiple tape partitions exists. In the CP1 1287 GB are waiting for premigration. This TS7700 has only 1 FC 5274 installed, so throttling was applied (PMT and CPYT).

```

TAPE VOLUME CACHE STATE V4.0
TS7700 VIRTUALIZATION ENGINE MODEL: TS7720 TAPE ATTACH
TOTAL INSTALLED/ENABLED GBS: 23859 / 23859
TOTAL ADJUSTED CACHE USED GBS: 2196
CACHE ENCRYPTION STATUS: CAPABLE
OVERCOMMITTED CACHE PARTITIONS: NONE
PRIMARY CACHE RESIDENT ONLY INFORMATION
PRIVATE CACHE USED GBS: 3848
SCRATCH CACHE USED GBS: 0
CP  ALLOC  USED   PIN   PKP   PRM   COPY  CPYT
0   8859   2012   0    2012  0     0     0
PRIMARY TAPE MANAGED PARTITIONS
CP  ALLOC  USED   PGO   PG1  PMIGR D_PMIGR  COPY  PMT  CPYT
1   8000   1756  1287  469  1287  0     0     41  41
2   4000    79    0     79   0     0     0     41  41
3   2000    0     0     0     0     0     0     41  41
4    0     0     0     0     0     0     0     41  41
5   1000    0     0     0     0     0     0     41  41
6    0     0     0     0     0     0     0     41  41
7    0     0     0     0     0     0     0     41  41

```

Figure 11-21 EXAMPLE of LI REQ,distlib,CACHE

11.5.4 Tuning cache usage - Making your cache deeper

A deeper cache improves 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 that requires a recall from physical tape. The TS7700 statistics provide a cache hit ratio for read mounts that can be monitored to ensure that 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.
- ▶ For TS7740 or tape attached partitions in the TS7700T, use the Storage Class (SC) construct to use Preference Level 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 while possible, to be available for recalls. Many clients have effectively made their cache deeper by examining their jobs and identifying which of them are most likely not to be recalled. Use SC to assign these jobs to PG0.

- ▶ For TS7700D or a TS7700T CP0, set the SC constructs to use **prefer remove** for volumes that you do not expect to be mounted. Use **pinned** for those that you know you will be mounting and need fast access times, and **prefer keep** for the others.
- ▶ With the SC construct PG1, the volume on the selected TVC for I/O operations is preferred to be in the cache of that cluster. The copy that is made on the other clusters is preferred to be removed from cache. If the TS7700 is used for the copy, ensure that this default setting is not overridden by the **Host Console** command. The behavior can be set by using **SETTING, CACHE, COPYFSC**:
 - When disabled, logical volumes that are copied into cache from a Peer TS7700 are managed as PG0 (prefer to be removed from cache).

- When the **ENABLE** keyword is specified, the logical volumes that are copied into the cache from a peer TS7700 are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7700.

This setting works on a distributed library level. It needs to be specified on each cluster. For a deep cache, **DISABLE** is the preferred keyword.

- ▶ By default, logical volumes that are recalled into cache are managed as though they were newly created or modified. You can modify cache behavior by using the **SETTING** Host Console command: **SETTING, CACHE, RECLPGO**:
 - When enabled, logical volumes that are recalled into cache are managed as PG0 (prefer to be removed from cache). This overrides the actions that are defined for the SC that is associated with the recalled volume.
 - When the **DISABLE** keyword is specified, logical volumes that are recalled into cache are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7700.

This setting works on a distributed library level. It needs to be specified on each cluster. The preferred keyword depends on your requirements. **ENABLE** is the best setting if it is likely that the recalled logical volumes are used only once. With the setting **DISABLE**, the logical volume stays in cache for further retrieval if the SC is defined as PG1 in the cluster that is used for the I/O TVC.

11.5.5 Tuning cache usage - Management of unwanted copies

In rare cases, a copy of a logical volume on a cluster exists that should not exist on this cluster according to the Management Class (MC) definition. This situation can occur after a read, without the “retain copy mode option” or after migration scenarios. This copy was flagged as “E” for exist.

In previous releases, these copies were deleted at mount/demount time, if the copy was inconsistent. If the copy was consistent, it was kept in the cluster.

With R3.2, this command was enhanced, and now enables you to determine not only how unwanted copies are treated, but also when this command is run.

With R3.1, a new HCR command was introduced, with the following options:

LI REQ,distributed library,SETTING,EXISTDEL,CRITERIA,[STALE|ALL NONE]

- ▶ **STALE**: The “E” copy is deleted if this copy is inconsistent. This is the default.
- ▶ **ALL**: The “E” copy is deleted from the cluster, if all other non-“E” copy mode sites are consistent.
- ▶ **NONE**: The “E” copy will never be deleted.

Before Release 3.2, these settings were acknowledged only if the logical volume is mounted or unmounted. With Release 3.2, a new command has been introduced to determine when these settings are acknowledged:

LI REQ,distributed library,SETTING,EXISTDEL,WHEN,[ACTCLOSE|AUTO]

- ▶ **ACTCLOSE**: The “E” copy is deleted by mount /dismount processing. This is the same behavior before R3.2. This is the default.
- ▶ **AUTO**: The “E” copy is deleted by periodic checks. The check runs once per 24 hours, and deletes 100 “E” copies during that process.

The deletion of an “E” copy can be processed only if all clusters in the grid are available. That is necessary because the status of all copies needs to be determined.

These commands are only available if all clusters are on R3.2 or higher.

11.6 Cache throughput / Cache bandwidth

The TS7700 cache has a finite bandwidth, depending on your actual environment (number of installed cache drawers and models). Other influences are compression ratio and block sizes used. For more information, see the performance white paper, *IBM TS7700 R4 (TS7760) Performance*, WP102652. Always use the most recently published performance white papers available on the Techdocs website at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

This TVC bandwidth is shared between the host I/O (compressed), copy activity (from and to other clusters), premigration activity, recalls for read and remote writes from other clusters. The TS7700 balances these tasks by using various thresholds and controls to prefer host I/O.

This section gives you some information how to determine if your cache bandwidth is a limitation, for what actions the cache bandwidth is used, and tuning actions that you can consider.

11.6.1 Interpreting Cache throughput: Performance graph

As described in “Historical Summary” on page 656, the Historical Summary throughput chart give you an overview about the cache bandwidth. Due to the nature of the MI, you can see only the throughput of maximum 1 day, and only one cluster at a time. However, it is a good starting point if you are looking to an actual performance problem to identify, if the cache bandwidth is exhausted.

11.6.2 Interpreting cache throughput: VEHSTATS HOURFLOW

In the VEHSTATS reports, you find a report called “HOURFLOW.” This report gives you per cluster, on a 15-minute or 1-hour interval, the overall compressed throughput, and which task is using the cache bandwidth. Figure 11-22 shows such a report.

REPORT=HOURFLOW (16090)		DATA FLOW IN MiB/sec BY CLUSTER										RUN ON 09MAY2016 @ 14:12:40		PAGE 1								
F DIST_LTB_ID=00		VE_CODE_LEVEL= 32.02.0001																				
Day	Time	Avg CPU Util	Max CPU Util	Avg Disk Util	Max Disk Util	MiB/s Total Xfer	MiB/s To_TVC Dev_Wr	MiB/s Fr_TVC Dev_Rd	MiB/s To_TVC Recv	MiB/s Fr_TVC Sent	MiB/s To_TVC Recall	MiB/s Fr_TVC PreMig	MiB/s By_GGM	Queue GiB_to PreMig	Queue GiB_to Copy	Queue GiB_to Recv	Write Throt Impac%	Copy Throt Impac%	Avg mSec	MiB/s To_TVC RMT_WR	MiB/s Fr_TVC RMT_RD	Intvl Sec
Sun	16:15:00	6	11	1	7	18.6	0.0	0.0	0.0	0.0	18.6	0.0	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	16:30:00	7	16	4	11	60.2	0.0	0.1	0.0	0.0	56.8	3.1	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	16:45:00	21	34	15	33	115.7	68.3	0.3	0.0	0.0	41.4	5.6	0.0	0	55	19	0.00	0.00	0	0.0	0.0	900
Sun	17:00:00	26	45	25	70	107.2	50.0	2.1	0.0	0.0	55.0	0.0	0.0	0	32	80	0.00	0.00	0	0.0	0.0	900
Sun	17:15:00	20	38	23	71	103.7	20.0	0.0	0.0	0.0	70.4	13.3	0.0	0	23	24	0.00	0.00	0	0.0	0.0	900
Sun	17:30:00	23	43	31	80	497.0	0.0	0.0	0.0	0.0	57.0	439.9	0.0	0	0	11	0.00	0.00	0	0.0	0.0	900
Sun	17:45:00	15	35	15	48	174.0	0.3	0.0	0.0	0.0	44.3	129.3	0.0	0	0	4	0.00	0.00	0	0.0	0.0	900
Sun	18:00:00	10	37	7	19	89.8	1.0	0.0	0.0	0.0	67.1	21.6	0.0	0	0	38	0.00	0.00	0	0.0	0.0	900
Sun	18:15:00	14	29	11	48	113.0	2.1	0.0	0.0	0.0	54.5	56.3	0.0	0	3	20	0.00	0.00	0	0.0	0.0	900
Sun	18:30:00	10	27	9	39	78.2	3.4	0.0	0.0	0.0	47.0	27.8	0.0	0	0	2	0.00	0.00	0	0.0	0.0	900
Sun	18:45:00	9	15	5	13	56.2	4.5	0.0	0.0	0.0	51.7	0.0	0.0	0	7	0	0.00	0.00	0	0.0	0.0	900
Sun	19:00:00	8	17	4	12	52.9	4.5	0.0	0.0	0.0	27.0	21.3	0.0	0	0	3	0.00	0.00	0	0.0	0.0	900
Sun	19:15:00	8	14	2	9	14.5	0.0	0.0	0.0	0.0	0.0	14.5	0.0	0	0	5	0.00	0.00	0	0.0	0.0	900
Sun	19:30:00	6	12	2	6	9.0	0.0	0.0	0.0	0.0	0.0	9.0	0.0	0	0	3	0.00	0.00	0	0.0	0.0	900
Sun	19:45:00	7	13	2	8	11.4	0.0	0.0	0.0	0.0	0.0	11.4	0.0	0	0	7	0.00	0.00	0	0.0	0.0	900
Sun	20:00:00	5	11	1	7	8.7	0.3	0.0	0.0	0.0	0.0	8.4	0.0	0	0	3	0.00	0.00	0	0.0	0.0	900
Sun	20:15:00	6	10	2	7	4.1	0.0	0.0	0.0	0.0	0.0	4.1	0.0	0	3	7	0.00	0.00	0	0.0	0.0	900
Sun	20:30:00	7	12	3	10	23.1	0.0	0.0	0.0	0.0	0.0	23.1	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	20:45:00	4	10	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	21:00:00	8	14	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	21:15:00	9	18	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	21:30:00	5	10	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	21:45:00	8	19	1	6	4.8	2.6	0.0	0.0	0.0	0.0	2.2	0.0	0	7	0	0.00	0.00	0	0.0	0.0	900
Sun	22:00:00	9	14	5	10	29.4	0.3	0.0	0.0	0.0	0.0	29.1	0.0	0	0	2	0.00	0.00	0	0.0	0.0	900
Sun	22:15:00	11	17	3	10	5.7	2.0	0.0	0.0	0.0	0.0	3.7	0.0	0	4	4	0.00	0.00	0	0.0	0.0	900
Sun	22:30:00	7	14	3	14	40.1	0.0	0.0	0.0	0.0	0.0	40.1	0.0	0	0	2	0.00	0.00	0	0.0	0.0	900
Sun	22:45:00	7	12	1	6	6.2	0.0	0.0	0.0	0.0	0.0	6.2	0.0	0	0	3	0.00	0.00	0	0.0	0.0	900
Sun	23:00:00	8	18	1	12	4.4	0.0	0.0	0.0	0.0	0.0	4.4	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900
Sun	23:15:00	15	20	2	5	25.5	3.7	16.7	0.0	0.0	0.0	5.1	0.0	0	0	2	0.00	0.00	0	0.0	0.0	900
Sun	23:30:00	15	19	3	6	26.7	2.6	11.5	0.0	0.0	0.0	12.5	0.0	0	3	17	0.00	0.00	0	0.0	0.0	900
Sun	23:45:00	7	20	5	29	32.9	0.1	2.1	0.0	0.0	0.0	30.6	0.0	0	0	0	0.00	0.00	0	0.0	0.0	900

Figure 11-22 Hourflow Report of VEHSTATS

The MiB/s Total Xfer is an average value. Notice that some peaks maybe higher.

In this example, the maximum value is 497 MBps, but that is driven by the premigration task. If that would cause an performance issue, you should review the PMPRIOR and PMTHLVL setting for tuning.

11.6.3 Tuning Cache bandwidth: Premigration

To tune the usage of cache bandwidth, you have several possibilities.

- ▶ When premigration will be run (PMPRIOR and PMTHLVL)
- ▶ Amount of drives for premigration
- ▶ Use delayed premigration to either never premigrate data or delay the premigration to a more suitable time slot

Fast Host Write premigration algorithm

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 MiB/s, the CPU is more than 99% bus, and the total I/O rate (read and write) against the cache is greater than 200 MBps (not MiB/s).

The circle on the graph (Figure 11-23) illustrates this algorithm in effect. During the 16:15 to 16:45 intervals, the amount of premigrate activity is limited. During the next six intervals, the premigration activity is zero. After this period of intense host activity and CPU usage, the premigrate tasks are allowed to start again.




Figure 11-23 Fast Host Write premigration algorithm

11.6.4 Premigration and premigration throttling values

These two parameters are used to define actions, triggered by the amount of non-premigrated data in the cache. Both values are applicable only to the TS7700T and the TS7740.

Premigration Priority threshold (PMPRIOR)

When this threshold is crossed, premigration processes start and increase, and the host throughput tends to decrease from the peak I/O rate. The default value is 1600 GB. The TS7740 and TS7700T uses various criteria to manage the number of premigration tasks. The TS7700 looks at these criteria every 5 seconds to determine whether more premigration tasks must be added. Adding a premigration task is based on the following factors, among others:

- ▶ Host-compressed write rate
- ▶ CPU activity
- ▶ The amount of data that needs to be premigrated per pool
- ▶ The amount of data that needs to be premigrated in total

Premigration Throttling threshold (PMTHLVL)

When this threshold is crossed, the host write throttle and copy throttle are both started. The purpose is to slow incoming data to allow the amount of non-premigrated data to be reduced and not rise above this threshold. The default value is 2000 GB.

As stated before, the workload creating data (Host I/O, Copy, or remote write) in the TS7700T CP0 is not subject to throttling for these values.

Premigration Throttling threshold (PMTHLVL)

When this threshold is crossed, the host write throttle and copy throttle are both started. The purpose is to slow incoming data to allow the amount of non-premigrated data to be reduced and not rise above this threshold. The default value is 2000 GB.

As stated before, the workload creating data (Host I/O, Copy, or remote write) in the TS7700T CP0 is not subject to throttling for these values.

Determining the values for your environment: PMPRIOR and PMTHLVL

When you define the values for PMPRIOR and PMTHLVL, they have several dependencies and consequences. Especially after hardware replacements, you need to review these parameters to ensure that the parameters are adjusted.

There is no guideline about the values of PMPRIOR and PMTHLVL. The following aspects need to be considered:

- ▶ Installed number of FC for Cache enablement for a TS7740
- ▶ Installed number of FC 5274 for premigration queue size for a TS7700T
- ▶ Workload profile
- ▶ Requirement regarding how long data should stay in cache unpemigrated

You need to determine the balance. If throttling occurs, it can be monitored with the MI or the VEHSTATS reports. You should review the values periodically.

To adjust the parameters, use the Host Console Request command. When you try to define a not allowed value, the TS7700 automatically uses an appropriate value. Details about these settings are described 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 following keywords:

- ▶ SETTING, CACHE, PMPRIOR
- ▶ SETTING, CACHE, PMTHLVL

<http://www.ibm.com/support/techdocs/atmastr.nsf/Web/TechDocs>

PMPRIOR setting

If the value of PMPRIOR is crossed, the TS7740/TS7700T starts the premigration. This might decrease the resources available for other tasks in the TS7740/TS7700T and shorten the peak throughput period of a workload window.

Raising this threshold increases the timeline where the TS7740/TS7700T can run in Peak mode. However, the exposure is more for non-premigrated data in cache.

Having a low PMPRIOR causes data to be premigrated faster and avoids hitting the premigration throttling threshold.

The default is 1600 GB, and the maximum is the value of PMTHLVL minus 500 GB.

PMTHLVL setting

After the PMTHLVL is crossed, the Host I/O, remote writes, and copies into a TS7740 and TS7700T are throttled. In a TS7700T, only workload to the Tape partitions is subject to throttling. If the data continues to increase after you hit the PMTHLVL, the amount of delay for throttling will continue to increase.

Raising the threshold avoids the application of the throttles, and keeps host and copy throughput higher. However, the exposure is more for non-premigrated data in cache.

The default value is 2000 GB. The maximum is:

- ▶ TS7740: The number of installed TB cache enablement (FC 5276) times 1 TB minus 500 GB
- ▶ TS7700T: The number of installed premigration queue size (FC 5274) times 1 TB

How to determine the Premigration Queue Size Feature Codes or Tape Cache Enablement

The size of the installed feature codes can be determined in the following ways:

- ▶ The MI, in the window “Feature Code License”
- ▶ The VEHSTATS Report H30TVCx:
 - TVC_SIZE: For a TS7740, the enabled tape cache (not the installed size) is displayed. For a TS7700D or TS7700T the installed cache size is displayed.
 - PARTITION SIZE: For a TS7740 the enabled tape cache size is displayed again, for the TS7700T the partition size is reported.
 - PRE-MIG THROT VALUE: Shows the PMTHLVL value. In a TS7700T, the recommendation is to set the PMTHLVL equal to the amount of FC 5274. If that recommendation was used, this is the amount of FC installed.

11.7 TS7700 throughput: Host I/O increments

In a TS7700, the read/write throughput from the Host is generally limited by the number of Host I/O increments that are installed. This can be from 1 increment (100 MBps uncompressed data) up to 25 increments (unlimited) in a TS7700 with R3.2 and 8 GB Ficon installed. To understand how many MiB/s Host I/O a TS7700 can absorb as a maximum, the following aspects need to be considered:

- ▶ The configuration:
 - The amount and type of drawers installed
 - The FICON attachment

- ▶ The compression ratio
- ▶ The replication mode
- ▶ The blocksize of the I/O
- ▶ The read/write ratio

Especially in a TS7720T, the cache bandwidth can be a limit for the HOST I/O as well.

To determine information about the Host I/O, you have multiple choices.

11.7.1 HOST I/O in the performance graphs

In the metrics, you can select the Host I/O information, as shown in Figure 11-24.




Figure 11-24 Select the Host I/O metrics

Then the Host I/O is displayed in the performance graph, as shown in Figure 11-25.




Figure 11-25 Performance Graph with Host I/O

11.7.2 HOST I/O in the VEHSTATS

In the VEHSTATS report H20VIRT you find the statistic regarding the host I/O. Reported are the actual installed Host I/O increments, and the so called attempted throughput. As long the host I/O did not reach the limit, "LESS" is reported.

The most important fields on this page give you information regarding your Host I/O (throughput) behavior. The MAX THRPUT is the enabled Host Throughput limit with FC 5268.

The ATTMPT THRPUT is the amount of MBps the host wanted to deliver to the TS7700. In combination with the next three values, you can determine whether the installed FC 5268 is sufficient, or if an upgrade would provide a better performance. The DELAY MAX, DELAY AVG, and PCT of 15 Sec Intervals fields tell you how much delay is applied to each second during the 15-second sample interval. VEHSTATS reports thousandths of a second with a decimal point, which is the same as milliseconds if no decimal point is present. Statistics records are created every 15 minutes (900 seconds), so there are 60 of the 15-second intervals that are used to report the 15-minute interval values. Most especially, the PCT of 15 sec INTVLS tells you how often a delay occurred.

Our example in Figure 11-26 on page 674 shows a TS7700 that benefits from more host throughput increments to provide a better job workflow.

This calculation is only an indicator. If you want to enhance the number of host I/O increments, talk to your IBM representative for sizing help.

GRID	DIST_LIB_ID= 0				VNODE_ID= 0		NODE_SERIAL=		
RECORD	VIRTUAL_DRIVES--				MAX	ATTMPT	_THROUGHPUT_		PCT_OF
TIME	INST	MIN	AVG	MAX	THRPUT	THRPUT	Delay_/15Sec	MAX	15Sec
					R2.2	CALC		AVG	INTVLS
							<----R3.0.0063---->		
05:00:00	256	0	0	0	500	less	.000	.000	0
05:15:00	256	0	0	0	500	less	.000	.000	0
05:30:00	256	0	4	128	500	1291	.613	.025	5
05:45:00	256	128	174	185	500	1377	.637	.634	100
06:00:00	256	175	188	203	500	1385	.639	.635	100
06:15:00	256	173	178	185	500	1381	.638	.634	100
06:30:00	256	177	188	199	500	1547	.677	.635	100
06:45:00	256	165	186	197	500	1381	.638	.635	100
07:00:00	256	153	166	181	500	1381	.638	.634	100
07:15:00	256	144	157	173	500	1381	.638	.635	100
07:30:00	256	123	143	161	500	1381	.638	.634	100
07:45:00	256	21	83	123	500	1381	.638	.603	100
08:00:00	256	0	7	25	500	780	.359	.014	23
08:15:00	256	0	0	0	500	less	.000	.000	0
08:30:00	256	0	0	0	500	less	.000	.000	0

Figure 11-26 Throughput Host I/O in VEHSTATS

11.7.3 Host Throughput Feature Codes

If the TS7700 is equipped or has been upgraded with 8-Gb FICON cards, take into account the fact that more throughput increments might need to be considered to unleash the full data transfer capabilities.

The previous 4-Gb FICON system had a maximum of 10 throughput increments, any data rate above 1 gigabyte per second (GBps) were given for free. With the new 8-Gb cards (or after an upgrade occurs), the new throughput increments limit is 25 GBps.

If a cluster was able to achieve speeds faster than 1 GBps with 4-Gb FICON cards and 10 throughput increments in the past, that will no longer be true because the TS7700 limits them to exactly 1 GBps, supposing that 8-Gb FICON cards were installed and the same 10 throughput increments were licensed. Thus, consider purchasing enough throughput increments (up to 25) to allow TS7700 cluster to run at unthrottled speeds.

See Chapter 7, “Hardware configurations and upgrade considerations” on page 227 for more details.

Figure 11-27 shows the feature code license entry picture.




Figure 11-27 Feature Code license entry picture

Figure 11-28 shows the installed increments (Feature Code (FC) 5268). In this example, four increments are installed. The throughput is limited to 400 megabytes per second (MBps) because of number of the installed 100 MBps increments.




Figure 11-28 Feature Code licenses example

11.7.4 Tuning for HOST I/O

Tuning for the Host I/O is limited. If the increments are exhausted, there is no possibility in the TS7700 to solve this issue. The only possibility is to change your workload profile in the attached hosts.

If the Host I/O is limited because the cache bandwidth is on the limit, refer to “Tuning Cache bandwidth: Premigration” on page 670.

11.8 Grid link and replication performance

The following section gives an insight into what aspects and definitions influence the gridlink performance. It also provides information how to monitor and to tune the performance of the copy actions.

The grid link and replication performance depends on the following aspects:

- ▶ Installed grid link hardware
- ▶ Sufficient bandwidth and quality of the provided network
- ▶ Chosen replication mode
- ▶ Defined amount of concurrent copytasks
- ▶ Number of remote write/read operations

Remember that cache bandwidth is always an influencing factor, but was already described in the last paragraphs. So this information is not included in this section.

11.8.1 Installed grid link hardware: Mixing of different Grid link adapters

It is not supported to have different grid link adapter types in one single cluster. However, in a grid, there can be a situation in which some clusters are connected to the grid link with 10 GB adapters, and other clusters are connected with 1 GB adapters. That is especially true for migration or upgrade scenarios.

In the TS7700 grid, there is a 1:1 relationship between the primary and primary adapters, and the secondary and secondary adapters. Due to that reason, in a mixed environment of 2*10 GB and 4*1 GB adapters, the clusters with the 4*1 GB links cannot use the full speed of the installed grid link adapters.

Remember, that 4*10 GB can be installed only in a VEC. A VEB/C07 with R4.0 cannot be upgraded to use 4*10 GB.

11.8.2 Bandwidth and quality of the provided network

The network between the TS7700s must have sufficient bandwidth to account for the total replication traffic. If you are sharing network switches among multiple TS7700 paths or with other network traffic, the total 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, other key factors affect the throughput that the TS7700 can achieve. The following factors directly affect performance:

- ▶ 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, and performance

The TS7700s attempt to drive the network links at the full 1-Gb rate for the two or four 1-Gbps links, or at the highest possible load at the two 10-Gbps links, 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.

Important: When the system attempts to give the network more data than it can handle, it discards packets that it cannot handle. This process causes TCP to stop, resynchronize, and resend amounts of data, resulting in a 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 that is 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 wide-area LANs (WANs) capability.
- ▶ Flow control between the switches is also a potential factor to ensure that the switches are able to pace with each other's rate.
- ▶ Be sure that the performance of the switch can handle the data rates that are expected from all of the network traffic.

Latency between the sites is the primary factor. However, packet loss, because of bit error rates or because the network is not capable of the maximum capacity of the links, causes TCP to resend data, which multiplies the effect of the latency.

11.8.3 Selected replication mode

As already mentioned, the number of concurrent running copy tasks, and the replication mode has an influence to the overall performance of a TS7700 grid. The following shared resources are consumed:

- ▶ CPU cycles of the TS7700
- ▶ Grid bandwidth
- ▶ Cache bandwidth

Synchronous mode copy

Synchronous mode copy can have a positive effect to the cache bandwidth. In opposite to all other replication modes, where the data needs to be read from cache to produce the copy, the SYNC is written directly to the remote sync cluster.

In addition, the synchronous mode copy does also not adhere to the same rules as the other copy modes, as shown in Table 11-1.

Table 11-1 Synchronous mode rules comparison to Run or Deferred modes

Subject	Synchronous mode copy	Run or Deferred mode
Data direction	Data is pushed from the primary cluster.	Data is pulled from the secondary cluster.
Throttling	Synchronous mode copies will not be throttled.	These copies can be throttled.
Number of concurrent copies can be controlled by a setting.	No	Yes
Copies can be halted by a <code>disable gridlink</code> command.	No	Yes

Synchronous mode showing in the MI

In the MI the synchronous mode mounts are only visible in the Virtual Tape Device view, of the Host I/O cluster. The chosen second synchronous target cluster does not show any mount. That is true, because for the synchronous mount an internal device is used, which is not shown in the Virtual Tape drive panel, as shown in Figure 11-29.

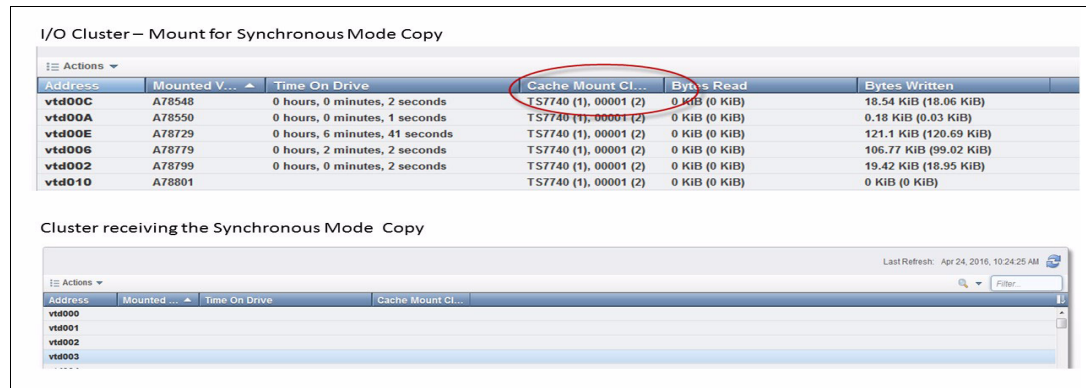


Figure 11-29 Virtual Tape Drive view for synchronous mode copy

Synchronous mode monitoring in the VEHSTATS

The information for synchronous mode are shown in H30TVCx, independent for each partition. As in the MI, all numbers that are related to the synchronous mode operation are only shown in the Host I/O cluster for these mounts. The “secondary” synchronous mode does not reflect synchronous mounts.

In the following picture, you see that in total 44 scratch mounts (FAST NUM MNTS) were made, all of them are scratch mounts. In addition, you see the same number in the SYNC NUM MNTS field, which means, that the same number of mounts to a remote cluster has been run.

The receiver of the synchronous copy reports nothing, as shown in Figure 11-30.

RECORD	AVG	MAX	AVG	MAX	PART	TOTAL	FAST_RDY	CACHE_HIT	CACHE_MIS	SYNC_MODE	P-MIG
END_TIME	CPU_UTIL	DISK_UTIL	HIT%	MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS
07:00:00	4	11	0	1	0	0	.00	0	.00	R2.1	R1.5
07:15:00	13	38	31	89	100	88	.71	44	.71	0	.00
07:30:00	23	33	67	95	100	194	.64	97	.64	0	.00
07:45:00	30	39	87	100	100	150	.85	75	.85	0	.00
08:00:00	23	29	64	91	100	128	.58	64	.58	0	.00
08:15:00	14	29	33	87	0	0	.00	0	.00	0	.00
08:30:00	11	34	23	99	100	64	.77	32	.77	0	.00
08:45:00	29	41	85	100	100	192	.81	96	.81	0	.00
09:00:00	25	33	68	100	100	76	.81	38	.81	0	.00
09:15:00	13	29	35	76	100	54	.62	27	.62	0	.00
09:30:00	14	27	42	86	100	94	.58	47	.58	0	.00
09:45:00	4	9	2	17	0	0	.00	0	.00	0	.00

RECORD	AVG	MAX	AVG	MAX	PART	TOTAL	FAST_RDY	CACHE_HIT	CACHE_MIS	SYNC_MODE	P-MIG
END_TIME	CPU_UTIL	DISK_UTIL	HIT%	MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS
07:00:00	4	10	0	1	0	0	.00	0	.00	R2.1	R1.5
07:15:00	13	31	23	69	0	0	.00	0	.00	0	.00
07:30:00	24	38	47	82	0	0	.00	0	.00	0	.00
07:45:00	28	35	60	85	0	0	.00	0	.00	0	.00
08:00:00	23	35	44	74	0	0	.00	0	.00	0	.00
08:15:00	13	24	22	55	0	0	.00	0	.00	0	.00
08:30:00	10	28	14	49	0	0	.00	0	.00	0	.00
08:45:00	25	34	46	64	0	0	.00	0	.00	0	.00
09:00:00	24	39	49	87	0	0	.00	0	.00	0	.00
09:15:00	16	37	30	78	0	0	.00	0	.00	0	.00
09:30:00	18	24	27	40	0	0	.00	0	.00	0	.00
09:45:00	5	9	3	18	0	0	.00	0	.00	0	.00

Figure 11-30 Synchronous mode copies shown in VEHSTATS

There is no further information (such as the number of concurrent sync copies, inconsistency at interval, and so on) for the synchronous mode available, because none of them is applicable.

Only if the synchronous mode copy could not be processed, and sync-deferred was activated, are reports written. However, then these copies are reported with “DEFERRED” and there is no possibility for a further drill down.

Run copies monitoring in the VEHSTATS

The information about the RUN copies are contained in the H33GRID report in the receiving cluster, as shown in Figure 11-31 on page 680. To see the whole grid performance, you need to look at each cluster individually. To understand how much RUN copies (amount of Ivols and MB) were processed in an interval, look to LVOLS TO_TVC_BY_RUN_COPY.

To understand, if RUN copies were queued for processing in that interval, look at AV_RUN_QUEAGE -- MINUTES--. Having numbers in here means that RUN could not be processed accordingly. Having RUN Ivols waiting also means that the job cannot process further. If the report shows multiple indications for this behavior, take a closer look at the number of concurrent copy activities and the grid link usage.

You might want to consider increasing the number of concurrent RUN tasks. Also check if all receiving clusters were available, or if a cluster went to service or had an outage during that interval.

Figure 11-31 shows the H33GRID report.

08NOV15SU	LVOLS TO RECEIVE	MiB TO RECEIVE	AV_DEF QUEUE ---MINUTES---	AV_RUN QUEUE	#_LVOLS TIM_DLY CPY_QUE	LVOLS __TO_TVC_BY__ RUN_COPY	MB__	LVOLS __TO_TVC_BY__ DEF_COPY	MB__
01:00:00	21	3894	101	0	0	20	18172	30	11494
02:00:00	6	5945	5	0	0	20	19348	20	6203
03:00:00	21	15408	12	0	0	20	12074	14	9884
04:00:00	1	792	0	0	0	57	37145	33	24100
05:00:00	9	2291	7	0	0	56	34895	12	8762
06:00:00	37	28216	22	0	0	74	50541	16	9668
07:00:00	2	1490	0	0	0	48	28204	48	34513
08:00:00	10	4039	12	0	0	36	24274	14	10863
09:00:00	1	686	0	0	0	23	7137	14	10336
10:00:00	21	15442	11	12	0	16	10758	10	7649
11:00:00	17	10982	3	0	0	51	27853	19	14407
12:00:00	146	84369	28	0	0	0	0	0	0
13:00:00	254	146289	65	0	0	0	0	0	0
14:00:00	531	305714	72	0	0	0	0	0	0
15:00:00	13	8010	164	0	0	0	0	0	0
16:00:00	9	5633	186	0	0	0	0	0	0
17:00:00	9	5633	201	0	0	0	0	0	0

Figure 11-31 H33Grid report to see copy behavior for RUN and Deferred copies

Deferred copies monitoring in the VEHSTATS

For deferred copies you have in report H33GRID only the AV_DEF QUEUE -- MINUTES --- Usually the deferred copy queue is constantly increasing during batch processing. The reason is that deferred copies are throttled by the sending cluster, when specific workload values are reached.

To understand, if that throttling applies and is the reason for the increase, look in the H30TVCx report. The H30TVC1 report contains the information of CP0, and the H30TVC2-8 report contains the CP1- CP7 information.

The deferred copy throttling is for all partitions identically, so it is sufficient to look into one of the H30TVCx reports. As shown in Figure 11-32 on page 681, the H30TVC report contains detail information about the occurred throttling.

```

-----DEFER_COPY_THROTTLING-----
  NUM   NUM   AVG
15MIN 30SEC   SEC BASE
INTVL SMPLS /INTVL SECS REASN
-----R1.5----- R3.0
  1     1   .001 .125 x0000
  2     9   .009 .125 x0000
  4    63   .065 .125 x0000
  4    34   .035 .125 x0000
  2    13   .013 .125 x0000
  3    20   .020 .125 x0000
  4    33   .034 .125 x0000
  3     5   .005 .125 x0000
  3    16   .016 .125 x0000
  3    19   .019 .125 x0000
  3    10   .010 .125 x0000
  2     3   .003 .125 x0000
  2     3   .003 .125 x0000
  4    86   .089 .125 x0000
  4   110   .114 .125 x0000
  4   120   .125 .125 x0000
  4    89   .092 .125 x0000
  4    32   .033 .125 x0000
  3    14   .014 .125 x0000
  4    14   .014 .125 x0000
  3     9   .009 .125 x0000
  4    17   .017 .125 x0000
  3    14   .014 .125 x0000
  1     1   .001 .125 x0000

```

Figure 11-32 Deferred copy throttling

- ▶ NUM 15 MIN INTVL: Shows the number of intervals in the hour the deferred throttling occurred. A 4 means that in every interval of the 1-hour report, deferred throttling occurred. A 1 means, that only in one interval deferred copy throttling happened.
- ▶ NUM 30 SEC SMPLS: Shows in how many 30-second samples in the reported intervals the throttling occurred. That means, that in an hour report, you have a maximum of 120 samples (60 minutes * 2 samples of 30 seconds each). If 1 is reported, only in 30 seconds of the whole our a deferred throttling occurred.
- ▶ AVG SEC INTVL: Shows the actual amount of penalty in seconds given for each deferred copy action in this interval.

Looking to the Figure 11-32, you find one interval where deferred copy throttling occurred in all 120 samples. This results in the maximum value of .125 s penalty for each copy operation. However, in the report shown, there was no interval where no deferred copy throttling occurred, but in some intervals were limited throttling measured.

Be aware, that depending on your network, a throttling higher than 20 ms normally results in little or no deferred copy action. For further information how to influence the deferred copy throttling, refer to 11.8.5, "Tuning possibilities for copies: Deferred Copy Throttling" on page 682.

11.8.4 Tuning possibilities for copies: COPYCOUNT Control

There can be several reasons for tuning the counts of the number of concurrent copy jobs over the grid links.

Values can be set for the number of concurrent RUN copy threads and the number of Deferred copy threads. The allowed values for the copy thread count are 5 - 128. The default value is 20 for clusters with two 1-Gbps Ethernet links, and 40 for clusters with four 1-Gbps or two 10-Gbps Ethernet links. Use the following parameters with the **LIBRARY** command:

- ▶ **SETTING, COPYCNT, RUN**
- ▶ **SETTING, COPYCNT, DEF**

Increase the copy count

Increasing the copy count can be beneficial, if the following conditions occur:

- ▶ The gridlinks are not saturated
- ▶ The number of very small logical volumes is high
- ▶ An upgrade from 2 to 4 grid links was made - and the number of copy tasks were not adjusted.

In this case, an increase of the copy count might reduce the RPO.

Be aware, that usually one gridlink with 1 Gbps can be saturated by 10 copies running in parallel. If the logical volumes are small, you might see gaps in the grid link usage, when only a few copies are running, because some volumes finished, and new lvols were selected for copy processing. In this situation, it might be beneficial to have more copies running concurrently.

Take into account that if too many copies are running concurrently, you will overflow the grid link. That can result in package loss and retries, and can lead overall to a lower performance of the grid link environment.

If you want to increase the number, do that in smaller steps (5 or 10) to get experience with the new setting. In addition, do not define values that are too high, especially if you use synchronous mode copy concurrently to the RUN and Deferred traffic.

Decrease the copy count

Decreasing the copy count can be beneficial in the following situations:

- ▶ If you have limited network bandwidth (for example, less than the 100 MiB/s).
- ▶ If the bandwidth is limited, running too many copies in parallel prolongs the single copy time. This can result in time outs. When you cross this threshold, the system switches from RUN to IMMED-Deferred. For a deferred copy, the copy is deleted from the actual copy tasks and will be scheduled back to the copy queue.
- ▶ If packet loss is reported and hardware issue or the grid link quality is not the reason.
- ▶ If too many copies are running in parallel, a single copy stream might run into a packet time out, which is reported as packet loss.

11.8.5 Tuning possibilities for copies: Deferred Copy Throttling

The deferred copy throttle (DCT) value is used to regulate outgoing deferred copies to other clusters to prefer host throughput. For some, host throughput is more important than the deferred copies, but for others, deferred copies are as important. Adjusting the DCT value and threshold can enable you to tune the performance of the deferred copies.

Deferred Copy throttle value

When the DCT threshold is reached, the TS7700 adds a delay to each block of deferred copy data that is sent across the grid links from a cluster. The larger the delay, the slower the overall copy rate becomes.

The performance of the grid links is also affected by the latency time of the connection. The latency has a significant influence on the maximum grid throughput. For example, with a one-way latency of 20 - 25 milliseconds (ms) on a 2 x 1 Gb grid link with 20 copy tasks on the receiving cluster, the maximum grid bandwidth is approximately 140 MBps. Increasing the number of copy tasks on the receiving cluster increases the grid bandwidth closer to 200 MBps.

The default DCT is 125 ms. The effect on host throughput as the DCT is lowered is not linear. Field experience shows that the knee of the curve is at approximately 30 ms. As the DCT value is lowered toward 30 ms, the host throughput is affected somewhat and deferred copy performance improves somewhat. At and below 30 ms, the host throughput is affected more significantly, as is the deferred copy performance.

If the DCT needs to be adjusted from the default value, try an initial DCT value of 30 - 40 ms. Favor the value toward 30 ms if the client is more concerned with deferred copy performance, or toward 40 ms if the client is concerned about sacrificing host throughput.

After you adjust the DCT, monitor the host throughput and Deferred Copy Queue to see whether the wanted balance of host throughput and deferred copy performance is achieved. Lowering the DCT improves deferred copy performance at the expense of host throughput.

A DCT of "0" eliminates the penalty completely, and deferred copies are treated equally as host I/O. Depending on your RPO requirements that is also a feasible setting.

The DCT value can be set by using 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 by using the keywords SETTING, THROTTLE, and DCOPYT:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

Deferred Copy throttle value threshold

This value is used to determine the average host I/O rate at which 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 *DCTAVGTD – DCT 20-Minute Average Threshold* looks at the 20-minute average of the compressed host read and write rate. The threshold defaults to 100 MBps. The *Cache Write Rate – Compressed writes to disk cache* includes host write, recall write, grid copy-in write, and cross-cluster write to this cluster. The threshold is fixed at 150 MBps. *Cluster Utilization* looks at both the CPU usage and the disk cache usage. The threshold is when either one is 85% busy or more.

DCT is applied when *both* of the following conditions are true:

- ▶ Cluster utilization is greater than 85% or the cache write rate is more than 150 MBps.
- ▶ The 20-minute average compressed host I/O rate is more than DCTAVGTD.

The preceding algorithm was added in R2.0. The reason to introduce the cache write rate at R2.0 was due to the increased CPU power on the IBM POWER7 processor. The CPU usage is often below 85% during peak host I/O periods.

Before R2.0, the cache write rate was not considered. Use the following parameters with the **LIBRARY** command to modify the DCT value and the DCTAVGDT:

- ▶ **SETTING, THROTTLE, DCOPT**
- ▶ **SETTING, THROTTLE, DCTAVGDT**

Note: The recommendation is to not change DCTAVGDT and instead use for the tuning the DCOPTY only. Changing DCTAVGDT might not reduce the throttling as expected.

Application of Deferred Copy Throttle

The next two charts illustrate the use of DCT. In Figure 11-33, the amount of data that is being copied out is small because the DCT is being applied. DCT is applied because the compressed host I/O is above the DCT threshold, which is set to the default of 100 MBps.

Figure 11-34 on page 685 shows the compressed host I/O dropping below the 100 MBps threshold. As a result, the rate of deferred copies to other clusters is increased substantially.

Figure 11-33 shows the behavior when the DCT is used. The deferred copies are limited (light blue bars), and Host I/O (green bar) and Premigration (yellow bar) are preferred.

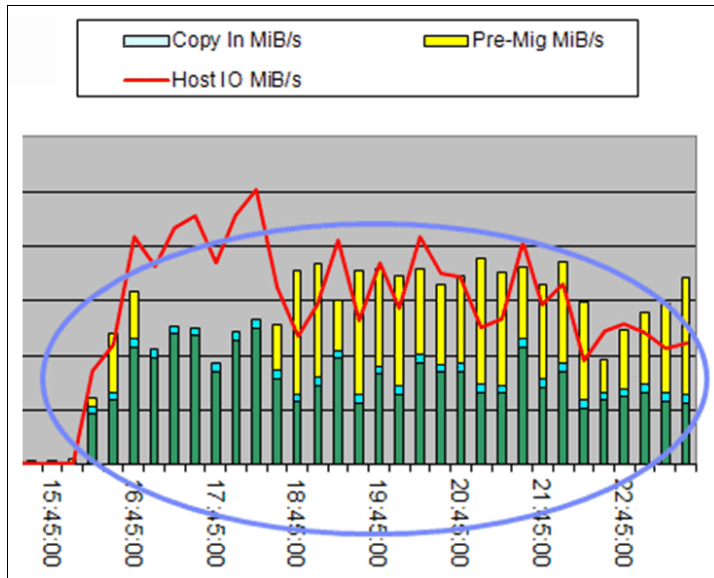


Figure 11-33 DCT being applied

In Figure 11-34, you see the effect when DCT is “turned off” because the host throughput drops under 100 MBps (green bar). The number of deferred copy writes in MBps increases (light blue bar).




Figure 11-34 DCT turned off

11.8.6 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 the grid links is monitored periodically, and if one link is performing worse than the other link by an IBM Service Support Representative (IBM SSR)-alterable value, a warning message is generated and sent to the host. The purpose of this warning is to alert you 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.

For example, a setting of 60% means that if one link is running at 100%, the remaining links are marked as degraded if they are running at less than 60% of the 100% link. The grid link performance is available with the Host Console Request function, and on the TS7700 MI. The monitoring of the grid link performance by using the Host Console Request function 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 and GRIDLINK keywords:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

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 5-minute intervals that link degradation was detected before reporting an attention message. The default value is 9.
- ▶ Generate Call Home iterations: The number of consecutive 5-minute intervals that link degradation was detected before generating a Call Home. The default value is 12.

The default values are set to 60% for the threshold, nine iterations before an attention message is generated, and 12 iterations before a Call Home is generated. Use the default values unless you are receiving intermittent warnings and support indicates that the values need to be changed. If you receive intermittent warnings, let the SSR change the threshold and iteration to the suggested values from support.

For example, suppose that clusters in a two-cluster grid are 2000 miles apart with a round-trip latency of approximately 45 ms. The normal variation that is seen is 20 - 40%. In this example, the threshold value is at 25% and the iterations are set to 12 and 15.

11.9 Considerations for the backend TS7740 / TS7700T

In the backend, we advise you to monitor the two different resources, backend drives and backend cartridges. To determine if you have sufficient backend resources, how they are actually used, and how the use increases over time, you have several possibilities. The most difficult question is if you have sufficient backend drives.

11.9.1 Amount of Back-end drives

It is important to ensure that enough back-end drives are available. If not enough backend drives are available, you might see the following issues:

- ▶ Recalls are too slow, because no free backend drive was available
- ▶ Premigration cannot be processed sufficiently, and therefore throttling occurred due to reaching the limit of premigration queue sizes
- ▶ More backend cartridges are needed, because no drives were available to run the reclaim

If there are insufficient back-end drives, the performance of the TS7740/TS7700T diminishes.

In a TS7740, there was a direct dependency between Host I/O Increments throughput and the number of backend drives. This was understandable, because the TS7740 had a limited cache and all data that was written to the cache also needed to be premigrated to backend cartridges.

This strict dependency does not exist any more in a TS7700T. First of all, part of the Host I/O can be written in the CP0, and this data will never be premigrated. Second, a part of the data might be expired in cache, even if they were written to CPx using the Delay Premigration parameter.

Therefore, such a strict relationship does not exist any more.

As a guideline for the TS7740, use the ranges of back-end drives that are listed in Table 11-2 based on the host throughput that is configured for the TS7740. The lower number of drives in the ranges is for scenarios that have few recalls. The upper number is for scenarios that have numerous recalls. Remember, these are guidelines, not rules.

Table 11-2 Performance increments versus back-end drives

Throughput (MiB/s)	Back-end drives
Up to 400	4 - 6
400 - 800	6 - 8
800 - 1200	8 - 12
1200 - 1600	10 - 16
1600 or higher	16

So, if the Host increments cannot be used for guidance anymore, the question is how to determine the number of needed backend drives.

As described previously, there is no overall rule. Here are some general statements:

- ▶ The more physical backend pools you use, the more physical drives you need. Each physical pool is treated independently (premigration, reclaim).

- ▶ Depending on the used physical cartridge type and the reclaim value, the amount of still valid data can be very high. Therefore, a reclaim has to copy a high amount of data from one physical tape to another tape. This uses two drives, and the more data that has to be transferred, the longer these drives will be occupied. Reclaim is usually a low priority task, but if not enough reclaims can be run, the number of necessary tapes increases.
- ▶ Tape drives are also needed for Copy Export and Secure data overwrite
- ▶ Data expires in cache without premigration and data in CP0 does not need tape drives
- ▶ Low hit ratio requires more recalls from the backend

Installing the correct number of back-end drives is important, along with the drives being available for use. *Available* means that they are operational and might be idle or in use. The Host Console Request function can be used to set up warning messages for when the number of available drives drops. Setting 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 these keywords:

- ▶ SETTING, ALERT, PDRVLOW
- ▶ SETTING, ALERT, PDRVCRT

Use the following website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

11.9.2 Monitor Backend drives in the MI

In the MI you can see the number and status of the backend drives, the actual mount situation, and in the performance section you find a snapshot of mount times.

11.9.3 Monitor Backend drives in the VEHSTATS

The report H32TDU12 gives you an overview about the usage of the backend drives. Example 11-1 shows:

- ▶ How many physical tape drives were installed, and how many 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 volumes back into cache
 - MIG: Premigration of logical volumes from cache to physical tape
 - RCM: Reclamation
 - SDE: Secure Data Erase

Example 11-1 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 following fields are the most important fields in this report:

- ▶ **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, you 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.

Be aware, that the number of tape drives might be misleading. The report does not recognize the "IDLE" state. Idle means that the tape drive is mounted, but not in use. Therefore, you might see a maximum - or even an average - usage that is equal to the installed drives. That might or might not be a performance issue. To identify if that really is a bottleneck, it is necessary to have a closer look at the overall situation.

To do so, first review which mounts are run. If a reclaim is still processed, there was no performance issue (except if a panic reclaim occurred).

If no reclaim was run, have a closer look at how long in average a physical cartridge could be mounted. To calculate this, take the total mounts and divide it by the number of installed drives and the amount of intervals sample. That shows how long a physical tape cartridge can be mounted on a physical tape drive. If this value is lower than 10 minutes, further investigations should be done. If this value is lower than 4 minutes, a performance issue is likely.

The H32GUPxx (General Pool Use) report is shown in Example 11-2. 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 11-2 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	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	LVOLS	GB	WRITTN	READ	SCR	PRIV	PCT	POOL	VOLS	VOLS	
UPD INT=>	-ON	THE	HOUR-		ON	THE	HR		-ON	THE	HOUR-		ON	THE	HR																
4:15:00	65079	18052	5412	0	2	56	25	01	00	00	0	0	0	0	0	0	0	25	02	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	0	25	02	00	00	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	0	25	02	00	00	0	0	0	0	0	25	02	00	00	
5:00:00	65630	18206	94721	0	2	57	25	01	00	00	0	0	0	0	0	0	0	25	02	00	00	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	0	25	02	00	00	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	0	25	02	00	00	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	0	25	02	00	00	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	0	25	02	00	00	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	0	25	02	00	00	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	0	25	02	00	00	0	0	0	0	0	25	02	00	00	

Check the ODERV12 statements from the BVIR jobs to select which types of cartridges are used in your environment. Only four different types of media can be reported at the same time.

11.9.4 Monitor Backend drives with a LI REQ command

The LI REQ,distlib,PDRIVE command gives you a snapshot about the physical tape drive environment. It shows the installed drives, the model, and the status. In addition, you can see what action is performed, for which pool the drive is working, and the actual mounted physical volume. If the status is not idle, also the actual logical volume in use is provided.

Figure 11-35 shows an output of this command:

```
LI REQ,DTS7720,PDRIVE
CBR1020I Processing LIBRARY command: REQ,DTS7720,PDRIVE.
CBR1280I Library DTS7720 request. 523
Keywords: PDRIVE
-----
PHYSICAL DRIVES V2 .1
  SERIAL NUM   TYPE  MODE  AVAIL  ROLE  POOL   PVOL   LVOL
0000078D1224 3592E07          Y  IDLE  00
0000078D0BAA 3592E07          Y  IDLE  00
0000078DBC65 3592E08          Y  IDLE  00
0000078DBC95 3592E08  E08   Y  MIGR  01  R00011  A51571
0000078DBC5C 3592E08  E08   Y  MIGR  01  R00001  A66434
0000078DBC87 3592E08  E08   Y  MIGR  01  R00002  A66462
```

Figure 11-35 LI REQ PDRIVE command example

11.9.5 Tune the usage of Back-end drives

If you need to tune the backend drive usage, you have multiple possibilities:

- ▶ Review the usage of delay premigration for data with short lifecycle
- ▶ Review the number of physical pools
- ▶ Review reclaim operations
- ▶ Review the amount of premigration drives
- ▶ Review Copy Export operation

Delay premigration usage

To eliminate the necessity to premigrate and reclaim data, you might consider using delay premigration for some of your data. To determine if this is a viable solution, you need to analyze the information from your tape management system. This shows if any data has such a short lifetime cycle, and how much cache space would be needed to do so.

Contact your IBM representative if you need further assistance to do so.

Amount of physical pools

As mentioned, each physical pool is treated independently regarding the backend drive usage. That is true for premigration and reclaim. Reducing the number of active pools can be helpful if the backend drive environment shows bottlenecks.

Reclaim operations

Reclaim operations use two drives per reclaim task. Reclaim operations also use CPU MIPs, but does not use any cache bandwidth resources, because the data will be copied from physical tape to physical tape directly. If needed, the TS7740/TS7700T can allocate pairs of idle drives for reclaim operations, making sure to leave one drive available for recall.

Reclaim operations affect host performance, especially during peak workload periods. Tune your reclaim tasks by using both the reclaim threshold and Inhibit Reclaim schedule.

Reclaim threshold

The reclaim threshold directly affects how much data is moved during each reclaim operation. The default setting is 35% for each pool. Clients tend to raise this threshold too high because they want to store more data on their stacked volumes. The result is that reclaim operations must move larger amounts of data and use drive resources that are needed for recalls and premigration. After a reclaim task is started, it does not free its back-end drives until the volume being reclaimed is empty.

Table 11-3 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. This change cuts the time and resources that are needed for reclaim in half. However, it raises the needed number of backend cartridges and slots in the library.

Table 11-3 Reclaim threshold by cartridge capacity

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
4000 GB	400 GB	800 GB	1200 GB	1600 GB
10000 GB	1000 GB	2000 GB	30200 GB	4000 GB

Inhibit Reclaim schedule

Use the Inhibit Reclaim schedule to inhibit reclaims during your busy periods, leaving back-end drives available for recalls and premigrates tasks. Generally, start the inhibit 60 minutes before 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 use two back-end drives per task, and CPU cycles as well. 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. For example, moderate host I/O during the uninhibited period and reclaim might use too many back-end drives, CPU cycles, or both.

With the **Host Library Request** command, you can limit the number of reclaim tasks in the TS7740/TS7700T. The second keyword RECLAIM can be used along with the third keyword of RCLMMAX. This expansion applies only to the TS7740/TS7700T. Also, the Inhibit Reclaim schedule is still acknowledged.

The maximum number of reclaim tasks is limited by the TS7740/TS7700T, based on the number of available back-end drives, as listed in Table 11-4. These values have changed during the evolution of the product, and might be different in previous releases.

Table 11-4 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 the number of premigration drives (maximum drives)

Each storage pool enables you to define the maximum number of back-end drives to be used for premigration tasks. Several triggers can cause the TS7740/TS7700T 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 must wait until a premigration task is complete for a back-end drive to be freed. 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 window. For a V06, the maximum number of premigration drives per pool must not exceed 6. Extra drives do not increase the copy rate to the drives. For a V07, TS7720T, or TS7760T, premigration can benefit from having 8 - 10 drives available for premigration, the default value is 10. There is no benefit to have more than 10 running premigration.

The limit setting is in the TS7740/TS7700T MI. For Copy Export pools, set the maximum number of premigration drives. If you are exporting a small amount of data each day (one or two cartridges' worth of data), 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.

To determine the number of drives to set, consider MB/GB written to a pool, compute MiB/s, compute maximum and average, and the number of premigration drives per pool. Base the number of drives by using 50 - 70 MBps per drive for models up to the TS1140, approximately 100 MBps for a TS1140, and approximately 140 MBps for a TS1150. These values are different from the native data rate because of the resources used by a TS7700 (for example, database updates and logical volume selection for premigration).

Avoiding 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 run the operation during a slower workload time.

11.9.6 Amount of Back-end cartridges

The number of needed back-end cartridges is determined not only by the amount of data being stored on the cartridges. Some other parameters can influence the number of cartridges you need.

- ▶ Reclaim value
- ▶ number of physical pools and pool properties
- ▶ Delete Expire setting

To monitor your actual situation, but also the trend of the cartridge usage, you have several possibilities.

11.9.7 Monitor the usage of Back-end cartridges on the MI

To view the active pools, select them from the **Physical volumes** → **Active data distribution** menu. Figure 11-36 shows the active pools and correspondent data distribution (number of cartridges by occupancy percentage range).

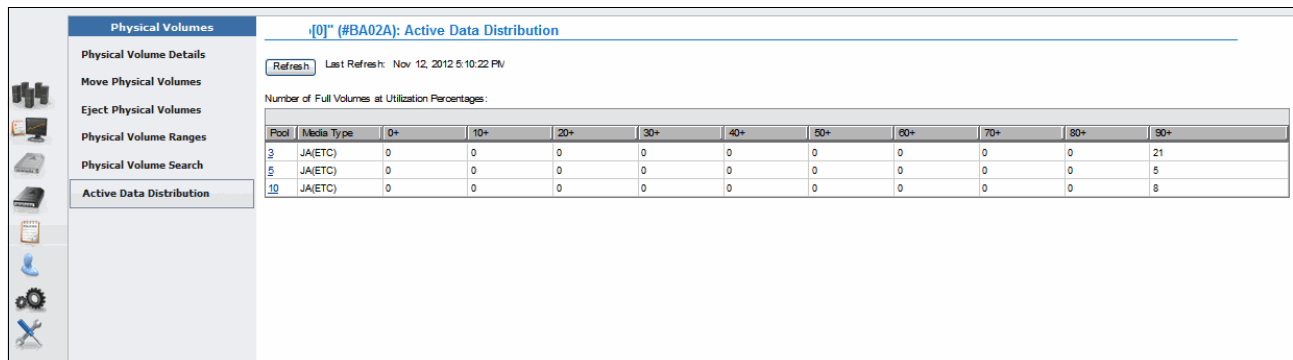


Figure 11-36 Pool Active Data Distribution

Click a pool link. Information for the pool is displayed, as shown in Figure 11-37.

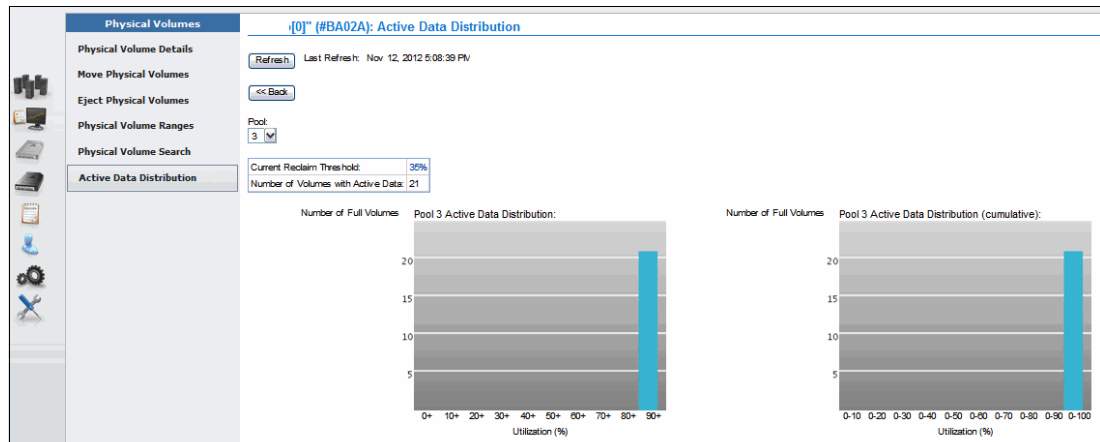


Figure 11-37 Information Display of selected Pool

Review your Active Data Distribution. A low utilization percentage results in a higher number of stacked volumes. Also, ensure that you monitor the number of empty stacked volumes to avoid an “out of stacked volumes” condition. If you have defined multiple physical pools, you might need to check this on a per pool basis, depending on your Borrow/Return policies. In this example, Pool 3 has the **borrow, return** parameter enabled.

11.9.8 Monitor the usage of Back-end cartridges with VEHSTATS

As explained in REF NEEDED, you can use the H32GUPxx report to view the cartridges on a per pool base.

In addition, to see the trend of empty cartridges, you should use the H32CSP report. This report provides an overview of the empty cartridges in the common scratch pool on a cartridge type basis, as presented in Example 11-3.

Example 11-3 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	0	0	0
4:30:00	0	0		42	0	0	0	0	0	0
4:45:00	0	0		42	0	0	0	0	0	0
5:00:00	0	0		41	0	0	0	0	0	0

Remember that it is not sufficient to check only the scratches in the common scratch pool. In addition, you need to check that all pools can borrow from the CSP and will return the empty cartridges to the CSP. If a pool is set to **no borrow**, you need to ensure that always enough empty cartridges are in inside this pool. This number is reflected in the H32GUPxx reports.

Keep in mind that in a heterogeneous environment, backlevel cartridges (JA/JB) can only be used for read and not for write purposes.

11.9.9 Tuning of the usage of Back-end cartridges with VEHSTATS

As stated, you have several possibilities to influence the number of used backend cartridges.

Reclaim value

As explained in the backend cartridge section, you can change the reclaim value to gain more empty cartridges. The lower the percentage of the reclaim value, the more cartridges are needed. The higher this value is, the more valid data needs to be transferred and physical tape drives are required.

To find a good balance, review the active data distribution. Some times, it is sufficient to change to a slightly higher reclaim value.

Amount of physical pool and pool properties

For each physical pool usually two empty cartridges are kept inside the pool. Especially for smaller configurations with JD cartridges, that might increase the need for more cartridges.

In addition, the pool properties should be reviewed. **No borrow/Keep** has a negative influence.

Amount of physical pool and pool properties

The delete expire value in the scratch categories defines how long a logical volume and the data on it will still be treated as valid data. For more information about delete expire, see “Logical Volume Delete Expire Processing versus previous implementations” on page 117.

Keep in mind, that a short delete expire value might reduce your cartridge usage, but a short value does not enable you to rescue any unintentional deleted data. We suggest not to use a value below 5 days. A best practise is to use 7 days.

11.10 Throttling the TS7700

As explained previously, *Throttling* is the mechanism adopted to control and balance several tasks that run at the same time within the TS7700, prioritizing certain tasks over others.

Generally speaking, we distinguish three different types of throttling:

- ▶ Host Write Throttling: Host I/O will be throttled
- ▶ Copy Throttling: RUN copies pulled from other clusters will be throttled
- ▶ Deferred Copy Throttling: Deferred copies pulled from other clusters will be throttled

The reasons for using DCT and how to tune the DCT have already been explained in “Tuning possibilities for copies: Deferred Copy Throttling” on page 682.

11.10.1 Monitoring throttling with the MI

In the grid summary and the cluster summary, there is a throttling indicator on the head of the cluster. Hover over the indicator to see the throttling type and the throttling impact. This is only a snapshot view.

In addition, you can see the throttling on the performance graph, as explained in Figure 11-17 on page 662.

11.10.2 Monitoring throttling with VEHSTATS

In the H30TVCx report for each interval the throttling is reported for each of the different throttling types. To interpret these values, you need to review the throttling reasons in the following figures. Figure 11-38 shows the reasons for host write throttling.

Value	Description
x00	No throttling during the interval
x01	Premigration steady state (PMTHLVL)
x02	Low on cache free space
x04	Immediate copy throttling
x08	Excess cached content for copy
x10	Grid premigration steady state (throttling outbound copies because the target cluster is premigration throttling)
	All other values are reserved

Figure 11-38 Host Write Throttling reasons

Figure 11-39 shows the reasons for copy throttling.

Value	Description
x00	No throttling during the interval
x01	Premigration steady state (PMTHLVL)
x02	Low on cache free space
	All other values are reserved

Figure 11-39 Copy Throttling reasons

If a value of “x03” is shown in the H30TVC, that means reason X01 and X02 are applicable at the same time.

To understand if the throttling is a real performance issue, analyze in how many samples of the interval throttling happened, and the relative throttling impact value (%RLTV IMPAC VALUE). Even if throttling occurred, this might be only in a few samples during the interval, which means that the real impact to the write might or might not influence the overall production runtime.

11.10.3 Tuning to avoid the throttling

For the different types and reasons for throttling several tuning possibilities can be considered.

Some of them are parameter changes (PMTHLVL adjustments, ICOPYT), while other issues can only be solved by providing higher bandwidth or more resources (cache, drives).

While adjusting a parameter might be beneficial for a specific situation, it might have another impact to other behaviors in the grid. Therefore, we recommend to discuss such tuning measurements with your IBM representative up front.

However, the next section describe a common tuning action.

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 (R1.7). Since Release 2.0, the value was changed to 20 or more minutes. 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) on the host. When a copy is changed to immediate-deferred, the RUN task is completed, and the immediate copy becomes a high priority deferred copy. See “Immediate-copy set to immediate-deferred state” on page 715 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 is immediate, better overall performance has been seen when the host write throttle because of immediate copies is turned off. You are trading off host I/O for length of time that is 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 keywords SETTING, THROTTLE, ICOPYT:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

11.11 Adjusting parameters in the TS7700

As described in the previous section, the LI REQ command gives you various adjustment possibilities. They can be subparameters of the following elements:

- ▶ SETTING
- ▶ SETTING2
- ▶ Independent parameters

Some of them have changed the cluster behavior, while others have an influence on the grid allocation behavior (for example, SAA, DAA, LOWRANK).

Although these settings can be modified by using z/OS or the MI, we suggest that you first check the parameter settings in case of any issue, and determine if they have been modified. Remember, that most of the settings are persistent, so after a Service or a Power Off of the cluster, they are still active.

Important: Library commands change the behavior of the whole cluster. If a cluster is attached to multiple LPARs from the same client, or to a multi-tenant environment, the change that is run from one LPAR influences all attached LPARs.

If you have a shared TS7700, consider restricting the usage of the **Library** command.

11.12 Monitoring after service or outage

Use this window (Figure 11-40) to view the pending updates for the IBM 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.

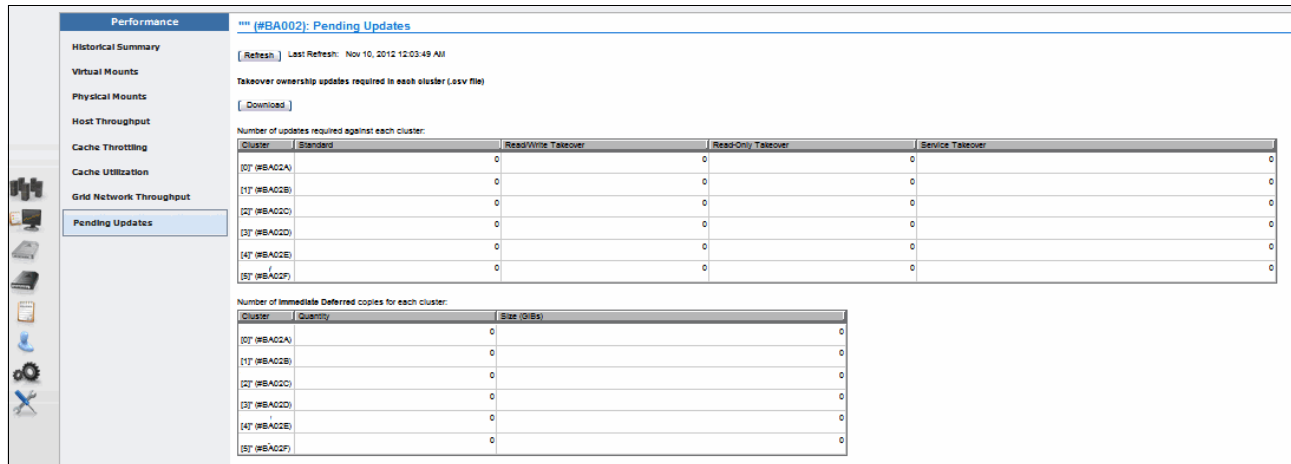


Figure 11-40 Grid Pending Updates window

This window provides a summary of the number of outstanding updates for each cluster in an IBM TS7700 Grid. You can also use this window to monitor the progress of pending immediate-deferred copies, which, like pending updates, normally result from changes that are made while a cluster is Offline, in service prep mode, or in service mode.

Remember: Pending immediate-deferred copies need to be avoided. They might be a result of overload or grid network problems.

With Release 3.2, the download section also includes the tape with a *hot token*. Hot tokens are volumes that have been changed during an unavailability of the cluster and now need a reconciliation. The reconciliation is run during the cluster online setting.

11.13 Performance evaluation tool: Plotting cache throughput from VEHSTATS

To change the “knobs” settings to alter 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.

Support is available at the Techdocs website:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101465>

Figure 11-41 shows the cache throughput plotted from VEHSTATS.




Figure 11-41 Cache throughput plotted from VEHSTATS

When evaluating performance, a graph that reveals a significant amount of information succinctly is the *cache throughput* for a cluster graph.

There are performance tools available on Techdocs that take 24 hours of 15-minute VEHSTATS data, seven days of 1-hour VEHSTATS data, or 90 days of daily summary data and create a set of charts for you.

The following material does not contain any information about the specifics of a Tape Attach model. However, the other information is still valuable, and has not changed, especially how to create the excel spreadsheet and the charts.

See the following Techdocs site for the performance tools, and the class replay for detailed information about how to use the performance tools:

► Tools:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4717>

► Class replay:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4872>

The 24-hour, 15-minute data spreadsheets include the cache throughput chart. The cache throughput chart has two major components: The uncompressed host I/O line and a stacked bar chart that shows the cache throughput.

The cache throughput chart includes the following components (all values are in MiB/s):

- Compressed host write: This is the MiB/s of the data that is written to cache. This bar is hunter green.
- Compressed host read: This is the MiB/s of the data read from cache. This bar is lime green.
- Data that is copied out 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 value that is 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 following descriptions are for plotting Cluster 0's cache throughput. Use the appropriate columns when plotting other clusters. These bars are cyan.

- ▶ Data that is 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 same rules for DCT apply for *data copied out*. These bars are light blue.
- ▶ Compressed 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 yellow.
- ▶ Compressed data recalled from tape to cache: This is the rate at which data is being read from tape into cache for a mount that requires a recall. This bar is dark blue.
- ▶ Compressed remote reads from this cluster: This is the rate that other clusters use this TVC as I/O cache for read. This bar is orange.
- ▶ Compressed remote writes to this cluster: This is the rate of synchronous copies. This bar is burnt orange.

This tool contains spreadsheets, data collection requirements, and a 90-day trending evaluation guide to assist you in the evaluation of the TS7700 performance. Spreadsheets for a 90-day, 1-week, and 24-hour evaluation are provided.

One 90-day evaluation spreadsheet can be used for one-cluster, two-cluster, three-cluster, or four-cluster grids and the other evaluation spreadsheet can be used for five-cluster and six-cluster grids. There is an accompanying data collection guide for each. The first worksheet in each spreadsheet has instructions for populating the data into the spreadsheet. A guide to help with the interpretation of the 90-day trends is also included.

There are separate one-week spreadsheets for two-cluster, three-cluster, four-cluster, five-cluster, and six-cluster grids. The spreadsheets use the one-hour interval data to produce charts for the one-week period. There is also a data collection guide.

There are separate 24-hour spreadsheets for two-cluster, three-cluster, four-cluster, five-cluster, and six-cluster grids. The spreadsheets use the 15-minute interval data to produce charts for the 24-hour period. There is also a data collection guide.

These spreadsheets are intended for experienced TS7700 users. A detailed knowledge of the TS7700 is expected, and familiarity with using spreadsheets.

11.14 Bulk Volume Information Retrieval

With the potential to support hundreds of thousands of logical volumes in a TS7700 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 subsystem that allow it to virtualize a tape volume also allow for a simple and effective method to transfer the information to a requesting application.

The TS7700 converts the format and storage conventions of a tape volume into a standard file that is managed by a file system within the subsystem. With BVIR, you are able to obtain information about all of the logical volumes that are managed by a TS7700.

The following data is available from a TS7700:

- ▶ Volume Status Information
- ▶ Cache Contents Information
- ▶ Physical Volume to Logical Volume Mapping Information
- ▶ Point-in-Time Statistics
- ▶ Historical Statistics
- ▶ Physical Media Pools
- ▶ Physical Volume Status
- ▶ Copy Audit

For more information, see the *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide* at the following URL:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101094>

11.14.1 Overview of the BVIR function

The TS7700 converts the format and storage conventions of a tape volume into a standard file that is managed by a file system within the subsystem. It 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 because standard IBM utilities, such as IEBGENER, provide the function that is 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 that is being requested. The records are in human-readable form, containing 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 server recognizes it as a request volume and “primes” the subsystem for the next step.

Remember: Some information that is 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 grid configuration with multiple clusters, use an MC 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 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 has completed appending to the data set, the host is notified that the mount is complete. The requested data can then be accessed like any other tape data set. In a job entry subsystem 2 (JES2) environment, the job control language (JCL) to complete the two steps can be combined into a single job. However, in a JES3 environment, they must be run in separate jobs because the volume will not be unmounted 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 do not affect its contents. Write accesses to the logical volume overwrite its contents. The logical volume can be returned to SCRATCH status and reused by any application.

Note: Due to the two-step approach, BVIR volumes cannot be written with LWORM specifications. You need to assign a Data Class without LWORM for BVIR volumes.

Figure 11-42 shows the process flow of BVIR.

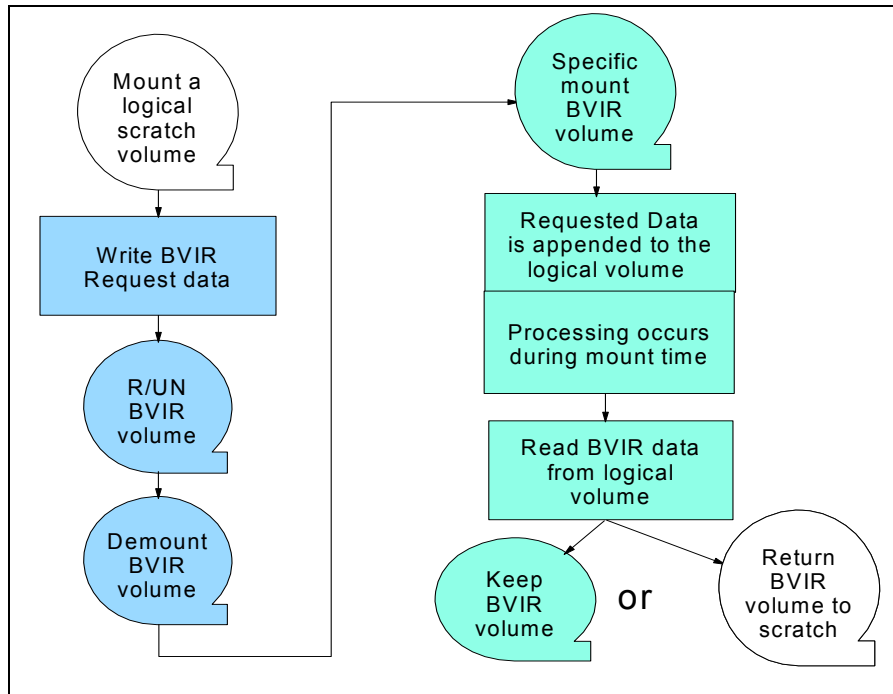


Figure 11-42 BVIR process flow

The building of the response information requires a small amount of resources from the TS7700. Do not use the BVIR function to “poll” for a specific set of information and only issue one request at a time. Certain requests, for example, the volume map, might take several minutes to complete.

To prevent “locking” out another request during that time, the TS7700 is designed to handle two concurrent requests. If more than two concurrent requests are sent, they are processed as previous requests are completed.

Although the requested data is always in a human-readable format, depending on the request, the data that is returned from the TS7700 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 11-4.

Example 11-4 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

Clarification: When records are listed in this chapter, there is an initial record showing “1234567890123...” This record does not exist, but it is provided to improve readability.

Record 0 is identical for all requests, and it is not part of the output; it is for support for records 1 - 5 only. Records 6 and higher 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 both the hardware serial number and the distributed library ID of the TS7700.
- ▶ Record 5 contains all blanks.

Records 6 - *N* and higher contain the requested data. The information is described in general terms. Detailed information about these records is in the *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide* at the following URL:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101094>

11.14.2 Prerequisites

Any logical volume that is defined to a TS7700 can be used as the request/response volume. Logical volumes in a TS7700 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 that is 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 can be used. Likewise, when the response volume's data is no longer needed, the logical volume must 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).

11.14.3 Request data format

Several types of data can be requested. The type of data that is 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” (fixed block) and a logical record size of 80 bytes in uncompressed data format (TRTCH=NOCOMP). 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 complete the record.

Important:

- ▶ The request fields must be as shown. Not beginning with the first character position of the record, or by using extra blanks between words, results in a failed request.
- ▶ The file must be written in uncompressed format to have it correctly interpreted by the TS7700.

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 that is used to read them is the Undefined (U) format. See Appendix E, “Sample job control language” on page 871 for more information.

In a multi-site TS7700 grid configuration, the request volume must be created on the cluster for which the data is being requested. The MC 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 is listed in the following sections.

Record 1

Record 1 must contain the command exactly as shown in Example 11-5.

Example 11-5 BVIR request record 1

```
1234567890123456789012345678
VTS BULK VOLUME DATA REQUEST
```

The format for the request's data set records is shown in Table 11-5.

Table 11-5 BVIR request record 1

Record 1: Request identifier		
Bytes	Name	Contents
1 - 28	Request identifier	VTS 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 11-6.

Table 11-6 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-aligned, 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 ABC000 - ABC999 have been defined to the cluster, a request of VOLUME STATUS ABC1_0 returns 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 that is being requested. Optionally, -yyy can also be specified and indicates that historical statistics from xxx through yyy are being requested. Valid days are 001 - 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 that are requested, that will be indicated in the response record.

This can occur if a request is made for a day before 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 TS7700's 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 that are 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 that are assigned to the specified pool, that is indicated in the response record. Data can be requested for pools 0 - 32.

For point-in-time and historical statistics requests, any additional characters that are 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 from which the data is to be obtained.

Use a specific MC that has a copy policy that is defined to indicate that only the wanted 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 an MC 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.

11.14.4 Response data format

When the request data set has been written to the volume and then closed and unmounted, when mounted again, the TS7700 validates the contents of the request volume and appends the requested data records to the data set.

Human-readable appended records can vary in length, depending on the reports that are requested and can be 80 - 640 bytes. Binary data appended records can be variable in length of up to 24,000 bytes. The data set is now a response data set. The appropriate block counts in the end of file (EOF) records are updated to reflect the total number of records written to the volume.

These records contain the specific response records based on the request. If the request cannot be understood or was invalid, that is indicated. The record length is fixed; the record length of each response data is listed in Table 11-7.

Table 11-7 Record length of response data

BVIR request	Record length in bytes
VOLUME STATUS <i>vvvvvv</i>	64
CACHE CONTENTS	80
VOLUME MAP	80
POINT IN TIME STATISTICS	24000
HISTORICAL STATISTICS FOR <i>xxx-yyy</i>	24000
PHYSICAL MEDIA POOLS	80
PHYSICAL VOLUME STATUS VOLUME <i>zzzzzz</i>	440
PHYSICAL VOLUME STATUS POOL <i>xx</i>	440
COPY AUDIT COPYMODE INCLUDE/EXCLUDE <i>libids</i>	80

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 are not valid, one or more error description records are appended to the data set or the data set is unmodified before 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 complete the record. All binary records are variable in length and are not padded.

Tips:

- ▶ In the response records, the dates and times that are presented are all based on the internal clock of the TS7700 handling the request. The internal clock of a TS7700 is not synchronized to the host, but it is synchronized with all other TS7700s.
- ▶ The host and the TS7740 can be synchronized to a Network Time Protocol (NTP) server, but they use a different NTP server with a different timing protocol. Slight time differences are still possible when NTP is used.

The response data set contains both request records that are described in 11.14.3, “Request data format” on page 702, and the response data set contains three explanatory records (Records 3 - 5) and starting with Record 6, the actual response to the data request.

The detailed description of the record formats of the response record is in the following white papers:

- ▶ *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101094>
- ▶ *IBM Virtualization Engine TS7700 Series Statistical Data Format white paper:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829>

The response data set has this general format:

- ▶ 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.
- ▶ Record 4
This record contains both the five-character hardware serial number of the TS7700, and the five-character distributed library sequence number of the cluster that generated the response.
- ▶ Record 5
This record contains all blank characters.
- ▶ Record 6 - *N* and Record 7
These records contain the specific response records based on the request. If the request cannot be understood or was invalid, that is indicated.

11.14.5 Interpreting the BVIR response data

This section explains how to interpret each BVIR Response Data Set for the specific request information, such as the following information:

- ▶ Volume Status Information
- ▶ Cache Contents Information
- ▶ Physical Volume to Logical Volume Mapping Information
- ▶ Point in Time Statistics
- ▶ Historical Statistics
- ▶ Physical Media pools

- ▶ Physical Volume Status
- ▶ Copy Audit

Clarification: 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 TS7700 cluster that contains information that is related to the management of the logical volumes on the cluster and copy and resynchronization processes when the TS7700s 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.

The volume status information that is returned represents the status of the volume on the cluster the requested volume was written to. In a TS7700 grid configuration, separate requests must be sent to each cluster to obtain the volume status information for the individual clusters. Using the volume serial number mask that is 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 that are defined as described in the white paper. Fields are presented in the order that is defined in the table and are comma-separated. The overall length of each record is 640 bytes with blank padding after the last field, as needed. The first few fields of the record that is returned for VOLSER ABC123 are shown in Example 11-6.

Example 11-6 BVIR volume status information

```
1234567890123456789012345678901234567890123456789012345678901234567890123
ABC123,0,2009-04-22-11.56.45.871263,0,0,32,0,N,2548,N,8719,N...
```

Important information is derived from the records:

- ▶ Data Inconsistent

This field indicates whether the cluster has a valid version of the data. If it indicates that the data on the logical volume is not valid, the same volume on another TS7700 in the grid has been modified and it has not yet been copied.

Suppose that you use the deferred Copy Consistency Point (which is typically when there is significant distance between the TS7700s in the grid configuration). In this situation, some number of volumes will not be consistent between the TS7700s at any point in time.

If a situation occurs that renders the site inoperable where the source data is, by sending the Volume Status request to an operable TS7700, this field can be used to identify the volumes that were not copied before the situation so that appropriate recovery steps can be run for them.

- ▶ MES Volume

This field indicates that the logical volume was created in the TS7700 Cluster or even created within a VTS before being merged into a grid configuration. Volumes that existed in a TS7700 cluster before being included in a grid configuration are not automatically copied to the other TS7700 clusters in the configuration until they have been accessed and closed.

This field can be used to determine which volumes in each TS7700 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 can 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 is provided in 11.18.1, “VEHSTATS tool overview” on page 725. You can also access the TAPETOOL FTP site at the following URL:

<ftp://ftp.software.ibm.com/storage/tapetool/>

► Copy Required for Cluster *n*

This field indicates that a copy to another TS7700 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 have 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, 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. Two types of ownership takeover are available:

- 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 that are 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 that are 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 that the volume is assigned to within the TS7700 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 Cluster, you can specify that after a certain time after being returned to scratch, the contents of a volume can be deleted. This field indicates whether the data that is associated with the volume has been deleted on the cluster.

► **Removal State**

As part of normal processing in a TS7700 Grid configuration where a mixture of both TS7740 and TS7700D clusters exists, a data removal or migration process occurs where data is removed from TS7700D clusters to prevent TS7700D clusters from overrunning their TVC. This field, and the removal time stamp, can be used to determine whether the data that is associated with the volume has been removed.

► **Hot**

This field represents the cluster's view of which clusters have obsolete 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.

Cache content information

This report provides a list of all volumes that are currently kept in cache. 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 sent to each cluster to obtain the cache contents.

The response records are written in 80-byte fixed block (FB) format.

Remember:

- 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.
- The contents of the cache typically are all private volumes. However, some might have been returned to SCRATCH status soon after being written. The TS7700 does not filter the cache contents based on the private or SCRATCH status of a volume.

Physical volume to logical volume mapping information

The TS7700 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 results when a logical volume is being moved from one physical volume to another. For a while, the volume is shown on more than one physical volume. This can result in a few logical volumes that are reported as being on physical volumes that they were on in the past, but are not presently on.

Even with inconsistencies, the mapping data is useful if you want to design jobs that recall data efficiently from physical volumes. If the logical volumes that are reported on a physical volume are recalled together, the efficiency of the recalls is increased. If a logical volume with an inconsistent mapping relationship is recalled, it recalls correctly, but an extra amount of a separate physical volume might be required.

The physical volume to logical volume mapping that is associated with the physical volumes managed by the specific cluster to which the request volume is written is returned in the response records. In a TS7700 grid configuration, separate requests must be sent to each cluster to obtain the mapping for all physical volumes.

The response records are written in 80-byte FB format.

This report is only available for the TS7740 and TS7700T.

Tip: 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.

Physical media pools

The TS7740/TS7700T supports separating the physical volumes that it manages into pools. The supported pools include a pool that contains scratch (empty) volumes that are common, and up to 32 pools that can 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 0 (common scratch pool), because it contains only empty volumes, 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 being filled, and have been marked as full, is returned. The count for empty includes physical volumes that have been 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) is 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 that are marked as full for a pool.

Information is returned for the common pool and all other pools that are defined and have physical volumes that are associated with them.

The physical media pool information that is managed by the specific cluster to which the request volume is written is returned in the response records. In a TS7700 grid configuration, separate requests must be sent to each cluster to obtain the physical media pool information for all clusters.

The response records are written in 80-byte FB format. Counts are provided for each media type associated with the pool (up to a maximum of eight).

Physical volume status

A database is maintained on each TS7740/TS7700T cluster that contains information that is 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 sent 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 that is specified in the request that exists in the cluster. A response record consists of the database fields that are defined as shown in Example 11-7 for Volume A03599. The overall length of each record is 400 bytes with blank padding after the last field, as needed.

Example 11-7 Sample response record for VOLSER A03599

A03599,2,FULL,READ-WRITE,2007-05-05-06.40.08.030061,2007-05-04-13.45.15.918473,...

Tip: The generation of the response might take several minutes to complete depending on the number of volumes that is requested and how busy the TS7700 cluster is at the time of the request.

Copy audit

A database is maintained on each TS7700 cluster that contains status information about the logical volumes that are defined to the grid. Two key pieces of information are whether the cluster contains a valid copy of a logical volume and whether the copy policy for the volume indicates that it must have a valid copy.

This request runs an audit of the databases on a set of specified TS7700 distributed libraries to determine whether any volumes do not have a valid copy on at least one of them.

If no further parameter is specified, the Audit checks whether a logical volume has a copy on the specified cluster or not. There is no validation if that cluster should have a copy or not. To take the copy modes into account, you need to specify a second parameter **COPYMODE**.

Using COPYMODE

If the COPYMODE option is specified, whether the volume is supposed to have a copy on the distributed library is considered 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 “S”, “R”, “D”, or “T”, that cluster is considered during the audit.

If the copy policy for a volume on a specific cluster is “N”, 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 clusters specified as part of the request.

The specified clusters might not have a copy for several reasons:

- ▶ The copy policy that is associated with the volume did not specify that any of the clusters that are 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 that are used in a disaster recovery test need to be only on the disaster recovery TS7700 and not on the production TS7700s. If the request specified only the production TS7700 tape drives, all of the volumes that are used in the test are returned in the list.

- ▶ The copies have not yet been made from a source TS7700 to one or more of the specified clusters. This can be because the source TS7700 or the links to it are unavailable, or because a copy policy of Deferred was specified and a copy has not been completed when the audit was run.
- ▶ Each of the specified clusters contained a valid copy at one time, but has since removed it as part of the TS7700 hybrid automated removal policy function. Automatic removal can take place on TS7700D or TS7700T clusters in all configuration scenarios (hybrid or homogeneous). In a TS7700T, only data in CP0 is subject for autoremoval.

The Copy Audit is intended to be used for the following situations:

- ▶ A TS7700 is to be removed from a grid configuration. Before its removal, you want to ensure that the TS7700s that are to remain in the grid configuration have a copy of all the important volumes that were created on the TS7700 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 TS7700s in a grid configuration is no longer available and you want to determine which, if any, volumes on the remaining TS7700s do not have a valid copy.

In the Copy Audit request, you need to specify which TS7700 clusters are to be audited. The clusters are specified by using their associated distributed library ID (this is the unique five-character library sequence number that is defined when the TS7700 Cluster was installed). If more than one distributed library ID is specified, they are separated by a comma. The following rules determine which TS7700 clusters are to be included in the audit:

- ▶ When the INCLUDE parameter is specified, all specified distributed library IDs are included in the audit. All clusters that are associated with these IDs must be available or the audit fails.
- ▶ When the EXCLUDE parameter is specified, all specified distributed library IDs are excluded from the audit. All other clusters in the grid configuration must be available or the audit fails.
- ▶ 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 fails and the response indicates the IDs that are considered invalid.
- ▶ Distributed library IDs must be specified or the Copy Audit fails.
- ▶ Here are 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 that is 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 that is associated with distributed library ID DA01A.
 - COPY AUDIT INCLUDE DA01B,DA01C: Audits the copy status of volumes on the clusters that are 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 that is 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 11-8. The response records are written in 80-byte FB format.

Example 11-8 BVIR message when Copy Audit is requested

```
123456789012345678901234567890123456789012
L00001,R,D,D,N,N,N,N,N,R,N,R,N,N,N,N,N,R
```

Tips:

- ▶ The output for Copy Audit includes Copy Consistency Points for up to eight TS7700 clusters. This is to provide for future expansion of the number of clusters that are supported in a TS7700 Grid to the designed 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 tape drives are at the time of the request.

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 indicates in Record 6 that the request is unknown, as shown in Example 11-9.

Example 11-9 BVIR message when an unknown or invalid request is submitted

```
12345678901234567890
UNKNOWN REQUEST TYPE
```

If the request file contains the correct number of records, the first record is correct, and the second is recognized but includes a variable that is not within the range that is supported for the request, the response file indicates in record 6 that the request is invalid, as shown in Example 11-10.

Example 11-10 BVIR message when an invalid variable is specified

```
12345678901234567890123456
INVALID VARIABLE SPECIFIED
```

11.15 Alerts and exception and message handling

The following section provides an overview about user-defined alerts in a TS7700, possible exceptions in a TS7700 environment, and upcoming messages. These messages can be used as input for an automation system and, depending on the user requirements, they should be sent with the automation to an email or alarm system.

11.15.1 Alerting of specific events

The TS7700 offers you a broad variety of threshold and timeout alerts. This section includes a brief introduction. For a detailed description, see the following white paper:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101091>

Most threshold alerts allow two thresholds per setting. The first one issues the “alarm” message that the threshold is now crossed. A second, lower limit that informs you when the condition is resolved, and that the amount of data has fallen beyond the threshold.

Before you adjust the values of the thresholds alerts, evaluate appropriate values for your installation. Use the performance evaluation tool in advance. Also, review the values after implementation and review them periodically (for example, twice a year) or after installation changes.

Values for the alerting that are too low lead to unnecessary messages and operator actions, Values that are too high might not give you enough time in a critical situation to react.

General statement for alerts

All of the following alerts send messages to the host system log if the pending inbound copy backlog indicates that the amount of uncopied data has exceeded the low or the critical warning limit specified (in GB). All parameters can be set independently, but the values have some dependencies that must be acknowledged. For all alerts, a message is created if the values are exceeded.

The default value for each parameter is 0, which indicates that no warning limit is set and messages are not generated.

Inbound backlog value (PCPYLOW and PCPYCRIT)

These values specify the threshold in GB of volumes waiting in the incoming copy queue. The PCPYLOW value defines the first level of warning. The PCPYCRIT represents the critical state of inbound copy backlog in GB.

These alerts are applicable for all TS7700s, and can be set for each distributed library independently.

Change the value with the following commands:

```
LI REQ,distributed library,SETTING,ALERT,PCPYLOW,value
LI REQ,distributed library,SETTING,ALERT,PCPYCRIT,value
```

Uncopied data in cache (COPYLOW and COPYRIT)

These values specify the threshold in GB of volumes waiting to be copied to another cluster in the TS7700 grid configuration. The COPYLOW value defines the first level of warning. The COPYCRIT represents a critical state of inbound copy backlog (in GB). These alerts are applicable for all TS7700s, and can be set for each distributed library independently.

Change the value with the following commands:

```
LI REQ,distributed library,SETTING,ALERT,COPYLOW,value
LI REQ,distributed library,SETTING,ALERT,COPYCRIT,value
```

Data in cache (RESLOW and RESHIGH)

These values specify the amount of data in cache. The RESLOW value defines the first level of warning. The RESHIGH represents a critical state of data in cache.

The following values are measured:

- ▶ TS7700D: All data in cache
- ▶ TS7740: Resident data in the cache that has not yet been copied to the back end
- ▶ TS7700T: Resident data in CP0

Change the value with the following commands:

```
LI REQ,distributed library,SETTING,ALERT,RESLOW,value
LI REQ,distributed library,SETTING,ALERT,RESHIGH,value
```

Important: The monitored values are different for each TS7700 model.

For the TS7740, if RESDLOW/RESHIGH was exceeded, check the number of available tape drives and empty cartridges in the TS7740, especially if you have multiple physical pools. Also, review the number of premigration tasks and reclaim schedules.

For a TS7700D and TS7700T CP0, if RESDLOW/RESHIGH was exceeded, check the amount of Pinned data, and the amount of data that is subject to auto removal. If the data is filled up with Pinned data and the data will not be expired by the host in the near future, you might run out of cache.

Data in cache (RSDTLOW and RSTDHIGH)

These alerts are similar to those described previously, but now for the TS7700T tape partitions. The same actions apply as for the TS7740 when the thresholds are exceeded. The RESTDLOW value defines the first level of warning. The RESTDHIGH represents a critical state of data in cache of non-premigrated data from all tape partitions.

Change the value with the following commands:

```
LI REQ,distributed library,SETTING,ALERT,RSDTLOW,value
LI REQ,distributed library,SETTING,ALERT,RSDHIGH,value
```

Physical drive availability (PDRVLOW and PDRVCRIT)

These values specify the number of available backend drives in a TS7700T or TS7740. The PDRVLOW value defines the first level of warning. The PDRVCRIT represents a critical state. They are applicable only for TS7700T and TS7740 and can be set for each distributed library independently.

Change the value with the following commands:

```
LI REQ,distributed library,SETTING,ALERT,PDRVLOW,value
LI REQ,distributed library,SETTING,ALERT,PDRVCRIT,value
```

Physical scratch availability (PSCRLOW and PSCRCRIT)

These values specify the number of available backend cartridges in a TS7700T or TS7740. The PSCRLOW value defines the first level of warning. The PSCRCRIT represents a critical state.

Change the value with the following commands:

```
LI REQ,distributed library,SETTING,ALERT,PSCRLOW,value
LI REQ,distributed library,SETTING,ALERT,PSCRCRIT,value
```

11.15.2 Handling Replication Exceptions

In certain conditions, the selected replication mode cannot be run. This is true for Synchronous mode copy and Immediate (RUN) copies.

To ensure, that the production is not impacted, a switch from the origin replication mode to a “Sync-Deferred” or “Immed-Deferred” is possible. For the synchronous mode copy, the user can define whether the synchronous mode copy will change to a “SYNC-Deferred” state, or if the job will fail.

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 before 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 enters the immediate-deferred state.

The volume remains in an immediate-deferred state until all of the requested RUN consistency points contain a valid copy. The immediate-deferred volume replicates 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 that are targeted to receive an immediate copy are not available. Independently of why a volume might enter the immediate-deferred state, 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.

The reasons why a volume moves to the immediate-deferred state are contained in the Error Recovery Action (ERA) 35 sense data. The codes are divided into unexpected and expected reasons. From a z/OS host view, the ERA is part of message IOS000I (Example 11-11).

Example 11-11 Message IOS000I

```
IOS000I 1029,F3,EQC,0F,0E00,,**,489746,HADRMMBK 745
.50408035000000206011(B31000011C005800)2182(50000FFF)CE1F0720EED40900
IMMEDIATE MODE COPY - EXPECTED FAILURE - REASON CODE = 82
COPY COUNT - 1 OF 2 COMPLETED SUCCESSFULLY
```

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 verified as successful before surfacing Sense Status Information (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 that is targeted for a RUN copy is not available (unexpected outage).
 - 0x03 - 40 minutes have elapsed and one or more copies have timed out.
 - 0x04 - Is reverted to immediate-deferred because of health/state of RUN target clusters.
 - 0x05 - Reason is unknown.
 - Expected - 0x80 to 0xFF: The reasons are based on the configuration or a result of planned outages:
 - 0x80 - One or more RUN target clusters are out of physical scratch cache.
 - 0x81 - One or more RUN target 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 that are explicitly disabled through the Library Request operation.
 - 0x84 - The volume cannot be reconciled and is “Hot” against peer clusters.

The additional data that is 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 before sending the CCW(RUN) command, it cannot act on the failed status. However, future jobs can be suspended or other custom operator actions can be taken by using the information that is provided within the sense data.

Handling of the Immed-Deferred related messages

For each volume that cannot be replicated in the immed-Deferred mode, an IOS000I can be reported in the system log. Although that allows a message automation on the host to detect possible problems, the thousands of messages that are generated might overflow the host log. This is especially true in maintenance situations. A Host Console Request (HCR) command enables you to handle all OS000I triggered by “IMMED-Deferred.” The following options are available:

LI REQ,distributed library,SETTING,COPY,IMMSNS,[ALL|NONE|UNEXP]

- ▶ ALL: All messages are presented to the host.
- ▶ NONE: All messages are not presented to the host, except if no valid copy source is available. This is the default.
- ▶ UNEXP: Only unexpected failures will be presented to the host. All messages during a maintenance situation or during “Gridlink Disable” will not be presented because they are the result of a customer action.

Synchronous mode copy set to synchronous deferred

The default behavior of Synchronous mode copy (SMC) is to fail a write operation if both clusters with the “S” copy policy are not available or become unavailable during write operations. You can enable the Synchronous Deferred on Write Failure (SDWF) option to permit update operations to continue to any valid consistency point in the grid. If there is a write failure, the failed “S” locations are set to a state of “synchronous-deferred.”

After the volume is closed, any synchronous-deferred locations are updated to an equivalent consistency point through asynchronous replication. If the SDWF option is not selected (default) and a write failure occurs at either of the “S” locations, host operations fail and you must view only content up to the last successful sync point as valid.

For example, imagine a three-cluster grid and a copy policy of Sync-Sync Deferred (SSD), Sync Copy to Cluster 0 and Cluster 1, and a deferred copy to Cluster 2. The host is connected to Cluster 0 and Cluster 1. With this option disabled, both Cluster 0 and Cluster 1 must be available for write operations. If either one becomes unavailable, write operations fail. With the option enabled, if either Cluster 0 or Cluster 1 becomes unavailable, write operations continue. The second “S” copy becomes a synchronous-deferred copy.

In the previous example, if the host is attached to Cluster 2 only and the option is enabled, the write operations continue even if both Cluster 0 and Cluster 1 become unavailable. The “S” copies become synchronous-deferred copies.

The synchronous-deferred volume replicates with a priority greater than immediate-deferred and standard-deferred copies.

For more information about Synchronous mode copy, see the *IBM Virtualization Engine TS7700 Series Best Practices - Synchronous Mode Copy* white paper on Techdocs.

Handling of composite status change due to replication issues

When a cluster detects a “SYNC-Deferred” or an “Immed-Deferred” condition, a degradation of the composite library is reported. Although these conditions might have an impact to your disaster recovery capability, the production jobs might not be impacted at all. This is especially true if the cluster is not available due to maintenance purposes.

Therefore, a Host Console Request Command is provided, which enables you to define whether these conditions report the composite library as degraded or not.

LI REQ,composite library,SETTING,ALERT,DEFDEG, [ENABLE|DISABLE]

Grid link exception handling

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 the grid links is monitored periodically, and if one link is performing worse than the other link by an IBM Service Support Representative (IBM SSR)-alterable value, a warning message is generated and sent to the host. The purpose of this warning is to alert you 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. Here is the warning message format:

CBR3750I, G0030, Library xx Primary, Alternate, Primary2, and Alternate2 grid links are degraded. The degraded grid link port is reported.

A second message with a variable text “EA480E” is reported in the syslog.

For example, a setting of 60% means that if one link is running at 100%, the remaining links are marked as degraded if they are running at less than 60% of the 100% link. The grid link performance is available with the Host Console Request function, and on the TS7700 MI. The monitoring of the grid link performance by using the Host Console Request function 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 and GRIDLINK keywords:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

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 5-minute intervals that link degradation was detected before reporting an attention message. The default value is 9.
- ▶ Generate Call Home iterations: The number of consecutive 5-minute intervals that link degradation was detected before generating a Call Home. The default value is 12.

The default values are set to 60% for the threshold, nine iterations before an attention message is generated, and 12 iterations before a Call Home is generated. Use the default values unless you are receiving intermittent warnings and support indicates that the values need to be changed. If you receive intermittent warnings, let the SSR change the threshold and iteration to the suggested values from support.

For example, suppose that clusters in a two-cluster grid are 2000 miles apart with a round-trip latency of approximately 45 ms. The normal variation that is seen is 20 - 40%. In this example, the threshold value is at 25% and the iterations are set to 12 and 15.

11.16 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. Before describing VEHSTATS and the reports that you can obtain by running the VEHSTATS jobs, a general overview of the IBM Tape Tools is provided to guide you through the installation of these tools in a z/OS environment.

11.16.1 Introduction to IBM Tape Tools

All the TS7700 monitoring and evaluating tools that are described are at the following FTP site:

<ftp://ftp.software.ibm.com/storage/tapetool/>

Example 11-12 lists the content of the `Readme.txt` file that provides basic information about the tape tools.

Example 11-12 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 that are distributed through 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 whether 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 compressed into IBMTOOLS.EXE.

IBMTOOLS.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, send an email 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.

Tools that might be interesting for you are presented in Table 11-8.

Table 11-8 Tape tools selection

Tool	Major use	Benefit	Inputs	Outputs
BADBLKSZ	Identify small VTS blocksizes	Improve VTS performance, make jobs run faster	Logrec MDR & CA1, TLMS, RMM, ZARA,CTLT	VOLSER, Jobname, and Dsname for VTS volumes with small blocksizes.
BVIRHSTx	Get historical stats from TS7700	Creates U, VB, SMF format	TS7700	Statistics file.
BVIRPOOL	Identify available scratch by pool	Reports all pools at the same time	BVIR file	Physical media by pool.
BVIRPRPT	Reclaim Copy Export volumes	Based on active GB, not %	BVIR file	Detailed report of data on volumes.
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 reports
BVPITRPT	Point in Time stats as write to operator (WTO)	Immediately available	TS7700	Point in Time stats as WTO.
COPYVTS	Copy lvols from old VTS	Recall lvols based on selected applications	BVIR data & CA1, TLMS, RMM, ZARA, CTLT	IEBGENER to recall lvols and copy to new VTS.
DIFFEXP	Identify multi-file volumes with different expiration dates	Prevent single file from not allowing volume to return to scratch	CA1, TLMS, RMM, ZARA, CTLT	List of files not matching file 1 expiration date.
FSRMATCH	Replace *.HMIGTAPE.DATASET in SMF 14 with actual recalled dsname	Allows TapeWise and other tools by using SMF 14/15 data to report actual recalled data set	FSR records plus SMF 14, 15, 21, 30, 40, and so on	Updated SMF 14s plus all other SMF records as they were.
GETVOLS	Get VOLSERs from list of dsns	Automate input to PRESTAGE	CA1, TLMS, RMM, ZARA, CTLT	VOLSERs for requested dsns.
IOSTATS	Report job elapsed times	Show runtime improvements	SMF 30 records	Job-step detailed reporting
MOUNTMON	Monitor mount pending and volume allocations	Determine accurate mount times and concurrent drive allocations	Samples tape UCBs	Detail, summary, distribution, hourly, TGROUP, and system reporting.
ORPHANS	Identify orphan data sets in Tape Management Catalog (TMC)	Cleanup tool	CA1, TLMS, RMM, ZARA, CTLT	Listing file showing all multiple occurrence generation data group (GDGs) that have not been created in the last <i>nn</i> days.
PRESTAGE	Recall lvols to VTS	Ordered and efficient	BVIR VOLUME MAP	Jobs that are submitted to recall lvols.

Tool	Major use	Benefit	Inputs	Outputs
SMFILTER	IFASMFDP exit or E15 exit	Filters SMF records to keep just tape activity. Generates "tape" records to simulate optical activity	SMF data	Records for tape activity plus optional TMM or optical activity.
TAPECOMP	Show current tape compression ratios	See how well data will compress in VTS	Logrec MDR or EREP history file	Shift and hourly reports showing current read and write compression ratios.
TAPEWISE	Identify tape usage improvement opportunities	Shows UNIT=AFF, early close, UNIT=(TAPE,2), multi-mount, DISP=MOD, recalls	SMF 14, 15, 21, 30, and 40	Detail, summary, distributions, hourly, TGROUP, and system reporting.
TCDBMCH	Identify tape configuration database (TCDB) versus Tape Catalog mismatches	List VOLSER mismatches	CA1, TLMS, RMM, ZARA, CTLT	ERRRPT with mismatched volumes.
TMCREUSE	Identify data sets with create date equal to last ref date	Get candidate list for VTS PG0	CA1, TLMS, RMM, ZARA, CTLTF	Filter list of potential PG0 candidates.
VEHGRXCL	Graphing package	Graphs TS7700 activity	VEHSTATS flat files	Many graphs of TS7700 activity.
VEHSCAN	Dump fields in historical statistics file	Individual field dump	BVIR stats file	DTLRPT for selected interval.
VEHSTATS	TS7700 historical performance reporting	Show activity on and performance of TS7700	BVIRHSTx file	Reports showing mounts, data transfer, and box usage.
VEPSTATS	TS7700 point-in-time statistics	Snapshot of last 15 seconds of activity plus current volume status	BVIRPIT data file	Reports showing current activity and status.
VESYNC	Synchronize TS7700 after new cluster added	Identify lvols that need copies	BVIR data and CA1, TLMS, RMM, ZARA, CTLT	List of all VOLSERS to recall by application.
VOLLIST	Show all active VOLSERS from TMC. Also, get volume counts by group, size, and media.	Used to get a picture of user data set naming conventions. See how many volumes are allocated to different applications.	CA1, TLMS, RMM, ZARA, CTLT	Dsname, VOLSER, create date, and volseq. Group name, counts by media type.

11.16.2 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 includes the VEHSTATS and VEPSTATS tools, which use the Bulk Volume Information Retrieval (BVIR) reports for comprehensive performance analysis.

Figure 11-43 shows several tools that are available from the FTP site.

Index of ftp://public.dhe.ibm.com/storage/tapetool/

Up to higher level directory

Name	Size	Last Modified	
.message	1 KB	6/15/2000	12:00:00 AM
A_License_Agreement_for_IBM_Tape_Tools.pdf	22 KB	1/27/2011	12:00:00 AM
TS7680_Statistics_Report_Install.doc	46 KB	1/19/2011	12:00:00 AM
TS7700.VEHSTATS.Decoder.V20.pdf	951 KB	3/20/2014	9:40:00 AM
TapeWise.PDF	325 KB	8/17/2005	12:00:00 AM
VEHGRXCL.exe	1807 KB	9/13/2013	12:00:00 AM
VEHGRXCL.txt	4 KB	8/5/2009	12:00:00 AM
VTSGRXCL.EXE	805 KB	2/28/2007	12:00:00 AM
VTSGRXCL.TXT	2 KB	12/4/2003	12:00:00 AM
badblksz.txt	2 KB	8/5/2004	12:00:00 AM
batcntl.xmi	996 KB	2/1/2013	12:00:00 AM
eximcalc.123	11 KB	3/2/2007	12:00:00 AM
export.txt	1 KB	8/5/2004	12:00:00 AM
findlrg.txt	2 KB	5/7/2002	12:00:00 AM
fsrtmm.txt	1 KB	5/7/2002	12:00:00 AM
ftpcust.txt	2 KB	3/21/2002	12:00:00 AM
grpdsn.txt	2 KB	8/5/2004	12:00:00 AM
ibmcntl.xmi	2686 KB	1/9/2014	2:40:00 PM
ibmjcl.xmi	1330 KB	3/18/2014	2:40:00 PM
ibmload.xmi	5432 KB	3/12/2014	10:40:00 PM
ibmpat.xmi	280 KB	4/16/2011	12:00:00 AM
ibmtools.txt	6 KB	5/7/2013	12:00:00 AM
ifasmfdp.txt	3 KB	8/5/2004	12:00:00 AM
libmangr.pdf	151 KB	1/21/2005	12:00:00 AM
mountmon.pdf	277 KB	8/15/2005	12:00:00 AM
mountmon.txt	2 KB	5/7/2002	12:00:00 AM

Figure 11-43 Tape tools catalog

The index is at the following web address:

<http://public.dhe.ibm.com/storage/tapetool/>

IBM employees can access IBM Tape Tools on the following website:

<http://w3.ibm.com/sales/support/ShowDoc.wss?docid=SGDK749715N06957E45&node=brands,B5000|brands,B8S00|clientset,IA>

IBM Business Partners can access IBM Tape Tools on the following website:

<http://www.ibm.com/partnerworld/wps/servlet/ContentHandler/SGDK749715N06957E45>

For most tools, a text file is available. In addition, each job to run a tool contains a detailed description of the function of the tool and parameters that need to be specified.

Important: 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 that are contained in the `ibmtools.exe` file.

The `ibmtools.exe` file is a self-extracting compressed file that is expanded into four separate files:

IBMJCL.XMI	Contains the execution JCL for current tape analysis tools.
IBMCNTL.XMI	Contains parameters that are needed for job execution, but that do not need to be modified by the user.
IBMLOAD.XMI	Contains the load library for executable load modules.
IBMPAT.XMI	Contains the data pattern library, which is only needed if you run the QSAMDRVR 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 that you complete the following steps:

1. Copy, edit, and submit `userid.IBMT00LS.JCL($$CPYLIB)` to create a new JCL library that has a unique second node (`&SITE` symbolic). This step creates a private JCL library for you from which you can submit jobs while leaving the original as is. `CNTL` and `LOAD` can then be shared by multiple users who are running jobs from the same system.
2. Edit and submit `userid.SITENAME.IBMT00LS.JCL($$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 that 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 is issued that is good for the next 12 months. When you download the current tools package any time during the year, you have at least nine months remaining on the `EXPIRE` value. New values are issued in the 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 current `IBMT00LS.JCL(EXPIRE)` member.

11.16.3 IBM Tape Tools for TS7700 monitoring

Several tape tools that can be used to help you better understand your tape processing regarding TS7700 operation and migration are described.

IOSTATS

IOSTATS tool is part of the `ibmtools.exe` file, which is available at the following URL:

<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 performance before and after configuration changes.

IOSTATS can be run for a subset of job names for a certain period 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 running in a multi-cluster grid configuration, IOSTATS can be used for the following purposes:

- ▶ To evaluate the effect of the multi-cluster grid environment and to 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 and to compare job execution times before and after upgrading components of the TS7700. For example, you might want to verify the performance impact of a larger 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:

- ▶ Tape activity analysis, including reads and Disp=mod 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

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 authorized program facility (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 tapes are allocated
- ▶ How many drives are being used at any time
- ▶ What is the most accurate report of concurrent drive usage
- ▶ Which jobs are allocating too many drives

11.17 Using Volume Mount Analyzer

The Volume Mount Analyzer (VMA) is part of the DFSMS suite and is integrated in the DFSMS code. Based on the SMF 14, 15, 21, and 30, you can produce several reports, similar to some of the functions of Tapewise:

- ▶ Mounts and MBs processed by hour
- ▶ Concurrent logical drives used
- ▶ Logical volume distribution (single-file or multifile, and single-volume or multivolume)

The VMA has a filter option, which can be used to create separated reports for the following specific items:

- ▶ SYSID
- ▶ Drive Types
- ▶ Device Ranges
- ▶ High-level qualifier
- ▶ Accounting codes

Also, some so-called “Top” reports can be produced to get an overview of the tape usage. For more information, see *z/OS DFSMS Using the Volume Mount Analyzer*, SC23-6895.

11.18 Using VEHSTATS and VEHGRXCL for monitoring and reporting

This section shows how to work with the binary reports for point-in-time and historical statistics after you use the BVIR functions that are described in Appendix E, “Sample job control language” on page 871. 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 statistics 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 11.15, “Alerts and exception and message handling” on page 713.

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 11.16.2, “Tools download and installation” on page 721 for specifics about where to obtain these tools. Details about using BVIR are 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 at the following address:

<http://www.ibm.com/support/techdocs/atmastr.nsf/Web/TechDocs>

With the record layout of the binary BVIR response data, you can decode the binary file or you can use the record layout to program your own tool for creating statistical reports.

Some of the VEHSTATS reports were already discussed in previous sections. This section contains only information that is not yet covered.

11.18.1 VEHSTATS tool overview

The TS7700 activity is recorded in the subsystem. There are two types of statistics:

- ▶ 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 Appendix E, “Sample job control language” on page 871).

Because both types of statistical data are delivered in binary format from the BVIR functions, you must convert the content into a readable format.

You can do this task manually by using the information that is provided in the following documents:

- ▶ *IBM Virtualization Engine TS7700 Series Statistical Data Format white paper Version 2.0:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829>
- ▶ *IBM Virtualization Engine TS7700 Series VEHSTATS Decoder:*
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105477>

Alternatively, you can use an existing automation 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. Guidance for interpretation of the reports is available in 11.18.3, “VEHSTATS reports” on page 728.

You can use VEHSTATS to monitor TS7700 virtual and physical tape drives, and TVC activity to do trend analysis reports, which are based on BVIR binary response data. The tool summarizes TS7700 activity on a specified time basis, up to 90 days in time sample intervals of 15 minutes or 1 hour, depending on the data reported.

Figure 11-43 on page 722 highlights three files that might be helpful in reading and interpreting VEHSTATS reports:

- ▶ The TS7700.VEHSTATS.Decoder file contains a description of the fields that are listed in the various VEHSTATS reports.
- ▶ The VEHGRXCL.txt file contains the description for the graphical package that is 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 that are created with VEHSTATS. Follow the instructions in the VEHSTATS_Model.xls file to create these graphs.

11.18.2 Running the VEHSTATS jobs

You have several output options for VEHSTATS, and you must submit separate jobs, depending on your requirements. The IBMT00LS.JCL member VEHSTATS (Example 11-13) provides guidance about which job to choose.

Example 11-13 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. With these jobs, you can control where the historical statistics are accumulated for long-term retention.

The TS7700 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

BVIR volumes cannot be written with LWORM attributes. Ensure that the BVIR logical volumes have a Data Class without LWORM specification.

The VEHSTATS reporting program accepts any or all of the various formats of BVIR input. Define which input is to be used through a data definition (**DD**) statement in the VEHSTATS job. The three input DD statements are optional, but at least one of the statements that are shown in Example 11-14 must be specified.

Example 11-14 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 that are kept by the user. The **SMFNUM** parameter defines which record number is processed when you specify the **STATSMF** statement.

The fields that are shown in the various reports depend on which **ORDER** member in IBMT00LS.JCL is being used. Use the following steps to ensure that the reports and the flat file contain the complete information that you want 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 as indicated by an asterisk in the first column. Example 11-15 shows sample active and inactive definitions in the **ORDERV12** member of IBMT00L.JCL. The sample statements define whether you want the amount of data in cache to be displayed in MB or in GB.

Example 11-15 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 by using the graphics package from the IBM Tape Tools FTP site, 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.

If your cluster numbers do not start with 0, or the numbering has gaps, you need to use another parameter adjustment in the JCL. With “Define Distributed Library” = DEFDL, you can specify the order of your clusters, as shown in the following example:

```
DEFDL= H3833 1;  
DEFDL= H6395 2;
```

11.18.3 VEHSTATS reports

VEHSTATS can be used to monitor TS7700 drive and TVC activity, and to run trend analysis to see where the performance bottlenecks are. Also, comparative analysis can be used to determine whether an upgrade, such as adding more physical tape drives, might improve the overall performance of the TS7740. VEHSTATS is not a projection tool, but it provides the basis for an overall health check of the TS7700.

VEHSTATS gives you a huge amount of information. The following list shows the most important reports available for the TS7700, and the results and analysis that can help you understand the reports better:

- ▶ H20VIRT: Virtual Device Historical Records
- ▶ H21ADP0x: 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:
 - For a TS7700D and TS7740, this report represents the cache.
 - For a TS7700T, multiple TVCs (TVC2 and TVC3) are presented. TVC1 contains the data from CP0, TVC2 contains the data of CP1, and so on.
- ▶ H31IMEX: hNode Export/Import Historical Activity
- ▶ H32TDU12: hNode Library Historical Drive Activity
- ▶ H32CSP: hNode Library Hist Scratch Pool Activity
- ▶ H32GUPXX: General Use Pools 01/02 through General Use Pools 31/32
- ▶ H33GRID: hNode Historical Peer-to-Peer (PTP) Activity
- ▶ AVGRDST: Hrs Interval Average Recall Mount Pending Distribution
- ▶ DAYMRY: Daily Summary
- ▶ MONMRY: Monthly Summary
- ▶ COMPARE: Interval Cluster Comparison
- ▶ HOURFLAT: 15-minute interval or 1-hour interval
- ▶ DAYHSMRY: Daily flat file

Tip: Be sure that you have a copy of TS7700 VEHSTATS Decoder and the TS7700 Statistical white paper available when you familiarize yourself with the VEHSTATS reports.

Virtual Device Activity

Example 11-16 on page 729 shows the report for Virtual Device Activity. This report gives you an overview, per 15-minute interval, of the relevant time frame and shows the following information:

- ▶ The minimum, average, or maximum (MIN, AVG, or MAX) mounted virtual drives
- ▶ The amount of channel blocks written based on blocksize

Clarification: The report is provided per cluster in the grid. The report title includes the cluster number in the DIST_LIB_ID field.

Example 11-16 VEHSTATS report for Virtual Drive Activity - first half

```

1(C) IBM   REPORT=H2OVIRT (15102)           VNODE VIRTUAL DEVICE
GRID#=00186  DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=CL2H6395
03JUN15WE -VIRTUAL_DRIVES-                _THROUGHPUT_ PCT_OF _
RECORD      --MOUNTED--      MAX ATTMP  Delay_/15Sec 15Sec _
      TIME  INST MIN  AVG  MAX  THRPUT THRPUT   MAX   AVG INTVLS
              R2.2   CALC <-----R3.0.0063-----> <
06:00:00    256   0   6  12    100   507   .803   .200    69
07:00:00    256   5   8  12    100   502   .801   .257    85
08:00:00    256   3   8  12    100   497   .799   .230    81
09:00:00    256   3   8  12    100   515   .806   .256    86
10:00:00    256   0   1   9    100   432   .769   .190    48
11:00:00    256   0   0   0    100  less   .000   .000     0
12:00:00    256   0   0   0    100  less   .000   .000     0
13:00:00    256   0  10  25    100   684   .854   .413    67

```

Example 11-17 shows the second half of the report.

Example 11-17 VEHSTATS report for Virtual Drive Activity - second half

```

HISTORICAL RECORDS          RUN ON 13JUL2015 § 13:40:49    PAGE 24
VE_CODE_LEVEL=008.033.000.0025          UTC NOT CHG
  CLUSTER VS FICON CHANNEL
  AHEAD  AHEAD  BEHIND  BEHIND  -----CHANNEL_BLOCKS_WRITTEN_F
  MAX    AVG    MAX    AVG    <=2048    <=4096    <=8192
  -----R3.1.0073+----->
7585    3551    1064    242    17540     0     25650
7626    4600    1239    326    22632     0     32400
7638    4453     958    325    21943     0     31350
7491    4553     974    353    22913     0     32700
7664    2212    1564    387    14092     0     19500
  0      0      0      0      0         0     0
  0      0      0      0      0         0     0
8521    4534     713    108    19101     0     32063

```

With R3.1, new information was included. CLUSTER VS FICON CHANNEL shows you whether the TS7700 can take more workload (called *ahead*), or if the FICON tries to deliver more data than the TS7700 can accept at this specific point in time (called *behind*). You will normally see that numbers in both columns are shown.

Use the ratio of both numbers to understand the performance of your TS7700. In our example, the TS7700 is behind only 8% of the time in an interval. The TS7700 can handle more workload than is delivered from the host.

In addition, the report shows the CHANNEL BLOCKS WRITTEN FOR BLOCKSIZES. In general, the largest number of blocks are written at 65,546 or higher blocksize, but this is not a fixed rule. For example, DFSMSHsm writes a 16,384 blocksize, and DB2 writes blocksizes up to 256,000. The report contains more differences for blocksizes, but are not shown in the example. From an I/O point of view, analysis of the effect 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), the vnode adapter Historical Activity Combined report is sufficient to provide 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 11-18 shows a sample report for Adapter 3 of Cluster 0.

Example 11-18 Adapter 3 sample report

```

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 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 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 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 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 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 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 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 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 0 0 0

```

The following fields are the most important fields in this report:

- ▶ GBS_RTE: This field shows the actual negotiated speed of the FICON Channel.
- ▶ RD_MB and WR_MB: The amount of uncompressed data that is transferred by this FICON Channel.

The host adapter activity is summarized per adapter, and as a total of all adapters. This result is also shown in the vnode adapter Throughput Distribution report shown in Example 11-19.

Example 11-19 Extract of the Adapter Throughput Distribution report

```

1(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

```

The provided example is an extract of the report and summarizes the overall host throughput and shows how many one-hour intervals have shown which throughput.

For example, 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.

Cache Partition Activity

This report provides details of Cache Partitions Activity in the TS7700. For a TS7740 and a TS7700D, only one H30TVC1 is provided. For a TS7700T, multiple H30TVCx can be shown. The H30TVC1 for a TS7700T contains the information of the resident partition. If you have defined two Tape Partitions, you also get H30TVC2 and H30TVC3. You can identify the following information for each 15-minute interval:

- ▶ CPU and Disk Utilization
- ▶ The number of scratch (Fast Ready) mounts, cache hits, cache misses, and Sync mounts
- ▶ The value that is defined for PMTHLVL
- ▶ The percentage of read, write, and deferred copy throttling, and the throttling reasons
- ▶ The capacity and number of logical volumes by preference group (0 or 1) in cache, and an indicator for the residency time

The report also shows information about the Preference Groups. The following fields are the most important fields in this report:

- ▶ The ratio between FAST_RDY_MOUNTS, CACHE_HIT_MOUNTS, and CACHE_MISS_MOUNTS. In general, a high number of CACHE_MISSES might mean that additional cache capacity is needed, or cache management policies need to be adjusted. For a TS7700T, you might also reconsider the Tape Partition sizes.
- ▶ FAST_RDY_AVG_SECS and CACHE_HIT_AVG_SECS need to show only a few seconds. CACHE_MIS_AVG_SECS can list values higher than a few seconds, but higher values (more than 2 - 3 minutes) might indicate a lack of physical tape drives. For more information, see Example 11-20.

Example 11-20 List of Values

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

Peer-to-Peer Activity

The Peer-to-Peer Activity report that is shown in Example 11-21 provides various performance metrics of grid activity. This report can be useful for installations working in Deferred copy mode. This report enables, for example, the analysis of subsystem performance during peak grid network activity, such as determining the maximum delay during the batch window.

Example 11-21 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	MiB	AV_DEF	AV_RUN	MiB_TO	CALC	V_MNTS	MiB_XFR	MiB_XFR	1-->0	CALC	1-->2	CALC	1-->3	CALC
TO	TO	QUEAGE	QUEAGE	TVCBY	MiB/	DONE_BY	FR_DL	TO_DL	TVC_BY	MiB/	TVC_BY	MiB/	TVC_BY	MiB/	TVC_BY
RECEIVE	RECEIVE	---	MINUTES---	COPY	SEC	OTHR_DL	RMT_WR	RMT_RD	COPY	SEC	COPY	SEC	COPY	SEC	COPY
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	0.0
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

For the time 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 (in MB)

- ▶ The average age of copy jobs on the deferred and immediate copy queue
- ▶ The amount of data (in MB) to and from the TVC driven by copy activity
- ▶ The amount of data (in MB) copied from other clusters (inbound data) to the cluster on which the report was run

Tip: Analyzing the report that is shown in Example 11-21, 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.

The following fields are the most important fields in this report:

- ▶ **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 that is achieved when replicating data between the clusters that are 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.

11.18.4 VEHGRXCL tool overview

VEHGRXCL is a tool that can be downloaded from the IBM Tape Tools and used as the graphical interface for the records that are provided by VEHSTATS. The VEHGRXCL.EXE file contains VEHSTATS_Model.ppt and VEHSTATS_Model.xls. You can use these files to create graphs on cluster activity based on the flat files that are created with VEHSTATS. Detailed instructions about how to include your data in the tool are described in the first worksheet in the VEHSTATS_Model.xls file that is created as part of the installation procedure.

The following steps describe the sequence of actions in general to produce the graphs of your grid environment:

1. Run the BVIRHSTV program to collect the TS7700 BVIR History data for a selected period (suggested 31 days). Run the VEHSTATS program for the period to be analyzed (a maximum of 31 days is used).
2. Select one day during the analysis period to analyze in detail, and run the VEHSTATS hourly report for that day. You can import the hourly data for all days and then select the day later in the process. You also decide which cluster will be reported by importing the hourly data of that cluster.
3. File transfer the two space-separated files from VEHSTATS (one daily and one hourly) to your workstation.
4. Start Microsoft Excel and open this workbook, which must be in the directory C:\VEHSTATS.
5. Import the VEHSTATS daily file into the "Daily data" sheet by using a special parsing option.

6. Import the VEHSTATS hourly file into the “Hourly data” sheet, by using a special parsing option. Copy 24 hours of data for your selected day and cluster and paste it into the top section of the “Hourly data” sheet.
7. Open the accompanying VEHSTATS_MODEL.PPT Microsoft PowerPoint presentation and ensure that automatic links are updated.
8. Save the presentation with a *new name* so as not to modify the original VEHSTATS_MODEL.PPT.
9. Verify that the PowerPoint presentation is correct, or make any corrections necessary.
10. Break the links between the workbook and the presentation.
11. Edit or modify the saved presentation to remove blank or unneeded charts. Save the presentation with the links broken.

The following examples of PowerPoint slides give an impression of the type of information that is provided with the tool. You can easily update these slides and include them in your own capacity management reports.

Figure 11-44 gives an overview of all of the sections included in the PowerPoint presentation.

Agenda

- This presentation contains the following sections: In PowerPoint, right click on the section name and then “Open Hyperlink” to go directly to the beginning of that section.
 - [Overview](#)
 - [Data transfer](#)
 - [Virtual mounts](#)
 - [Virtual mount times](#)
 - [Virtual Drive and Physical Drive usage](#)
 - [Physical mounts](#)
 - [Physical mount times](#)
 - [Data compression ratios](#)
 - [Blocksizes](#)
 - [Tape Volume Cache performance](#)
 - [Throttling](#)
 - [Multi cluster configuration \(Grid\)](#)
 - [Import/Export Usage](#)
 - [Capacities: Active Volumes and GB stored](#)
 - [Capacities: Cartridges used](#)
 - [Pools \(Common Scratch Pool and up to 4 Storage Pools \)](#)

Figure 11-44 Sample VEHGRXCL - Agenda

Figure 11-45 gives an overview of the reported period.

Customers Grid February	
TS7700 Serial #	CL1
Grid #	ACEF1
First day of analysis	1-Feb-11
Last day of analysis	28-Feb-11
Number of days	28
TVC size (GB)	13744
Overall average mount time (secs)	10.0
Overall cache miss %	14.4
Max daily cache miss %	41.0
Max physical drives mounted	16
Max virtual drives mounted	69
Max total virtual mounts per day	3553
Max scratch virtual mounts per day	2208
Max Read GB per day	1503
Max Write GB per day	6168
Max 15-minute Read MB per sec	161
Max 15-minute Write MB per sec	381

Figure 11-45 Sample VEHGRXCL - Overview

Figure 11-46 is an example throughput, expressed in MBps.




Figure 11-46 Sample VEHGRXCL - Maximum and average throughput

Figure 11-47 is an example of physical mounts.




Figure 11-47 Sample VEHGRXCL - All physical mounts

11.18.5 VEHAUDIT overview

VEHAUDIT is also based on the BVIR information and a part of the tapetool suite. It has several independent reports,

Copy Distribution report (CPYDST)

This report shows long copies taken from one cluster to the other clusters in the grid. Family relationships can be reflected as well. This report can be used to identify the RPO, and how long in general the creation of copies takes.

Detail Report (DTLRPT)

The DTLRPT provides an audit functionality, and provides information about which logical volumes have a deviation from the requested copy consistency policy on the specified cluster. In opposition to the BVIRAUD function, it does not compare several clusters to find out, if copies are missing. Instead, an audit of a specific cluster is performed, and the output report shows every logical volume that is not in the appropriate state.

The report contains not only missing copies, but also inconsistent data level, data corruption of a logical volume (only detected at read), and existing (but unwanted) copies.

For each logical volume reported, it also reports all of the following information:

- ▶ Type of issue
- ▶ Constructs,
- ▶ Creation, last reference and expire date
- ▶ Category
- ▶ Size

- ▶ Data set name
- ▶ Physical volume placement (only TS7740/TS770T)
- ▶ information from the tape management system (if input is provided)

MES Error Report (MESERRP)

This report shows a list of MES records of corrupted volumes. A volume is corrupted if the "Read_Error" flag is set.

Multi Copy (MLTCPY)

This report shows if more copies exist than requested from the copy consistency policy perspective.

STALE

This report contains logical volumes that are expired, but not yet deleted. This report should usually be empty.

TOSYNC

If fewer copies exist for a logical volume than requested in the management class, this report should be generated. The tool is able to eliminate the scratches. The list produced in this report can then be used as input for copy refresh processing.

11.19 IBM z/OS commands for monitoring

In addition to the previously introduced methods and options for monitoring the TS7700, the following extra points offer further subsystem monitoring.

11.19.1 DISPLAY SMS commands

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 11-22, separated by a dashed line.

Example 11-22 DISPLAY SMS command responses

```

D SMS,OAM
F OAM,D,OAM,L=DENEKA-Z
CBR1100I OAM status: 171
TAPE TOT  ONL  TOT  TOT  TOT  TOT  TOT  ONL  AVL  TOTAL
      LIB  LIB  AL  VL  VCL  ML  DRV  DRV  DRV  SCRATCH
      2   2   0   0   2   0   528  256  256   12
There are also 3 VTS distributed libraries defined.
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.
-----
D SMS,LIB(ALL),DETAIL
F OAM,D,LIB,L=DENEKA-Z
CBR1110I OAM library status: 174
TAPE  LIB  DEVICE  TOT  ONL  AVL  TOTAL  EMPTY  SCRATCH  ON  OP
LIBRARY  TYP  TYPE  DRV  DRV  DRV  SLOTS  SLOTS  VOLS
DTS7720  VDL  3957-VEB  0   0   0    0     0     0   Y   Y
HYDRAE   VDL  3957-V07  0   0   0   185   133   0   Y   Y

```

```

HYDRAG VCL GRID 512 256 256 0 0 12 Y Y
HYDRAO VDL 3957-V06 0 0 0 400 340 0 Y Y
HYDV06 VCL 3957-V06 16 0 0 0 0 0 Y Y

```

```

-----
D SMS,LIB(ALL),STATUS
IGD002I 13:47:31 DISPLAY SMS 176

```

```

LIBRARY CLASS SYSTEM= 1
DTS7720 TAPE +
HYDRAE TAPE +
HYDRAG TAPE +
HYDRAO TAPE +
HYDV06 TAPE +

```

```

***** LEGEND *****
. THE LIBRARY IS NOT DEFINED TO THE SYSTEM
+ THE LIBRARY IS ONLINE
- THE LIBRARY IS OFFLINE
P THE LIBRARY IS PENDING OFFLINE

```

```

-----
D SMS,LIB(HYDRAG),DETAIL
F OAM,D,LIB,HYDRAG,L=DENEKA-Z
CBR1110I OAM library status: 179

```

TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON	OP
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS	Y	Y
HYDRAG	VCL	GRID	512	256	256	0	0	12	Y	Y

```

-----
MEDIA TYPE SCRATCH COUNT SCRATCH THRESHOLD SCRATCH CATEGORY
MEDIA1 7 0 0011
MEDIA2 5 0 0012

```

```

-----
DISTRIBUTED LIBRARIES: HYDRAE DTS7720

```

```

-----
LIBRARY ID: 00186
OPERATIONAL STATE: AUTOMATED
ERROR CATEGORY SCRATCH COUNT: 0
CORRUPTED TOKEN VOLUME COUNT: 0

```

```

-----
Library supports outboard policy management.
Library supports logical WORM.

```

```

-----
D SMS,LIB(HYDRAE),DETAIL
F OAM,D,LIB,HYDRAE,L=DENEKA-Z
CBR1110I OAM library status: 168

```

TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON	OP
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS	Y	Y
HYDRAE	VDL	3957-V07	0	0	0	185	133	0	Y	Y

```

-----
COMPOSITE LIBRARY: HYDRAG

```

```

-----
LIBRARY ID: 01052
CACHE PERCENTAGE USED: 0
OPERATIONAL STATE: AUTOMATED
SCRATCH STACKED VOLUME COUNT: 12
PRIVATE STACKED VOLUME COUNT: 5

```

```

-----
Library supports import/export.
Library supports outboard policy management.
Library supports logical WORM.
Convenience I/O station installed.
Convenience I/O station in Output mode.

```

Bulk input/output not configured.

```
-----  
D SMS,VOL(A13052)  
CBR1180I OAM tape volume status: 143  
VOLUME MEDIA STORAGE LIBRARY USE W C SOFTWARE LIBRARY  
TYPE GROUP NAME ATR P P ERR STAT CATEGORY  
A13052 MEDIA1 SGG00001 HYDRAG P N NOERROR PRIVATE  
-----  
RECORDING TECH: 36 TRACK COMPACTION: YES  
SPECIAL ATTRIBUTE: NONE ENTER/EJECT DATE: 2014-02-12  
CREATION DATE: 2014-02-12 EXPIRATION DATE: 2014-02-13  
LAST MOUNTED DATE: 2014-02-12 LAST WRITTEN DATE: 2014-02-12  
SHELF LOCATION:  
OWNER: DENEKA  
LM SG: SGG00001 LM SC: SC00000K LM MC: MNSSN068 LM DC: D100N006  
LM CATEGORY: 001F  
-----
```

Logical volume.
Volume is cache resident.
Valid copy in each distributed library.

```
-----  
D SMS,VOL(A13051)  
CBR1180I OAM tape volume status: 146  
VOLUME MEDIA STORAGE LIBRARY USE W C SOFTWARE LIBRARY  
TYPE GROUP NAME ATR P P ERR STAT CATEGORY  
A13051 MEDIA1 *SCRTCH* HYDRAG S N N NOERROR SCRMED1  
-----  
RECORDING TECH: 36 TRACK COMPACTION: YES  
SPECIAL ATTRIBUTE: NONE ENTER/EJECT DATE: 2014-02-12  
CREATION DATE: 2014-02-12 EXPIRATION DATE:  
LAST MOUNTED DATE: 2014-02-12 LAST WRITTEN DATE: 2014-02-12  
SHELF LOCATION:  
OWNER:  
LM SG: LM SC: LM MC: LM DC:  
LM CATEGORY: 0011  
-----
```

Logical volume.

For more information, see Chapter 10, “Host Console operations” on page 601 and *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

11.19.2 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 11-23 on page 739 shows the output of the **LIBRARY DD** command that you can use to verify whether all virtual drives are available.

Example 11-23 Sample response for LI DD,libname command

```

LI DD,ATVIGA
RESPONSE=EGZB
CBR1220I Tape drive status: 338
DRIVE  DEVICE  LIBRARY  ON  OFFREASN  LM  ICL  ICL  MOUNT
NUM    TYPE    NAME      LI  OP  PT  AV  CATEGRY  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 10, “Host Console operations” on page 601 and *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

11.20 What to look for and where

This chapter describes tools, provides considerations, and gives you important information to help you monitor and understand the performance indicators of your TS7700 grid. Table 11-9 summarizes where you can find information and what observations you can make.

Table 11-9 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
TS7700 health check	TS7700 MI	Display each composite library	Each shift	Report any offline or degraded status

Information	Source	Tracking	Reporting interval	Observation
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
Virtual volumes managed	VEHSTATS	Rolling weekly trend	Daily	Capacity planning: maximum one million per grid
MB stored	VEHSTATS	Rolling weekly trend	Daily	Capacity planning and general awareness
Back-end 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 and general awareness
Available slot count	D SMS, LIB(ALL), DETAIL	Rolling weekly trend	Daily	Capacity planning and general awareness

Information	Source	Tracking	Reporting interval	Observation
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

Most checks that you need to 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 is intended as a basis for monitoring. You can tailor this information to best fit your needs.

11.21 Virtual Device Allocation in z/OS with JES2

z/OS (JES2 only) allocation characteristics in general are described and how allocation algorithms are being influenced by z/OS allocation parameter settings EQUAL and BYDEVICES, by the TS7700 Copy Consistency Points, by Override Settings, and by the Allocation Assistance functions is shown. Carefully plan for your device allocation requirements because improper use of the parameters, functions, and settings can have unpredictable results.

Various scenarios are presented, showing the influence of the algorithms involved in virtual device allocation. A configuration with two 3-cluster grid configurations, named GRID1 and GRID2, is used. Each grid has a TS7700D (Cluster 0) and a TS7740 (Cluster 1) at the primary Production Site, and a TS7740 (Cluster 2) at the Disaster Site. The TS7700D Cluster 0 in the Production Site can be considered a deep cache for the TS7740 Cluster 1 in the scenarios that are described next.

Figure 11-48 gives a general overview of the configuration.

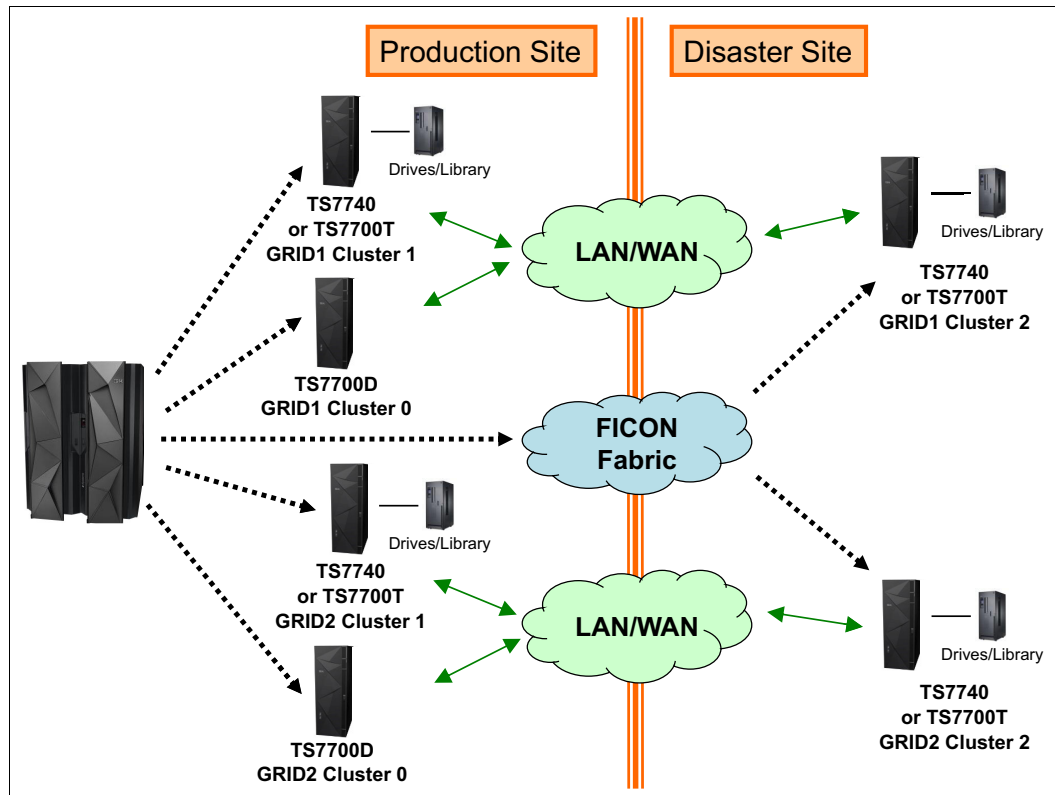


Figure 11-48 Grid configuration overview

In Figure 11-48, the host in the Production Site has direct access to the local clusters in the Production Site, and has access over the extended FICON fabric to the remote clusters in the Disaster Site. The extended FICON fabric can include dense wavelength division multiplexing (DWDM) connectivity, or can use FICON tape acceleration technology over IP. Assume that connections to the remote clusters have a limited capacity bandwidth.

Furthermore, there is a storage management subsystem (SMS) Storage Group (SG) per grid. The groups are defined in the SMS SG routine as GRID1 and GRID2. SMS equally manages SGs. The order in the definition statement does not influence the allocations.

The following scenarios are described. Each scenario adds functions to the previous scenario so that you can better understand the effects of the added functions:

- ▶ **EQUAL allocation:** Describes the allocation characteristics of the default load-balancing algorithm (EQUAL) and its behavior across the sample TS7700 configuration with two grids. See 11.21.1, “EQUAL allocation” on page 743.
- ▶ **BYDEVICES allocation:** Adds the new BYDEVICES algorithm to the configuration. It explains how this algorithm can be activated and the differences from the default EQUAL algorithm. See 11.21.2, “BYDEVICES allocation” on page 745.
- ▶ **Allocation and Copy Consistency Point setting:** Adds information about the effect of the Copy Consistency Point on the cache data placement. The various TS7700 Virtualization override settings influence this data placement. See 11.21.3, “Allocation and Copy Consistency Point setting” on page 747.

- ▶ Allocation and device allocation assistance (DAA): DAA is activated, and the effects on allocation are described. The unavailability of cluster and device information influences the allocation when DAA is enabled. DAA is enabled, by default. See 11.21.4, “Allocation and device allocation assistance” on page 749.
- ▶ Allocation and scratch allocation assistance (SAA): SAA is activated, and its effects are described. A sample workload in this scenario is presented to clarify SAA. The advantages of SAA and the consequences of the unavailability of clusters and devices are explained. SAA is disabled, by default. See 11.21.5, “Allocation and scratch allocation assistance” on page 752.

11.21.1 EQUAL allocation

For non-specific (scratch) allocations, by default, MVS device allocation will first randomize across all eligible libraries and then, after a library is selected, will randomize on the eligible devices within that library. In terms of the TS7700, *library* refers to a composite library because the MVS allocation has no knowledge of the underlying clusters (distributed libraries). The default algorithm (EQUAL) works well if the libraries under consideration have an equal number of online devices.

For example, if two libraries are eligible for a scratch allocation and each library has 128 devices, over time, each library will receive approximately half of the scratch allocations. If one of the libraries has 128 devices and the other library has 256 devices, each of the libraries still receives approximately half of the scratch allocations. The allocations are independent of the number of online devices in the libraries.

Remember: With EQUAL allocation, the scratch allocations are randomized across the libraries. EQUAL allocation is not influenced by the number of online devices in the libraries.

In this first scenario, both DAA and SAA are assumed to be disabled. With the TS7700, you can control both assistance functions with the **LIBRARY REQUEST** command. DAA is **ENABLED** by default and can be **DISABLED** with the command. SAA is **DISABLED** by default and can be **ENABLED** with the command. Furthermore, none of the TS7700 override settings are used.

Assuming that the MC for the logical volumes has a Copy Consistency Point of [R,R,R] in all clusters and that the number of available virtual drives are the same in all clusters, the distribution of the allocation across the two grids (composite libraries) is evenly spread. The multi-cluster grids are running in **BALANCED** mode, so there is no preference of one cluster above another cluster.

With the default algorithm EQUAL, the distribution of allocations across the clusters (in a multiple cluster grid) depends on the order in which the library port IDs were initialized during IPL (or input/output definition file (IODF) activate). The distribution of allocations across the clusters also depends on whether the library port IDs in the list (returned by the **DEVSERV QLIB,composite-library-id** command) randomly represent each of the clusters.

Alternatively, the distribution depends on if the library port IDs in the list tend to favor the library port IDs in one cluster first, followed by the next cluster, and so on. The order in which the library port IDs are initialized and appear in this **DEVSERV** list can vary across IPLs or IODF activates, and can influence the randomness of the allocations across the clusters.

So with the default algorithm EQUAL, there might be times when device randomization within the selected library (composite library) appears unbalanced across clusters in a TS7700 that have online devices. As the number of eligible library port IDs increases, the likelihood of this imbalance occurring also increases. If this imbalance affects the overall throughput rate of the library, consider enabling the BYDEVICES algorithm described in 11.21.2, “BYDEVICES allocation” on page 745.

Remember: Exceptions to this can also be caused by z/OS JCL backward referencing specifications (UNIT=REF and UNIT=AFF).

With z/OS V1R11 and later, and z/OS V1R8 through V1R10 with APAR OA26414 installed, it is possible to change the selection algorithm to BYDEVICES. The algorithm EQUAL, which is the default algorithm that is used by z/OS, can work well if the libraries (composite libraries) under consideration have an equal number of online devices and the previous cluster behavior is understood.

The non-specific (scratch) allocation distribution is shown in Figure 11-49.

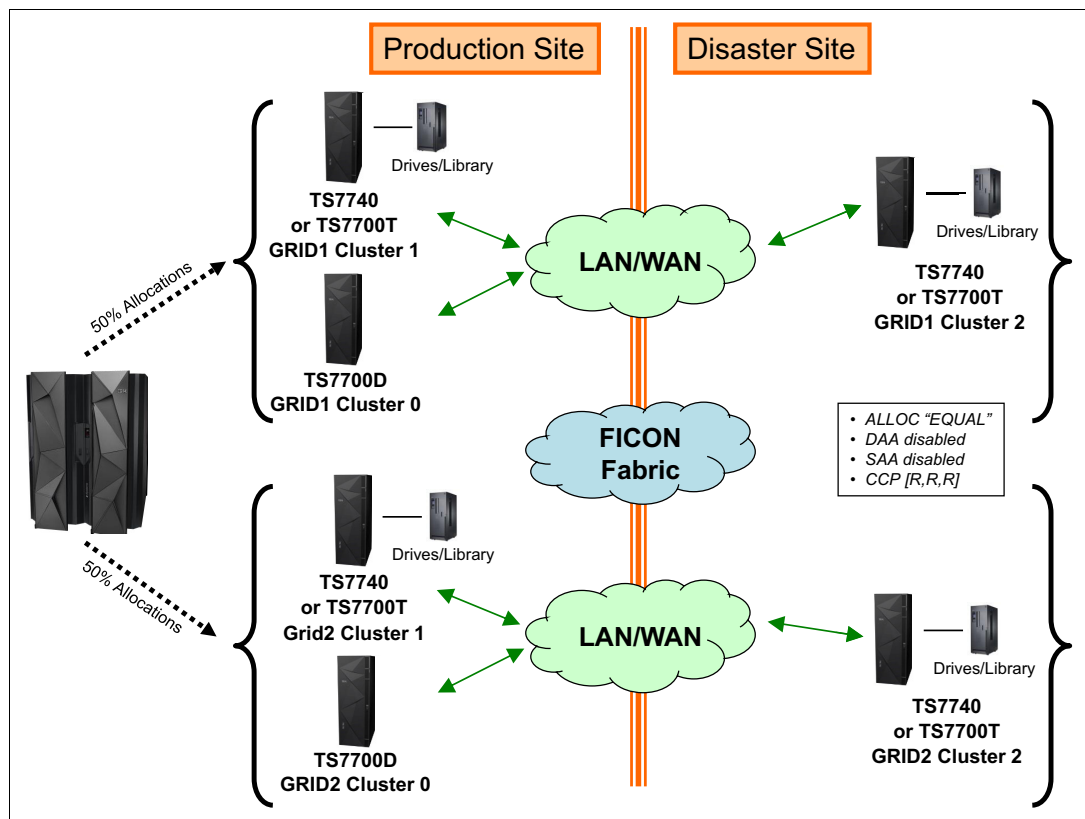


Figure 11-49 ALLOC EQUAL scratch allocations

For specific allocations (DAA DISABLED in this scenario), it is first determined which of the composite libraries, GRID1 or GRID2, has the requested logical volume. That grid is selected and the allocation can go to any of the clusters in the grid. If it is assumed that the logical volumes were created with the EQUAL allocation setting (the default), it can be expected that specific device allocation to these volumes will be distributed equally among the two grids.

However, how well the allocations are spread across the clusters depends on the order in which the library port IDs were initialized, and whether this order was randomized across the clusters.

In a TS7740 multi-cluster grid configuration, only the original copy of the volume stays in cache, normally in the mounting cluster's TVC for a Copy Consistency Point setting of [R,R,R]. The copies of the logical volume in the other clusters are managed as a TVC Preference Level 0 (PG0 - remove from cache first) unless an SC specifies Preference Level 1 (PG1 - stay in cache) for these volumes.

A number of possibilities can influence the cache placement:

- ▶ You can define an SC for the volume with Preference Level 0 (PG0). The logical volume does not stay in the I/O TVC cluster.
- ▶ You can set the CACHE COPYFSC option, with a **LIBRARY REQUEST, GRID [1] / [2], SETTING, CACHE, COPYFSC, ENABLE** command. When the **ENABLE** keyword is specified, the logical volumes that are copied into the cache from a peer TS7700 cluster are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7740 cluster that is receiving the copy.

Therefore, a copy of the logical volume also stays in cache in each non-I/O TVC cluster where an SC is defined as Preference Level 1 (PG1). However, because the TS7700D is used as a deep cache, there are no obvious reasons to do so.

In the hybrid multi-cluster grid configuration that is used in the example, there are two cache allocation schemes, depending on the I/O TVC cluster selected when creating the logical volume. Assume an SC setting of Preference Level 1 (PG1) in the TS7740 Cluster 1 and Cluster 2.

- ▶ If the mounting cluster for the non-specific request is the TS7700D Cluster 0, only the copy in that cluster stays. The copies in the TS7740 Cluster 1 and Cluster 2 will be managed as Preference Level 0 (PG0) and will be removed from cache after placement of the logical volume on a stacked physical volume. If a later specific request for that volume is directed to a virtual device in one of the TS7740s, a cross-cluster mount from Cluster 1 or Cluster 2 occurs to Cluster 0's cache.
- ▶ If the mounting cluster for the non-specific request is the TS7740 Cluster 1 or Cluster 2, not only the copy in that cluster stays, but also the copy in the TS7700D Cluster 0. Only the copy in the other TS7740 cluster will be managed as Preference Level 0 (PG0) and will be removed from cache after placement of the logical volume on a stacked physical volume.

Cache preferencing is not valid for the TS7700D cluster. A later specific request for that logical volume creates only a cross-cluster mount if the mount point is the vnode of the TS7740 cluster that is not used at data creation of that volume.

With the EQUAL allocation algorithm that is used for specific mount requests, there are always cross-cluster mounts when the cluster where the device is allocated is not the cluster where the data is. Cache placement can limit the number of cross-cluster mounts but cannot avoid them. Cross-cluster mounts over the extended fabric are likely not acceptable, so vary the devices of Cluster 2 offline.

11.21.2 BYDEVICES allocation

The alternative algorithm BYDEVICES randomizes scratch allocations across all devices. For example, if two libraries are eligible for a scratch allocation and each library has 128 devices, over time each library will receive approximately half of the scratch allocations, similar to the EQUAL algorithm. Again, in terms of the TS7700, "library" refers to a composite library because MVS allocation has no knowledge of the underlying clusters (distributed libraries).

However, if one of the libraries has 128 devices and the other library has 256 devices, over time, the library that has 128 devices will receive 1/3 of the scratch allocations and the library that has 256 devices will receive approximately 2/3 of the scratch allocations. This is different compared to the default algorithm EQUAL, which does not take the number of online devices in a library into consideration.

Clarification: With BYDEVICES, the scratch allocation randomizes across all devices in the libraries, and is influenced by the number of online devices.

With z/OS V1R11 and later, and z/OS V1R8 through V1R10 with APAR OA26414 installed, it is possible to influence the selection algorithm. The BYDEVICES algorithm can be enabled through the ALLOCxx PARMLIB member by using the SYSTEM TAPELIB_PREF(BYDEVICES) parameter, or it can be enabled dynamically through the SETALLOC operator command by entering **SETALLOC SYSTEM,TAPELIB_PREF=BYDEVICES**.

The alternative BYDEVICES algorithm can be replaced by the default EQUAL algorithm by specifying EQUAL through the **SETALLOC** command or the ALLOCxx PARMLIB member in a similar manner. Before enabling the new load balancing support, care must be taken to ensure that the wanted results are achieved. This is important if the libraries are being shared across multiple systems and the systems are at different levels of support, or if manual actions have already been taken to account for the behavior of algorithms that were used in previous releases.

Consideration: The **SETALLOC** operator command support is available only in z/OS V1R11 or later releases. In earlier z/OS releases, BYDEVICES must be enabled through the ALLOCxx PARMLIB member.

Assume that GRID1 has a total of 60 virtual devices online and GRID2 has 40 virtual devices online. For each grid, the distribution of online virtual drives is 50% for Cluster 0, 25% for Cluster 1, and 25% for Cluster 2.

The expected distribution of the scratch allocations is as shown in Figure 11-50.

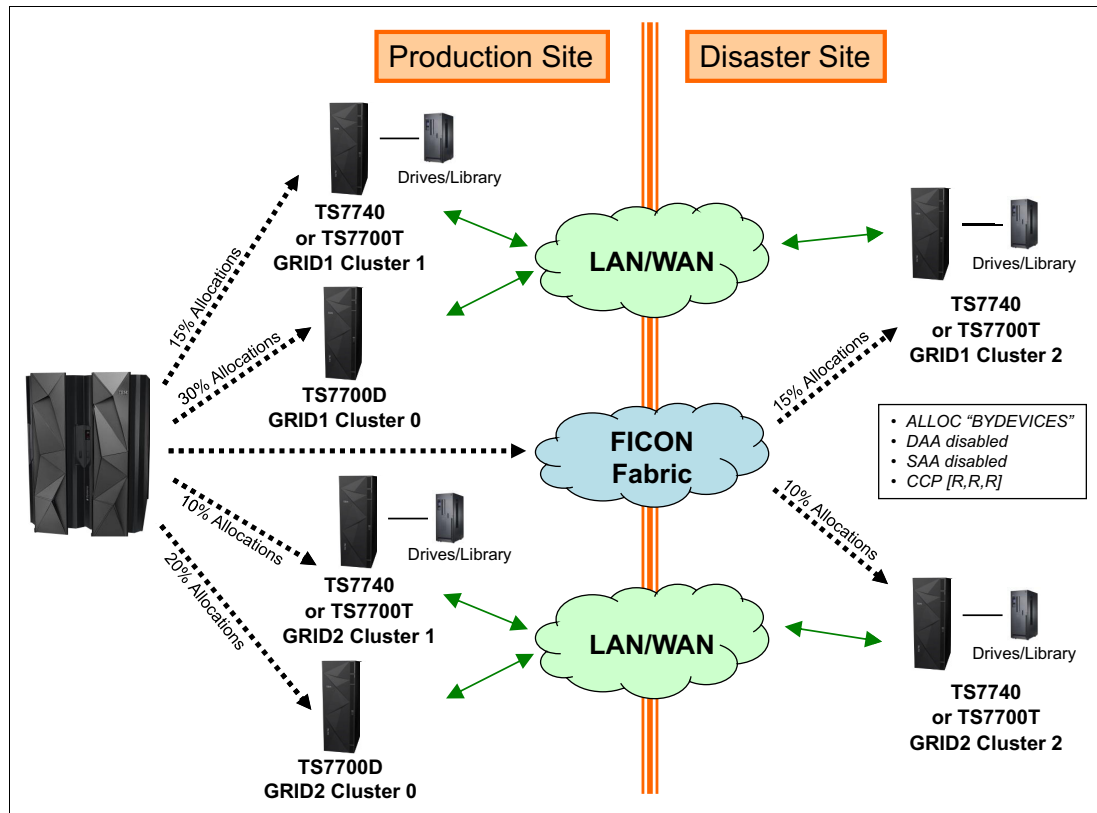


Figure 11-50 ALLOC BYDEVICES scratch allocations

As stated in 11.21.1, "EQUAL allocation" on page 743, DAA is ENABLED by default and was DISABLED by using the **LIBRARY REQUEST** command. Furthermore, none of the TS7700 override settings are activated.

For specific allocations (DAA DISABLED in this scenario), it is first determined which one of the composite libraries, GRID1 or GRID2, has the requested logical volume. That grid is selected, and the allocations can go to any cluster in the grid and are proportionately distributed based on the number of online devices in each cluster. The logical volume cache placement possibilities and the two allocation schemes, both described in 11.21.1, "EQUAL allocation" on page 743, are also applicable for the BYDEVICES allocation.

With the BYDEVICES allocation algorithm that is used for specific mount requests, there are always cross-cluster mounts when the cluster where the device is allocated is not the cluster where the data is. Cache placement can limit the number of cross-cluster mounts but cannot avoid them. Cross-cluster mounts over the extended fabric are likely not acceptable, so vary the devices of Cluster 2 offline.

11.21.3 Allocation and Copy Consistency Point setting

By defining the Copy Consistency Point, you control if and how a volume will be placed in a determined cluster of the grid. If you plan to use the TS7700D Cluster 0 as a deep cache, you probably prefer to define the MC Construct as [R,D,D]. By defining this, Cluster 0 is the primary placeholder of the data. At job completion time, only this cluster has a valid copy of the data. The other cluster creates a deferred copy of that logical volume afterward.

For more information about Copy Consistency Points, see the *IBM TS7700 Best Practices - Synchronous Mode Copy* and *IBM TS7700 Best Practices - Copy Consistency Point* white papers:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102098>

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101230>

It is further assumed that the allocation characteristics apply as described in 11.21.2, “BYDEVICES allocation” on page 745. Both DAA and SAA are DISABLED in this scenario, too, and none of the TS7700 override settings are used.

For non-specific (scratch) allocations, the BYDEVICES algorithm randomizes across all devices, resulting in allocations on all three clusters of each grid. Subsequently, i/O TVC selection assigns the TVC of Cluster 0 as the I/O TVC due to the Copy Consistency Point setting. Many factors can influence this selection, as explained in 2.2.2, “Tape Volume Cache” on page 31.

Normally, the cluster with a Copy Consistency Point of R(un) gets preference over other clusters. As a consequence, the TVC of Cluster 0 is selected as the I/O TVC and cross-cluster mounts are issued from both Cluster 1 and Cluster 2.

By activating the override setting “Prefer Local Cache for Fast Ready Mount Requests” in both clusters in the Disaster Site, cross-cluster mounts are avoided but the copy to Cluster 0 is made before the job ends, caused by the R(un) Copy Consistency Point setting for this cluster. Therefore, by further defining a family for the Production Site clusters, Cluster 1 retrieves its copy from Cluster 0 in the Production Site location, avoiding using the remote links between the locations.

The method to prevent device allocations at the Disaster Site, implemented mostly today, is just varying offline all the remote virtual devices. The disadvantage is that in losing a cluster in the Production Site, an operator action is required to vary online manually the virtual devices of Cluster 2 of the grid with the failing cluster. With the TS7700 R2.0, an alternative solution is using scratch allocation assistance (SAA), which is described in 11.21.5, “Allocation and scratch allocation assistance” on page 752.

For specific allocations, the algorithm that is described in 11.21.2, “BYDEVICES allocation” on page 745 applies when DAA is disabled. It is first determined which of the composite libraries, GRID1 or GRID2, has the requested logical volume. Subsequently, that grid is selected and the allocation over the clusters is randomized. It can be assumed that, if the requested logical volumes were earlier created with the BYDEVICES allocation scenario, these logical volumes are spread over the two grids. The allocation distribution within the grid over the three clusters is determined by the number of the online devices in each of the clusters.

Cluster 0 is likely to have a valid copy of the logical volume in the cache due to the Copy Consistency Point setting of [R,D,D]. If the vnodes of Cluster 1 and Cluster 2 are selected as mount points, it results in cross-cluster mounts. It might happen that this volume has been removed by a policy in place for TS7700D Cluster 0, resulting in the mount point TVC as the I/O TVC.

In the TS7700, activating the *Force Local TVC to have a copy of the data* override first results in a recall of the virtual volume from a stacked volume. If there is no valid copy in the cluster or if it fails, a copy is retrieved from one of the other clusters before the mount completes. Activating the *Prefer Local Cache for non-Fast Ready Mount Requests* override setting recalls a logical volume from tape instead of using the grid links for retrieving the data of the logical volume from Cluster 0. This might result in longer mount times.

With the TS7700, an alternative solution can be considered by using device allocation assistance (DAA) that is described in 11.21.4, "Allocation and device allocation assistance" on page 749. DAA is enabled by default.

Figure 11-51 shows the allocation results of specific and non-specific allocations when the devices of the remote clusters in the Disaster Site are not online. Allocation BYDEVICES is used. GRID1 has a total of 60 devices online and GRID2 has 40 devices online. For each grid, the distribution of online devices is 75% for Cluster 0 and 25% for Cluster 1.

Cross-cluster mounts might apply for the specific allocations in Cluster 1 because it is likely that only the TS7700D Cluster 0 will have a valid copy in cache. The red arrows show the data flow as result of these specific allocations.

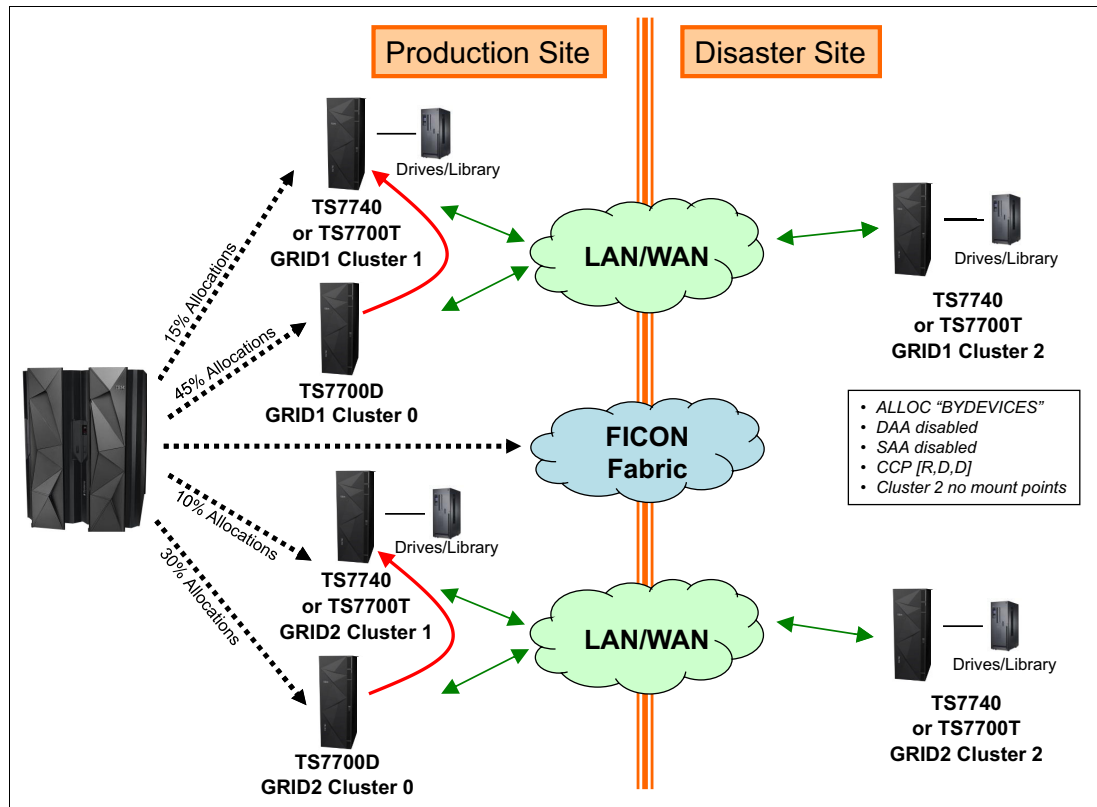


Figure 11-51 Allocation and Copy Consistency Points set at R,D,D

11.21.4 Allocation and device allocation assistance

Device allocation assistance (DAA) allows the host to query the TS7700 to determine which clusters must be preferred for a private (specific) mount request before the actual mount is requested. DAA returns to the host a ranked list of clusters (the preferred cluster is listed first) where the mount must be run. If DAA is enabled, it is for the composite library, and it is used by all z/OS JES2 or JES3 LPARs having the proper level of supporting software (z/OS V2R1 or above).

The selection algorithm orders the clusters first by those having the volume already in cache, then by those having a valid copy on tape, and then by those without a valid copy. Later, host processing attempts to allocate a device from the first cluster that is 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 enables the host to direct the mount request to the cluster that will result in the fastest mount, typically the cluster that has the logical volume resident in cache.

If the mount is directed to a cluster without a valid copy, a cross-cluster mount results. Thus, in special cases, even if DAA is enabled, cross-cluster mounts and recalls can still occur.

For JES2, if the default allocation algorithm EQUAL is used, it supports an ordered list for the first seven library port IDs returned in the list. After that, if an eligible device is not found, all of the remaining library port IDs are considered equal. The alternative allocation algorithm BYDEVICES removes the ordered library port ID limitation.

With the TS7700, install the additional APAR OA30718 before enabling the new BYDEVICES algorithm. Without this APAR, the ordered library port ID list might not be acknowledged correctly, causing specific allocations to appear randomized.

In the scenario that is described in 11.21.3, “Allocation and Copy Consistency Point setting” on page 747, if you enable DAA (this is the default) by entering the command **LIBRARY REQUEST, GRID[1] / [2], SETTING, DEVALLOC, PRIVATE, ENABLE**, it influences the specific requests in the following manner. The Copy Consistency Point is defined as [R,D,D]. It is assumed that there are no mount points in Cluster 2.

It is further assumed that the data is not in the cache of the TS7740 Cluster 1 anymore because this data is managed as TVC Preference Level 0 (PG0), by default. It is first determined which of the composite libraries, GRID1 or GRID2, has the requested logical volume. Subsequently, that grid is selected and the allocation over the clusters is determined by DAA. The result is that all allocations select the TS7700D Cluster 0 as the preferred cluster.

You can influence the placement in cache by setting the **CACHE COPYFSC** option with the **LIBRARY REQUEST, GRID[1] / [2], SETTING, CACHE, COPYFSC, ENABLE** command. When the **ENABLE** keyword is specified, the logical volumes that are copied into the cache from a peer TS7700 cluster are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7740 cluster receiving the copy.

Therefore, a copy of the logical volume also stays in cache in each non-I/O TVC cluster where an SC is defined as Preference Level 1 (PG1). However, because the TS7700D is used as a deep cache, there are no obvious reasons to do so.

There are two major reasons why Cluster 0 might not be selected:

- ▶ No online devices are available in Cluster 0, but are in Cluster 1.
- ▶ The defined removal policies in the TS7700D caused Cluster 0 to not have a valid copy of the logical volume anymore.

In both situations, DAA selects the TS7740 Cluster 1 as the preferred cluster:

- ▶ When the TS7740 Cluster 1 is selected due to lack of online virtual devices on Cluster 0, cross-cluster mounts might happen unless the TS7700 override settings, as described in 11.21.3, “Allocation and Copy Consistency Point setting” on page 747, are preventing this from happening.
- ▶ When the TS7740 Cluster 1 is selected because the logical volume is not in the TS7700D Cluster 0 cache anymore, its cache is selected for the I/O TVS and because the Copy Consistency Point setting is [R,D,D], a copy to the TS7700D Cluster 0 is made as part of successful RUN processing.

Even when DAA is enabled, there might be specific mounts for which the device affinity call is not made. For example, DFSMSshm goes to allocation when appending a volume, requiring that a scratch volume be mounted. Then, when a device is allocated and a volume is to be mounted, it selects from the list of HSM-owned volumes. In this case, because the allocation started as a scratch request, the device affinity is not made for this specific mount.

The MARKFULL option can be specified in DFSMSshm to mark migration and backup tapes that are partially filled during tape output processing as full. This enforces a scratch tape to be selected the next time that the same function begins.

Figure 11-52 shows the allocation result of specific allocations. The devices of the remote clusters in the Disaster Site are not online. GRID1 has in total 60% of specific logical volumes and GRID2 has 40% of the specific logical volumes. This was the result of earlier BYDEVICES allocations when the logical volumes were created.

The expected distribution of the specific allocations is as shown. Cross-cluster mounts might apply in situations where DAA selects the vnode of Cluster 1 as the mount point. The red arrows show the data flow for both the creation of the copy of the data for scratch allocations and for specific allocations.

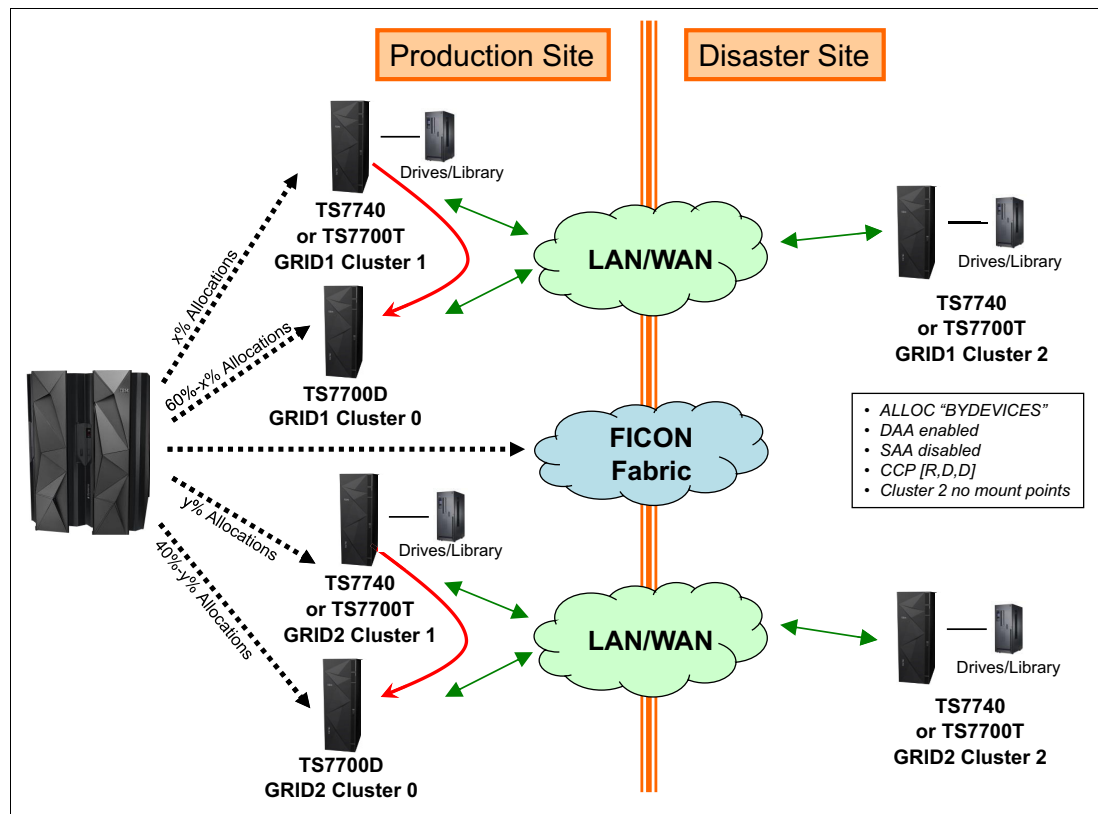


Figure 11-52 Allocation and DAA

With DAA, you can vary the devices in the Disaster Site Cluster 2 online without changing the allocation preference for the TS7700D cache if the logical volumes exist in this cluster and this cluster is available. If these conditions are not met, DAA manages local Cluster 1 and remote Cluster 2 as equal and cross-cluster mounts over the extended fabric are issued in Cluster 2.

A new copy of the logical volume is created due to the MC setting [R] for Cluster 0. This is likely not an acceptable scenario and so, even with DAA ENABLED, vary the devices of Cluster 2 offline.

If you plan to have an alternative MC setup for the Disaster Site (perhaps for the Disaster Test LPARs), you must carefully plan the MC settings, the device ranges that must be online, and whether DAA is enabled. You will probably read production data and create test data by using a separate category code.

If you do not want the grid links overloaded with test data, vary the devices of Cluster 0 and Cluster 1 offline on the disaster recovery (DR) host only and activate the TS7700 Override Setting "Force Local TVC" to have a copy of the data. A specific volume request enforces a mount in Cluster 2 even if there is a copy in the deep cache of the TS7700D Cluster 0.

11.21.5 Allocation and scratch allocation assistance

The scratch allocation assistance (SAA) function can be used when there is a need for a method to have z/OS allocate to specific clusters (candidate clusters) for a workload. For example, DFSMSshm Migration Level 2 (ML2) migration can be directed to a TS7700D cluster with its deep cache, and the archive workload needs to be directed to a TS7740 cluster within the same grid configuration.

SAA is an extension of the DAA function for scratch mount requests. SAA filters the list of clusters in a grid to return to the host a smaller list of candidate clusters that are designated as scratch mount candidates. By identifying a subset of clusters in the grid as sole candidates for scratch mounts, SAA optimizes scratch mounts to a TS7700 grid.

When a composite library supports/enables the SAA function, the host sends an SAA handshake to all SAA-enabled composite libraries and provide the MC that will be used for the upcoming scratch mount. A cluster is designated as a candidate for scratch mounts by using the Scratch Mount Candidate option on the MC construct, which is accessible from the TS7700 MI, as shown in Figure 11-53. By default, all clusters are considered candidates. Also, if SAA is enabled and no SAA candidate is selected, all clusters are considered as candidates.

Name:	*Test													
Secondary Pool:	0													
Description:	Test Management Class													
Retain Copy Mode:	<input type="checkbox"/>													
<table border="1"> <thead> <tr> <th>Clusters</th> <th>Copy Mode</th> <th>Scratch Mount Candidate</th> </tr> </thead> <tbody> <tr> <td>"Zoro[0]" (#BA39A)</td> <td>Rewind Unload (RUN)</td> <td><input checked="" type="checkbox"/></td> </tr> <tr> <td>"Juno[1]" (#BA39B)</td> <td>Deferred</td> <td><input type="checkbox"/></td> </tr> <tr> <td>"Diana[2]" (#BA39C)</td> <td>Deferred</td> <td><input type="checkbox"/></td> </tr> </tbody> </table>			Clusters	Copy Mode	Scratch Mount Candidate	"Zoro[0]" (#BA39A)	Rewind Unload (RUN)	<input checked="" type="checkbox"/>	"Juno[1]" (#BA39B)	Deferred	<input type="checkbox"/>	"Diana[2]" (#BA39C)	Deferred	<input type="checkbox"/>
Clusters	Copy Mode	Scratch Mount Candidate												
"Zoro[0]" (#BA39A)	Rewind Unload (RUN)	<input checked="" type="checkbox"/>												
"Juno[1]" (#BA39B)	Deferred	<input type="checkbox"/>												
"Diana[2]" (#BA39C)	Deferred	<input type="checkbox"/>												

Figure 11-53 Scratch Mount Candidate definition

The targeted composite library uses the provided MC definition and the availability of the clusters within the same composite library to filter down to a single list of candidate clusters. Clusters that are unavailable or in service are excluded from the list. If the resulting list has zero clusters present, the function then views all clusters as candidates.

In addition, if the filtered list returns clusters that have no devices that are configured within z/OS, all clusters in the grid become candidates. The candidate list is not ordered, meaning that all candidate clusters are viewed as equals and all clusters that are excluded from the list are not candidates.

Because this function introduces system burden into the z/OS scratch mount path, a new **LIBRARY REQUEST** option is introduced to globally enable or disable the function across the entire multi-cluster grid. SAA is disabled, by default. When this option is enabled, the z/OS JES software obtains the candidate list of mount clusters from a given composite library.

Use the **LIBRARY REQUEST, GRID[1] / [2], SETTING, DEVALLOC, SCRATCH, ENABLE** command to enable SAA. All clusters in the multi-cluster grid must be at R2.0 level before SAA is operational. A supporting z/OS APAR OA32957 is required to use SAA in a JES2 environment of z/OS. For JES3, the minimum supported release is z/OS R2.1. Any z/OS environment with earlier code can exist, but it continues to function in the traditional way regarding scratch allocations.

Assume that there are two main workloads. The *application* workload consists of logical volumes that are created and then retrieved on a regular, daily, weekly, or monthly basis. This workload can best be placed in the TS7700D deep cache. The *backup* workload is normally never retrieved and can best be placed directly in the TS7740 Cluster 1. SAA helps direct the mount point to the most efficient cluster for the workload:

- ▶ The application workload can best be set up in the following manner. In the MC construct, the MC is defined with a Copy Consistency Point of [R,D,D]. Cluster 0 is selected in all clusters as Scratch Mount Candidate. In Cluster 1, the SC can best be set as TVC Preference Level 1. This is advised because in cases where Cluster 0 is not available or no online devices are available in that cluster, Cluster 1 can be activated as the mount point. Cluster 2 can set Preference Level 0.

You can control the placement in cache per cluster by setting the **SETTING CACHE COPYFSC** option. When the **ENABLE** keyword is specified, the logical volumes that are copied into the cache from a peer TS7700 cluster are managed by using the actions that are defined for the SC construct associated with the volume as defined at the TS7740 cluster receiving the copy. The SC in Cluster 0 needs to have a Volume Copy Retention Group of Prefer Keep. Logical volumes can be removed from the TS7700D deep cache if more space is needed.

- ▶ The Backup workload can best be set up in the following manner. In the MC construct, the MC is defined with a Copy Consistency Point of [D,R,D] or [N,R,D]. Cluster 1 is selected in all clusters as Scratch Mount Candidate. In Cluster 1 and Cluster 2, the SC can best be set as TVC Preference Level 0. There is no need to keep the data in cache.

The SC in Cluster 0 can have a Volume Retention Group of Prefer Remove. If Cluster 0 is activated as mount point because of the unavailability of Cluster 1 or because there are no online devices in that cluster, the logical volumes with this MC can be removed first when cache removal policies in the TS7700D require the removal of volumes from cache.

With these definitions, the scratch allocations for the application workload are directed to TS7700D Cluster 0 and the scratch allocations for the Backup workload are directed to TS7740 Cluster 1. The devices of the remote clusters in the Disaster Site are not online. Allocation "BYDEVICES" is used. GRID1 has in total 60 devices online and GRID2 has 40 devices online. For each grid, the distribution of online devices is now not determined within the grid by the number of online devices, as in the scenario BYDEVICES, but by the SAA setting of the MC.

The expected distribution of the scratch allocations is shown in Figure 11-54.

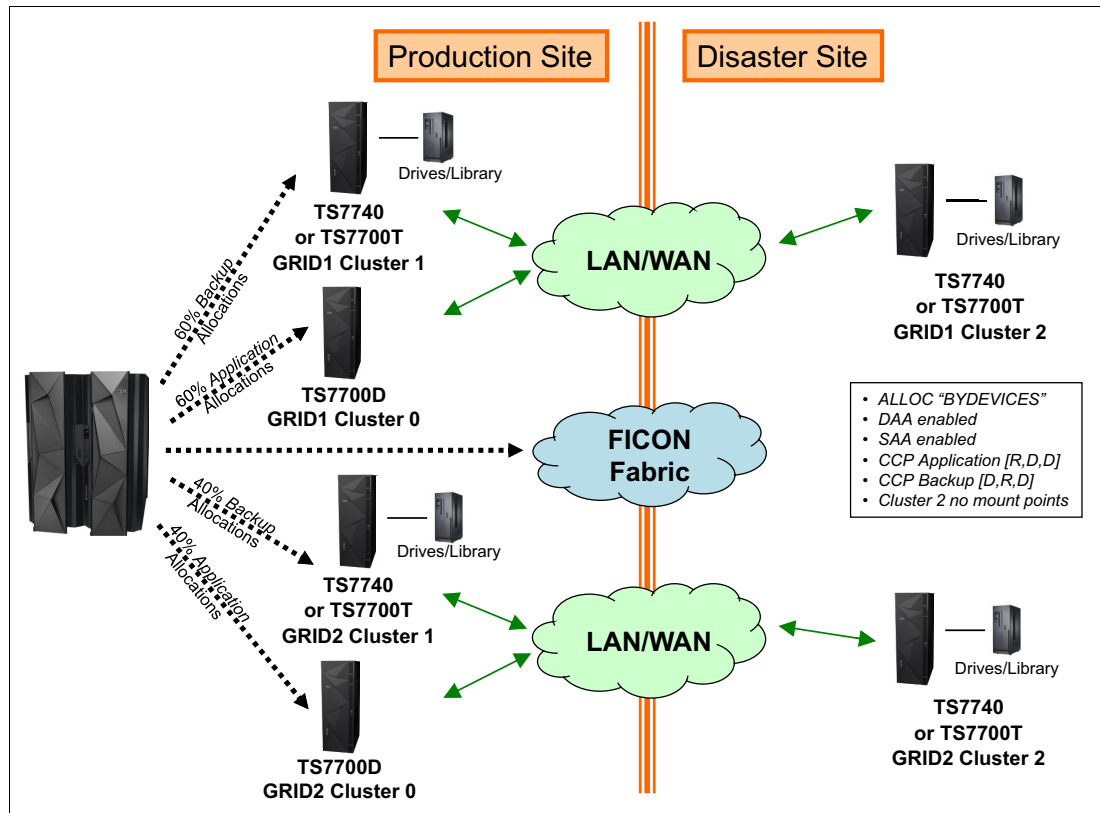


Figure 11-54 Allocation and SAA

Clusters not included in the list are *never* used for scratch mounts unless those clusters are the only clusters that are known to be available and configured to the host. If all candidate clusters have either all their devices varied offline to the host or have too few devices varied online, z/OS will not revert to devices within non-candidate clusters. Instead, the host goes into allocation recovery. In allocation recovery, the existing z/OS allocation options for device allocation recovery (WTOR | WAITHOLD | WAITNOH | CANCEL) are used.

Any time that a service outage of candidate clusters is expected, the SAA function needs to be disabled during the entire outage by using the **LIBRARY REQUEST** command. If left enabled, the devices that are varied offline can result in zero candidate devices, causing z/OS to enter the allocation recovery mode. After the cluster or clusters are again available and their devices are varied back online to the host, SAA can be enabled again.

If you vary offline too many devices within the candidate cluster list, z/OS might have too few devices to contain all concurrent scratch allocations. When many devices are taken offline, first disable SAA by using the **LIBRARY REQUEST** command and then re-enable SAA after they have been varied back on.

If you plan to have an alternative MC setup for the Disaster Site (perhaps for the Disaster Test LPARs), carefully plan the MC settings, the device ranges that need to be online, and whether SAA will be used. Read production data and create test data that uses a separate category code. If you use the same MC as used in the Production LPAR and define in Cluster 2 the MC with SAA for Cluster 2 and not for Cluster 0 or 1 (as determined by the type of workload), Cluster 2 might be selected for allocations in the Production LPARs.

SAA randomly selects one of the clusters for determining the scratch mount candidate clusters in the MC constructs. Therefore, the devices in Cluster 2 must not be made available to the Production LPARs and the devices in the clusters in the Production Site must not be made available in the Disaster Site.

Furthermore, the Copy Consistency Point for the MCs in the Disaster Site can be defined as [D,D,R] or even [N,N,R] if it is test only. If it is kept equal with the setting in the Production Site, with an [R] for Cluster 0 or Cluster 1, cross-cluster mount might occur. If you do not want the grid links overloaded with test data, update the Copy Consistency Point setting or use the TS7700 Virtualization override setting “Prefer Local Cache for Fast Ready Mount Requests” in Cluster 2 in the Disaster Site.

Cross-cluster mounts are avoided, but the copy to Cluster 0 or 1 is still made before the job ends, caused by the Production R(un) Copy Consistency Point setting for these clusters. By further defining a family for the Production Site clusters, the clusters source their copies from the other clusters in the Production Site location, optimizing the usage of the remote links between the locations.

Be aware that SAA is only influencing the mount processing. If you have multiple clusters defined as SAA, but the cluster not available is defined as the only TVC cluster (for example, [R,N,N] or [N,N,R]) then the job is not able to run. The mount will be processed, but the job will hang, because no TVC can be selected.



Copy Export

This chapter explains the Copy Export and Copy Export Recovery functions and how to use them.

This chapter includes the following sections:

- ▶ Copy Export overview and considerations
- ▶ Implementing and running Copy Export
- ▶ Using Copy Export Recovery

12.1 Copy Export overview and considerations

Copy Export provides a function to enable a copy of selected logical volumes that is written to the IBM TS7700 to be removed and taken offsite for disaster recovery purposes. In addition, because the data is a copy of the logical volumes, the volumes remain intact and are still accessible by the production system.

Control of Copy Export

Storage Group (SG) and Management Class (MC) constructs are defined to use separate pools for the primary and secondary copies of the logical volume. The existing MC construct, which is part of Advanced Policy Management (APM), is used to create a secondary copy of the data to be Copy Exported. The MC actions are configured through the TS7700 Management Interface (MI).

An option on the MI window enables designation of a secondary pool as a Copy Export pool. As logical volumes are written, the secondary copy of the data is written to stacked volumes in the Copy Export pool.

Workflow of a Copy Export process

Typically, you run the Copy Export operation on a periodic basis. Because the purpose is to get a copy of the data offsite for disaster recovery (DR) purposes, performing it soon after the data is created minimizes the time for the recovery point objective (RPO). When the time comes to initiate a Copy Export, a Copy Export job is run from the production host.

The TS7700 pre-migrates any logical volumes in the Copy Export pool that have not been pre-migrated. Any new logical volumes that are written after the Copy Export operation is initiated are not included in the Copy Export set of physical volumes. These volumes will be copy exported in the next run, because the Copy Export is an incremental process. Therefore you need all Copy Export physical volumes from all Copy Export operations to do a full recovery. In each Copy Export session, the TS7700 writes a complete TS7700 database to each of the physical volumes in the Copy Export set. It is possible to select to write the database backup to all of the physical volumes, or to limited number of physical volumes. For recovery, you should use the database from the last Copy Export session.

During a Copy Export operation, all of the physical volumes with active data on them in a specified secondary pool are candidates to be exported. Only the logical volumes that are valid on that TS7700 are considered during the running of the operation. Logical volumes that are currently mounted during a Copy Export operation are excluded from the export set, as are any volumes that are not currently in the Tape Volume Cache (TVC) of the export cluster.

The host that initiates the Copy Export operation first creates a dedicated *export list volume* on the TS7700 that runs the operation. The export list volume contains instructions about the execution of the operation, and a *reserved file* that the TS7700 uses to provide completion status and export operation information.

As part of the Copy Export operation, the TS7700 creates response records in the reserved file. These records list the logical volumes that are exported and the physical volumes on which they are located. This information can be used as a record for the data that is offsite. The TS7700 also writes records in the reserved file on the export list volume that provide the status for all physical volumes with a state of Copy Exported.

The Copy Export job can specify whether the stacked volumes in the Copy Export set must be ejected immediately or placed into the export-hold category. When Copy Export is used with the export-hold category, you need to manually request the ejection of the export-hold volumes.

The choice to eject as part of the Copy Export job or to eject them later from the export-hold category is based on your operational procedures. The ejected Copy Export set is then transported to a disaster recovery site or vault. Your RPO determines the frequency of the Copy Export operation.

In heterogeneous drive configurations, the previous generation of drives is normally used for read-only operations. However, the Copy Export operation uses previous generation of 3592 tape drives to append the DB backup to physical volumes so that previous generation of cartridges can also be exported.

12.1.1 General considerations for Copy Export

Consider the following information when you are planning to use the Copy Export function for disaster recovery:

- ▶ Both TS7740 and TS7700T support the Copy Export function. If a Copy Export Recovery is needed, the Copy Export sets can be used to restore data at a location that has equal or newer TS7700 microcode with physical tape drives that can read a Copy Export set of physical volumes. A TS7700T Copy Export set can be restored into both TS7740 and TS7700T. A TS7740 Copy Export set can also be restored into both TS7740 and TS7700T.
- ▶ When using the Copy Export acceleration (LMTDBPVL) option, the database backup is appended only to the first two and the last two volumes that are exported. These corresponding tapes with the database backup are selected and listed in the alphabetical order of the physical tape VOLSER. If the LMTDBPVL option was set, and there is a failure appending the DB backup, a different physical volume is selected to contain the database backup so that four physical volumes have the DB backup. The LMTDBPVL option can be specified through either of the following JCL:
 - **OPTIONS1,COPY,LMTDBPVL** (volumes will be marked as export hold)
 - **OPTIONS1,COPY,EJECT,LMTDBPVL** or **OPTIONS1,COPY,LMTDBPVL,EJECT** (volumes will be ejected from the library)
- ▶ The Copy Export operation might fail depending on the combination of installed tape drives, media types, or recording formats of physical volumes in the secondary pool, and the existence of LMTDBPVL option. Exportable physical volumes are explained in detail in Table 12-1 on page 761. If unexportable physical volumes are included in the Copy Export set physical volumes, the Copy Export operation fails and returns message CBR3856I.
- ▶ Specific logical volumes are not specified as part of a Copy Export operation. Instead, all valid logical volumes on the physical volumes in the specified secondary pool are considered for export. After the first time that Copy Export is performed for a pool, the logical volumes that will be exported are the ones for that pool that have been newly written or modified since the last export began.

Previously exported volumes that have not been changed are not exported. For recovery, all exported physical volumes that still contain active data from a source TS7700 need to be included because not all of the logical volumes that are created are going to be on the last set exported.

- ▶ The primary copy of the logical volumes that is exported remains in the inventory of the TS7700 grid. Exported volumes are *always* copies of volumes still in the TS7700.

- ▶ Only those logical volumes that are assigned to the secondary pool that is specified in the export list volume that is resident on a physical volume of the pool or in the cache of the TS7700 performing the export operation are considered for export. For a grid configuration, if a logical volume is to be copied to the TS7700 that is performing the Copy Export operation, but that copy had not yet completed when the export is initiated, it is not 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, are copied to stacked volumes in the secondary pool as part of the Copy Export operation.
- ▶ If the logical volumes are assigned to CP0 in TS7700T, they are resident only and never copied to any physical volumes. Therefore, logical volumes in CP0 cannot be exported.
- ▶ Any logical volume that is assigned to the specified secondary pool in the TS7700 after the Copy Export operation is initiated is not part of the export, and is written to a physical volume in the pool but is not exported. This includes host-sourced and copy-sourced data.
- ▶ Logical volumes that are currently mounted cannot be Copy Exported.
- ▶ Only one Copy Export operation can be performed at a time.
- ▶ Only one secondary physical volume pool can be specified per export operation, and it must have been previously defined as a Copy Export pool.
- ▶ The export list volume cannot be assigned to the secondary pool that is specified for the operation. If it is, the Copy Export operation fails.
- ▶ During the execution of a Copy Export operation, if the TS7700 cannot access the primary copy and the secondary copy exists in a pool that is defined for the Copy Export, that secondary version is made inaccessible and the mount fails. This occurs regardless of whether that secondary pool is involved in the current Copy Export operation.

The library that is associated with the TS7700 running the Copy Export operation must have an I/O station feature for the operation to be accepted. Empty the I/O station before running Copy Export and prevent it from going to the full state.

- ▶ A minimum of four physical tape drives must be available to the TS7700 for the Copy Export operation to be performed. The operation is terminated by the TS7700 when fewer than four physical tape drives are available. In heterogeneous drive configurations, the operation is terminated when fewer than four 3592-E08 drives are available. Processing for the physical stacked volume in progress when the condition occurred is completed, and the export status file records reflect what was completed before the operation was terminated.
- ▶ Copy Export and the insertion of logical volumes are mutually exclusive functions in a TS7700 or grid.
- ▶ 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 waits 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 ended.
- ▶ During execution, if the TS7700 determines that a physical volume that is assigned to the specified secondary pool contains one or more primary logical volumes, that physical volume and any secondary logical volumes on it are excluded from the Copy Export operation.

- ▶ To minimize the number of physical volumes that are used for Copy Export, use the highest capacity media and physical drive format that is compatible with the recovery TS7700. You might also want to reduce the number of concurrent tape devices that the TS7700 uses when copying data from cache to the physical volumes in secondary pool used for Copy Export. You can change it using the Maximum Devices in Pool Properties in the MI.
- ▶ 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. The default limit is 2000 per TS7700 to ensure that they all fit into the receiving library. This value can be adjusted to a maximum of 10,000 volumes by using Copy Export Settings in MI.
- ▶ The recovery TS7700 must have physical tape drives that can read the physical volumes from a source TS7700. If a source TS7700 writes the volumes by using the native E08 format, the recovery TS7700 must also have 3592-E08 drives running in native format mode. If a source TS7700 writes the JB volumes by using the native E07 format, the recovery TS7700 must also have 3592-E07 drives. If a source TS7700 writes the JC or JK volumes by using the native E07 format, the recovery TS7700 must have 3592-E07 or E08 drives.

If the exporting pool on the source TS7700 is set up to encrypt the data, the recovery TS7700 must also be set up to handle encrypted volumes and have access to the IBM Encryption Key Manager with replicated keys from the production site. If the source TS7700 writes the volumes in J1A or emulated J1A mode, the recovery TS7700 must have 3592-E07 or the previous generations 3592 model drives.

- ▶ There is an available service offering where a customer can merge data from more than one source TS7700 Copy Export backup. However, basically, the recovery TS7700 cannot contain any previous data, and a client-initiated recovery process cannot merge data from more than one source TS7700 together. As a part of the Copy Export Recovery, an option is provided to erase any previous data on the TS7700. This enables a TS7700 that is used for disaster recovery testing to be reused for testing of a different source TS7700's data.
- ▶ For the secondary pool that is used for Copy Export, the designated reclaim pool must not be the same pool as the primary pool or its reclaim pool.

Note: If the reclaim pool for the Copy Export pool is the same as either the Copy Export primary pool or its reclaim pool, the primary and backup copies of a logical volume can exist on the same physical tape.

Table 12-1 shows the exportable physical volumes based on tape drives, export format, and the LMTDBPVL option.

Table 12-1 Exportable physical volumes based on tape drives, export format, and LMTDBPVL option

Installed tape drives	Export format of copy export pool	LMTDBPVL enabled?	Exportable physical volumes
E08 only	Default	Yes or no	JC/JK in E07 format and JC/JK/JD/JL in E08 format
	E08		JC/JK/JD/JL in E08 format
	E07		JC/JK in E07 format

Installed tape drives	Export format of copy export pool	LMTDBPVL enabled?	Exportable physical volumes
E08 and E07	Default	Yes	Any media types in any recording format
	E08		JC/JK/JD/JL in E08 format
	E07		JB/JC/JK in E07 format
	Default	No	JB in E06 format, JB/JC/JK in E07 format, and JC/JK/JD/JL in E08 format
	E08		JC/JK/JD/JL in E08 format
	E07		JB/JC/JK in E07 format
E08 and E06	Default	Yes	JA/JJ in J1A format, JA/JJ/JB in E05/E06 format, JC/JK in E07 format, and JC/JK/JD/JL in E08 format
	E08		JC/JK/JD/JL in E08 format
	E07		JC/JK in E07 format
	Default	No	JA/JJ/JB in E05/E06 format, JC/JK in E07 format, and JC/JK/JD/JL in E08 format
	E08		JC/JK/JD/JL in E08 format
	E07		JC/JK in E07 format
E08 and E05	Default	Yes or no	JA/JJ in J1A format, JA/JJ/JB in E05 format, JC/JK in E07 format, and JC/JK/JD/JL in E08 format
	E08		JC/JK/JD/JL in E08 format
	E07		JC/JK in E07 format
E08 and J1A	Default	Yes or no	JA/JJ in J1A format, JC/JK in E07 format, and JC/JK/JD/JL in E08 format
	E08		JC/JK/JD/JL in E08 format
	E07		JC/JK in E07 format
E07 only	Default	Yes or no	JB in E06 format and JB/JC/JK in E07 format
	E07		JB/JC/JK in E07 format
	E06		JB in E06 format
E06 only	Default	Yes or no	JA/JJ/JB in E05/E06 format
E05 only	Default	Yes or no	JA/JJ in J1A format and JA/JJ/JB in E05 format
J1A only	Default	Yes or no	JA/JJ in J1A format

12.1.2 Copy Export grid considerations

Copy Export is supported in both grid and stand-alone environments. You need to remember several considerations that are unique to the grid environment.

Performing Copy Export

The first consideration relates to performing Copy Export. In a grid configuration, a Copy Export operation is performed against an individual TS7700, not across all TS7700 clusters. Set up Copy Export in a grid plan based on the following guidelines:

- ▶ Decide which TS7700 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 an MC name. On the TS7700 that is used to export that data, define a secondary physical volume pool for that MC name and also ensure that you indicate that it is an export pool. Although you need to define the MC name on all TS7700s in the grid configuration, specify only the secondary physical volume pool on the one 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 exceptional useful case of this approach is if you want one of the TS7700s in the grid configuration to have a secondary copy of the data for the case the primary copies on other TS7700s are inaccessible.

- ▶ While you are defining the MC name for the data, also ensure that the TS7700 to perform the export operation has a copy policy that specifies that it is to have a copy.
- ▶ When the Copy Export operation is run, the export list volume must be valid only on the TS7700 that is performing the operation. You must define a unique MC name to be used for the export list volume. For that MC name, you must define its copy policy so that a copy is only on the TS7700 that is to perform the export operation. If the VOLSER that is specified for the export list volume when the export operation is initiated is on more than one TS7700, the Copy Export operation fails.

Tip: If the MC specified for the Copy Export operation is defined to more than one cluster, the Copy Export fails and the following CBR message is displayed:

```
CBR3726I FUNCTION INCOMPATIBLE ERROR CODE 32 FROM LIBRARY XXX FOR VOLUME  
xxxxxx.
```

```
X'32' There is more than one valid copy of the specified export list volume in  
the TS7700 grid configuration.
```

Consider this Copy Export example:

- a. A Copy Export with the export list volume EXP000 is initiated from a host that is 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 that is attached to C0, a copy of EXP000 is created both on the C0 and C1.

- d. The 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.

In the previous 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 need to set up an MC for the data that is to be exported so that it has a copy on Cluster 1 and a secondary copy in pool 8. To ensure that the data is on that cluster and is consistent with the close of the logical volume, you want to have a copy policy of Rewind Unload (RUN). You define the following information:

- ▶ Define an MC, for example, MCCEDATA, on Cluster 1:

```

Secondary Pool      8
Cluster 0 Copy Policy  RUN
Cluster 1 Copy Policy  RUN
  
```

- ▶ Define this same MC on Cluster 0 without specifying a secondary pool.
- ▶ To ensure that the export list volume is written to Cluster 1 and exists only there, define an MC, for example, MCELFVOL, on Cluster 1:

```

Cluster 0 Copy Policy  No Copy
Cluster 1 Copy Policy  RUN
  
```

- ▶ Define this MC on Cluster 0:

```

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 is internally routed to the TS7700 that has the valid copy of the specified export list volume. Operational and completion status is broadcast to all hosts attached to all of the TS7700s in the grid configuration.

It is assumed that Copy Export is performed regularly, and logical volumes whose copies were not complete when a Copy Export was initiated will be exported the next time that 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 before initiating the operation by using the Volume Status function of the Bulk Volume Information Retrieval (BVIR) facility. You can then be sure that all critical volumes are exported during the operation.

Performing Copy Export Recovery

The next consideration relates to how Copy Export Recovery is performed. Although there is an available service offering where a customer can merge data from more than one source TS7700 Copy Export backup, Copy Export Recovery is basically to a stand-alone empty TS7700. As part of a client-initiated 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 has volume ownership of all volumes.

It is possible that one or more logical volumes might 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 before the start of the operation. Each copy-exported physical volume remains under the management of the TS7700 from which it was exported.

Normally, you return the empty physical volumes to the library I/O station that associated with the source TS7700. They are then 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 define for that TS7700 match the VOLSERs of the physical volumes that you want to move.
- ▶ Have the original TS7700 stop managing the copy-exported volumes by entering the following command from the host:

```
LIBRARY REQUEST, libname, COPYEXP, volser, DELETE
```

12.1.3 Reclaim process for Copy Export physical volumes

The physical volumes that are exported during a Copy Export operation continue to be managed by the source TS7700 regarding 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 decreases until the physical volume becomes eligible for reclamation based on your provided criteria for its pool.

Figure 12-1 on page 766 shows how the Reclaim Threshold Percentage is set in Physical Volume Pool Properties. If the ratio between active data size and total bytes written to the physical volume is lower than the Reclaim Threshold Percentage, the physical volume becomes eligible for reclamation. The ratio between active data size and media capacity is not used for the comparison with Reclaim Threshold Percentage.

Exported physical volumes that are to be reclaimed are not brought back to the source TS7700 for processing. Instead, a new secondary copy of the remaining valid logical volumes is made by using the primary logical volume copy as a source. It is called *Offsite Reclaim*. Offsite Reclaim does not start while Copy Export is running, and follows Inhibit Reclaim Schedule. If there is more than one volume eligible for Offsite Reclaim, it tries to make the exported physical volumes EMPTY one by one in the ascending order of their active data size.

There is another kind of Offsite Reclaim called Priority Offsite Reclaim. It is Offsite Reclaim for user-specified exported physical volumes. Users can make an exported physical volume eligible for Priority Offsite Reclaim by issuing the following library request from the host: **LIBRARY REQUEST, *libname*, COPYEXP, *volser*, RECLAIM**. Priority Offsite Reclaim processing is the same as for normal Offsite Reclaim, but it runs in priority over normal Offsite Reclaim, and does not follow Inhibit Offsite Reclaim Schedule.

Figure 12-1 shows the Reclaim Threshold Percentage for a normal Offsite Reclaim.

The screenshot shows the configuration window for Physical Volume Pools. The title bar indicates the cluster is #BA99E (Cluster 5) and the current view is Physical Volume Pools. A dropdown menu shows 'Pool 9' is selected. Below this, the 'Pool Properties' section contains a table of settings:

Media Class:	3592
First Media (Primary):	Any 3592
Second Media (Secondary):	None
Borrow Indicator:	Borrow, Return
Reclaim Pool:	9
Maximum Devices:	All Compatible Devices
Export Pool:	Not Defined
Export Format:	Default
<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
Sunset Media Reclaim Threshold Percentage (%):	35

At the bottom of the window are 'OK' and 'Cancel' buttons.

Figure 12-1 Reclaim Threshold Percentage is set in Physical Volume Pool Properties

The next time that the Copy Export operation is performed, the physical volumes with the new copies are also exported. After the Copy Export completes, the physical volumes that were reclaimed (which are offsite) are no longer considered to have valid data (empty), and can be returned to the source TS7700 to be used as new scratch volumes.

Tip: If a physical volume is in Copy Export hold state and becomes empty, it is automatically moved back to the common scratch pool (or the defined reclamation pool) when the next Copy Export operation completes.

Monitoring for Copy Export data

The 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.

12.1.4 Copy Export process messages

During the execution of the Copy Export operation, the TS7700 sends informational messages to its attached hosts. These messages are in the syslog and are shown in Table 12-2.

Note: All messages are prefaced with CBR3750I.

Table 12-2 SYSLOG messages from the library

Message description	Action needed
E0000 EXPORT OPERATION STARTED FOR EXPORT LIST VOLUME XXXXXX This message is generated when the TS7700 begins the Copy Export operation.	None.
E0002 OPENING EXPORT LIST VOLUME XXXXXX FAILED This message is generated when opening the export list volume failed during the Copy Export operation.	Check whether the export list volume or cache file system is in a bad state.
E0005 ALL EXPORT PROCESSING COMPLETED FOR EXPORT LIST VOLUME XXXXXX This message is generated when the TS7700 completes an export operation.	None.
E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL IN EXPORT-HOLD This message is generated during Copy Export operations when an exported stacked volume 'YYYYYY' has been assigned to the export-hold category. The 'LLLLLLLL' field is replaced with the distributed library name of the TS7700 performing the export operation.	None.
E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL IN EJECT This message is generated during Copy Export operations when an exported stacked volume 'YYYYYY' has been assigned to the eject category. The physical volume is placed in the convenience I/O station. The 'LLLLLLLL' field is replaced with the distributed library name of the TS7700 performing the export operation.	Remove ejected volumes from the convenience I/O station.
E0013 EXPORT PROCESSING SUSPENDED, WAITING FOR SCRATCH VOLUME This message is generated every 5 minutes when the TS7700 needs a scratch stacked volume to continue export processing and there are none available.	Make one or more physical scratch volumes available to the TS7700 performing the export operation. If the TS7700 does not get access to a scratch stacked volume in 60 minutes, the operation is ended.
E0014 EXPORT PROCESSING RESUMED, SCRATCH VOLUME MADE AVAILABLE This message is generated when, after the export operation was suspended because no scratch stacked volumes were available, scratch stacked volumes are again available and the export operation can continue.	None.
E0015 EXPORT PROCESSING TERMINATED, WAITING FOR SCRATCH VOLUME This message is generated when the TS7700 ends the export operation because scratch stacked volumes were not made available to the TS7700 within 60 minutes of the first E0013 message.	The operator must make more TS7700 stacked volumes available, perform an analysis of the export status file on the export list volume, and reissue the export operation.

Message description	Action needed
<p>E0016 COPYING LOGICAL EXPORT VOLUMES FROM CACHE TO STACKED VOLUMES</p> <p>This message is generated when the TS7700 begins, and every 10 minutes during, the process of copying logical volumes that are only resident in the TVC to physical volumes in the specified secondary physical volume pool.</p>	None.
<p>E0017 COMPLETED COPY OF LOGICAL EXPORT VOLUMES TO STACKED VOLUMES</p> <p>This message is generated when the TS7700 completes the copy of all needed logical volumes from cache to physical volumes in the specified secondary physical volume pool.</p>	None.
<p>E0018 EXPORT TERMINATED, EXCESSIVE TIME FOR COPY TO STACKED VOLUMES</p> <p>The export process has been ended because one or more cache resident-only logical volumes that are needed for the export were unable to be copied to physical volumes in the specified secondary physical volume pool within a 10-hour period from the beginning of the export operation.</p>	Call for IBM support.
<p>E0019 EXPORT PROCESSING STARTED FOR POOL XX</p> <p>This message is generated when the TS7700 export processing for the specified secondary physical volume pool XX.</p>	None.
<p>E0020 EXPORT PROCESSING COMPLETED FOR POOL XX</p> <p>This message is generated when the TS7700 completes processing for the specified secondary physical volume pool XX.</p>	None.
<p>E0021 DB BACKUP WRITTEN TO STACKED VOLUMES, PVOL01, PVOL02, PVOL03, PVOL04</p> <p>(Where PVOL01, PVOL02, PVOL03, and PVOL04 are the physical volumes to which the database backup was appended).</p> <p>This message is generated if the Copy Export acceleration (LMTDBPVL) option was selected on the export.</p>	None.
<p>E0022 EXPORT RECOVERY STARTED</p> <p>The export operation has been interrupted by a TS7700 error or a power off condition. When the TS7700 is restarted, it attempts recovery of the operation.</p>	None.
<p>E0023 EXPORT RECOVERY COMPLETED</p> <p>The recovery attempt for interruption of an export operation has been completed.</p>	Perform an analysis of the export status file on the export list volume and reissue the export operation, if necessary.
<p>E0024 XXXXXX LOGICAL VOLUME WITH INVALID COPY ON LLLLLLLL</p> <p>This message is generated when the TS7700 performing the export operation has determined that one or more (XXXXXX) logical volumes that are associated with the auxiliary storage pool that is specified in the export list file do not have a valid copy resident on the TS7700. The 'LLLLLLLL' field is replaced by the distributed library name of the TS7700 performing the export operation. The export operation continues with the valid copies.</p>	When the export operation completes, perform an analysis of the export status file on the export list volume to determine the logical volumes that were not exported. Ensure that they have completed their copy operations and then perform another export operation.

Message description	Action needed
<p>E0025 PHYSICAL VOLUME XXXXXX NOT EXPORTED, PRIMARY COPY FOR YYYYYY UNAVAILABLE</p> <p>This message is generated when the TS7700 detected a migrated-state logical volume 'YYYYYY' with an unavailable primary copy. The physical volume 'XXXXXX' on which the secondary copy of the logical volume 'YYYYYY' is stored was not exported.</p> <p>This message is added at code level R1.7.</p>	<p>The logical volume and the physical volume will be eligible for the next Copy Export operation after the logical volume is mounted and unmounted from the host. An operator intervention is also posted.</p>
<p>E0026 DB BACKUP WRITTEN TO ALL OF STACKED VOLUMES</p> <p>This message is generated when the Copy Export acceleration (LMTDBPVL) option is <i>not</i> selected.</p>	<p>None.</p>
<p>R0000 RECLAIM SUCCESSFUL FOR EXPORTED STACKED VOLUME YYYYYY</p> <p>This message is generated when the TS7700 has successfully completed reclaim processing for an exported stacked volume that was exported during a previous copy export operation.</p> <p>Note: A copy exported physical volume can become eligible for reclaim based on the reclaim policies that are defined for its secondary physical volume pool, or through the host console request command.</p>	<p>The exported physical volume no longer contains active data and can be returned from its offsite location for reuse. If it is placed in export-hold, it should be returned when the next copy export is completed.</p>

When a stacked volume that is associated with a Copy Export operation is ejected from a library (placed in export-hold or is physically ejected from the library), you see status message E0006, which is sent by the library (see Table 12-2 on page 767). Removable Media Management (RMM) intercepts this message and performs one of these actions:

- ▶ If the stacked volume is predefined to RMM, RMM marks the volume as ejected or in-transit and sets the movement/store date that is associated with the stacked volume.
- ▶ If the stacked volume is not predefined to RMM and the STACKEDVOLUME(YES) option in RMM is specified, RMM automatically adds the stacked volume to its control data set (CDS).

To have DFSMSrmm policy management manage the retention and movement for volumes that are created by Copy Export processing, you must define one or more volume vital record specifications (VRSs). For example, assume that all Copy Exports are targeted to a range of volumes STE000 - STE999. You can define a VRS as shown in Example 12-1.

Example 12-1 VRS definition

```
RMM AS VOLUME(STE*) COUNT(99999) LOCATION(location)
```

As a result, all matching stacked volumes that are set in AUTOMOVE 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 listed, a copy-exported stacked volume can become eligible for reclamation based on the reclaim policies that are 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 offsite location for reuse.

For users that use DFSMSrmm, when you have stacked volume support that is enabled, DFSMSrmm automatically handles and tracks the stacked volumes that are created by Copy Export. However, there is no way to track which logical volume copies are on the stacked volume. Retain the updated export list file, which you created and the library updated, so that you have a record of the logical volumes that were exported, and on what exported stacked volume they are.

For more information and error messages that are related to the Copy Export function in RMM, see the *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874.

12.2 Implementing and running Copy Export

Implementing and running Copy Export are described. For more information and error messages that relate to the Copy Export function, see the *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide*, which is available at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092>

12.2.1 Setting up data management definitions

To set up the data management definitions, perform the following steps:

1. Decide the MC construct name (or names).

As part of the plan for using the Copy Export function, you must decide on at least one MC construct name. A preferred practice is to make the name meaningful and relate to the type of data to be on the pool or the location where the data is sent.

For example, if the pool is used to send data to the primary disaster recovery site in Atlanta, a name like MCPRIATL can be used. "MC" indicates MC, "PRI" indicates that it is for the primary recovery site, and "ATL" indicates Atlanta. Up to an eight-character name can be defined.

2. Define the MC names to DFSMS.

After the MC names are selected, the names must be defined to DFSMS and to the TS7700. For details about defining the MC in DFSMS, see *z/OS DFSMSdfp Storage Administration*, SC23-6868.

None of the settings are used for system-managed tape. All settings that are associated with an MC name are defined through the TS7700, not the DFSMS windows.

3. Define the MC names to the TS7700.

You must also define the MC names on the TS7700 because you are not using the Default MC settings for Copy Export volumes. Define a Secondary Pool for the copies to be exported.

For details about how to add an MC, see "Management Classes window" on page 439.

4. Define the VOLSER ranges for the 3592 media.

You must define the VOLSER range (or ranges) for the physical volumes to use for Copy Export if you plan to use a specific VOLSER range. Ensure that you define the same pool that you used in the MC definition as the Home Pool for this VOLSER range.

Tip: 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 can simplify the task of knowing the volumes to use in recovery, and of returning a volume that no longer has active data on it to the TS7700 that manages it.

For details about how to define the VOLSER ranges, see “Defining VOLSER ranges for physical volumes” on page 529.

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 MC definition, and, optionally, in the VOLSER range definition, select **Copy Export** in the Export Pool field.

For more information about how to change the physical volume pool properties, see “Defining physical volume pools in the TS7700T” on page 531.

6. Code or modify the MC automatic class selection (ACS) routine.

Add selection logic to the MC ACS routine to assign the new MC name, or names.

7. Activate the new construct names and ACS routines.

Before new allocations are assigned to the new MC, the Source Control Data Set (SCDS) with the new MC definitions and ACS routines must be activated by using the **SETSMS SCDS** command.

12.2.2 Validating before activating the Copy Export function

Before the logical volumes are exported, you must perform several general validations. Before you initiate the operation, check that the TS7700 has the required physical drives and scratch physical volume resources. Verify that the TS7700 is not near the limit of the number of physical volumes that can have a status of Copy Exported and modify the value, if required. Depending on your production environment, you might want to automate these validation steps.

Follow these validation steps:

1. Check whether data is in an older format. If you migrated from a B10 or B20 VTS to the TS7700 by using the outboard migration method, you might have data that is still in the older VTS format. The TS7700 cannot export data in the old format, so you must check whether any of the data to export was written with the old format.
2. Validate that the TS7700 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 that are attached to the TS7700. If fewer than the required numbers of physical drives are available, you must call for service to repair drives before you perform the Copy Export operation.

See Example 12-2 for the output of the **PDRIVE** request. This command is only valid when run against a distributed library.

Example 12-2 Data that is returned by the PDRIVE request

```

LI REQ,BARR03A,PDRIVE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR03A,PDRIVE.
CBR1280I LIBRARY BARR03A REQUEST. 768
KEYWORDS: PDRIVE
-----
PHYSICAL DRIVES V2 .1
SERIAL NUM      TYPE  MODE  AVAIL  ROLE  POOL    PVOL    LVOL
0000078DAD6A   3592E07                Y  IDLE   00
0000078DAD9B   3592E07                Y  IDLE   00
0000078DBA58   3592E08      E08    Y  MIGR   01  JD0402  S00006
0000078DB88A   3592E08      E08    Y  MIGR   02  JC0863  S00102
0000078DB887   3592E08                Y  IDLE   00
0000078DB89C   3592E08                Y  IDLE   00

```

In the response that is shown in Example 12-2, you can see the following information:

- Four E08 drives and two E07 drives are defined.
 - All nine drives are available (AVAIL=Y).
 - The ROLE column describes which drive is performing. The following values can be indicated:
 - IDLE: The drive is not in use for another role or is not mounted.
 - SECE: The drive is being used to erase a physical volume.
 - MIGR: The drive is being used to copy a logical volume from the TVC to a physical volume. In this display, logical volume SO0006 is being copied to physical volume JD0402.
 - RECA: The drive is being used to recall a logical volume from a physical volume to the TVC.
 - RCLS: The drive is being used as the source of a reclaim operation.
 - RCLT: The drive is being used as the target of a reclaim operation.
3. Check that the pool to be exported has sufficient scratch physical volumes and that the TS7700 is under the volume limit for copy-exported volumes in all pools. The limit by default is a total of 2,000 volumes, but this limit can be modified in the **SETTINGS** option of the TS7000 MI with a maximum of 10,000 volumes. You can use the Library Request host console command that specifies the **POOLCNT** request. See Example 12-3 for the response to the **LI REQ, <library-ID>, POOLCNT** command.

Example 12-3 Data that is returned from the POOLCNT command

```

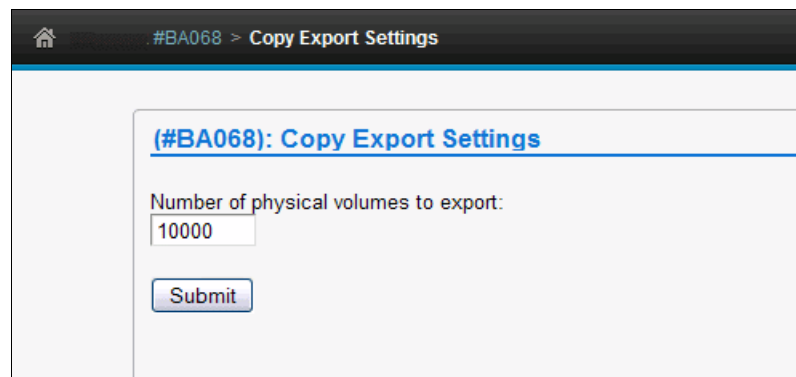
LI REQ,BARR68A,POOLCNT
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,POOLCNT.
CBR1280I LIBRARY BARR68A REQUEST. 919
KEYWORDS: POOLCNT
-----
PHYSICAL MEDIA COUNTS V2
POOL MEDIA  EMPTY  FILLING  FULL  ERASE  ROR  UNAVAIL  CXPT
0   JA    164
0   JJ    38
1   JA     2     6    12     0     0     1     0
9   JJ     0     4    22     0     0     0    45

```

Pool 0 is the Common Scratch Pool. Pool 9 is the pool that is used for Copy Export in this example. Example 12-3 shows the command **POOLCNT**. The response is listed per pool:

- The media type used for each pool
- The number of empty physical volumes that are available for Scratch processing
- 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 (ROR) 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)

Use the MI to modify the maximum-allowed number of volumes in the copy-exported state (Figure 12-2).



The screenshot shows a web-based interface for configuring Copy Export settings. At the top, there is a header bar with a home icon, the text "#BA068 > Copy Export Settings", and a title "#BA068): Copy Export Settings". Below the title, there is a label "Number of physical volumes to export:" followed by a text input field containing the number "10000". A "Submit" button is located below the input field.

Figure 12-2 Maximum allowable number of volumes in copy-exported state

You must determine when you usually want to start the Copy Export operation. Thresholds might be the number of physical scratch volumes or other values that you define. These thresholds can 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 10.1.3, “Host Console Request function” on page 608.

12.2.3 Running the Copy Export operation

To begin the Copy Export process, create an export list volume that provides the TS7700 with information about which data to export and the options to use during the operation (Figure 12-3 on page 775).

If you use a multi-cluster grid, be sure to create the export list volume only on the same TS7700 that is used for Copy Export, but not on the same physical volume pool that is used for Copy Export. If more than one TS7700 in a multi-cluster grid configuration contains the export list volume, the Copy Export operation fails.

Ensure that all volumes that are subject for copy export are in the TVC of the TS7700 where the copy export will be run. If there are copies from other clusters that have not been processed, you can promote them in the copy queue. Use a host console request (HCR) command with the **COPY,KICK** option to do so:

```
LI REQ,distributed library,LVOL,A08760,COPY,KICK
```

Complete these steps to run the Copy Export operation:

1. Create the export list volume JCL (Example 12-4).

Example 12-4 Sample JCL to create an export list volume of Pool 9

```
//*****  
//* 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,LMTDBPVL  
/*  
/* Remove LMTDBPVL to not use accelerate  
  
//*****  
//* 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 that is 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 step 1 in Example 12-4 on page 774. In this example, you are exporting Pool 09. Option EJECT in record 2 tells the TS7700 to eject the stacked volumes upon completion.

With only `OPTIONS1,COPY` (without `EJECT`), the physical volumes are placed in the export-hold category for later handling and left in the library by an operator.

- A reserved file, as created in step 2 in Example 12-4 on page 774. This file is reserved for future use.
 - An export status file, as created in step 3 in Example 12-4 on page 774. In this file, the information is stored from the Copy Export operation. You must keep this file because it contains information related to the result of the Export process and must be reviewed carefully.
- All records must be 80 bytes.
 - The export list file must be written without compression. Therefore, you must assign a Data Class (DC) that specifies `COMPACTION=NO` or you can overwrite the DC specification by coding `TRTCH=NOCOMP` in the JCL.

Important: Ensure that the files are assigned an MC that specifies that only the local TS7700 has a copy of the logical volume. You can either have the ACS routines assign this MC, or you can specify it in the JCL. These files need to have the same expiration dates as the longest of the logical volumes you export because they must be kept for reference.

Figure 12-3 shows the setting of an MC on the MI for the export list volume in a multi-cluster grid configuration. RN means one copy locally at RUN (R) and no copy (NN) on the other cluster.




Figure 12-3 Management Class settings for the export list volume

2. The Copy Export operation is initiated by running the **LIBRARY EXPORT** command. In this command, *logical VOLSER* is a variable, and is the logical volume that is used in creating the export list file.

The command syntax is shown in Example 12-5.

Example 12-5 Library export command

```
LIBRARY EXPORT,logical VOLSER
```

3. The host sends a command to the composite library. From there, it is routed to the TS7700 where the export list volume is.
4. The running TS7700 validates the request, checking for required resources, and if all is acceptable, the Copy Export continues.
5. Logical volumes that are related to the exported pool that still are only in cache can delay the process. They are copied to physical volumes in the pool as part of the Copy Export run.
6. Messages about the progress are sent to the system console. All messages are in the format that is shown in Example 12-6. See Table 12-2 on page 767 for an explanation of Library Message Text.

Example 12-6 Library message format

```
CBR3750I Message from library library-name: message text.
```

After a successful completion, all physical tapes that are related to the export pool are ejected if the EJECT option was specified. The operator can empty the I/O station and transport the tapes to another location.

To obtain a list of the virtual volumes that were exported during the COPY EXPORT operation, use the Physical Volumes Details selection in the MI. Specify the volume or volumes that were written to during the EXPORT. Those VOLSERs are listed in the CBR3750I messages on the syslog. Click **Download List of Virtual Volumes**.

Figure 12-4 shows the physical volume details.

The screenshot shows the 'Physical Volume Details' page for volume YJB012. The page title is '"asika[0]" (#BA97A): Physical Volume Details'. The left sidebar contains a 'Physical Volumes' menu with options: Physical Volume Details (selected), Move Physical Volumes, Eject Physical Volumes, Physical Volume Ranges, Physical Volume Search, and Active Data Distribution. The main content area displays the volume ID 'YJB012' and a 'Get Details' button. Below this is a table of physical volume details.

Physical Volume Details:	
Volser:	YJB012
Type:	JB(ETCL)
Recording Format:	E05
Volume State:	Copy Exported
Capacity State:	Full
Key Label 1:	-
Key Label 2:	-
Encrypted Time:	-
Home Pool:	0
Current Pool:	12
Mount Count:	2
Virtual Volumes Contained:	100
Pending Actions:	-
Copy Export Recovery:	Yes
Database Backup:	20141008094816

At the bottom of the main content area, there is a 'Download List of Virtual Volumes' button.

Figure 12-4 Physical volume details selection for list of exported volumes

Note: The copy export can also be initiated through JCL using the CBRXLCS FUNC=EXPORT programming interface. Refer to SAMPLE members CBRSPPLCS and CBRSPX03, provided in SYS1.SAMPLIB, and documented in the OAM Planning, Installation, and Storage Administration Guide for Tape Libraries:

http://www.ibm.com/support/knowledgecenter/SSLTBW_2.2.0/com.ibm.zos.v2r2.idao300/toc.htm

Sample member CBRSPX03, writes the three required files on the export list volume using a private volume and export list format 03 and has a 4th step (STEP4) that starts CBRSPPLCS to initiate the copy export. CBRSPPLCS is an example program that starts the CBRXLCS programming interface and would need to be modified to suit your business needs. When modified, it would need to be assembled and link-edited on your system, for it to be usable through JCL.

12.2.4 Canceling a Copy Export operation

Examine the export status file records to see what has been processed before 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 canceling 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 and the extra effect of running the export operation is undesirable.
- ▶ 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.

A request to cancel an export operation can be initiated from any host that is attached to the TS7700 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 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 that is attached to the TS7700 that can run the **CANCEL** command, you can cancel the operation through the TS7700 MI. After confirming the selection, a cancel request is sent to the TS7700 that is processing the Copy Export operation.

Regardless of whether the cancellation originates from a host or the MI, the TS7700 can process it in the following manner:

- ▶ 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 is placed in the export-hold or eject category before the cancel processing can continue. The export 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 subsystem processing.
- ▶ A completion message is sent to all hosts that are attached to the TS7700 indicating that the export was canceled by a host request. The message contains information about how much export processing completed before the execution of the cancellation request.

12.2.5 Host completion message

At the completion of the Copy Export operation, a completion message is broadcast to all hosts that are attached to the TS7700. 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 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 that is shown in Example 12-7.

Example 12-7 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 that is shown in Example 12-8.

Example 12-8 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, examine the export status file to determine what errors or exceptions were encountered.

Either of the completion messages provides statistics about what was processed during the operation. The following statistics are reported:

- ▶ Requested-number: This is the number of logical volumes that are associated with the secondary volume pool that is specified in the export list file. Logical volumes that are 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 that is specified in the export list file and it has a valid copy on the TS7700 performing the export. Logical volumes that are associated with the specified secondary volume pool that were previously exported are not considered to be resident in the TS7700.
- ▶ 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.

Clarification: The number of megabytes (MB) exported is the sum of the MB integer values of the data that is 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 R1.4 and later, this value is 0.

It is possible that multiple physical cartridges are written to during the COPY EXPORT even if a small amount of data was exported. This is primarily due to the optimization of the operation by using multiple available drives that are configured for the Copy Export pool.

12.3 Using Copy Export Recovery

The recovery process can be done in a test mode for DR testing purposes. This enables a test restore without compromising the contents of the Copy Export sets. An example of how to use a Copy Export Recovery process is provided.

Consideration: *Clients can run a Copy Export Recovery process only in a stand-alone cluster. After the recovery process completes, you can create a multi-cluster grid by joining the grid with another stand-alone cluster. However, there is an IBM service offering to recover to an existing grid.*

The following instructions for how to implement and run Copy Export Recovery also apply if you are running a DR test. If it is a test, it is specified in each step.

12.3.1 Planning and considerations for testing Copy Export Recovery

You must consider several factors when you prepare a recovery TS7700 for the Copy Export volumes. Copy Export Recovery can be run in various ways. The planning considerations for Copy Export Recovery are described.

Copy Export Recovery can be used to restore previously created and copy-exported tapes to a new, empty TS7700 cluster. The same subset of tapes can be used to restore a TS7700 in an existing grid if the new empty restore cluster replaces the source cluster that is no longer present.

This enables data that might have existed only within a TS7740 or TS7700T in 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.

Client-initiated Copy Export Recovery

Client-initiated recovery restores copy-exported tapes to a stand-alone TS7700 for DR testing or as a recovery site. The considerations for Copy Export Recovery to a stand-alone TS7700 cluster, which can be prepared in advance, are described. The TS7700 and associated library that are to be used for recovery of the copy-exported logical volumes must meet the following requirements:

- ▶ The recovery TS7700 must have physical tape drives that match the capabilities of the source TS7700, including encryption capability if the copy-exported physical volumes are encrypted.
- ▶ If the source copy-exported volumes are encrypted, the recovery TS7700 must have access to a key manager that has the EKs for the data.
- ▶ There must be enough library storage slots in the library that is associated with the recovery TS7700 to hold all of the copy-exported physical volumes from the source TS7700.
- ▶ Only the copy-exported volumes from a single source TS7700 can be used in the recovery process.
- ▶ The recovery TS7700 cannot be part of a grid configuration.
- ▶ The recovery TS7700 must be configured as Cluster 0.
- ▶ The recovery TS7700 and its associated MI must be configured, have code that is 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 as the source TS7700.
- ▶ If the recovery TS7700 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 system is emptied of data.
- ▶ If another TS7700 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 following results occur:
 - The VOLSERS that conflict cannot be recovered.
 - A warning message is displayed in the recovery status window on the recovery TS7700 MI.
 - You cannot use the same library for both the source and recovery TS7700.
- ▶ Other than the physical drive compatibility requirements listed, the source and recovery TS7700 can have different configuration features, such as different cache capabilities and performance enablement features.
- ▶ You must add scratch physical volumes to the recovery TS7700 even if you are going to be only reading data. A minimum of two scratch volumes per defined pool in the recovery TS7700 are needed to prevent the recovery TS7700 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 TS7700, do so only after the recovery has been run and the recovery TS7700 is ready to be brought online to its attached hosts. Otherwise, their inventory records are 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 TS7700 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 that are defined for the storage management constructs that were restored during the recovery are the actions that you want to continue to use.

12.3.2 Performing Copy Export Recovery

Perform the following steps:

1. With the TS7700 and library in an online state, log in to the MI and select **Service** → **Copy Export Recovery**.

You see only the Copy Export Recovery menu item if you have been given Administrator-level or Manager-level access by the overall system administrator on the TS7700. The Copy Export Recovery menu item is not displayed if the TS7700 is configured in a grid configuration. Contact your IBM Service Support Representative (IBM SSR) if you must recover a TS7700 that is a member of a grid.

2. If the TS7700 determines that data or database entries exist in the cache, Copy Export Recovery cannot be performed until the TS7700 is empty.

Figure 12-5 shows the window that opens to inform you that the TS7700 contains data that must be erased.



Figure 12-5 Copy Export Recovery window with erase volume option

3. Ensure that you are logged in to the correct TS7700. Then, select **Erase all existing volumes before the recovery** and click **Submit**. A window opens that provides you with the option to confirm and continue the erasure of data on the recovery TS7700 or to abandon the recovery process. It describes the data records that 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**. The TS7700 begins the process of erasing the data and all database records. As part of this step, you are logged off from the MI.

4. After waiting about 1 minute, log in to the MI. Select **Settings** → **Copy Export Recovery Status** to follow the progress of the Copy Export Recovery.

The following tasks are listed in the task detail window as the erasure steps are being performed:

- a. Taking the TS7700 offline.
- b. The existing data in the TS7700 database is being removed.
- c. The existing data in the TS7700 cache is being removed.
- d. Cleanup (removal) of existing data.
- e. Requesting the TS7700 go online.
- f. Copy Export Recovery database cleanup is complete. After the erasure process is complete, the TS7700 returns to its online state.

Note: If an error occurs during the erasure process, the task detail window provides a list of errors that occurred and indicates the reason and any action that needs to be taken.

5. Starting with an empty TS7700, you must perform several setup tasks by using the MI that is associated with the recovery TS7700. For many of these tasks, you might have to verify only that the settings are correct because 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 TS7700 must know the VOLSER ranges that it owns. This step is done through the MI that is associated with the recovery TS7700.
 - b. If the copy-exported physical volumes were encrypted, set up the recovery TS7700 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 write data to the recovery TS7700, 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 running 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 TS7700 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.

6. With the TS7700 in its online state, but with all virtual tape drives varied offline to any attached hosts, log in to the MI and select **Service** → **Copy Export Recovery**.

The TS7700 determines that it is empty and enables the operation to proceed. Load the copy-exported physical volumes into the library. Multiple sets of physical volumes have likely been exported from the source TS7700 over time. All of the exported stacked volumes from the source TS7700 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.

Important:

- ▶ Before continuing the recovery process, be sure that all of the copy-exported physical volumes have been added. Any volumes that are not known to the TS7700 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 TS7700.
- ▶ Do not add any physical scratch cartridges now. You can do that after the Copy Export Recovery operation has completed and you are ready to bring the recovery TS7700 online to the hosts.

7. After you add all of the physical volumes into the library and they are now known to the TS7700, enter the volume serial number of one of the copy-exported volumes from the last set that was exported from the source TS7700. It contains the last database backup copy, which is used to restore the recovery TS7700 database. The easiest place to find a volume to enter is from the export status file on the export list volume from the current Copy Export operation.

Remember: If you specified the LMTDBPVL option when performing the export, only a subset of the tapes that were exported have a valid database backup that can be used for recovery. If a tape that is selected for recovery does not have the backup, the user gets the following error: “The database backup could not be found on the specified recovery volume”.

If you are using the Copy Export Recovery operation to perform a disaster recovery test, keep the **Disaster Recovery Test Mode** check box selected. The normal behavior of the TS7700 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 was written to while it was mounted, the actions defined by the storage management constructs are taken.

If the volume was not modified, actions are only taken if there has been a change in the definition of the storage management constructs since the last time that the volume was unloaded. For example, suppose that a logical volume is assigned to an SG, which had last had the volume written to pool 4. Furthermore, either the SG was not explicitly defined on the recovery TS7700 or it specified a different pool.

In this case, on the unload of the volume, a new copy of it is written to the pool determined by the new SG definition, even though the volume was only read. If you are merely accessing the data on the recovery TS7700 for a test, you do 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 merely running a test is reclamation. Running reclamation while performing a test will require scratch physical volumes and enable the copy-exported volumes to be 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 that are used for recovery maintain their status of Copy Exported so that they cannot be reused or used in a subsequent Copy Export operation. If you are using Copy Export Recovery because of a real disaster, clear the check box.

Enter the volume serial number, select the check box, and then click **Submit**.

8. A window opens and 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**.
9. The TS7700 begins the recovery process. As part of this step, you are logged off from the MI.
10. After waiting about 1 minute, log in to the MI and select **Settings** → **Copy Export Recovery Status** 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 the overall status.

The following tasks are listed in the task detail window as the Copy Export Recovery steps are performed:

- a. The TS7700 is taken offline.
- b. The requested recovery tape XXXXXX is being mounted on device YYY.
- c. The database backup is being retrieved from the specified recovery tape XXXXXX.
- d. The requested recovery tape is being unmounted following the retrieval of the database backup.
- e. The database backup that is retrieved from tape is being restored on the TS7700.
- f. The restored database is being updated for this hardware.
- g. The restored database volumes are being filtered to contain the set of logical volumes that were Copy Exported.
- h. Token ownership is being set to this cluster from the previous cluster.

- i. The restored database is being reconciled with the contents of cache, XX of YY complete.
- j. Logical volumes are being restored on the Library Manager, XX of YY complete.
- k. Copy Export Recovery is complete.
- l. Copy Export Recovery from physical volume XXXXXX.
- m. The request is made for the TS7700 to go online.
- n. The recovered data is loaded into the active database.
- o. The process is in progress.

After the Copy Export Recovery process completes successfully, the MI returns to its full selection of tasks.

11. Now, add scratch physical volumes to the library. At least 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 the VOLSER ranges that it owns. The steps are described in “Defining VOLSER ranges for physical volumes” on page 529.

12. If you ran Copy Export Recovery because of a real disaster (you cleared the **Disaster Recovery Test Mode** check box), verify that the defined storage management construct actions will manage the logical and physical volumes in the manner that is needed.

During Copy Export Recovery, the storage management constructs and their actions are restored to the storage management constructs, and their actions are defined on the source TS7700. If you want the actions to be different, change them through the MI that is associated with the recovery TS7700.

You can now view the completed results of the Copy Export Recovery in Figure 12-6.

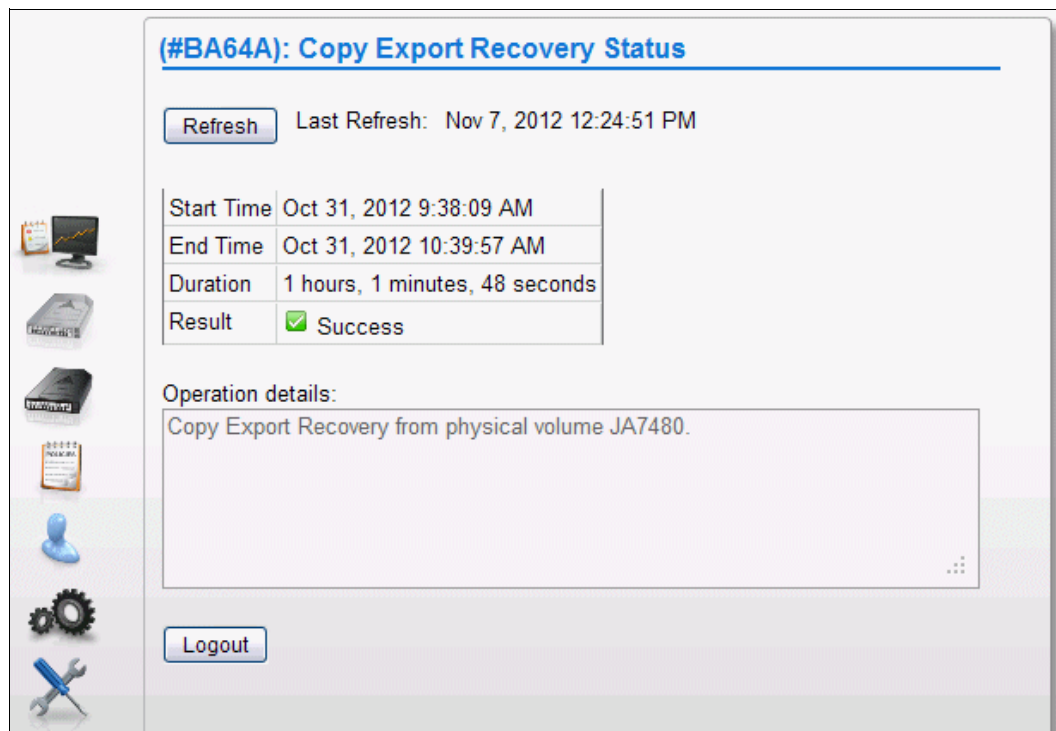


Figure 12-6 Copy Export Recovery Status

If an error occurs, various possible error texts with detailed error descriptions can help you solve the problem. For more information and error messages that are 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 URL:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092>

If everything is completed, you can vary the virtual devices online, and the tapes are ready to read.

Tip: For more general considerations about DR testing, see Chapter 5, “Disaster recovery” on page 189.

12.3.3 Restoring the host and library environments

Before you can use the recovered logical volumes, you must restore the host environment also. The following steps are the minimum steps that you need to continue the recovery process of your applications:

1. Restore the tape management system (TMS) CDS.
2. Restore the DFSMS data catalogs, including the tape configuration database (TCDB).
3. Define the I/O gen by using the Library ID of the recovery TS7700.
4. Update the library definitions in the source SCDS with the Library IDs for the recovery TS7700 in the composite library 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 MI for the grid and cluster to match the library names defined to DFSMS. That way, the names that are shown on the MI windows match those names that are used at the host for the composite library and distributed library.

To set up the composite name that is used by the host to be the grid name, complete the following steps:

1. Select **Configuration** → **Grid Identification Properties**.
2. In the window that opens, enter the composite library name that is used by the host in the grid nickname field.
3. You can optionally provide a description.

Similarly, to set up the distributed name, complete the following steps:

1. Select **Configuration** → **Cluster Identification Properties**.
2. In the window that opens, enter the distributed library name that is used by the host in the Cluster nickname field.
3. You can optionally provide a description.

These names can be updated at any time.



Disaster Recovery Testing

This chapter describes disaster recovery (DR) testing in a TS7700 grid configuration.

This chapter includes the following sections:

- ▶ DR Testing Overview
- ▶ DR Testing Methods
- ▶ DR Testing General Considerations
- ▶ DR for FlashCopy Concepts and Command Examples
- ▶ DR Testing Methods Examples
- ▶ Expected Failures During a DR Test

13.1 DR Testing Overview

In a perfect world, there would be no need for disaster recovery testing. However, in reality, there are innumerable factors that could lead to a disaster that prevents the usage of one or more of your production TS7700 clusters in a grid environment. Therefore, it is important to prepare and test your environment for such a scenario.

The fortunate thing is that, in many cases, recovering from a disaster is easier and requires fewer steps than having to simulate a disaster and then clean up your disaster environment as if the simulation had never happened. While Chapter 12 discussed disaster recovery concepts in general, this chapter focuses on concepts related to disaster recovery testing specifically, providing examples where needed and includes step-by-step walkthroughs for 4 methods that clients may use to accomplish DR testing in a TS7700 grid environment. Those methods are:

1. DR Testing using FlashCopy
2. DR Testing using Write Protect Mode on DR cluster(s)
3. DR Testing without Using Write Protect Mode on DR cluster(s)
4. DR Testing by Breaking the Grid Links connections to DR cluster(s)

All of these methods have their advantages and disadvantages, so it is important that before you decide which method to use, that you weigh the advantages and disadvantages of each method against your environment and resources and then choose which method best fits your DR testing needs and ability.

The description of each method makes the assumption that you are familiar with the DR concepts presented in Chapter 12. The end of this chapter contains a step-by-step list on how to perform a DR test using each method. While it may be tempting to jump right to these lists, it is recommended that you review this chapter in its entirety before DR testing to ensure that you are familiar with the concepts and options available for DR testing in a TS7700 grid environment.

13.2 DR Testing Methods

This section describes four different methods that can be used to test disaster recovery in a TS7700 grid environment.

13.2.1 Method 1: DR Testing using FlashCopy

This method of DR testing uses the FlashCopy functionality that was introduced in Release 3.1. This function enables a DR host to perform testing against a point in time consistency snapshot while production operations and replication continue. With FlashCopy, production data continues to replicate during the entire DR test and the same logical volume can be mounted at the same time by a DR host and a production host.

With FlashCopy and the underlying Write Protect Mode for DR testing, DR test volumes can be written to and read from while production volumes are protected from modification by the DR host. All access by a DR host to write-protected production volumes is provided by using a snapshot in time (a flash) of the logical volumes. Because of this, a DR host continues to have read access to any production volumes that have been returned to scratch while the FlashCopy is active.

During a DR test, volumes might need to be mounted from both the DR and production hosts. Before FlashCopy for DR testing, these mounts were serialized such that one host access received an IN USE exception. This was especially painful when the true production host was the instance that fails the mount.

FlashCopy enables logical volumes to be mounted in parallel to a production host and a DR host. Production hosts can scratch volumes, reuse volumes, or modify volumes without affecting the copy of the production data that is used by the DR host while the FlashCopy is active. This method has the following advantages and disadvantages:

► Advantages:

- Once the FlashCopy has been enabled, all read activity against volumes that are included in the FlashCopy (those in write-protected categories on one or more DR clusters) are from that point-in-time copy. This very closely simulates a real disaster scenario where one or more production clusters are no longer accessible and the disaster clusters have access to the production data from a point-in-time. Volumes belonging to categories that are excluded from write-protection on the DR clusters can continue to have data written to them during the DR test.
- Data written from a production host to the production cluster(s) can continue to be copied to the disaster cluster(s) without the risk of a disaster host accessing the live data. While the FlashCopy is active, the disaster host can only access the point-in-time copy of the production data already present on the disaster cluster(s) at the time of the FlashCopy.

► Disadvantages:

- Your disaster clusters must be composed of at least one TS7720 or TS7760 in order to use the FlashCopy functionality.
- The FlashCopy on the DR clusters ensures that if a logical volume is changed by a production host on a production cluster and that change is propagated to the DR clusters, a copy of the previous data is still kept in the DR clusters. This leads to a higher cache utilization on the DR clusters.
- The Write Protect Mode and Write Protect Exclude categories must be configured correctly in order for any data to be able to be written to the DR clusters during a DR test. If they are configured incorrectly (for example, by defining production categories as being excluded from write-protect), production data might be overwritten.
- Release 3.1 and newer release levels are required to use the FlashCopy functionality.

13.2.2 Method 2: DR Testing using Write Protect Mode on DR clusters

This method uses the Write Protect Mode functionality in TS7700 cluster(s) to prevent all write activity or volume attribute changes to the hardware categories that are NOT in the Exclude-from-write-protect list in the DR clusters. The only categories that should be in this list are those categories that will be used by the DR host to read and write from DR volumes that were processed by host cartridge entry on the DR clusters. All other categories (such as the categories that production volumes belong to), will be write-protected on the DR clusters.

This method has the following advantages and disadvantages:

Advantages:

- By enabling Write Protect Mode on the disaster clusters, even if a job on the DR host tries to mount a production volume for write on a disaster cluster, the cluster will prevent the write at the hardware level.

- Production data can still be written to the production clusters, and those clusters can still copy data to the disaster clusters so that in the event of a real disaster, the data on the disaster clusters will be more up-to-date than if the copying did not occur.

Disadvantages:

- The Write Protect Mode and Write Protect Exclude categories must be configured correctly in order for any data to be able to be written to the DR clusters during a DR test. If they are configured incorrectly (for example, by defining production categories as being excluded from write-protect), production data might be overwritten.
- There is no point-in-time simulation. The data on the volumes used during a DR test can change if those volumes are written to by a production system on a production cluster and those changes are propagated to the disaster clusters. Jobs running on the DR host that are reading data from production volumes on the DR clusters might fail if they do not account for this possibility.

If you determine that FlashCopy for DR is not suitable to your DR environment, using this method is the recommended alternative.

13.2.3 Method 3: DR testing without using Write Protect Mode on DR clusters

This is similar to the previous method, except instead of using the Write Protect Mode functionality in the DR clusters to prevent any writes issued from the DR host to a production volume on the DR clusters, this method relies on the ability (and correct configuration) of the TMS on a DR host to prevent volumes in the production volume range from being written to by the DR host.

This method has the following advantages and disadvantages:

Advantages:

- Production data can still be written to the production clusters and those clusters can still copy data to the disaster clusters so that in the event of a real disaster, the data on the disaster clusters will be more up-to-date than if the copying did not occur.

Disadvantages:

- There is no hardware-enabled write protection that would prevent a DR host from writing to a production volume. The TMS on the disaster host **MUST** be configured to prevent any writes directed toward production volumes.
- There is no point-in-time simulation. The data on the volumes used during a DR test can change if those volumes are written to by a production system on a production cluster and those changes are propagated to the disaster clusters. Jobs running on the DR host that are reading data from production volumes on the DR clusters might fail if they do not account for this possibility.
- Return-to-scratch processing might need to be suspended on the production hosts during the DR test. Refer to the Appendix 13.3.13, “Returning to scratch without using Selective Write Protect” on page 799 for additional information.

If your choice is between using Write Protect Mode and not using Write Protect Mode, it is suggested to use Write Protect Mode (Method 2), to provide an additional level of write-protection in case the TMS on the DR host is not configured correctly to prevent writes to the production volumes.

13.2.4 Method 4: Breaking the interconnects between the TS7700 grid

The final method discussed that you can choose to simulate that of a real disaster is to break the grid links between the production clusters and DR clusters in a TS7700 grid.

As with the previous methods, this method has its advantages and disadvantages:

Advantages:

- After the grid links have been broken, you are assured that any production data that is accessed from a DR cluster by the DR host is data that had been copied to the DR cluster before the grid links were broken.
- Return-to-scratch processing initiated by a production host again production volumes on production clusters does not affect the copy of the volumes on the DR clusters. The copy on the DR clusters can continue to be accessed for read by the DR host.
- DR volumes that are created for use during the DR test are not copied to the production clusters.

Disadvantages:

- If a real disaster occurs while the DR test is in progress, data that was created by the production site after the grid links were broken is lost.
- The disaster clusters must be allowed to takeover read-only volume ownership from the production clusters. Normally, the takeover function is only used in the event of a real disaster.
- Breaking the grid links must be done by your CE (SSR). Do not only disable a grid link with the Library Request command to run this method. Disabling the grid link with the command does not stop synchronous mode copies and the exchange of status information.

The concern about losing data in a real disaster during a DR test is the major drawback to using this DR method. Because of this, if it is possible to use one of the DR methods described earlier (using FlashCopy or Write Protect Mode), it is advised to use one of those methods.

Important: 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, probably during your batch window. To understand throttling, see 11.3.7, “Throttling in the TS7700” on page 646.

After the test ends, you might have many virtual volumes in the pending copy status. When TS7700 grid links are restored, communication is restarted, and the first task that the TS7700 runs is to make a copy of the volumes that are created during your links broken window. This task can affect the TS7700 performance.

If your DR test runs over several days, you can minimize the performance degradation by suspending copies by using the **GRIDCNTL** Host Console command. After your DR test is over and your CE has brought back the grid links, you can enable the copy again during a low activity workload to avoid or minimize performance degradation. See 10.1.3, “Host Console Request function” on page 608 for more information.

13.3 DR General Considerations

As you design a test that involves the TS7700 grid configuration, there are several capabilities that are designed into the TS7700 that you must consider.

13.3.1 The z/OS test environment represents a point in time

The test environment is typically a point in time, which means that at the beginning of the test, the catalog, TCDB, and TMS 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 effect is that it is possible that a volume that has been returned to SCRATCH status by the production system is used in a test. The test system's catalogs and TMS do not reflect that change. If the "Ignore fast ready characteristics of write protected categories" option is selected when Write Protect Mode is enabled on the DR clusters, the data can still be accessed, regardless if the logical volume is defined as scratch or not.

13.3.2 The data that is available in the DR cluster

In a real disaster, the data that is available in the clusters in your remaining site might not be consistent with the content in your TMS catalog. This situation depends on the selected Copy Modes, and if the copies are already processed.

During your DR test, production data is updated on the remaining production clusters. Depending on your selected DR testing method, this updated data can be copied to the DR clusters. Also, it depends on the DR testing method if this updated data is presented to the DR host, or if a FlashCopy from a Time Zero is available.

Without the FlashCopy option, both alternatives (updating the data versus not updating the data) have advantages and disadvantages. For more information, see 13.5.4, "Method 4: Breaking the grid link connections" on page 818.

Also, the DR host might create some data in the DR clusters. For more information, see "Creating data during the disaster recovery test from the DR host: Selective Write Protect" on page 794.

13.3.3 Write Protect Mode

This function enables clients to emulate DR events by running test jobs at a DR location within a TS7700 grid configuration, enabling volumes only within specific categories to be manipulated by the test application. This function prevents any changes to production-written data, which is accomplished by excluding up to 32 categories from the cluster's write-protect enablement.

When a cluster is write-protect-enabled, all volumes that are protected cannot be modified or have their category or storage construct names modified. As in the TS7700 write-protect setting, the option is grid partition scope (a cluster) and configured through the MI. Settings are persistent, except for DR FLASH, and saved in a special repository.

Also, the new function enables any volume that is assigned to one of the categories that are contained within the configured list to be excluded from the general cluster's write-protect state. The volumes that are assigned to the excluded categories can be written to or have their attributes modified.

In addition, those scratch categories that are not excluded can optionally have their Fast Ready characteristics ignored, including Delete Expire and hold processing, enabling the DR test to mount volumes as private that the production environment has since returned to scratch (they are accessed as read-only).

One exception to the write protect is those volumes in the insert category. To enable a volume to be moved from the insert category to a write-protect-excluded category, the source category of insert cannot be write-protected. Therefore, the insert category is always a member of the excluded categories.

Be sure that you have enough scratch space when Expire Hold processing is enabled to prevent the reuse of production scratched volumes when planning for a DR test. Suspending the volumes' return-to-scratch processing during the DR test is also advisable.

Because selective write protect is a cluster-wide function, separated DR drills can be conducted simultaneously within one multi-cluster grid, with each cluster having its own independent client-configured settings. Again, DR FLASH is the exception to this statement.

13.3.4 Protection of your production data

In a real disaster this is not an issue because the remaining systems become your production environment.

However, during a DR test you need to ensure that the actions on the DR site do not have an influence on the data from production. Therefore, the DR host must not have any connections to the clusters in production. Ensure that all devices that are attached to the remaining production clusters are offline (if they are FICON attached to the DR site).

The Write Protect mode prevents any host action (write data, host command) sent to the test cluster from creating new data, modifying existing data, or changing volume attributes such as the volume category. The Write Protect mode still enables logical volumes to be copied from the remaining production clusters to the DR cluster.

As an alternative to the Write Protect Mode or if you would like an additional safeguard, if you want to prevent overwriting production data, you can use the TMS on the DR host to enable only read-access to the volumes in the production VOLSER ranges. For more information, see 13.3.12, "Considerations for DR tests without Selective Write Protect mode" on page 797.

13.3.5 Separating production and disaster recovery hosts: Logical volumes

The DR host is an isolated LPAR that needs to be segregated from the production environment. To avoid any interference or data loss, complete these optional steps:

1. Define host-specific media categories for Media1/2, Error, and Private. These categories must be unique from any category that is used by a production system to ensure separation of production volumes from DR (disaster) volumes.
2. Limit the usage of logical volumes by using the TMS.
3. Define separate logical volume serial ranges (insert process).

To ensure that the inserted volume ranges are not accepted by the production systems, you need to perform the following steps:

- ▶ Changes on production systems:
 - Use the RMM **REJECT ANYUSE(TST*)**, **PRITITION VOLUME(TST*) TYPE(NORMM) SMT(IGNORE) NOSMT(IGNORE)** or **OPENRULE VOLUME(TST*) TYPE(RMM) ANYUSE(REJECT)** parameter, which means to *not* use VOLSERS named TST* here.

- ▶ Changes on the DR test systems:
 - Use the RMM `VLPOOL PREFIX(TST*) TYPE(S)` parameter to enable use of these volumes for default scratch mount processing.
 - Change `DEVSUPxx` to point to other categories, which are the categories of the TST* volumes.

Figure 13-1 shows the process to insert cartridges in a DR site to perform a DR test.

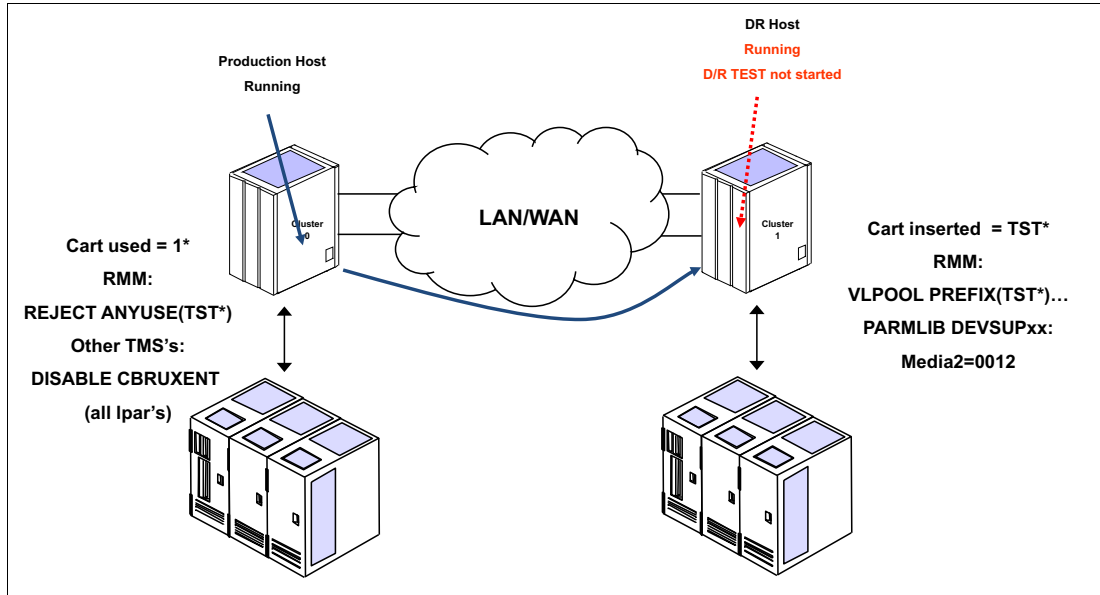


Figure 13-1 Insertion considerations in a disaster recovery test

After these settings are done, insert the new TST* logical volumes. It is important that the DR volumes that are inserted by using the MI are associated with the DR host so that the TS7700 at the DR site has ownership of the inserted volumes. The DR host must be running before the insertion is run.

Important: Ensure that one logical unit has been or is online on the test system before entering logical volumes.

Any new allocations for output that are performed by the DR host use only the logical volumes that are defined for the DR test. At the end of the DR test, the volumes can be returned to SCRATCH status and left in the library. Or, if you prefer, they can be deleted by using the EJECT command in ISMF on the DR host.

13.3.6 Creating data during the disaster recovery test from the DR host: Selective Write Protect

During the DR test, you might want to write data from the DR host to the DR clusters. These DR tests typically include running a batch job cycle that creates data on DR volumes.

This test can be handled in two ways:

- ▶ Have a different cluster 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 involves 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 DR cluster before the start of the test.

Important: The test volumes that are inserted by using the MI must be associated with the cluster that is used as DR cluster so that cluster has ownership of the inserted volumes.

If you require that the DR host be able to write new data, you can use the Write Protect Mode for DR testing function that enables you to write to volumes belonging to certain categories during DR testing. With Selective Write Protect, you can define a set of volume categories on the TS7700 that are excluded from the Write Protect Mode. This configuration enables the test host to write data onto a separate set of logical volumes without jeopardizing normal production data, which remains write-protected.

This requires that the DR host use a separate scratch category or categories from the production environment. If DR volumes also must be updated or if you want to run a TMS housekeeping process that is limited to the DR volumes, the DR 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 separate, not yet used categories on the DR host by using the DEVSUPxx member. Be sure that you define a minimum of four categories in the DEVSUPxx member: MEDIA1, MEDIA2, ERROR, and PRIVATE.

In addition to the DR host specification, you must also define on the DR clusters those volume categories that you are planning to use on the DR host and that need to be excluded from Write-Protect mode.

For more information about the necessary definitions for DR testing with a TS7700 grid that uses Selective Write Protect, see 13.5.2, "Method 2: Using Write Protect Mode on DR clusters" on page 814.

The Selective Write Protect function enables you to read production volumes and to write new volumes from the beginning of tape (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.

At the end of the DR test, be sure to clean up the data that was written to DR volumes during the DR test.

13.3.7 Creating data during the disaster recovery test from the disaster recovery host: Copy policies

If you are using the same MCs used in production, the data that is 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 MCs differently at the DR clusters than at the production clusters.

Using a copy mode of No Copy for the production clusters prevents the DR clusters from making a copy of the DR test data. It does not interfere with the copying of production data.

Remember to set the content of the MCs back to the original contents during the cleanup phase of a DR test.

13.3.8 Restoring the DR host from a production host

When the DR methods at the end of this chapter discuss restoring a DR environment from a production environment, what is meant is that, at a minimum, the following need to be obtained from a point-in-time copy from a production host and restored and activated on the DR host:

1. The tape management system (TMS) CDS.
2. The DFSMS data catalogs, including the tape configuration database (TCDB).
3. The input/output definition file (IODF)
4. The SMS source control data set (SCDS)

13.3.9 Scratch runs during the disaster recovery test from the production host

If return-to-scratch processing runs on a production host for a production volume, that volume can no longer be read by the production host while it is in scratch status. However, it can still be read by a DR host from a DR cluster that has Write Protect Mode active (either with or without DR FlashCopy) if the category the volume is in is being write-protected and the cluster has "Ignore fast read characteristics of write protected categories." enabled.

During DR testing, you might want to either turning off return-to-scratch processing on the production hosts or configure a long expire-hold time for the production tapes that can be scratched to ensure that the data can still be accessed during the DR test.

For scratch processing run during the DR test from the production host without using Selective Write Protect, see 13.3.12, "Considerations for DR tests without Selective Write Protect mode" on page 797.

13.3.10 Scratch runs during the disaster recovery test from the DR host

Depending on the selected method, a return-to-scratch procedure that is run on the DR host should be carefully considered. If Write Protect Mode is enabled and the production category is set to Write Protect Excluded, you can run a scratch procedure on the DR host. It is advised to limit the scratch procedure to the DR volume serial range inserted on the DR host.

If you choose not to use Write Protect or define the production categories as excluded from write protect, a return-to-scratch procedure that is run on a DR host might lead to data loss. If possible, it is best to avoid running any housekeeping process during a DR test.

13.3.11 Cleanup phase of a disaster recovery test

When a DR test is complete, you should clean up the DR environment so that it is in the same condition as before you started the DR test. During this process, you should delete the data from the DR clusters that was written by the DR host.

If this data is not deleted (set to scratch and EJECTed via ISMF) after the DR test, this unneeded data will continue to occupy cache or tape space. Because the volumes this data resides on remain in a PRIVATE category, they will never expire and will continue to occupy space indefinitely.

For this reason, be sure to return to scratch those DR volumes that are written to (converted from SCRATCH to PRIVATE) during the DR test and, at the very least (if you don't want to delete the volumes), ensure that the scratch category that they are assigned to has an expiration time specified in the TS7700 MI. Otherwise, space on the TS7700 will continue to be wasted because these logical volumes will not be overwritten.

Ownership takeover

If you perform the DR 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 a production cluster, that cluster has volume ownership.

If you attempt to mount one of those volumes from the DR host without ownership takeover enabled, the mount fails because the DR cluster cannot request ownership transfer from the production cluster. By enabling ROT, the test host can mount the production logical volumes and read their contents.

The DR host is not able to modify the production site-owned volumes or change their attributes. The volume appears to the DR host as a write-protected volume. Because the volumes that are going to be used by the DR host for writing data were inserted through the MI that is associated with the DR cluster, that DR cluster already has ownership of those volumes. The DR host has complete read and write control of these volumes.

Important: Never enable Write Ownership Takeover mode for a test. WOT mode must be enabled only during a loss or failure of the production TS7700.

If you are not going to break the links between the sites, normal ownership transfer occurs whenever the DR host requests a mount of a production volume.

13.3.12 Considerations for DR tests without Selective Write Protect mode

The TS7700 contains several features that can be used to prevent production volumes from being written to at the hardware level during a DR test, namely Write Protect Mode (which is also enabled during FlashCopy) and (in the case of a DR test accomplished by breaking the grid links) read-only takeover mode. As an alternative to using either of these methods (or in addition to, if wanted), you can use the TMS on the DR host to enable only read-only access to the volumes in the production VOLSER ranges.

For example, with DFSMSrmm, you can insert these extra statements into the EDGRMMxx parmlib member on the DR host:

- ▶ For production volumes in a range of A00000 - A09999, add this statement:
REJECT OUTPUT(A0*)
- ▶ For production volumes in a range of ABC000 - ABC999, add this statement:
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 correctly. If the product is DFSMSShsm, tapes that are filling (seen as not full) from the test system control data set (CDS) must be modified to full by running commands. If DFSMSShsm then later needs to write data to tape, it requires a scratch volume that is related to the test system's logical volume range.

As a result of recent changes in DFSMSrmm, it now is easier to manage this situation:

- ▶ In z/OS V1R10, the new commands **PRITITION** and **OPENRULE** provide for flexible and simple control of mixed system environments as an alternative to the REJECT examples used here. These new commands are used in the EDGRMMxx member of parmlib.
- ▶ You can specify extra EXPROC controls in the **EDGHSKP SYSIN** file to limit the return-to-scratch processing to specific subsets of volumes. So, you can just EXPROC the DR volumes on the DR host and the PROD volumes on the PROD host. You can still continue to run regular batch processing, and also run expiration on the DR host.

Figure 13-2 helps you understand how you can protect your tapes in a DR test while your production system continues running.

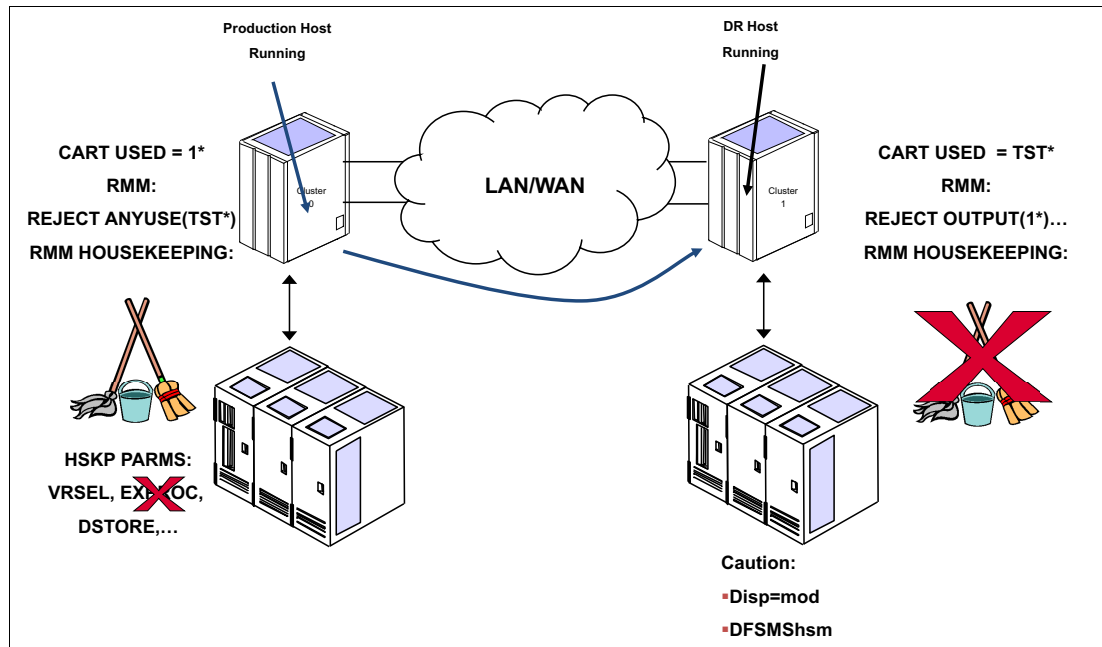


Figure 13-2 Work process in a disaster recovery test

Clarification: The term **HSKP** is used because this term is typically the job name that is used to run the RMM EDGHSKP utility that is used for daily tasks, such as vital records processing, expiration processing, and backup of control and journal data sets. However, it can also see the daily process that must be done with other TMSs. This publication uses the term HSKP to mean the daily process on RMM or any other TMSs.

This includes stopping any automatic short-on-scratch process, if enabled. For example, RMM has one emergency short-on-scratch procedure.

To illustrate the implications of running the HSKP task in a DR host, see the example in Table 13-1, which displays the status and definitions of a production volume in a normal situation.

Table 13-1 VOLSER AAAAAA before returned to scratch from the disaster recovery site

Environment	DEVSUP	TCDB	RMM	MI	VOLSER
PROD	0002	Private	Master	000F	AAAAAA
DR	0012	Private	Master	000F	AAAAAA

In this example, volume AAAAAA is the master in both environments. However, due to a procedural error, it is returned to scratch by the DR host. You can see its status in Table 13-2.

Table 13-2 VOLSER AAAAAA after returned to scratch from the disaster recovery site

Environment	DEVSUP	TCDB	RMM	MI	VOLSER
PROD	0002	Private	Master	0012	AAAAAA
DR	0012	Scratch	Scratch	0012	AAAAAA

Volume AAAAAA is now in scratch category 0012. This presents two issues:

- ▶ If you need to access this volume from a production host, you need to change its status to master (000F) using ISMF ALTER from SCRATCH to PRIVATE on the Prod host before you can access it. Otherwise, you will lose the data on the volume, which can have serious consequences, for example, 1000 production volumes are accidentally returned to scratch by the DR host.
- ▶ On the DR host, RMM is set to reject using the production volumes for output. If this volume is mounted in response to a scratch mount on the DR host, it will be rejected by RMM. Imagine the scenario where the TS7700 must mount 1,000 scratch volumes before the TS7700 mounts a volume that RMM does not reject. This would not be a desirable situation.

To provide maximum protection from a system operator perspective, perform these tasks to protect production volumes from unwanted return-to-scratch processing:

- ▶ Ensure that the RMM HSKP procedure is not running on any host during the test window of the DR host. There is a real risk of data loss if the DR host returns production volumes to scratch and you have defined in the TS7700 that the expiration time for the corresponding category is 24 hours. After this time, volumes can become unrecoverable.
- ▶ Ensure that the RMM short-on-scratch procedure does not start. The results can be the same as running an HSKP.

If you are going to perform the test with the site-to-site links broken, you can use the ROT mode to prevent the test system from modifying the production site's volumes. For more information about ownership takeover, see 2.3.34, "Autonomic Ownership Takeover Manager" on page 90.

In addition to the protection options that are described, 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.

13.3.13 Returning to scratch without using Selective Write Protect

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 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 more EXPROC controls in the EDGHSKP SYSIN file to limit the return-to-scratch processing to specific subsets of volumes. So, you can 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.

If a volume is returned to a scratch (Fast Ready) category during a DR test by a production host, mounting that volume through a specific mount does not recall the previously written data (even though the DR host sees it as a private volume).

The TS7700 always mounts a blank volume from a scratch (Fast Ready) category. It might be possible to recover the data on the volume by assigning the volume back to a private (non-Fast Ready) category, or (only in pre-Release 3.0) taking that category out of the scratch (Fast Ready) list and trying the mount again.

Even if the number of volumes in the list is larger than the number of volumes that are needed per day times the number of days of the test, you still need to take steps to make it unlikely that a volume that is needed for test is 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 at the following URL:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101281>

2. Suspend only the return-to-scratch processing for the production volume that is needed for the test. For RMM, this can be done by using policy management through vital record specifications (VRSs). A volume VRS can be set up that covers each production volume so that this overrides any existing policies for data sets.

For example, assume that 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 is 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 that is required for test.

After the DR test is finished, you have a set of volumes in the TS7700 that belong to DR test activities. You need to decide what to do with these tapes. As a test ends, the RMM database and VOLCAT will probably be destaged (along with all of the data that is used in the DR test). However, until an action is taken, the volumes remain defined in the MI database.

- One is in master status.
- The others are in SCRATCH status.

If the volumes are not needed anymore, manually release the volumes and then run EXPROC to return the volumes to scratch under RMM control. If the tapes will be used for future test activities, manually release these volumes. The cartridges remain in the SCRATCH status and ready for use. Remember to use a Scratch category with expiration time to ensure that no space is wasted.

Important: Although volumes in the MI remain ready to use, you must ensure that the next time that you create the DR test environment that these volumes are defined to RMM and the TCDB. Otherwise, you cannot use them.

13.4 DR for FlashCopy Concepts and Command Examples

When enabled, FlashCopy allows two instances of a volume to exist on the same DR cluster. The DR host accesses the content of a logical volume from a point zero, while at the same time an active copy of the logical volume can be updated with new copies pulled from the production cluster. You do not need a break of the grid link to ensure that only data from time zero is available to the DR host.

For a detailed technical description, see *IBM Virtualization Engine TS7700 Series Best Practices - FlashCopy for Disaster Recovery Testing*, which is available at the Techdocs website (search for the term **TS7700**):

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

The following terms apply to FlashCopy:

- ▶ **Live Copy:** A real-time instance of a virtual tape within a grid that can be modified and replicated to peer clusters. This is the live instance of a volume in a cluster that is the most recent version of the volume on that cluster. If the Live Copy is also consistent relative to the grid, it can be altered by a production host or from a DR host when it is in the exclusion list of write protect.
- ▶ **FlashCopy:** A snapshot of a live copy at time zero. The content in the FlashCopy is fixed and does not change even if the original copy is modified or if replication events occur. A FlashCopy might not exist at a particular cluster if a live volume was not present within that cluster at time zero. In addition, a FlashCopy does not imply consistency because the live copy might have been down level to the grid, or simply incomplete at time zero. An active FlashCopy indicates that Write Protect Mode is active.
- ▶ **DR Family:** A set of TS7700 clusters (most likely those at the DR site) that serve the purpose of DR. One to seven clusters can be assigned to a DR family. The DR family is used to determine which clusters should be affected by a flash request or write-protect request by using a host console request command (HCR). A DR Family needs at least one TS7760 or TS7720.
- ▶ **Write Protect Mode (existing function):** When Write Protect Mode is enabled on a cluster, host commands fail if they are sent to logical devices in that cluster and attempt to modify a volume's data or attributes and that volume is not excluded from write protect. The FlashCopy is created on a cluster when it is in the write protect mode only. Also, only write-protected virtual tapes are flashed. Virtual tapes that are assigned to the excluded categories are not flashed.
- ▶ **Time Zero:** The time when the FlashCopy is taken within a DR family. The time zero mimics the time when a real disaster happens. Customers can establish the time zero using a host console request command.

Basic requirements and concepts

All clusters in the grid must be running with R3.1 or higher microcode level to enable this function.

The FlashCopy for DR testing function is supported on TS7700 Grid configurations where at least one TS7760 or TS7720 cluster exists within the DR location. The function cannot be supported under TS7740-only grids or where a TS7740 is the only applicable DR cluster. A TS7740 might be present and used as part of the DR test if at least one TS7760 or TS7720 is also present in the DR site.

The Write Protect exclusion categories are not a subject for the flash. For these categories only, a Live Copy exists.

During an enabled Flash, the autoremoval process is disabled for the TS7760/TS7720 members of the DR Family. A TS7760/TS7720 within a DR location requires extra capacity to accommodate the reuse of volumes and any DR test data that is created within an excluded category. Volumes that are not modified during the test require no additional TS7760/TS7720 disk cache capacity. The extra capacity requirement must be considered when planning the size of the TS7760/TS7720 disk cache.

If you are using Time Delay Replication Policy, also check the cache usage of the remaining production cluster TS7760/TS7720. Volumes can be removed from the TS7760/TS7720 only when the T copies are processed (either in the complete grid, or in the family).

DR Family

In R4.0, one DR Family can be defined. A DR Family can be defined, modified, and deleted with the Library Request command. After a flash is enabled, a DR Family cannot be modified.

At least one TS7760 or TS7720 must be part of the DR Family. You can optionally include one or more TS7740s. The TS7740 does not have the same functions in a DR Family that the TS7760/TS7720 has. The Write Protect excluded media categories needs to be consistent on all clusters in a DR Family. If they are not consistent, the FlashCopy cannot be enabled.

Creating a DR Family or adding a cluster to the DR Family

A DR Family can be created, or a cluster may be added to a previously created DR Family by using the following command (Example 13-1):

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, ADD, <CLUSTER ID>
```

Example 13-1 Create a DR Family and add a cluster

```
-LI REQ,HYDRAG,DRSETUP,DRFAM01,add,2
CBR1020I Processing LIBRARY command: REQ,HYDRAG,DRSETUP,DRFAM01,ADD,1.
CBR1280I Library HYDRAG request. 939
Keywords: DRSETUP,DRFAM01,ADD,1
-----
DRSETUP V1 0.0
DR FAMILY DRFAM01 WAS NEWLY CREATED

CLUSTER 1 WAS ADDED TO DR FAMILY DRFAM01 SUCCESSFULLY
```

Checking the Current Settings of a DR Family

After using any DRSETUP command for a particular DR Family, it is good to check the current status of the DR Family to ensure that it matches what you expect.

The settings for a DR Family can be checked by using the following command (Example 13-2):

```
LI REQ, <COMPOSITE>,DRSETUP, SHOW, <FAMILYNAME>
```

Example 13-2 Check the DR Family Settings

```
LI REQ,HYDRAG,DRSETUP,SHOW,DRFAM01
CBR1020I Processing LIBRARY command: REQ,HYDRAG,DRSETUP,SHOW,DRFAM01.
CBR1280I Library HYDRAG request. 302
Keywords: DRSETUP,SHOW,DRFAM01
-----
DRSETUP V1 0.0
DR FAMILY VIEW
```

```

ID  FAM NAME      FLASH      FLASH TIME (UTC)  LCOPY  MEMBER CLUSTERS
1   DRFAM01  INACTIVE                          N/A  FAMILY  - 1 2 - - - - -
-----
FAMILY MEMBER WRITE PROTECT STATUS VIEW
CLUSTER  WRT-PROTECT  EXCATS-NUM  IGNORE-FR  ENABLED-BY
CLUSTER1  DISABLED      3           TRUE       N/A
CLUSTER2  DISABLED      3           TRUE       N/A
-----
CATEGORIES EXCLUDED FROM WRITE PROTECTION WITHIN DR FAMILY  DRFAM01
CLUSTER  ACTIVE EXCLUDED CATEGORIES
CLUSTER1  0092 009F 3002
CLUSTER1  0092 009F 3002

```

13.4.1 Livecopy enablement in a DR Family

A DR Family must contain at least one TS7760 or TS7720. If a TS7740, TS7760T, or TS7720T is present within a DR Family, an option is available allowing the “live” copy on the TS7740/TS7760T/TS7720T cluster to be accessed if it is a completed replication or was otherwise consistent within the TS7740/TS7760T/TS7720T before Time Zero of the DR test. This is applicable if the TS7760/TS7720 removed its copy or if the TS7740 was the only target of the volume. This option is called LIVECOPY.

The purpose of LIVECOPY is to allow read access from a DR host to production volumes that were consistent before time zero of a FlashCopy and do not exist in cache on the FLASHed TS7760/TS7720 but do exist on a physical backend tape that is attached to a TS7700 or is in the cache of a TS7740. If a volume in this state is accessed from a DR host and LIVECOPY is enabled, the mount is satisfied. If a volume is in this state and LIVECOPY is NOT enabled, the mount fails. To ensure that during a DR test only data from Time Zero are used, all mounts need to be run on the TS7760/TS7720.

Important: Use the TS7740 in a DR Family only for remote mounts. Do not vary online the TS7740 devices directly to the DR host.

The option is disabled by default. If you choose to enable this functionality, you must explicitly enable the option using the library request command with “LIVECOPY” keyword as follows (Example 13-3):

```
LI REQ,<clib_name>,DRSETUP,<family_name>,LIVECOPY,FAMILY
```

Example 13-3 Enable the LIVECOPY option

```

LI REQ,HYDRAG,DRSETUP,DRFAM01,LIVECOPY,FAMILY
CBR1020I Processing LIBRARY command: REQ,HYDRAG,DRSETUP,DRFAM01,LIVECOPY
FAMILY.
CBR1280I Library HYDRAG request. 154
Keywords: DRSETUP,DRFAM01,LIVECOPY,FAMILY
-----
DRSETUP V1 0.0
LIVE COPY USAGE HAS BEEN UPDATED TO FAMILY SUCCESSFULLY

```

To disable the livecopy option, you must run the following command (Example 13-4):

```
LI REQ, <clib_name>, DRSETUP, <family_name>, LIVECOPY, NONE
```

Example 13-4 Disable the LIVECOPY option

```
LI REQ, HYDRAG, DRSETUP, DRFAM01, LIVECOPY, NONE
CBR1020I Processing LIBRARY command: REQ, HYDRAG, DRSETUP, DRFAM01, LIVECOPY
NONE.
CBR1280I Library HYDRAG request. 154
Keywords: DRSETUP, DRFAM01, LIVECOPY, NONE
-----
DRSETUP V1 0.0
LIVE COPY USAGE HAS BEEN UPDATED TO NONE SUCCESSFULLY
```

The livecopy setting is persistent. Disabling the FlashCopy does not change the setting. Only a complete deletion of the DR Family can change the setting. You can verify the current LIVECOPY setting using the **DRSETUP, SHOW** command. The output of this command contains a column titled "LCOPY". If the value under LCOPY is FAMILY, this indicates that LIVECOPY is active for the DR Family. If the value under LCOPY is NONE, LIVECOPY is not enabled for the DR Family. Example **DRSETUP, SHOW** output that shows a DR Family where LIVECOPY is enabled can be found in Example 13-2 on page 802.

Write Protect and FlashCopy enablement / disablement

The FlashCopy is based on a Write Protect Mode. You can enable the Write Protect Mode first and the FlashCopy later, or you can enable them together. If you want to disable the FlashCopy, you need first to disable the FlashCopy and later on the Write Protect Mode. Both of these actions can be run with a single command.

Note: A FlashCopy cannot be enabled if Write Protect Mode was enabled from the MI.

Do not enable the FlashCopy if production hosts with tape processing have device allocations on the clusters where the Flash will be enabled. Failures might occur because the read-only mode does not enable subsequent mounts.

Starting FlashCopy and Write Protect Mode for a DR Family

After a DR Family has been created and you are ready to initiate Write Protect Mode and the FlashCopy simultaneously, you can issue the following command to do so (Example 13-5):

```
LI REQ, <COMPOSITE>, DRSETUP, <FAMILYNAME>, DOALL, ENABLE
```

Example 13-5 Enable the FlashCopy

```
LI REQ, HYDRAG, DRSETUP, DRFAM01, DOALL, ENABLE
CBR1020I Processing LIBRARY command: REQ, HYDRAG, DRSETUP, DRFAM01, DOALL
ENABLE.
CBR1280I Library HYDRAG request. 154
Keywords: DRSETUP, DRFAM01, DOALL, ENABLE
-----
DRSETUP V1 0.0
WRITE PROTECT STATUS HAS BEEN ENABLED SUCCESSFULLY
FlashCopy HAS BEEN CREATED SUCCESSFULLY
```

13.4.2 Stopping FlashCopy and Write Protect Mode for a DR Family

After the cleanup from a DR test is complete, you can disable FlashCopy and Write Protect Mode for the DR Family by using the following command (Example 13-6):

Example 13-6 Disable the Write Protect and FlashCopy

```
LI REQ,HYDRAG,DRSETUP,DRFAM01,DOALL,DISABLE
CBR1020I Processing LIBRARY command: REQ,HYDRAG,DRSETUP,DRFAM01,DOALL
DISABLE.
CBR1280I Library HYDRAG request. 765
Keywords: DRSETUP,DRFAM01,DOALL,DISABLE
-----
DRSETUP V1 .0
WRITE PROTECT STATUS HAS BEEN DISABLED SUCCESSFULLY
FlashCopy HAS BEEN DELETED SUCCESSFULLY
```

Commands to check volume status during a DR test

During a DR test, you might want to check the status of these logical volumes that are involved in the DR test:

- ▶ Newly produced volumes from production
- ▶ Updated volumes from production
- ▶ Newly produced volumes from DR

You can use the following commands to identify if a FlashCopy exists for a specific volume, and the status from the livecopy and the FlashCopy.

To do so, use the **D SMS,VOL(xxxxxx)** and the **D SMS,VOL(xxxxxx),FLASH** commands. If the livecopy volume is identical to the FlashCopy volume, the status is **ACTIVE**. Only if the logical volume was updated from production, and a second instance exists, the status changes to **CREATED** (Example 13-7).

Example 13-7 Display of a logical volume after modification from production - Livecopy

```
LI REQ,HYDRAG,LVOL,A08760
CBR1020I Processing LIBRARY command: REQ,HYDRAG,LVOL,A08760.
CBR1280I Library HYDRAG request. 883
Keywords: LVOL,A08760
-----
LOGICAL VOLUME INFORMATION V3 0.0
LOGICAL VOLUME:          A08760
MEDIA TYPE:              ECST
COMPRESSED SIZE (MB):    2763
MAXIMUM VOLUME CAPACITY (MB): 4000
CURRENT OWNER:          cluster1
MOUNTED LIBRARY:
MOUNTED VNODE:
MOUNTED DEVICE:
TVC LIBRARY:             cluster1
MOUNT STATE:
CACHE PREFERENCE:       PG1
CATEGORY:               000F
LAST MOUNTED (UTC):     2014-03-11 10:19:47
LAST MODIFIED (UTC):    2014-03-11 10:18:08
LAST MODIFIED VNODE:    00
```

```

LAST MODIFIED DEVICE:      00
TOTAL REQUIRED COPIES:      2
KNOWN CONSISTENT COPIES:   2
KNOWN REMOVED COPIES:      0
IMMEDIATE-DEFERRED:        N
DELETE EXPIRED:            N
RECONCILIATION REQUIRED:    N
LWORM VOLUME:              N
FlashCopy:                 CREATED

```

LIBRARY	RQ	CACHE	PRI	PVOL	SEC	PVOL	COPY	ST	COPY	Q	COPY	CP
cluster1	N	Y	-----	-----			CMPT		-			RUN
cluster2	N	Y	-----	-----			CMPT		-			RUN

Example 13-8 shows the flash instance of the same logical volume.

Example 13-8 Display of a logical volume after modification from production - Flash volume

```

LI REQ,HYDRAG,LVOL,A08760,FLASH
CBR1020I Processing LIBRARY command: REQ,HYDRAG,LVOL,A08760,FLASH
CBR1280I Library HYDRAG request. 886
Keywords: LVOL,A08760,FLASH

```

```

LOGICAL VOLUME INFORMATION V3 0.0
FlashCopy VOLUME:          A08760
MEDIA TYPE:                 ECST
COMPRESSED SIZE (MB):       0
MAXIMUM VOLUME CAPACITY (MB): 4000
CURRENT OWNER:              cluster2
MOUNTED LIBRARY:
MOUNTED VNODE:
MOUNTED DEVICE:
TVC LIBRARY:                cluster1
MOUNT STATE:
CACHE PREFERENCE:           ---
CATEGORY:                   000F
LAST MOUNTED (UTC):         1970-01-01 00:00:00
LAST MODIFIED (UTC):        2014-03-11 09:05:30
LAST MODIFIED VNODE:
LAST MODIFIED DEVICE:
TOTAL REQUIRED COPIES:       -
KNOWN CONSISTENT COPIES:    -
KNOWN REMOVED COPIES:      -
IMMEDIATE-DEFERRED:        -
DELETE EXPIRED:            N
RECONCILIATION REQUIRED:    N
LWORM VOLUME:               -

```

LIBRARY	RQ	CACHE	PRI	PVOL	SEC	PVOL	COPY	ST	COPY	Q	COPY	CP
cluster2	N	Y	-----	-----			CMPT		-			RUN

Only the clusters from the DR Family are shown (in this case only a TS7720 was defined in the DR Family). This information is also available on the MI.

In Example 13-9 on page 808, you see a copy with an active, created FlashCopy. That means that the logical volume is not only in a write-protected category and part of the flash, but also that the logical volume was updated during the DR test. Therefore, the flash instance was created. The detail for last access by a host is the information from the livecopy (even on the DR Cluster).

To see the information from the created FlashCopy instance, select the FlashCopy CREATED field. This opens a second view, as shown in Figure 13-3.

Volser of Virtual Volume: [Get Details](#)

Virtual Volume Summary:

Virtual

Physical

Cluster [1] (#01052)

Cluster [2] (#00001)

Virtual volume details:

Volser	A10046
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	3,992.7 MiB
Maximum Volume Capacity (Device)	4,000 MiB
Current Owner	[1] (#01052)
Currently Mounted	No
vNode	-
Virtual Drive	-
Cached Copy Used for Mount	[2] (#00001)
Mount State	-
Cache Management Preference Group	1
Last Accessed by a Host	Feb 23, 2014, 12:15:14 PM
Last Modified	Feb 23, 2014, 12:14:32 PM
Category	001F
Storage Group	SGG00001
Management Class	MNNDN040
Storage Class	SC00000K
Data Class	D000N004
Volume Data State	Active
Flash Copy	Created
Earliest Deletion On	-
Logical WORM	No

Cluster-specific Virtual Volume Properties:

Figure 13-3 Display of a logical volume with an active FlashCopy

Figure 13-4 shows the next view, which is opened by clicking **Created**.

UR Family Name: DRH-AM01	
Flash Copy Time: Feb 23, 2014, 11:59:19 AM	
Flash Copy details:	
Volser	A10046
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	2,717.2 MiB
Maximum Volume Capacity (Device)	4,000 MiB
Current Owner	"[2]" (#00001)
Currently Mounted	No
vNode	-
Virtual Drive	-
Cached Copy Used for Mount	"[2]" (#00001)
Mount State	-
Cache Management Preference Group	1
Last Accessed by a Host	Feb 23, 2014, 11:59:19 AM
Last Modified	Feb 23, 2014, 11:59:05 AM
Category	001F
Storage Group	SGG00001
Management Class	MNNDN040
Storage Class	SC00000K
Data Class	D000N004
Volume Data State	Active
Earliest Deletion On	-

Figure 13-4 Display of the FlashCopy information of a logical volume

During the execution of a DR test, you may monitor the cache usage of your TS7760/TS7720 clusters. For the TS7760/TS7720 cluster used as DR, you have two new possibilities.

The following HCR command provides you information about the space that is used by the FlashCopy on the bottom of the output. See Example 13-9.

LI REQ,<distributed library name>,CACHE

Example 13-9 Cache Consumption FlashCopy

```
LI REQ,distributed library name,CACHE
CBR1280I Library VTSDIST1 request.
Keywords: CACHE
```

```
-----
```

TAPE VOLUME CACHE STATE V3 0.0									
PRIMARY TAPE MANAGED PARTITIONS									
INSTALLED/ENABLED GBS 0/ 0									
CACHE ENCRYPTION STATUS:									
PARTITION	ALLOC	USED	PG0	PG1	PMIGR	COPY	PMT	CPYT	
0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0

```

PRIMARY CACHE RESIDENT ONLY INFORMATION
INSTALLED/ENABLED GBS 95834/ 95834
ADJUSTED CACHE USAGE 5172
CACHE ENCRYPTION STATUS: CAPABLE
ALLOCATED    USED    PIN    PKP    PRM    COPY    CPYT
  95834    5151     0    5150     0     0     0
FlashCopy INFORMATION
INDEX  ENABLED  SIZE
  1     YES    252
  2     NO     0
  3     NO     0
  4     NO     0
  5     NO     0
  6     NO     0
  7     NO     0
  8     NO     0

```

You can find the same information on the MI as well. You can select the following display windows:

- ▶ Monitor
- ▶ Performance
- ▶ Cache Usage

Figure 13-5 is an example of Cache Utilization output.

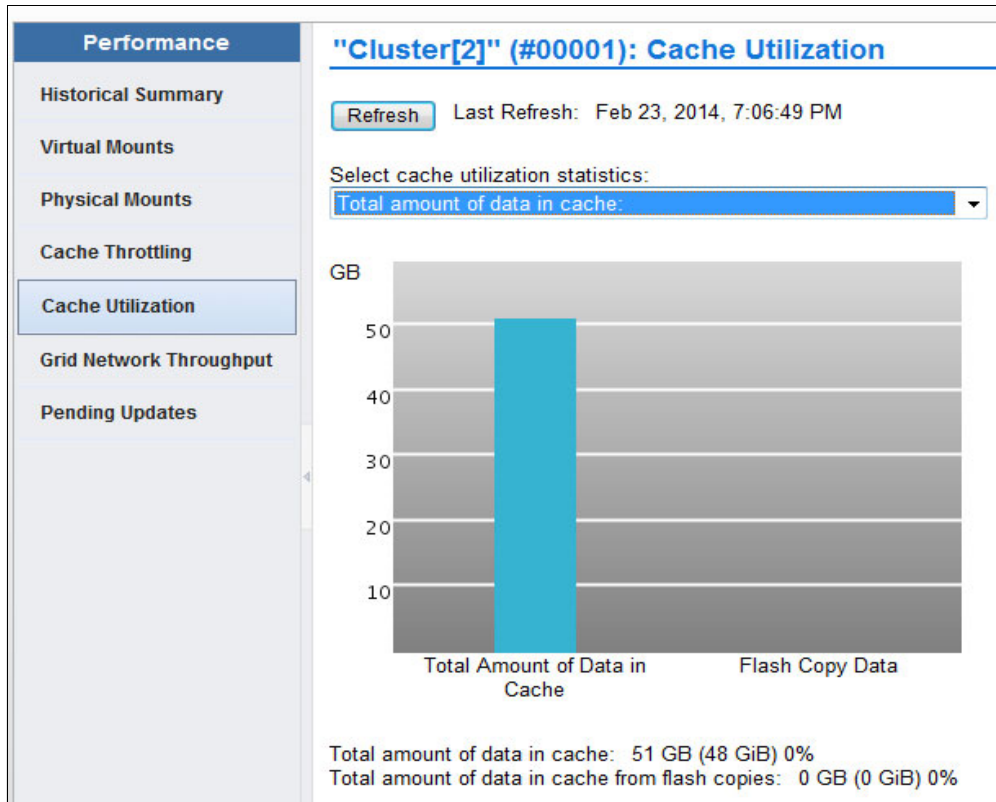


Figure 13-5 Cache usage of FlashCopy data

Also, you can control the usage of your virtual drives. You can select these displays on the MI:

- ▶ Virtual
- ▶ Virtual Tape Drives

Figure 13-6 is an example of virtual tape drive output.

Address	Mounted Volume	State	Time On Drive	Cache Mount Cluste	Mount Type
vtd00	ZKP002	idle - Write Protected	0 hours, 20 minutes, 42 seconds	#BA92C (2)	Flash Copy
vtd01	ZKD000	idle	0 hours, 19 minutes, 54 seconds	#BA92C (2)	Live Copy
vtd02	ZKP003	idle - Write Protected	0 hours, 19 minutes, 43 seconds	#BA92C (2)	Flash Copy
vtd03	ZKD005	idle	0 hours, 19 minutes, 29 seconds	#BA92C (2)	Live Copy
vtd04	ZKP004	idle - Write Protected	0 hours, 19 minutes, 3 seconds	#BA92C (2)	Flash Copy
vtd05	ZKP005	idle - Write Protected	0 hours, 18 minutes, 54 seconds	#BA92C (2)	Flash Copy
vtd06	ZKP006	idle - Write Protected	0 hours, 18 minutes, 47 seconds	#BA92C (2)	Flash Copy
vtd07	ZKP007	idle - Write Protected	0 hours, 18 minutes, 39 seconds	#BA92C (2)	Flash Copy
vtd08	ZKP008	idle - Write Protected	0 hours, 18 minutes, 32 seconds	#BA92C (2)	Flash Copy
vtd09	ZKP009	idle - Write Protected	0 hours, 18 minutes, 25 seconds	#BA92C (2)	Flash Copy
vtd0A					
vtd0B					
vtd0C					
vtd0D					
vtd0E					
vtd0F					
vtd10					
vtd11					

Figure 13-6 Virtual Tape Drive window during a FlashCopy for disaster recovery test

Considerations

DR tests have the following restrictions:

- ▶ There is no autoremoval of data from a TS7720 if the Flash is enabled.
- ▶ Do not perform the DR testing by using the FlashCopy function when a cluster in the grid is unavailable. An attempt to enable a FlashCopy in this situation results in a failure. You can perform the DR testing by using the FlashCopy function if all clusters in the grid are powered on (they can be in service/offline state).
- ▶ To perform the FlashCopy function, all clusters in a grid must be reachable through the grid links. Otherwise, host console commands to enable write protect mode or flash copy fail with an internal error.

13.5 DR Testing Methods Examples

Each method that is described in the following sections can be used as a step-by-step guide to running a DR test in a TS7700 grid environment. While it might be tempting to skip right to these lists, we advise that you review this chapter in its entirety before DR testing to ensure that you are familiar with the concepts and options available for DR testing in a TS7700 grid environment.

All of these methods have their advantages and disadvantages, so it is important that before you decide which method to use, that you weigh the advantages and disadvantages of each method against your environment and resources and then choose which method best fits your DR testing needs and ability.

Note: Each method assumes an independent DR site (DR host and at least one DR cluster). That is, it is assumed that no production hosts have had any devices online to the disaster clusters to read/write production data on those clusters.

13.5.1 Method 1: DR Testing using FlashCopy

The first method that you can choose to simulate a real disaster in your TS7700 grid uses FlashCopy on your disaster clusters.

The next section describes the steps that can be used to run DR testing using the FlashCopy functionality. For a detailed description of all commands, see *IBM Virtualization Engine TS7700 Series Best Practices - FlashCopy for Disaster Recovery Testing*, which is available at the Techdocs website (search for the term TS7700):

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

These steps were written in a checklist format to provide a reference of the steps that are needed to accomplish this method. It is advised that you review all of these steps before an actual DR exercise, as well as during the DR exercise. Because the steps were written to apply to more than one TS7700 grid configuration, make sure that before running each step that you understand each step and how it applies to your environment.

Method 1: DR Testing using FlashCopy: Steps

To perform DR using FlashCopy

1. Determine which MEDIA1, MEDIA2, ERROR, and PRIVATE categories will be used by the DR host during the DR test. These categories must be unique from any category that is used by a production system to ensure separation of production volumes from DR (disaster) volumes.
2. Using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that are chosen to the TS7700. This is done by selecting **Virtual** → **Categories** → **Add Scratch Category**.
3. Using the TS7700 MI, add the four new categories that are defined to the Exclude from Write Protect list in each cluster that will be in the DR Family. This is done by selecting **Settings** → **Cluster Settings** → **Write Protect Mode** → **Category Write Protect Properties** → **Add**.
 - a. If exclusion category counts are limited, MEDIA2 and PRIVATE are the most important to define. ERROR only needs to be added in order to allow a volume to be moved out of an ERROR state. MEDIA1 is only needed if MEDIA1 is used or any ACS (Automatic Class Selection) routine can result in using a default data class in which MEDIA1 is included (even if not used).
4. Issue the following command to determine the current status of the DR configuration within the TS7700 grid:

```
LI REQ, <COMPOSITE>,DRSETUP, SHOW
```

If there is already a DR Family that is defined, you must choose to either use it or delete it and start with a new configuration. To delete it, you must remove each cluster (remove a TS7760 or TS7720 cluster last) and when the last cluster has been removed, the DR Family will be automatically deleted. The following command can be used to remove a cluster:

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, REMOVE, <CLUSTER ID>
```

Wait for the command response that confirms that the cluster was removed before continuing. After the last cluster has been removed, the command response confirms that the DR Family was deleted because no members exist.

The steps that follow assume that you do not have a DR Family defined.

5. Create the DR Family that will be used for the DR test and add a cluster to the DR Family. If the DR Family will be composed of multiple clusters, add the TS7760/TS7720 first. This is done using the following command:

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, ADD, <CLUSTER ID>
```

Wait for the command response confirming that the DR Family was created and the cluster was added before continuing.

6. After the DR Family has been created, this command can be used repeatedly to add additional clusters to the DR Family.
7. Enable Write Protect Mode for the clusters in the DR Family by issuing the following command:

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, WP,ENABLE
```

Wait for the command response confirming that write protect has been enabled successfully.

8. Verify that Write Protect Mode is active for the clusters in the DR Family by issuing the following command:

```
LI REQ, <COMPOSITE>,DRSETUP, SHOW, <FAMILYNAME>
```

The above steps can often be completed in advance of a DR test. In fact, they can be set up once and left enabled indefinitely. If done far in advance, the only item to consider is that the Write Protect Mode would need to be disabled in the DR Family clusters in the event of a true DR event as part of the DR sequence.

9. Choose the volume serial ranges that will be created and used for input/output processing on the DR host. These are known as DR volumes. Update the TMS on the production hosts to ignore the ranges during host cartridge entry processing.
10. On each DR cluster, ensure that the Management Classes that will be used by the DR volumes on the DR hosts do not make copies to other clusters in the grid when the DR volumes are written to. If wanted, define a new management class on the DR clusters to be used explicitly for DR testing. On each DR cluster, set the 'Copy Mode' for the Management Class used to 'No Copy' for each non-DR cluster. This is done in **Constructs → Management Classes**.

11. IPL the DR host and restore the DR environment from the production environment.

12. Using the unique categories that are chosen in Step 1, define these MEDIA1, MEDIA2, ERROR and PRIVATE categories in the DEVSUPxx member on the DR host. These categories will be used by volumes created for use by the DR host during the DR test. After the categories are defined, IPL the DR host to ensure that the categories are used. Alternatively, the DS QL,CATS command can be used to dynamically set the categories without an IPL. If this alternative is used, be sure to update the DEVSUPxx as well with the new categories so that the categories will continue to be used if an IPL occurs.

13. If Livecopy usage is wanted, enable Livecopy by using the following command:

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, LIVECOPY,FAMILY
```

Wait for the command response confirming that the live copy usage is set to 'Family'.

14. Verify that the DR Family environment is as expected by using the following command:

```
LI REQ, <COMPOSITE>,DRSETUP, SHOW, <FAMILYNAME>
```


15. Update the TMS on the DR host to enable the new volume serial ranges to be accepted during host cartridge entry processing.
16. If a new management class was defined for each DR cluster in Step 10, modify the ACS routines on the DR host to direct new tape allocations to this management class. Activate the new SMS configuration on the DR host.
17. On the DR host, vary online the devices to the DR Family clusters that are either TS7760 or TS7720. DO NOT vary online devices in any TS7740 clusters.
18. Using the TS7700 MI, insert the new volume serial ranges. This can be done by selecting **Virtual → Virtual Volumes → Insert a new virtual volume range**. Verify that the DR Host has successfully processed the volumes during host cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes.
19. Change the Autoremoval Temporary Threshold on the TS7760/TS7720 used for DR testing to ensure that enough cache space is available for DR data and production data. This is only applicable for CPO and only if more than 10 TB is available in CPO. Wait until the removal process completes.
20. When you are ready to start the DR test, enable the FlashCopy using the following command:


```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, FLASH,ENABLE
```

 Wait for the command response confirming that the FlashCopy has been enabled successfully.
21. Verify that the DR Family environment is as expected by using the following command:


```
LI REQ, <COMPOSITE>,DRSETUP, SHOW, <FAMILYNAME>
```
22. Run the DR test.
23. When the DR test is complete, SCRATCH the volumes that were written to by the DR host. Be careful not to scratch production volumes. This can be done by using one of the following methods:
 - a. Run a TMS job on the DR host that scratches the volumes that are written to by the DR host (volumes that are converted from SCRATCH to PRIVATE during the DR test). The job must only include volumes that were inserted on the DR host (surfaced earlier in the CBR3610I messages).
 - b. Use ISMF to ALTER those volumes written to by the DR host to SCRATCH.
 - c. Use the CBRSP LCS SAMPLIB member to change the use attribute of each volume to SCRATCH
24. After all of the volumes that are processed by cartridge entry on the DR host are in SCRATCH status, they can be deleted from the TS7700 by using host EJECT processing. This can be done by using one of the following methods:
 - a. Run a TMS job to issue the EJECTs for these volumes
 - b. Use ISMF to EJECT each volume
 - c. Use the CBRSP LCS SAMPLIB member to eject each volume
25. Shutdown the DR host.
26. If the management class used on the DR cluster from Step 10 already existed before the DR test the 'Copy Mode' was updated for the DR test, change the 'Copy Mode' back to what it was before the DR test.
27. If you would like to keep the Write Protect Mode enabled and disable FlashCopy, use the following command:


```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, FLASH,DISABLE
```

Wait for the command response confirming that the FlashCopy has been deleted.

28. Alternatively, you can disable both Write Protect Mode and FlashCopy simultaneously by using the following command:

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, DOALL,DISABLE
```

Wait for the command response confirming that the FlashCopy has been deleted and write protect disabled.

29. Delete the DR Family. To delete it, you must remove each cluster (remove a TS7760 or TS7720 cluster last) and when the last cluster has been removed, the DR Family will be automatically deleted. The following command can be used to remove a cluster:

```
LI REQ, <COMPOSITE>,DRSETUP, <FAMILYNAME>, REMOVE, <CLUSTER ID>
```

Wait for the command response confirming that the cluster was removed before continuing. When the last cluster has been removed, the command response confirms that the DR Family was deleted because no members exist.

13.5.2 Method 2: Using Write Protect Mode on DR clusters

Another method that you can choose to use to simulate that of a real disaster is to use Write Protect Mode on your disaster clusters.

If you determine that FlashCopy for DR is not suitable to your DR environment, using the 'Write Protect Mode on DR clusters' method is the suggested alternative.

The following sections describe the steps that you can use to accomplish this method of DR testing. As with the previous method, these steps were written in a checklist format to provide a reference of the steps that are needed to accomplish this method. It is advised that you review all of these steps before an actual DR exercise, as well as during the DR exercise. Because the steps were written to apply to more than one TS7700 grid configuration, make sure that before running each step that you understand each step and how it applies to your environment.

Method 2: Using Write Protect Mode on DR clusters: Steps

To use Write Protect Mode on DR clusters, complete the following steps:

1. Determine which MEDIA1, MEDIA2, ERROR, and PRIVATE categories will be used by the DR host during the DR test. These categories must be unique from any category that is used by a production system to ensure separation of production volumes from disaster volumes.
2. Using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that were chosen to the TS7700. This is done by selecting **Virtual** → **Categories** → **Add Scratch Category**.
3. Using the TS7700 MI, add the four new categories that were defined to the Exclude from Write Protect list in each cluster that will be used as a DR cluster. This is done by selecting **Settings** → **Cluster Settings** → **Write Protect Mode** → **Category Write Protect Properties** → **Add**.

If exclusion category counts are limited, MEDIA2 and PRIVATE are the most important to define. ERROR only needs to be added in order to allow a volume to be moved out of an ERROR state. MEDIA1 is only needed if MEDIA1 is used or any ACS (Automatic Class Selection) routine can result in using a default data class in which MEDIA1 is included (even if not used).

4. Using the TS7700 MI, enable Write Protect Mode on each cluster that will be used as a DR cluster. This is done by selecting **SETTINGS** → **Cluster Settings** → **Write Protect Mode** → **Enable Write Protect Mode** → **Submit Changes**.

The previous steps can often be completed in advance of a DR test. In fact, they can be set up once and left enabled indefinitely. If done far in advance, the only item to consider is that the Write Protect Mode would need to be disabled in the DR clusters in the event of a true DR event as part of the DR sequence.

5. Choose the volume serial ranges that will be created and used for input/output processing on the DR host. Update the TMS on the production hosts to ignore the ranges during host cartridge entry processing.
6. On each DR cluster, ensure that the Management Classes that will be used by the DR volumes on the DR hosts do not make copies to other clusters in the grid when the DR volumes are written to. If wanted, define a new management class on the DR clusters to be used explicitly for DR testing. On each DR cluster, set the 'Copy Mode' for the Management Class used to 'No Copy' for each non-DR cluster. This is done in **Constructs → Management Classes**.
7. Restart the DR host and restore the DR environment from the production environment.
8. Using the unique categories that were chosen in Step 1, define these MEDIA1, MEDIA2, ERROR, and PRIVATE categories in the DEVSUPxx member on the DR host. These categories will be used by volumes created for use by the DR host during the DR test. After the categories are defined, IPL the DR host to ensure that the categories are used. Alternatively, the **DS QL,CATS** command can be used to dynamically set the categories without an IPL. If this alternative is used, be sure to update the DEVSUPxx as well with the new categories so that the categories will continue to be used if an IPL occurs.
9. Update the TMS on the DR host to enable the new volume serial ranges to be accepted during host cartridge entry processing.
10. If a new management class was defined for each DR cluster in Step 6, modify the ACS routines on the DR host to direct new tape allocations to this management class. Activate the new SMS configuration on the DR host.
11. On the DR host, vary online the devices to the DR clusters.
12. Using the TS7700 MI, insert the new volume serial ranges. This can be done by selecting **Virtual → Virtual Volumes → Insert a new virtual volume range**. Verify that the DR Host has successfully processed the volumes during cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes.
13. Run the DR test.
14. When the DR test is complete, SCRATCH the volumes that were written to by the DR host. Be careful not to scratch production volumes. This can be done using one of the following methods:
 - a. Run a TMS job on the DR host that scratches the volumes that are written to by the DR host (volumes that are converted from SCRATCH to PRIVATE during the DR test). The job must include only volumes that were inserted on the DR host (surfaced earlier in the CBR3610I messages).
 - b. Use ISMF to ALTER those volumes written to by the DR host to SCRATCH.
 - c. Use the CBRSP LCS SAMPLIB member to change the use attribute of each volume to SCRATCH.
15. When all of the volumes that are processed by cartridge entry on the DR host are in SCRATCH status, they can be deleted from the TS7700 by using host EJECT processing. This can be done using one of the following methods:
 - a. Run a TMS job to issue the EJECTs for these volumes.
 - b. Use ISMF to EJECT each volume.
 - c. Use the CBRSP LCS SAMPLIB member to eject each volume.

16. Shutdown the DR host.
17. If the management class used on the DR cluster from Step 6 already existed before the DR test and the 'Copy Mode' was updated for the DR test, change the 'Copy Mode' back to what it was before the DR test.
18. If you would like to keep the Write Protect Mode enabled on the DR clusters, it is a good precaution to take to prevent accidentally returning production tapes to SCRATCH by the DR host.
19. Alternatively, you can disable Write Protect Mode in each DR cluster by using the TS7700 MI. This is done by selecting **SETTINGS** → **Cluster Settings** → **Write Protect Mode** → **Disable Write Protect Mode** → **Submit Changes**.

13.5.3 Method 3: DR Testing without Write Protect Mode

Another method that you can choose to use to simulate that of a real disaster is to run your DR test from a DR host that is attached to one or more disaster clusters while the production hosts continue to write data across the grid.

If your choice is between using Write Protect Mode and not using Write Protect Mode, it is recommended to use Write Protect Mode (Method 2), to provide an additional level of write-protection in case the TMS on the disaster host is not configured correctly to prevent writes to the production volumes.

Described in the following sections are the steps that you can use to accomplish this method of DR testing. As with the previous method, these steps were written in a checklist format to provide a reference of the steps that are needed to accomplish this method. It is advised that you review all of these steps before an actual DR exercise, as well as during the DR exercise. Because the steps were written to apply to more than one TS7700 grid configuration, make sure that before running each step that you understand each step and how it applies to your environment.

Method 3: DR testing without Write Protect Mode: Steps

To perform DR testing without Write Protect Mode, complete the following steps:

1. Determine which MEDIA1, MEDIA2, ERROR, and PRIVATE categories will be used by the DR host during the DR test. These categories must be unique from any category used by a production system to ensure separation of production volumes from disaster volumes.
2. Using the TS7700 MI, add the MEDIA1 and MEDIA2 categories chosen to the TS7700. This is done by selecting **Virtual** → **Categories** → **Add Scratch Category**.
3. Choose the volume serial ranges that will be created and used for input/output processing on the DR host. Update the TMS on the production hosts to ignore the ranges during host cartridge entry processing.
4. On each DR cluster, ensure that the Management Classes that will be used by the DR volumes on the DR hosts do not make copies to other clusters in the grid when the DR volumes are written to. If wanted, define a new management class on the DR clusters to be used explicitly for DR testing. On each DR cluster, set the 'Copy Mode' for the Management Class used to 'No Copy' for each non-DR cluster. This is done in **Constructs** → **Management Classes**.
5. Restart the DR host and restore the DR environment from the production environment.
6. Using the unique categories that were chosen in Step 1, define these MEDIA1, MEDIA2, ERROR, and PRIVATE categories in the DEVSUPxx member on the DR hosts. These categories will be used by volumes created for use by the DR host during the DR test.

After the categories are defined, Restart the DR host to ensure that the categories are used. Alternatively, the DS QL,CATS command can be used to dynamically set the categories without an IPL. If this alternative is used, be sure to update the DEVSUPxx as well with the new categories so that the categories will continue to be used if an IPL occurs.

7. Update the TMS on the DR host to enable the new volume serial ranges to be accepted during host cartridge entry processing.
8. Update the TMS on the DR host to reject any output that is directed towards volumes in the production volume serial range. This is done as a safeguard to protect against production tapes being accidentally written to by the DR host.
9. If a new management class was defined for each DR cluster in Step 4, modify the ACS routines on the DR host to direct new tape allocations to this management class. Activate the new SMS configuration on the DR host.
10. On the DR host, vary online the devices to the DR clusters.
11. Using the TS7700 MI, insert the new volume serial ranges. This can be done by selecting **Virtual** → **Virtual Volumes** → **Insert a new virtual volume range**. Verify that the DR Host has processed the volumes during cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes.
12. Mark all DFSMSShsm ML2 volumes full by issuing the following command to DFSMSShsm:
 - a. F xxxxx,DELVOL MIGRATION(MARKFULL)
 - b. Run the F xxxxx,HOLD RECYCLE command
13. For maximum protection, ensure the following procedures DO NOT run:
 - a. RMM housekeeping activity at the DR site
 - b. Short-on-scratch RMM procedures at the DR site
14. Run the DR test
15. When the DR test is complete, SCRATCH the volumes that were written to by the DR host. Be careful not to scratch production volumes. This can be done using one of the following methods:
 - a. Run a TMS job on the DR host that scratches the volumes that are written to by the DR host (volumes that are converted from SCRATCH to PRIVATE during the DR test). The job must include only volumes that were inserted on the DR host (surfaced earlier in the CBR3610I messages).
 - b. Use ISMF to ALTER those volumes written to by the DR host to SCRATCH.
 - c. Use the CBRSP LCS SAMPLIB member to change the use attribute of each volume to SCRATCH.
16. After all the volumes that are processed by cartridge entry on the DR host are in SCRATCH status, they can be deleted from the TS7700 by using host EJECT processing. This can be done by using one of the following methods:
 - a. Run a TMS job to issue the EJECTs for these volumes
 - b. Use ISMF to EJECT each volume
 - c. Use the CBRSP LCS SAMPLIB member to eject each volume
17. Shutdown the DR host.
18. If the management class used on the DR cluster from Step 4 already existed before the DR test, and the 'Copy Mode' was updated for the DR test, change the 'Copy Mode' back to what it was before the DR test.

13.5.4 Method 4: Breaking the grid link connections

The final method that will be discussed that you can choose to simulate that of a real disaster is to break the grid links between the production clusters and disaster clusters in a TS7700 grid.

The concern about losing data in a real disaster during a DR test is the major drawback to using this DR method. Because of this, if it is possible to use one of the DR methods described earlier (using FlashCopy or Write Protect Mode), it is suggested to use one of those methods.

Important: Do not use logical drives in the DR site from the production site.

Described below are the steps that you may use to accomplish this method of DR testing. As with the previous methods, these steps were written in a checklist format to provide a reference of the steps needed to accomplish this method. It is suggested that you review all of these steps before an actual DR exercise, as well as during the DR exercise. As the steps were written to apply to more than one TS7700 grid configuration, make sure that before running each step that you understand each step and how it applies to your environment.

Method 4: Breaking the grid link connections: Steps

To break the grid link connections, complete the following steps:

1. Determine which MEDIA1, MEDIA2, ERROR, and PRIVATE categories will be used by the DR host during the DR test. These categories must be unique from any category that is used by a production system to ensure separation of production volumes from disaster volumes.
2. Using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that are chosen to the TS7700. This is done by selecting **Virtual** → **Categories** → **Add Scratch Category**.
3. Choose the volume serial ranges that will be created and used for input/output processing on the DR host. Update the TMS on the production hosts to ignore the ranges during host cartridge entry processing.
4. On each DR cluster, ensure that the Management Classes that will be used by the DR volumes on the DR hosts do not make copies to other clusters in the grid when the DR volumes are written to. If wanted, define a new management class on the DR clusters to be used explicitly for DR testing. On each DR cluster, set the 'Copy Mode' for the Management Class used to 'No Copy' for each non-DR cluster. This is done in **Constructs** → **Management Classes**.
5. Restart the DR host and restore the DR environment from the production environment.
6. Using the unique categories that were chosen in Step 1, define these MEDIA1, MEDIA2, ERROR, and PRIVATE categories in the DEVSUPxx member on the DR host. These categories will be used by volumes created for use by the DR host during the DR test. After the categories are defined, restart the DR host to ensure that the categories are used. Alternatively, the DS QL,CATS command can be used to dynamically set the categories without an IPL. If this alternative is used, be sure to update the DEVSUPxx as well with the new categories so that the categories will continue to be used if an IPL occurs.
7. Update the TMS on the DR host to enable the new volume serial ranges to be accepted during host cartridge entry processing.
8. Update the TMS on the DR host to reject any output that is directed towards volumes in the production volume serial range. This is done as a safe-guard to protect against production tapes being accidentally written to by the DR host.

9. If a new management class was defined for each DR cluster in Step 4, modify the ACS routines on the DR host to direct new tape allocations to this management class. Activate the new SMS configuration on the DR host.
10. On the DR host, vary online the devices to the DR clusters.
11. Using the TS7700 MI, insert the new volume serial ranges. This can be done by selecting **Virtual → Virtual Volumes → Insert a new virtual volume range**. Verify that the DR Host has successfully processed the volumes during cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes.
12. Engage your CE to break the grid link connections between the production clusters and the disaster clusters. As mentioned earlier, DO NOT disable a grid link with the Library Request command. Disabling the grid link with the command does not stop synchronous mode copies and the exchange of status information.
13. Use the TS7700 MI to enable read-only takeover mode on each disaster cluster to allow read-only access to volumes owned by each production cluster. This is done by selecting **Service → Ownership Takeover Mode**.
14. Mark all DFSMSShsm ML2 volumes full by issuing the following command to DFSMSShsm:
 - a. F xxxxx,DELVOL MIGRATION(MARKFULL).
 - b. Run the F xxxxx,HOLD RECYCLE command.
15. For maximum protection, ensure that the following procedures DO NOT run:
 - a. RMM housekeeping activity at the DR site
 - b. Short-on-scratch RMM procedures at the DR site
16. Run the DR test.
17. When the DR test is complete, SCRATCH the volumes that were written to by the DR host. Be careful not to scratch production volumes. Use one of the following methods:
 - a. Run a TMS job on the DR host that scratches the volumes that are written to by the DR host (volumes that are converted from SCRATCH to PRIVATE during the DR test). The job must include only volumes that were inserted on the DR host (surfaced earlier in the CBR3610I messages).
 - b. Use ISMF to ALTER those volumes written to by the DR host to SCRATCH.
 - c. Use the CBRSP LCS SAMPLIB member to change the use attribute of each volume to SCRATCH.
18. After all the volumes that are processed by cartridge entry on the DR host are in SCRATCH status, they can be deleted from the TS7700 by using host EJECT processing. This can be done using one of the following methods:
 - a. Run a TMS job to issue the EJECTs for these volumes.
 - b. Use ISMF to EJECT each volume.
 - c. Use the CBRSP LCS SAMPLIB member to eject each volume.
19. Shut down the DR host.
20. If the management class used on the DR cluster from Step 4 already existed before the DR test and the 'Copy Mode' was updated for the DR test, change the 'Copy Mode' back to what it was before the DR test.
21. Use the TS7700 MI to disable read-only takeover mode on each disaster cluster. This returns each disaster cluster to its normal state regarding takeover processing. Disable read/write takeover mode on the production clusters as well. This is done by using **Service → Ownership Takeover Mode**.
22. Engage your CE to re-establish the link connection between the production clusters and the disaster clusters.

13.6 Expected failures during a DR test

The next section covers some expected failures during a DR test.

The messages in Example 13-10 might appear if you try to read a logical volume that was not present at time zero in the DR Family.

Example 13-10 Expected failures during the disaster recovery test

```
IEF233A M 2500,A08759,,DENEKA1,STEP1,DENEKA.HG.TEST1.DUMP1
CBR4195I LACS retry possible for job DENEKA1: 399
IEE763I NAME= CBRLACS CODE= 140394
CBR4000I LACS WAIT permanent error for drive 2500.
CBR4171I Mount failed. LVOL=A08759, LIB=HYDRAG,
PVOL=??????,RSN=22
```

The message in Example 13-11 might also appear if you want to modify a volume that is in a write protect media category.

Example 13-11 Error message for volume in a write media category

```
IEF116I DENEKY6 STEP1 - MOUNT OF VOLUME PRIVAT ON DEVICE 2580 FAILED
IEE763I NAME= CBRLACS CODE= 14017E
CBR4000I LACS MOUNT permanent error for drive 2580.
CBR4126I Library HYDRAG drive is in read only mode.
IEF272I DENEKY6 STEP1 - STEP WAS NOT run
```

The message in Example 13-12 might occur if a job was running on the cluster while the FlashCopy was enabled.

Example 13-12 Message for job running on the cluster while FlashCopy was enabled

```
IEF233A M 2507,A10088,,DENEKA8,STEP2,DENEKA.HG.TEST1.DUMP1
IEC518I SOFTWARE ERRSTAT: WRITPROT 2507,A10088,SL,DENEKA8,STEP2
IEC502E RK 2507,A10088,SL,DENEKA8,STEP2
IEC147I 613-24,IFG0194F,DENEKA8,STEP2,AUS1,2507,,DENEKA.HG.TEST1.DUMP1
```

Appendixes

This part offers management and operational information for your IBM TS7700.

This part contains the following appendixes:

- ▶ Feature codes and RPQ
- ▶ IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments
- ▶ JES3 examples and information
- ▶ DEVSERV QLIB command
- ▶ Sample job control language
- ▶ Library Manager volume categories
- ▶ IBM TS7700 parameter examples
- ▶ Extra IODF examples
- ▶ Case study for logical partitioning of a two-cluster grid



A

Feature codes and RPQ

This appendix lists all feature codes (FC) and requests for price quotation (RPQs) that are related to the installation of IBM TS7700. It provides several quick-reference lists.

Exception: This appendix provides a general description of feature codes and to where they apply. Use the following link to the IBM Product IBM Knowledge Center for technical information related to the TS7700:

http://www.ibm.com/support/knowledgecenter/STFS69_4.0.0/ts7700_feature_codes_a11.html

This appendix includes the following sections:

- ▶ RPQ
- ▶ Feature code lists

RPQ

This section describes any RPQs that are available for the TS7700 according to component, machine type, and model.

3952 F06 RPQ

Available by RPQ is the ability for top exit as it relates to cables in the 3952 F06 frame. An example is RPQ **8B3670**.

Feature code lists

This section lists feature code descriptions for the TS7700 according to component, machine type, and model.

Clarification: The symbol “†” means that the specific feature has been withdrawn.

3952 F05 features

The F05 has the following features:

- ▶ FC 1903, Dual AC, power distribution unit
- ▶ FC 1904, Optional AC, switching power distribution unit
- ▶ FC 2704, Console Expansion 26 Port Ethernet switch/rackmount
- ▶ FC 2719, Console upgrade†
- ▶ FC 2725, TS3000 System Console Rackmount
- ▶ FC 2730, TS3000 System Console†
- ▶ FC 2732, TS3000 System Console†
- ▶ FC 2733, Internal modem†
- ▶ FC 2748, Optical drive
- ▶ FC 4743, Remove 3957-V06/VEA
- ▶ FC 5512, TS3000 System Console KVM keyboard, video, and mouse
- ▶ FC 5626, Plant install of 3957-VEA†
- ▶ FC 5627, Install 3957-VEB
- ▶ FC 5628, Plant install of 3957-V06†
- ▶ FC 5629, Install 3957-V07
- ▶ FC 5635, Plant install 3956-CS8†
- ▶ FC 5636, Plant install 3956-CS7†
- ▶ FC 5638, Plant install 3956-CC6†
- ▶ FC 5639, Plant install 3956-CC7†
- ▶ FC 5640, Plant install 3956-CC8†
- ▶ FC 5641, Plant install 3956-CX7†
- ▶ FC 5642, Field install 3956-CX7†
- ▶ FC 5646, Plant install 3956-XS7†
- ▶ FC 5647, Field install 3956-Xi
- ▶ FC 5648, Plant install 3956-CX6†
- ▶ FC 5649, Field install 3956-CX6†
- ▶ FC 5651, Plant install 3956-CS9
- ▶ FC 5652, Plant install 3956-CC9
- ▶ FC 5653, Plant install 3956-CX9
- ▶ FC 5654, Field install 3956-CX9

- ▶ FC 5655, Plant install 3956-XS9
- ▶ FC 5656, Field install 3956-XS9
- ▶ FC 5758, Integrated control path
- ▶ FC 5759, Integrated control path†
- ▶ FC 7312, TS7740 Base frame †
- ▶ FC 7322, TS7720 Base frame †
- ▶ FC 7323, TS7720 Storage expansion frame†
- ▶ FC 7330, TS7740 Encryption capable base frame
- ▶ FC 7331, TS7720 Encryption capable base frame
- ▶ FC 7332, TS7720 Encryption capable expansion frame
- ▶ FC 9110, Ship with R1.7 machine code†
- ▶ FC 9111, Ship with R2.0 machine code†
- ▶ FC 9112, Ship with R2.1 machine code†
- ▶ FC 9113, ship with R3.0 machine code†
- ▶ FC 9114, Ship with R3.1 machine code†
- ▶ FC 9115, Ship with R3.2 machine code†
- ▶ FC 9116, Ship with R3.3 machine code†
- ▶ FC 9323, Expansion frame attachment
- ▶ FC 9339, Replacing TS7740 System
- ▶ FC 9954, NEMA L6-30 power cord
- ▶ FC 9955, RS 9750 DP power cord
- ▶ FC 9956, IEC 309 power cord
- ▶ FC 9957, PDL 4.3 power cord
- ▶ FC 9958, Korean 4.3-m power cord
- ▶ FC 9959, Unterminated power cord
- ▶ FC 9966, Unterminated power cord (China)

3952 F06 features

The F06 has the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 1903, Dual AC, power distribution unit
- ▶ FC 1904, Optional AC, switching power distribution unit
- ▶ FC 2704, Console Expansion - 26 Port Ethernet switch/rackmount
- ▶ FC 2725, TS3000 System Console Rackmount
- ▶ FC 2748, Optical drive
- ▶ FC 5512, TS3000 System Console KVM keyboard, video, and mouse
- ▶ FC 5630, Install 3957 VEC
- ▶ FC 5657, Plant install 3956 CSA
- ▶ FC 5658, Plant Install 3956 XSA
- ▶ FC 5659, Field Install 3956 XSA
- ▶ FC 5758, Integrated Control Path
- ▶ FC 7333, TS7700 Encryption Capable Base Frame
- ▶ FC 7334, TS7700 Encryption Capable Expansion Frame
- ▶ FC 9323, Expansion Frame Attach
- ▶ FC 9339, Replacing TS7740 System
- ▶ FC 9954, NEMA L6-30 power cord
- ▶ FC 9955, RS 9750 DP power cord
- ▶ FC 9956, IEC 309 power cord
- ▶ FC 9957, PDL 4.3 power cord
- ▶ FC 9958, Korean 4.3-m power cord
- ▶ FC 9959, Unterminated power cord
- ▶ FC 9966, Unterminated power cord (China)
- ▶ FC AGK0, Ship with R4.0 Machine Code

Server features for 3957-V07, 3957-VEB and 3957-VEC

This section lists the server features for 3957-V07, 3957-VEB, and 3957-VEC.

3957-V07 Server features

The V07 has the following features:

- ▶ FC 0201, 9-micron LC/LC 31-meter
- ▶ FC 0203, 50-micron LC/LC 31-meter
- ▶ FC 1034, Enable dual port grid connection
- ▶ FC 1035, 10 Gb Grid optical LW connection
- ▶ FC 1036, 1 Gb grid dual port copper connection
- ▶ FC 1037, 1 Gb grid dual port optical sw connection
- ▶ FC 2714, Console expansion†
- ▶ FC 2715, Console attachment
- ▶ FC 3401, Enable 8 Gb FICON dual port
- ▶ FC 3438, 8 Gb FICON Short Wavelength Attachment
- ▶ FC 3439, 8 Gb FICON Long Wavelength Attachment
- ▶ FC 3441, FICON short-wavelength attachment
- ▶ FC 3442, FICON long-wavelength attachment
- ▶ FC 3443, FICON 10-km long-wavelength attachment
- ▶ FC 3462 Memory Upgrade
- ▶ FC 4015, Grid enablement
- ▶ FC 4016, Remove cluster from grid
- ▶ FC 4017, Cluster cleanup
- ▶ FC 5241, Dual port FC HBA
- ▶ FC 5267, 1 TB cache enablement
- ▶ FC 5268, 100 MiB/s increment
- ▶ FC 5270, Increased logical volumes
- ▶ FC 5271, Selective device access control
- ▶ FC 5272, Enable disk encryption - Local Key Management
- ▶ FC 5275, Additional virtual devices
- ▶ FC 5276, Enable disk encryption - External Key Management
- ▶ FC 5277, External Disk Encryption Certificate - Field
- ▶ FC 9000, Mainframe attachment
- ▶ FC 9111, Ship with R2.0 machine code†
- ▶ FC 9112, Ship with R2.1 machine code†
- ▶ FC 9113, Ship with R3.0 machine code†
- ▶ FC 9114, Ship with R3.1 machine Code†
- ▶ FC 9115, Ship with R3.2 machine Code†
- ▶ FC 9116, Ship with R3.3 machine Code†
- ▶ FC 9219, Attach to TS3500/TS4500
- ▶ FC 9277, External Disk Encryption Certificate - Plant
- ▶ FC 9350, Plant install TS7700 Server in 3952 F05†
- ▶ FC 9351, Field merge TS7700 Server in 3952 F05
- ▶ FC 9700, No factory cables†
- ▶ FC 9900, Tape Encryption configuration

3957-VEB Server features

The VEB has the following features:

- ▶ FC 0201, 9-micron LC/LC 31-meter
- ▶ FC 0203, 50-micron LC/LC 31-meter
- ▶ FC 1034, Enable dual port grid connection
- ▶ FC 1035, 10 Gb Grid optical LW connection

- ▶ FC 1036, 1 Gb grid dual port copper connection
- ▶ FC 1037, 1 Gb dual port optical SW connection
- ▶ FC 2714, Console expansion†
- ▶ FC 2715, Console attachment
- ▶ FC 3401, Enable 8 Gb FICON dual port
- ▶ FC 3438, 8 Gb FICON Short Wavelength Attachment
- ▶ FC 3439, 8 Gb FICON Long Wavelength Attachment
- ▶ FC 3441, FICON short-wavelength attachment
- ▶ FC 3442, FICON long-wavelength attachment
- ▶ FC 3443, FICON 10-km long-wavelength attachment
- ▶ FC 3462, Memory Upgrade
- ▶ FC 4015, Grid enablement
- ▶ FC 4016, Remove cluster from grid
- ▶ FC 4017, Cluster cleanup
- ▶ FC 5241, Dual port FC HBA
- ▶ FC 5268, 100 MiB/s increment
- ▶ FC 5270, Increased logical volumes
- ▶ FC 5271, Selective device access control DAC
- ▶ FC 5272, Enable disk encryption - Local Key Management
- ▶ FC 5273, TS7720/TS7760 Tape Attach enablement
- ▶ FC 5274, Enable 1 TB Pending Tape Capacity
- ▶ FC 5276, Enable disk encryption - External Key Management
- ▶ FC 5277, External Disk Encryption Certificate - Field
- ▶ FC 5275, Additional virtual devices
- ▶ FC 9000, Mainframe attachment
- ▶ FC 9111, Ship with R2.0 machine code†
- ▶ FC 9112, Ship with R2.1 machine code†
- ▶ FC 9113, Ship with R3.0 machine code†
- ▶ FC 9114, Ship with R3.1 machine code†
- ▶ FC 9115, Ship with R3.2 machine code†
- ▶ FC 9116, Ship with R3.3 machine Code†
- ▶ FC 9219, Attach to TS3500/TS4500
- ▶ FC 9268, Plant install 100 MiB/s throughput
- ▶ FC 9277, External Disk Encryption Certificate - Plant
- ▶ FC 9350, Plant install TS7700 Server in 3952 F5
- ▶ FC 9351, Field merge TS7700 Server in 3952 F05
- ▶ FC 9700, No factory cables
- ▶ FC 9900, Tape Encryption Enablement - optional for TS7720T

3957-VEC Server features

- ▶ FC 0201, 9-micron LC/LC 31-meter
- ▶ FC 0203, 50-micron LC/LC 31-meter
- ▶ FC 0983, TAA Compliance
- ▶ FC 1034, Enable dual port grid connection
- ▶ FC 1036, 1 Gb grid dual port copper connection
- ▶ FC 1037, 1 Gb dual port optical SW connection
- ▶ FC 1038, 10 Gb dual port optical LW connection
- ▶ FC 2715, Console attachment
- ▶ FC 3401, Enable 8 Gb FICON dual port
- ▶ FC 3438, 8 Gb FICON Short Wavelength Attachment
- ▶ FC 3439, 8 Gb FICON Long Wavelength Attachment
- ▶ FC 4015, Grid enablement
- ▶ FC 4016, Remove Cluster from Grid
- ▶ FC 4017, Cluster cleanup
- ▶ FC 5242, Dual Port 16 Gb Fibre Channel HBA

- ▶ FC 5268, 100 MBps increment
- ▶ FC 5270, Increased virtual volumes
- ▶ FC 5271, Selective device access control
- ▶ FC 5272, Enable disk encryption - Local Key Management
- ▶ FC 5273, TS7720/TS7760 Tape Attach enablement
- ▶ FC 5274, Enable 1 TB Pending Tape Capacity
- ▶ FC 5275, Additional virtual devices
- ▶ FC 5276, Enable disk encryption - External Key Management
- ▶ FC 9000, Mainframe attachment
- ▶ FC 9219, Attach to TS3500/TS4500
- ▶ FC 9268, Plant installation of 100 MB/s throughput
- ▶ FC 9350, Plant Install
- ▶ FC 9700, No factory cables
- ▶ FC 9900, Tape Encryption enablement
- ▶ FC AGK0, Ship with R4.0 Machine Code

Cache Controller features for 3956-CC9, 3956-CS9 and 3956-CSA

Feature codes that are available for the TS7700 3956 cache controller are listed by model.

Clarification: The symbol “†” means that specific feature has been withdrawn.

3956-CC9 Encryption Capable Cache Controller Drawer features

The CC9 has the following features:

- ▶ FC 7124, 13.2 TB SAS Storage
- ▶ FC 7404, Encryption plant or field
- ▶ FC 9352, Plant install a TS7740 Encryption Capable Cache Controller Drawer in a 3952 F05

3956-CS9 Encryption Capable Cache Controller Drawer features

The CS9 has the following features:

- ▶ FC 7115, 36 TB SAS Storage
- ▶ FC 7404, Encryption plant or field
- ▶ FC 9352, Plant install a TS7720 Encryption Capable Cache Controller Drawer in a 3952 F05

3956-CSA Encryption Capable Cache Controller Drawer features

The CSA has the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 7117, 48 TB SAS Storage
- ▶ FC 7404, Encryption
- ▶ FC 9352, Plant install

TS7700 Cache Drawer features 3956-CX9, 3956-XS9 and 3956-XSA

Feature codes available for the Cache Drawer are listed by model in this section.

Clarification: The symbol “†” means that specific feature has been withdrawn.

3956-CX9 Encryption Capable Cache Expansion Drawer features

The CX9 has the following features:

- ▶ FC 7124, 13.2 TB SAS Storage
- ▶ FC 7404, Encryption plant or field
- ▶ FC 9354, Plant install a TS7740 Encryption Capable Cache Expansion Drawer in a 3952 F05
- ▶ FC 9355, Field merge a TS7740 Encryption Capable Cache Expansion Drawer in a 3952 F05

3956-XS9 Encryption Capable Cache Expansion Drawer features

The XS9 has the following features:

- ▶ FC 7116, 33 TB SAS Storage
- ▶ FC 7404, Encryption plant or field
- ▶ FC 9354, Plant install a TS7720 Encryption Capable Cache Expansion Drawer in a 3952 F05
- ▶ FC 9355, Field merge a TS7720 Encryption Capable Cache Expansion Drawer in a 3952 F05

3956-XSA Encryption Capable Cache Expansion Drawer features

The XSA has the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 7117, 48 TB SAS Storage
- ▶ FC 7404, Encryption
- ▶ FC 9354, Plant Install
- ▶ FC 9355, Field Merge



IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments

This appendix documents the considerations for implementation and operation for IBM z/VM, IBM z/VSE, and IBM z/Transaction Processing Facility (IBM z/TPF) environments.

For more information, see the following documentation:

- ▶ *z/VM: DFSMS/VM Removable Media Services*, SC24-6185
- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *z/VSE V6R1.0 Administration*, SC34-2627

This appendix includes the following sections:

- ▶ Software requirements
- ▶ Software implementation in z/VM and z/VSE
- ▶ Software implementation in z/OS Transaction Processing Facility
- ▶ Implementing Outboard Policy Management for non-z/OS hosts

Software requirements

The following software products are required:

- ▶ IBM z/VM V5.4.0, or later

With z/VM, the TS7760, TS7740, and TS7720 are transparent to host software. z/VM V5R4 or later with the program temporary fixes (PTFs) for authorized program analysis report (APAR) VM64979 is required for both guest and native VM support. DFSMS/VM Function Level 221 with PTFs for Removable Media Services (RMS) APAR VM64773 and VM65005 is required for native VM tape library support. Environmental Record Editing and Printing (EREP) V3.5 plus PTFs are required.

- ▶ IBM z/VSE V4.3, or later

With z/VSE, TS7700 is transparent to host software. z/VSE supports the TS7760, TS7740, and TS7720 as a stand-alone system in transparency mode. z/VSE 5.1 supports both a single node or multi-cluster grid and Copy Export.

- ▶ IBM z/TPF V1.1, or later

With IBM z/TPF, the TS7760, TS7740, and TS7720 are supported in both a single node and a grid environment with the appropriate software maintenance. The category reserve and release functions are not supported by the TS7700.

Software implementation in z/VM and z/VSE

This section explains how to implement and run the TS7700 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.

General support information

Not all IBM Tape Libraries and TS7700 solutions are supported in all operating systems. Table B-1 shows a summary of several supported tape solutions for non-z/OS environments.

Table B-1 Supported tape solutions for non-z/OS platforms in z Systems environments

Platform/Tape system	IBM TS3500 tape library	TS7700	3592 drives
z/VM V5.4, V6.1, V6.2, and V6.3 native	Yes	Yes ^a	Yes
z/VSE V4.3 native ^b z/VSE V5.1 native ^c z/VSE V5.2 native z/VSE V6.1 native	Yes	Yes	Yes
z/VSE V4.3 under z/VM ^b z/VSE V5.1 under z/VM ^c z/VSE V5.2 under z/VM z/VSE V6.1 under z/VM	Yes	Yes ^a	Yes
zTPF	Yes	Yes	Yes

a. With restrictions: See “Considerations in all TS7700 environments” on page 833.

b. This platform and later includes support for logical Write Once Read Many (LWORM).

c. z/VSE V5.1 and later supports multi-cluster grid and Copy Export.

Note: z/VSE is only supported for releases 5.2 and later

Even if z/VM and z/VSE can use the TS7700, you must consider certain items. For information about support for TPF, see “Software implementation in z/OS Transaction Processing Facility” on page 839.

Note: An RPD that enables a TS4500 to be used by z/VM is available.

Considerations in all TS7700 environments

z/VSE cannot provide SMS constructs to TS7700. However, clients might be able to take advantage of some of the Outboard policy management functions if they predefine the constructs to the logical volumes when they are entered through the MI. Another possibility is to use dedicated physical pools in a TS7700 environment. After the insert processing of virtual volumes completes, you can define a default construct to the volume range as described in “Implementing Outboard Policy Management for non-z/OS hosts” on page 843.

TS7700 multi-cluster grid environments

z/VSE V5.1 and later supports multi-cluster grid and Copy Export.

Introduced by VM65789 is the ability for the RMS component of DFSMS/VM to use the COPY EXPORT functionality of a TS7700. COPY EXPORT allows a copy of selected logical volumes that are written on the backend physical tape that is attached to a TS7700 to be removed and taken offsite for disaster recovery purposes. For more information, review the memo bundled with VM65789 or see *z/VM: DFSMS/VM Removable Media Services*, SC24-6185.

For DR tests involving a TS7700 grid that is connected to hosts running z/VM or z/VSE, Release 3.3 of the TS7700 microcode introduced a new keyword on the DRSETUP command called SELFLIVE. This keyword provides a DR host the ability to access its self-created content that has been moved into a write-protected category when flash is enabled. For more information, see the *Library Request Command* white paper found at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101091>

z/VM native support that uses DFSMS/VM

DFSMS/VM Function Level 221 (FL221) is the only way for a z/VM system to communicate with a TS7700. DFSMS/VM FL221 is part of z/VM. The removable media services (RMS) function of DFSMS/VM FL221 provides TS7700 support in z/VM, as described in *DFSMS/VM Function Level 221 Removable Media Services*, SC24-6185.

Tape management

Although the RMS functions themselves do not include tape management system (TMS) services, such as inventory management and label verification, RMS functions are designed to interface with a TMS that can perform these functions. Additional information about third-party TMSs that support the TS7700 in the z/VM environment is in *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

Figure B-1 shows the z/VM native support for the TS7700.

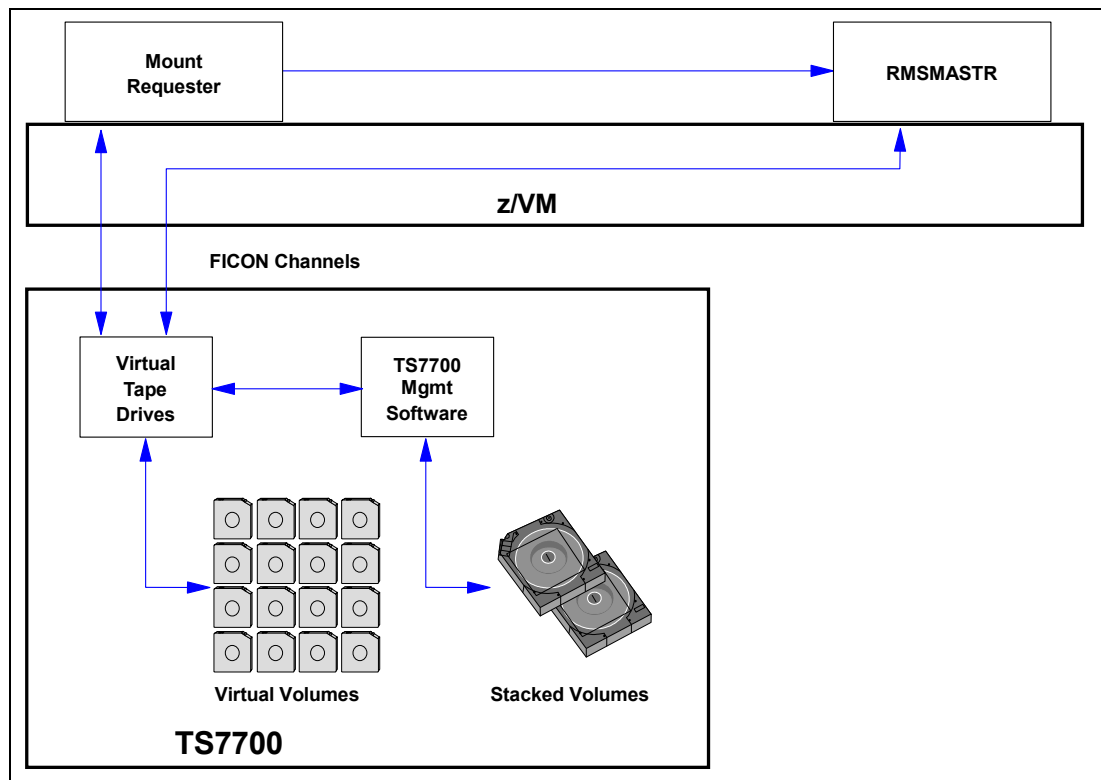


Figure B-1 TS7700 in a native z/VM environment using DFSMS/VM

When you use the TS7740, TS7720T, or TS7760T in a VM environment, consider that many VM applications or system utilities use specific mounts for scratch volumes. With specific mounts, every time a mount request is sent from the host, the logical volume might need to be recalled from the stacked cartridge if it is not on already in the Tape Volume Cache (TVC).

Instead, you might want to consider the use of a TS7760, TS7760T, TS7720, or TS7720T (cache resident partition CP0) for your VM workload. This keeps the data in the TVC for faster access. In addition, also consider that your VM backup must determine whether a TS7700 and its replication capabilities to remote sites provides what is needed or if physical tape is needed to move data offsite.

DFSMS/VM

After you define the new TS7700 tape library through HCD, you must define the TS7700 to DFSMS/VM if the VM system is to use the TS7700 directly. You define the TS7700 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 V6R2 DFSMS/VM Removable Media Services*, SC24-6185, for more information.

You have access to RMS as a component of DFSMS/VM. To enable RMS to run automatic insert bulk processing, you must create the RMBnnnn data file in the VMSYS:DFSMS CONTROL directory, where nnnn is the five-character tape library sequence number that is assigned to the TS7700 during hardware installation.

For more information 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 is shared by your VM system and other systems, more considerations apply. For more information, see *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409.

Native z/VSE

Native support is provided for the stand-alone grid TS7700 configuration in z/VSE Version 5.1 and later that support all IBM TS1150, TS1140, TS1130, TS1120, and 3592-J1A configurations without APARs in all automation offerings. This includes TS3500 Tape Library configurations.

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 enable the Tape Library to be supported through the IBM S/390® channel command interface commands. This function eliminates any XPCC/APPC communication protocol that is required by the old interface. The external interface (LIBSERV JCL and LIBSERV macro) remains unchanged.

Defining library support

First, define the type of support you are using by specifying the SYS ATL statement. You can define the following types:

- TLS** TLS Tape Library Support, which provides full VSE LPAR support.
- VSE** LCDD, which does not support TS1150/TS1140/TS1130/TS1120/3592 (only IBM 3490E and 3590), and does not support 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 VSE Guest server (VGS) and DFSMS are needed (see “z/VSE as a z/VM guest using a VSE Guest Server” on page 837).

For native support under VSE, where TS7700 is used only by z/VSE, select TLS. At least one tape drive must be permanently assigned to VSE.

Defining tape libraries

Next, define your tape library or libraries. This is done through a batch job as shown in Example B-1. Use skeleton member TLSDEF from ICCF Lib 59.

Example: B-1 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 goes through the LIBSERV JCL or macro interface. Example B-2 shows a sample job that uses LIBSERV to mount volume 123456 for write on device address 480 and, in a second step, to release the drive again.

Example: B-2 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
Dismount 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
Copy Export	LIBSERV COPYEX,VOL=123456,LIB=TAPELIB1

For more information, see *z/VSE System Administration Guide*, SC34-2627, and *z/VSE System Macros Reference*, SC34-2638.

VM/ESA and z/VM guest support

This section describes two host environments that enable you to use an IBM TS7700 while running it as a guest host system under z/VM.

Tip: When z/OS is installed as a 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 started for the first time:

```
STDEVOPT LIBRARY CTL
```


z/OS guests

The STDEV0PT 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 send 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 send commands to a tape library, which results in an I/O error (command reject) when MVS tries to send a command to the library. For more information about the STDEV0PT statement, see *z/VM V6.2 Resources*:

<http://www.vm.ibm.com/zvm620/>

z/VSE guests

Some VSE TMSs require VGS support and also DFSMS/VM RMS for communication with the TS7700.

If the VGS is required, define the LIBCONFIG file and FSMRMVGC EXEC configuration file on the VGS service system's A disk. This file 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 system. For more information, see 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 system.

z/VSE as a z/VM guest using a VSE Guest Server

When a z/VSE guest system uses a tape drive in the TS7700, the virtual tape drive must be attached to that system, and the virtual tape volume must be mounted on the drive. Because, as a virtual machine z/VSE cannot communicate with the TS7700 to request a tape mount, RMSMASTR (a z/VM system) 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 system, called the VGS, to communicate with RMSMASTR. VGS uses the standard facilities of RMS to interact with the TS7700 and the virtual drives.

Figure B-2 shows the flow and connections of a TS7700 in a z/VSE environment under a VM.

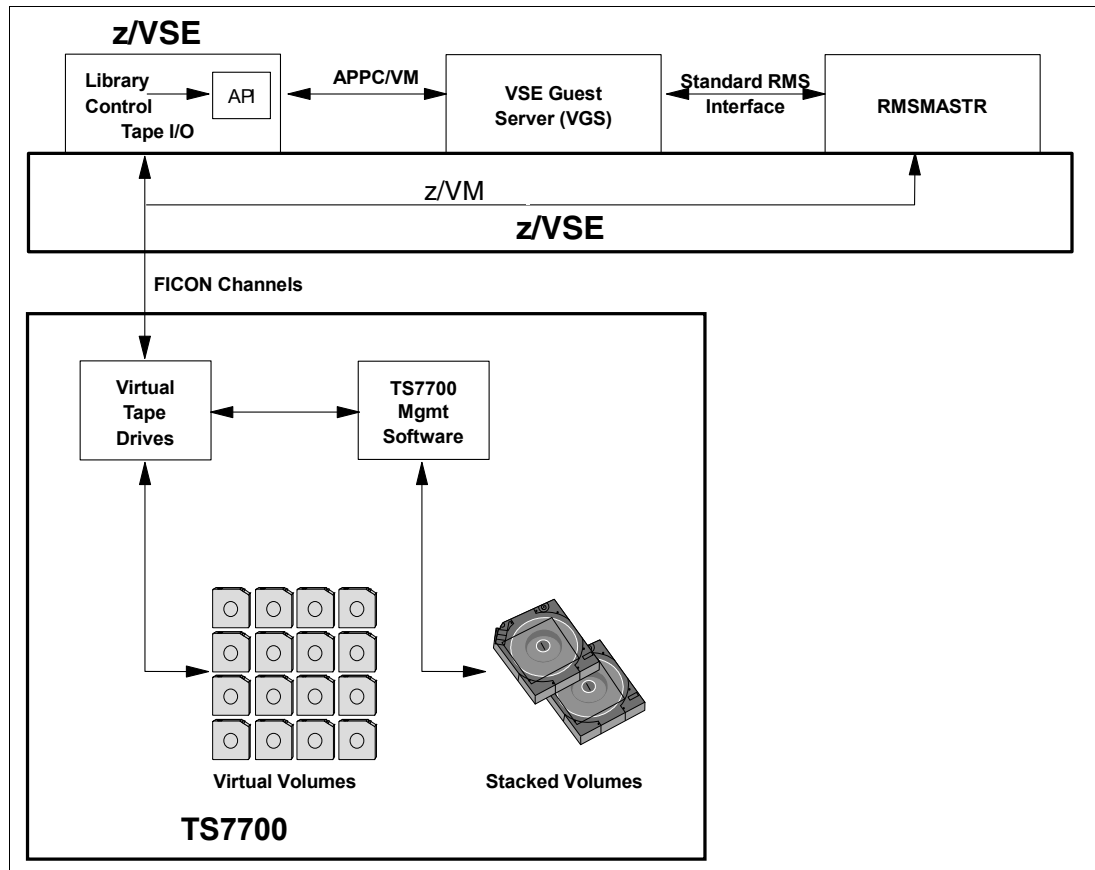


Figure B-2 TS7700 in a z/VSE environment as a VM guest

Tape management systems

As with the IBM VM/ESA native environment the TMS is responsible for keeping an inventory of volumes in the TS7700 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 B-3 shows CA-DYNAM/T VSE.




Figure B-3 TS7700 in a z/VSE environment as a VM guest (no VGS)

VSE uses original equipment manufacturer (OEM) tape management products that support scratch mounts. So, if you are using VSE under VM, you have the benefit of using the scratch (Fast Ready) attribute for the VSE library's scratch category.

For more information about z/VSE, see *z/VSE V6R1.0 Administration*, SC34-2627.

Software implementation in z/OS Transaction Processing Facility

This section describes the support for a TS7700 in a z/OS Transaction Processing Facility (z/TPF) environment with z/TPF V1.1. The z/TPF control program and several new and modified z/TPF E-type programs support the TS1150, TS7740, TS7720, and TS7760. The support is limited to a command-based interface.

Because z/TPF does not have a TMS or a tape catalog system, z/OS manages this function. In a z/TPF environment, most tape data is passed between the systems. In general, 90% of the tapes are created on z/TPF and read on z/OS, and the remaining 10% are created on z/OS and read on z/TPF.

Be sure to use the normal z/OS and TS7700 installation processes. For more information, see the following white paper, which describes some of the leading practices for implementing the TS7700 with z/TPF:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102117>

Usage considerations for TS7700 with z/TPF

z/TPF uses virtual volumes from the z/OS scratch pools and shares the TS7700 scratch categories with z/OS. The z/OS host runs the insert processing for these virtual volumes and continues to manage them based on the input obtained from z/TPF. z/TPF has a set of commands (**ztp1f**), which you use to load the volumes in z/TPF-allocated virtual drives.

After a volume is loaded into a z/TPF drive, 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 process it automatically.

On z/OS, you must update the TMS's catalog and the TCDB so that z/OS can process virtual volumes that are created by z/TPF. After the z/TPF-written volumes are added to the z/OS TMS catalog and the TCDB, normal expiration processing applies. When the data on a virtual volume expires and the volume is returned to scratch, the TS7700 internal database is updated to reflect the volume information maintained by z/OS.

Specifics for z/TPF and z/OS with a shared TS7700

From the virtual drive side, z/TPF must be allocated certain drive addresses. This information depends on the tape functions that are needed on z/TPF, and can vary with your set. Therefore, the TS7700 has tape addresses that are allocated to multiple z/TPF and z/OS systems, and can be shared by dedicating device addresses to other systems.

Tapes that are created on z/OS and read into z/TPF

Tapes that are created on z/OS and read into z/TPF use the same z/OS process for creating tapes. Now, when z/TPF wants to read this z/OS-created tape, it does a specific mount of the tape virtual server network (VSN) into a z/TPF-allocated drive by using the z/TPF (**ztp1f**) commands.

TS7700 performance for z/TPF

You can use the normal z/TPF Data Collection and Reduction reports that summarize read and write activity to the z/TPF-allocated drive. For TS7700 specific performance, use the normal TS7700 statistics that are offloaded to z/OS through the TS7700 Bulk Volume Information Retrieval (BVIR) function.

Support of large virtual volumes for z/TPF (2 GB and 4 GB)

z/TPF itself does not use functions, such as Data Class (DC), to control the logical volume size for specific mounts. User exits enable you to set construct names for a volume. If you are not using the user exits, you can set the default size through the TS7700 Management Interface (MI) during logical volume insertion, as described in "Implementing Outboard Policy Management for non-z/OS hosts" on page 843.

Consider the following information when you implement a TS7700 in a TPF environment:

- ▶ Reserving a tape category does not prevent another host from using that category. You are responsible for monitoring the use of reserved categories.
- ▶ Automatic insert processing is not provided in z/TPF.
- ▶ Currently, no IBM TMS is available for z/TPF.

Advanced Policy Management is supported in z/TPF through a user exit. The exit is called any time that a volume is loaded into a drive. Then, the user can specify, through the z/TPF user exit, whether the volume will inherit the attributes of an existing volume by using the clone VOLSER attribute. Or, the code can elect to specifically set any or all of the Storage Group (SG), Management Class (MC), Storage Class (SC), or DC construct names. If the exit is not coded, the volume attributes remain unchanged because the volume is used by z/TPF.

For z/TPF V1.1, APAR PJ31394 is required for this support.

Library interface

z/TPF has only one operator interface with the TS7700, which is a z/TPF functional message called **ZTPLF**. The various **ZTPLF** functions enable 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 z/TPF host does not keep a record of the volumes in the TS7700 tape library or manages 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.

Service information message and media information message presentation

Service information messages (SIMs) and media information messages (MIMs) report hardware-related problems to the operating system.

SIMs and MIMs are represented in z/TPF by EREP reports and the following messages:

- ▶ CEFRO354
- ▶ CEFRO355W
- ▶ CEFRO356W
- ▶ CEFRO357E
- ▶ CEFRO347W
- ▶ CDFRO348W
- ▶ CDFRO349E

Performance considerations for TS7700 multi-cluster grids with z/TPF

When clusters are operating within a TS7700 grid, they share information about the status of volumes and devices. Certain operations that are initiated by z/TPF require all the clusters in the grid to communicate with one another. Under normal conditions, this communication occurs without delay and no effect to z/TPF. In addition, if one cluster fails and the other clusters in the grid recognize that condition, the communication with that cluster is no longer needed.

The issue with z/TPF arises when the period that clusters wait before recognizing that another cluster in the grid failed exceeds the timeout values on z/TPF. This issue also means that during this recovery period, z/TPF cannot run any ZTPLF commands that change the status of a volume. This restriction includes loading 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 6 minutes. Attempting to send 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, z/TPF systems must be configured to avoid sending tape library commands to devices in a TS7700 grid along critical code paths within z/TPF. This task can be accomplished through the tape category change user exit in the segment CORU. To isolate z/TPF from timing issues, the category for a volume must not be changed if the exit is called for a tape switch. Be sure that the exit changes the category when a volume is first loaded by z/TPF and then not changed again.

To further protect z/TPF against periods in which a cluster is failing, z/TPF must keep enough volumes loaded on drives that are varied on to z/TPF so that the z/TPF system can operate without the need to load an extra volume on any drive in the grid until the cluster failure is recognized. z/TPF must have enough volumes that are loaded so that it can survive the 6-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. Eliminating the tape library commands from the critical paths in z/TPF helps z/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 relates 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, a volume that is owned by the cluster to which that device belongs is selected, if possible. z/TPF systems must always be configured so that any scratch category is populated with volumes that are owned by each cluster in the grid.

In this manner, z/TPF has access to a scratch tape that is owned by the cluster that was given the request for a scratch volume. If all of the volumes in a grid are owned by one cluster, a failure on that cluster requires a cluster takeover (which can take tens of minutes) before volume ownership can be transferred to a surviving cluster.

Guidelines

When z/TPF applications use a TS7700 multi-cluster grid that is 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 z/TPF devices and volumes on separate composite libraries (either single-cluster 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 z/TPF devices/volumes are 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 the **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 the segment CORU. To isolate z/TPF from timing issues, the category for a volume must never be changed if the exit is called for a tape switch. Be sure that the exit changes the category when a volume is first loaded by z/TPF, and then does not change it again.

- z/TPF must keep enough volumes loaded on drives that are varied on to z/TPF so that the z/TPF system can operate without the need to load extra volumes on any drive in the grid until a cluster failure is recognized and the cluster isolated. z/TPF must have enough volumes that are loaded so that it can survive the 6-minute period when a failing cluster prevents other devices in that grid from loading any new volumes.
 - z/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 method assures that during cluster failures, volumes on other clusters are available for use without having ownership transfers.
- Use the RUN Copy Consistency Point only for the cluster that is used as the z/TVC. All other clusters must be configured with the Deferred consistency point to avoid timeouts on the close of the volume.

Implementing Outboard Policy Management for non-z/OS hosts

Outboard Policy Management and its constructs are used only in DFSMS host environments where OAM can identify the construct names and dynamically assigns and resets them. z/VM, z/VSE, z/TPF, and other hosts cannot identify the construct names, and cannot change them. In addition, non-z/OS hosts use multiple Library Manager (LM) categories for scratch volumes, and can use multiple logical scratch pools on the Library Manager, as shown in Table B-2.

Table B-2 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'
Basic Tape Library Support (BTLS)	X'0FF2' - X'0FF8', X'0FFF'	8	X'FFFF'
Native VSE	X'00A0' - X'00BF'	32	X'FFFF'

Clarification: In a z/TPF environment, manipulation of construct names for volumes can occur when they are moved from scratch through a user exit. The user exit enables the construct names and clone VOLSER to be altered. If the exit is not implemented, z/TPF does not alter the construct names.

z/TPF use of categories is flexible. z/TPF enables each drive to be assigned a scratch category. Relating to private categories, each z/TPF has its own category to which volumes are assigned when they are mounted.

For more information about this topic, see the z/TPF IBM Knowledge Center:

<http://www.ibm.com/support/knowledgecenter/SSB23S/welcome>

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 the insert time of logical volumes, they are assigned construct names, also. Construct names can also be assigned later at any time.

Tip: In a z/OS environment, OAM controls the construct assignment and resets any static assignment that is made before using the TS7700 MI. 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, complete the following steps:

1. Define your pools and constructs.
2. Insert your logical volumes into groups through the TS7700 MI, as described in 9.3.3, “TS7700 definitions” on page 546. You can assign the required static construct names during the insertion as shown at the bottom part of the window in Figure B-4 on page 844.
3. On the left side of the MI, click **Virtual** → **Virtual Volumes** → **Insert Virtual Volumes**. The window in Figure B-4 opens. Use the window to insert virtual volumes. Select the **Set Constructs** check box and type the construct names.

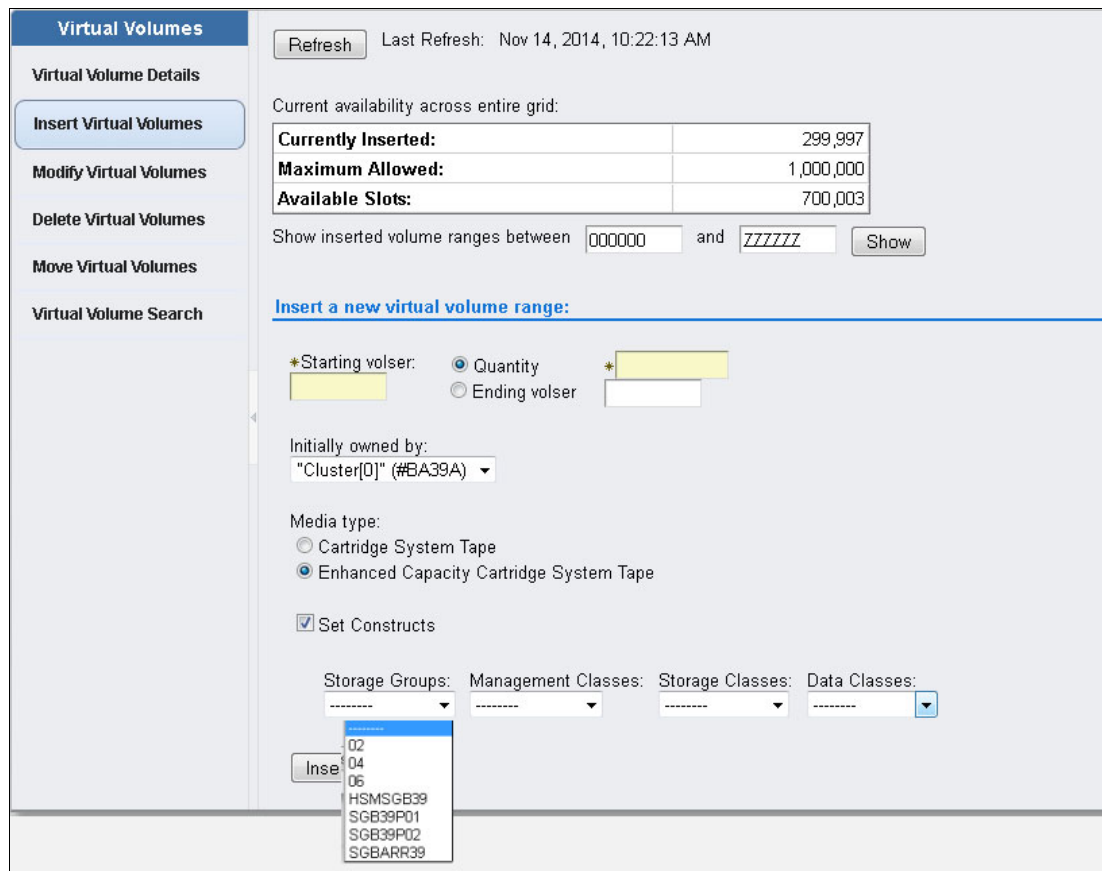


Figure B-4 Insert logical volumes by assigning static construct names

4. 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.

Define groups of logical volumes with the same construct names assigned and during insert processing, direct them to separate volume categories so that all volumes in one LM volume category have identical constructs assigned.

Host control is given by using the appropriate scratch pool. By requesting a scratch mount from a specific scratch category, the actions that are defined for the constructs that are assigned to the logical volumes in this category are run at the Rewind Unload (RUN) of the logical volume.



JES3 examples and information

This appendix provides several configuration examples where there are multiple tape libraries and tape devices, native devices in an IBM System Storage TS3500 tape library, and virtual devices in an IBM TS7700. This appendix describes the necessary parameters and considerations that must be installed in a job entry subsystem 3 (JES3) environment.

It provides 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, stand-alone cluster TS7700 Grid, and multi-cluster TS7700 Grid installed

This appendix includes the following sections:

- ▶ JES3 support for system-managed tape
- ▶ Example with two separate tape libraries
- ▶ Example with three Tape Libraries
- ▶ Processing changes

JES3 support for system-managed tape

JES3 tape library support with Data Facility System Management Subsystem (DFSMS) is described in the following sections. The primary purpose of this support is to maintain JES3 resource allocation and share tape devices. For detailed information, see *z/OS JES3 Initialization and Tuning Reference*, SA32-1005 and *z/OS JES3 Initialization and Tuning Guide*, SA32-1003.

DFSMS has support that provides JES3 allocation with the appropriate information to select a tape library device by referencing device strings with a common name among systems within a JES3 complex.

All tape library devices can be shared between processors in a JES3 complex. They must also be shared among systems within the same storage management subsystem complex (SMSplex).

Consideration: 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 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 can 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 enable 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 high-watermark setup name (HWSNAME) names in the JES3 initialization statements.

During converter/interpreter (CI) 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 HCD, they are not used in the job control language (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 is 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 that are associated with it. There is a fifth special esoteric name for a TS7700 distributed library if DAA or SAA is being used, which we describe separately:

- ▶ The *complex-wide name* is always LDGW3495. It enables you to address every device and device type in every library.
- ▶ The *library-specific name* is an eight-character string that is composed of LDG prefixing the five-digit library identification number. It enables you to address every device and device type in that specific library.
- ▶ The *complex-wide device type*, which is shown in Table C-1, defines the various device types that are used. It contains a prefix of LDG and a device type identifier. It enables you to address a specific device type in every tape library.

Table C-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
3592-E07 encryption-enabled	LDG359N

- ▶ 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 identification number, as shown in Table C-2.

Table C-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
3592-E07	LDN + library number	3592 Model E07 encryption-enabled in lib xx

It also enables you to address a specific device type in a specific tape library. In a stand-alone grid, or in a multiple cluster TS7700 grid, the previous references to the five-digit library identification number is to the composite library.

To set up a TS3500 tape library in a JES3 environment, complete the following steps:

1. Define LDGs. Prepare the naming conventions in advance. Clarify all the names for the LDGs that you need.
2. Include the esoteric names from step 1 in the hardware configuration definition (HCD) and activate the new Esoteric Device Table (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 **HWSNAME** statement.

Updating the JES3 INISH deck

To enable JES3 to allocate the appropriate device, you must code several definitions:

- ▶ **DEVICE** statements
- ▶ **SETNAME** statements
- ▶ **HWSNAME** (high-watermark setup) statements

These statements are described in detail.

DEVICE statement: Defining I/O devices for Tape Libraries

Use the **DEVICE** format to define a device so that JES3 can use it. A device statement (Figure C-1) must be defined for each string of tape library drives in the complex. **XTYPE** specifies a one-character to eight-character name, which is provided 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 and the same device type.*

The letters CA in the **XTYPE** definition indicate to you that this is a CARTRIDGE device, as shown in Figure C-1.

```

*/ Devices 3592-J1A, 3592-E05, 3592-E06, and 3592-E07 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
DEVICE,XTYPE=(LB13592N,CA),XUNIT=(0204,*ALL,,OFF),numdev=4

*/ Example for TS7700 ...../*
DEVICE,XTYPE=(LB3GRD1,CA),XUNIT=(3000,*ALL,,OFF),numdev=256

```

Figure C-1 DEVICE statement sample

Tape library drives cannot be used as support units by JES3 DSPs.

Exception: When Dump Job (DJ) is used with the **SERVER=YES** keyword. When this keyword is used, DJ uses MVS dynamic allocation to allocate the device, which uses **XUNIT**.

Therefore, do not specify **DTYPE**, **JUNIT**, and **JNAME** parameters on the **DEVICE** statements. No check is made during initialization to prevent 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 tape library fails because the DSP cannot allocate a 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 C-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 C-2 SETNAME rules

The **SETNAME** statement has these rules:

- ▶ 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.

Tip: 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

Use the **HWSNAME** statement to define which device names are subsets of other device names. You must specify all applicable varieties. The **HWSNAME** command has this syntax:

```
HWSNAME,TYPE=(groupname,{altname})
```

The variables specify the following information:

- ▶ The *groupname* variable: Specifies a device type valid for a high watermark setup.
- ▶ The *altname* variable: Specifies a list of valid user-supplied or IBM-supplied device names. These are alternative units to be used in device selection.

Consider the following example:

```
HWSNAME,TYPE=(LDGW3495,LDGF4001,LDGF4006,LDG359J,LDG359K,LDG359M,LDG359N,LDJF4001,  
LDKF4001,LDKF4006)
```

The LDG **HWSNAME** statements have the following rules:

- ▶ 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 tape library, the complex-device type name must also be included as an alternative 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 alternative.
- ▶ The library-specific device type name must be included. Alternative names can be specified in the following manner:
 - 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 alternative.
 - 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.

Ensure that all valid alternative names are specified.

Example with two separate tape libraries

The first example includes different native tape drives in two separate tape libraries. Figure C-3 shows a JES3 complex with two TS3500 tape libraries that are 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. In this example, the 3592-E05 drives are not encryption-enabled in either library.




Figure C-3 First JES3 configuration example

LDG definitions necessary for the first example

Table C-3 shows all of the LDG definitions that are needed in HCD. There are a total of eight esoterics to define.

Table C-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 that are needed for this configuration

These examples use a naming convention for XTYPE that contains the library (LB1, LB2) in the first three digits, and then the device type (Figure C-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=(LB23592K,CA),XUNIT=(2000,*ALL,,OFF),numdev=8
    
```

Figure C-4 First configuration example - device-type definition sample

SETNAME statements that are needed for this configuration

Figure C-5 includes all of 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=(LB23592K,CA),NAMES=(LDGW3495,LDGF4006,LDG359K,LDKF4006)
    
```

Figure C-5 First configuration example - SETNAME definition sample

For this example, you need three SETNAME statements for these reasons:

- ▶ One library with two different device types = Two SETNAME statements
- ▶ One library with one device type = One SETNAME statement

Tip: 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 that is needed for this configuration

The **HWSNAME** definition is tricky, so every statement shown in Figure C-6 is explained. 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 C-6 HWSNAME definition sample

The following numbers correspond to the numbers in Figure C-6:

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 has only 3592-E05s installed.
3. All 3592-E05s in library F4006 (LDKF4006) are a subset of library F4006. LDG359K is not specified because there are also 3592-E05s that are 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.

Example with three Tape Libraries

Figure C-7 on page 855 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 of a multiple cluster TS7700 grid (distributed library) installed. The multiple-cluster TS7700 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 installed with a composite library LIBRARY-ID of 13001. There are no native tape drives in that library.

Figure C-7 does not show the actual configuration for the TS7700 configurations regarding the numbers of frames, controllers, and the back-end drives. Only the drives and frames that are needed for the host definitions are displayed.

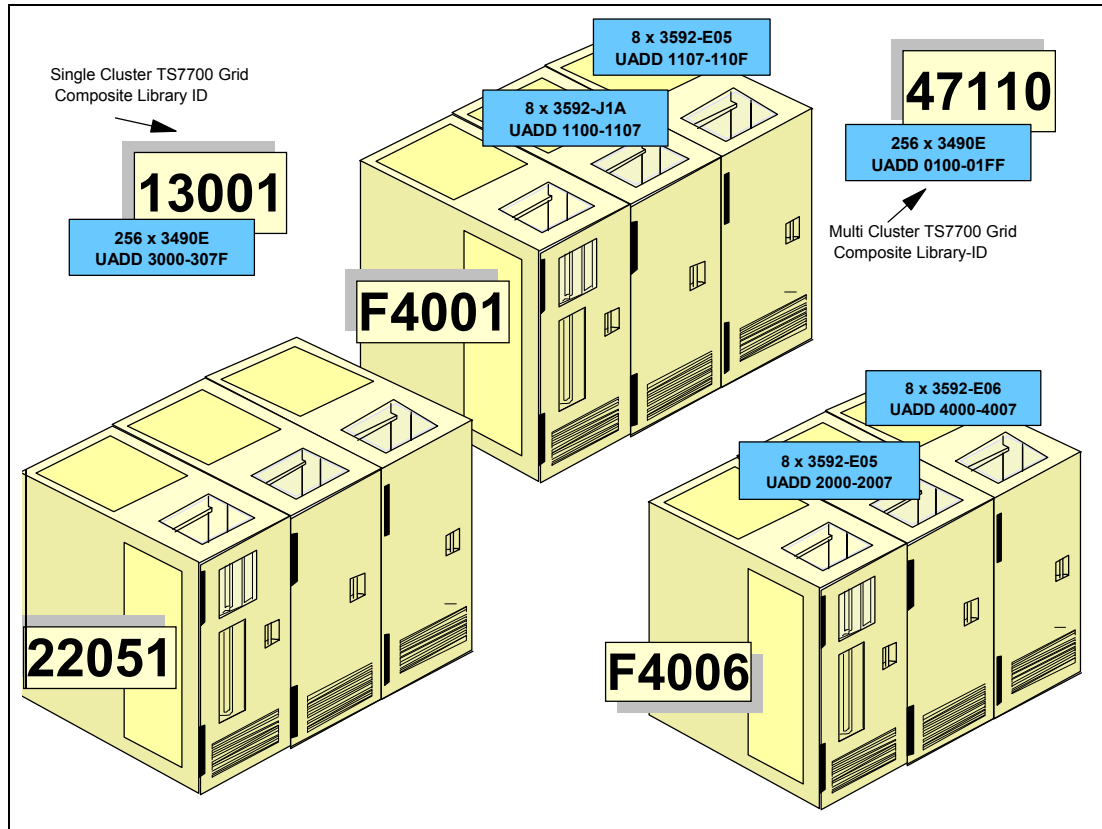


Figure C-7 Second JES3 configuration example

LDG definitions that are needed for the second configuration example

Table C-4 shows all the LDG definitions that are needed in the HCD of the second configuration example.

Table C-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 stand-alone cluster TS7700 grid. For a single cluster or multiple cluster TS7700 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 3490 devices in TS7700 ▶ Represents the 3592-J1A ▶ Represents the 3592-E05 ▶ Represents the 3592-E05 with Encryption ▶ Represents the 3592-E06

LDG definition	Value for LDG	Explanations
Library-specific device type		One definition for each device type in each library, except for the multi-cluster TS7700 grid:
	LDE13001	▶ Represents the virtual drives in the stand-alone cluster TS7700 grid in library 22051
	LDE47110	▶ Represents the virtual drives in the multicluster TS7700 grid in libraries F4001 and F4006
	LDJF4001	▶ Represents the 3592-J1A in library F4001
	LDKF4001	▶ Represents the 3592-E05 in library F4001
	LDLF4006	▶ Represents the encryption-enabled 3592-E05 in library F4006
	LDMF4006	▶ Represents the 3592-E06 in library F4006

Device statement needed

Figure C-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 Stand-alone 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 C-8 DEVICE statements for the second configuration example

Consideration: If you code **NUMDEV** in a peer-to-peer (PTP) IBM Virtual Tape Server (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

Figure C-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,,LDG3490E,,LDE13001)
SETNAME,XTYPE=LB12GRD,NAMES=(LDGW3495,LDG47110,LDG3490E,LDE47110)
```

Figure C-9 SETNAME statement values for the second example

High-watermark setup name statements

Figure C-10 shows the **HWSNAME** 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 C-10 High watermark setup statements for the second example

Additional examples

For additional examples, see the *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229 topic regarding JES3 sample initialization deck definition.

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 support module, as shown in Figure C-11 on page 859.
- ▶ 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 C-11 on page 859 shows the JES3 processing phases for CI and MDS. The processing phases include the support for system-managed direct access storage device (DASD) data sets.

The following major differences exist between TS3500 tape library deferred mounting and tape mounts for non-library drives:

- ▶ Mounts for non-library drives by JES3 are only for the first use of a drive. Mounts for the same unit are sent by IBM z/OS for the job. All mounts for TS3500 tape library drives are sent 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.

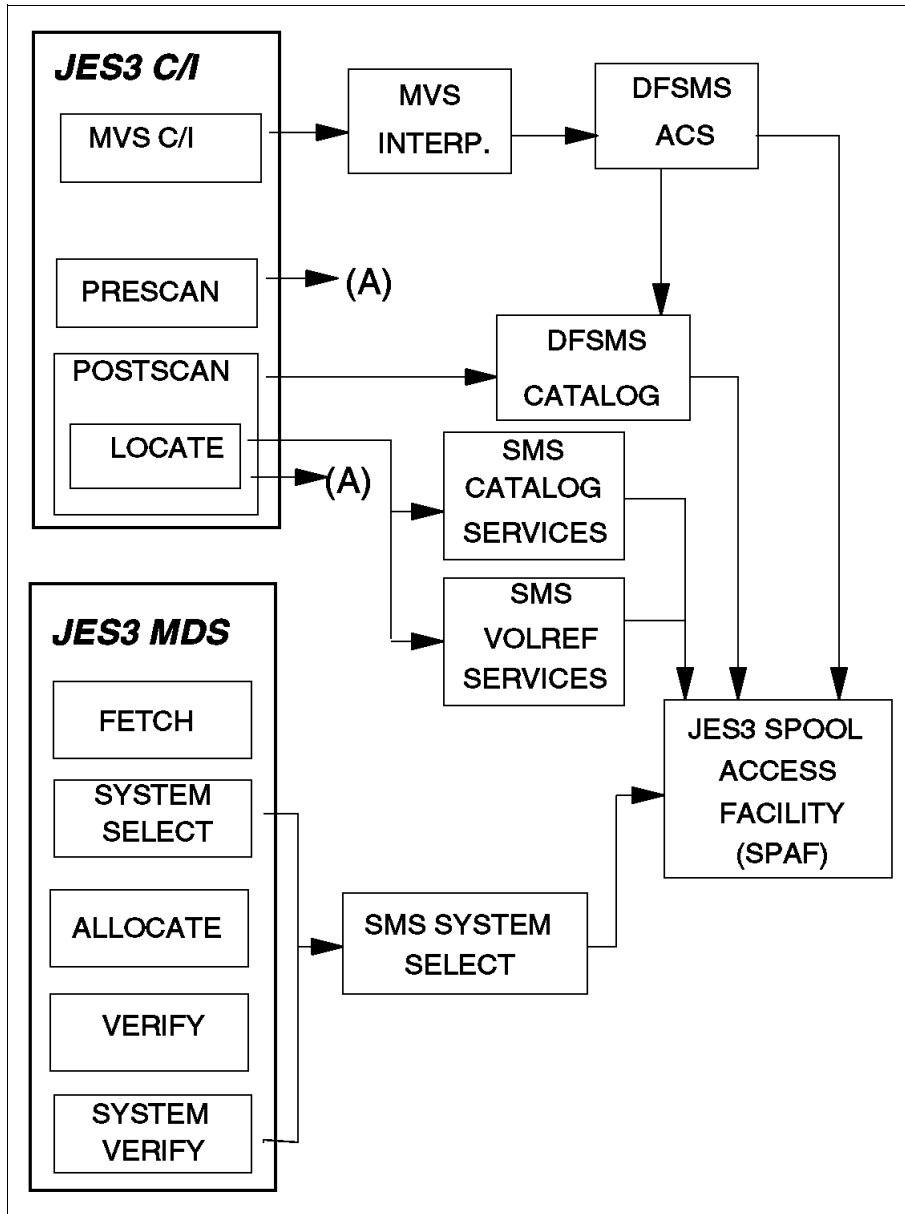


Figure C-11 JES3 CI and MDS processing phases

JES3/DFSMS processing

DFSMS is called by the z/OS interpreter to perform these functions:

- ▶ Update the scheduler work area (SWA) for DFSMS tape requests.
- ▶ Call ACS exits for construct defaults.

DFSMS system-managed tape devices are not selected by using the UNIT parameter in the JCL. For each data definition (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. This 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 the 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 the example, LDKF4001, and so on) is used.

New or modified data sets

For new data sets, ACS directs the allocation by providing Storage Group (SG), SC, and Data Class (DC). When the SG specified by ACS is defined in the active DFSMS configuration as a tape SG, 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 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 that are 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 are started during locate processing. It starts supervisor call (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 started 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. **VOL=REF=** now maps to the same SG 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

When tape library volumes are mounted and unmounted 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 that are requested in the JCL for each job to be run 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 running processor.

MDS processing also determines which processors are eligible to run 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 tape library allocation request through a subsystem interface (SSI) request to JES3 during step initiation when the job is running under the initiator. This support is not changed from previous releases.

DFSMS and z/OS provide all of the tape library support except for the interfaces to JES3 for MDS allocation and processor selection.

JES3 MDS continues to select tape units for the tape library. MDS no longer uses the **UNIT** parameter for allocation of tape requests for tape library requests. DFSMS determines the appropriate LDG name for the JES3 setup from the SG and DC assigned to the data set, and replaces the UNITNAME from the JCL with that LDG name. Because this action is 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.

Consideration: An LDG name that is specified as a UNITNAME in JCL can be used only to filter requests within the ACS routine. Because DFSMS replaces the externally specified UNITNAME, it cannot be used to direct allocation to a specific library or library device type unless SSMHONOR is specified on the unit parameter. For a description of SSMHONOR and changes that are needed for JES3 see *z/OS JES3 Initialization and Tuning Guide*, SA32-1003, and *z/OS MVS JCL Reference*, SA32-1005.

All components within z/OS and DFSMS request tape mounting and unmounting inside a tape library. They call a Data Facility Product (DFP) service, Library Automation Communication Services (LACS), rather than sending a write to operator (WTO), which is done by z/OS allocation, so all mounts are deferred until job run time. 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 is used to mount and verify a scratch tape. When the job finishes with the tape, either CLOSE or deallocation issues an unmount request through LACS, which removes the tape from the drive. MDS does normal breakdown processing and does not need to communicate with the tape library.

Multi-cluster grid considerations

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 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 might, 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 (MC).

- ▶ Copy Consistency Point

Copy Consistency Point has one of the largest influences on which cluster's cache is used for a mount. The Copy Consistency Point of Rewind Unload (R) takes precedence over a Copy Consistency Point of Deferred (D). For example, assuming that each cluster has a consistent copy of the data, if a virtual device on Cluster 0 is selected for a mount and the Copy Consistency Point is [R,D], the CL0 cache is selected for the mount. However, if the Copy Consistency Point is [D,R], CL1's cache is selected.

For workload balancing, consider specifying [D,D] rather than [R,D]. This more evenly distributes the workload to both clusters in the grid.

If the Copy Consistency Points for a cluster are [D,D], 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 cause the cluster that received the mount request to be the cache that is used to access the volume.

► Cluster families

If there are more than two clusters in the grid, consider defining *cluster families*. Especially in multisite configurations with larger distances between the sites, defining one cluster family per site can help reduce the grid link traffic between both sites.

► Copy Override settings

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 Copy Consistency Points defined in the MC can be different and tailored according to workload requirements.

You can find detailed information about these settings and other workload considerations in Chapter 6, "IBM TS7700 implementation" on page 213 and Chapter 11, "Performance and monitoring" on page 635.

Scratch allocation assistance and device allocation assistance

Unlike the system-managed tape support in the JES2 environment, the JES3 support relies on customer INISH deck setup and special tape-library-related esoteric names: complex-wide name library-specific name, library-specific device name, complex-wide device name, and a new library-specific distributed name for use with the allocation assist support. By default the allocation assist support is disabled in the JES3 environment. The following sections outline what is needed to use the device allocation assist support in a JES3 environment.

Note: This support is available starting with z/OS V2R1 plus APARs OA46747 and OA47041.

The first set of steps is common for device allocation assistance (specific mounts) and scratch allocation assistance (scratch mounts). Device allocation assistance can be used independent of the scratch allocation assistance support and vice versa:

complex-wide name Always LDGW3495. Indicates every device and device type in every library.

library-specific name

An eight character string that is composed of LDG prefixing the five-digit library identification number. Indicates every device and device type in that specific library (for example, LDG12345). In a TS7700, the library-specific name refers to the composite library.

library-specific device name

An eight character string that is composed of LDx prefixing the five-digit library identification number. Indicates every device with device type x in that specific library (for example, LDE12345, where "E" represents all 3490E devices in library 12345). In a TS7700, the library-specific device name refers to the composite library.

complex-wide device name

Contains a prefix of LDG and a device identifier that represents all devices of a particular system and model type in every tape library (for example, LDG3490E for 3490E devices).

library-specific distributed name

An eight character string that is composed of LDX prefixing the five-digit library identification number of the distributed library (or cluster) in a TS7700. Only for use with the TS7700, and only if the device allocation assist functions (DAA, SAA, or both) will be used by JES3.

The library-specific distributed name is used in addition to the previously listed esoteric names that are still needed. Define the LDXxxxx names only for distributed libraries (or clusters) that have devices that are connected to the host.

Specific allocation assistance enablement considerations

These installation steps must be followed to prevent job failures from occurring:

1. Ensure that all systems in the JES3plex are at z/OS V2R1 (this is needed because pre-execution and job execution can occur on different systems in the JES3plex). However, JES3 itself can be at a lower release level.
2. Make JES3 INISH deck changes as described in the following INISH deck example. The INISH deck changes define the library-specific distributed names that are associated with the distributed libraries clusters in a TS7700. All TS7700 tape drives that are used by JES3 (with devices that are connected to the host) should have the new esoteric names that are defined, regardless of whether the TS7700 is part of a single or a multi-cluster grid.
3. Roll out the JES3 INISH deck changes to all systems in the JES3plex (this roll-out can occur one system at a time). The new esoteric names are not passed to JES3 until the support is enabled through **DEVSUPxx**.
4. By default, the device allocation assist function is enabled at the library for all specific allocations. However, starting with Release 1.7 of the TS7700, this support can be disabled either by a tape hardware specialist (PFE) dialing into the library or starting with Release 2.0 of the TS7700, through the **MVS LIBRARY REQUEST** command. Verify that the DAA function is enabled at the library by using the **LIBRARY REQUEST, composite-libraryname, SETTING** command.
5. Lastly, enable the support to the host through the **DEVSUPxx PARMLIB** member by using the **JES3_ALLOC_ASSIST=YES** keyword (either at IPL or through the **SET DEVSUP=xx** operator command). The **SET DEVSUP=xx** operator command can be used to enable this support by routing the command to all systems in the JES3plex. After this support is enabled, the new library-specific distributed names can be returned to JES3. Ensure that steps 2 and 3 are completed before enabling this support.

Otherwise, job failures can occur if JES3 does not understand the new esoteric names being passed (because they were not defined in the JES3 INISH deck). If one of the systems in the JES3plex lags behind (in enablement of this support), all that might occur is that the device allocation assist preferred cluster list might not be acknowledged. JES3 and MVS allocation still see the same list of eligible devices.

Scratch allocation assistance enablement considerations

Complete the following installation steps to prevent job failures from occurring:

1. Ensure that all systems in the JES3plex are at z/OS V2R1 (this is needed because pre-execution and job execution can occur on different systems in the JES3plex). However, JES3 itself can be at a lower release level.
2. Make JES3 INISH deck changes as described in INISH deck example. The INISH deck changes define the library-specific distributed names that are associated with the distributed libraries clusters in a TS7700. All TS7700 tape drives used by JES3 (with devices connected to the host) should have the new esoteric names that are defined, regardless of whether the TS7700 is part of a single or a multi-cluster grid.
3. Roll out the JES3 INISH deck changes to all systems in the JES3plex (this roll-out can occur one system at a time). The new esoteric names are not passed to JES3 until the support is enabled through **DEVSUPxx**.
4. Enable the support to the host through the **DEVSUPxx PARMLIB** member by using the **JES3_ALLOC_ASSIST=YES** keyword (either at IPL or through the **SET DEVSUP=xx** operator command). The **SET DEVSUP=xx** operator command can be used to enable this support by routing the command to all systems in the JES3plex. After this support is enabled, the new library-specific distributed names can be returned to JES3.

Ensure that steps 2 and 3 are completed before enabling this support. Otherwise, job failures can occur if JES3 does not understand the new esoteric names being passed (because they were not defined in the JES3 INISH deck).

5. Then, unlike the specific allocation assistance support, the scratch allocation assistance support must be explicitly enabled at the library through the **LIBRARY REQUEST, composite-libraryname, SETTING, DEVALLOC, SCRATCH, ENABLE** command (disabled by default), and then policies must be set up at the library (on an MC basis) to request the support for a specific scratch allocation.

Before assigning an MC policy that uses the scratch allocation assistance support (specifies candidate clusters), ensure that step 4 is completed first. This helps ensure that the list of eligible devices that JES3 gets back matches the list of devices that MVS allocation got back during job run time. Even though MVS allocation has retry logic to try to circumvent ABEND05C-309, that retry logic is not guaranteed to succeed.

INISH deck example

Here is an example of an INISH deck for a TS7700 multi-cluster grid that has devices online in two clusters (other clusters whose devices are not connected to the host might exist for replication purposes). In this example, the composite library has library identification number X'12345' and the first distributed library in the grid has library identification number X'10001' and the second distributed library in the grid has library identification number X'10002'.

In this example, each distributed library in the grid has 256 devices for a total of 512. The changes that must be made to the INISH deck to use the optional allocation assist support in JES3 are shown in ***bold italic*** text. The INISH deck changes are needed only if the allocation assist functions are to be enabled by specifying **JES3_ALLOC_ASSIST=YES** in the **DEVSUPxx PARMLIB** member.

Before you enable the allocation assist functions, ensure that all TS7700 tape drives in the INISH deck are defined with the necessary **LDXxxxxx** names, even if the TS7700 is a stand-alone configuration consisting of one distributed library. In this example, rather than the device statement representing the composite library (as a whole), the device statements are defined at the distributed (or cluster) level and **LDXxxxxx** names are added (as needed) for each distributed library in a TS7700 that has devices that are connected to the JES3 host.

Device statements

Replace `DEVICE, XTYPE=(CLB12345, CA), XUNIT=(1100, *ALL, , OFF), NUMDEV=512` with the following statements:

```
DEVICE, XTYPE=(DLB10001, CA), XUNIT=(1100, *ALL, , OFF), NUMDEV=256
DEVICE, XTYPE=(DLB10002, CA), XUNIT=(1200, *ALL, , OFF), NUMDEV=256
```

These device statements are suggested examples that can be used. However, depending on the contiguous device ranges that are available, more than one device statement can be used to represent all of the devices in a composite library. Also, more than one device statement might be needed to represent the devices in a distributed library (and a device can occur in only one device statement). For example, if there are not 256 contiguous device addresses that start with 1100, the devices might be split as follows:

```
DEVICE, XTYPE=(DLB10001, CA), XUNIT=(1000, *ALL, , OFF), NUMDEV=128
DEVICE, XTYPE=(DLB10001, CA), XUNIT=(1100, *ALL, , OFF), NUMDEV=128
```

Also, one of the factors that are used by JES3 in selecting devices for volume mounting is the **ADDRSORT** parameter on the **SETPARAM** initialization statement. This parameter specifies that devices are either allocated in the same order as the **DEVICE** statement defining them (**ADDRSORT=NO**) or allocated by the order of their device numbers in ascending order (**ADDRSORT=YES**, which is the default).

In a multi-cluster grid environment today, customers might have used **ADDRSORT=NO** to distribute their work load across multiple clusters in the grid by defining each device individually and alternating devices across the clusters. With the allocation assist support enabled, because the goal is to direct allocation requests to specific distributed libraries clusters in the grid, **ADDRSORT=NO** is no longer needed. Within a distributed library (or cluster), it doesn't matter which device is used and the main purpose of the allocation assist support is to direct the allocation request to appropriate distributed libraries.

SETNAME statements

The following list illustrates the **SETNAME** statements:

- ▶ For the 3490E devices in *composite library 12345, distributed library (10001)*:
`SETNAME, XTYPE=DLB10001, NAMES=(LDGW3495, LDG12345, LDG3490E, LDE12345, LDX10001)`
- ▶ For the 3490E devices in *composite library 12345, distributed library (10002)*:
`SETNAME, XTYPE=DLB10002, NAMES=(LDGW3495, LDG12345, LDG3490E, LDE12345, LDX10002)`

High-watermark statements

The following list illustrates the **HWSNAME** statements:

```
HWSNAME, TYPE=(LDGW3495, LDG12345, LDG3490E, LDE12345, LDX10001, LDX10002)
HWSNAME, TYPE=(LDG12345, LDE12345, LDG3490E, LDX10001, LDX10002)
HWSNAME, TYPE=(LDE12345, LDG12345, LDG3490E, LDX10001, LDX10002)
HWSNAME, TYPE=(LDG3490E, LDE12345, LDG12345, LDX10001, LDX10002)
HWSNAME, TYPE=(LDX10001)
HWSNAME, TYPE=(LDX10002)
```

Note: The DLB10001 and DLB10002 device statement names are used here for illustration purposes. When defining the device statement names, any name (up to 8 characters) can be used.



DEVSERV QLIB command

The syntax and parameter explanations for the **DEVSERV QLIB** command are explained.

For a full explanation of these commands with examples, see *z/OS MVS System Commands*, SA22-7627.

The **DEVSERV QLIB** command can be used to complete these tasks:

- ▶ Request a list of tape library subsystems that are defined to the host. The libraries are listed by library-id.
- ▶ Request a list of devices within a library. The devices are listed by device number and the library port for each device is displayed.
- ▶ Request a list of the outstanding library mount orders (MOUNT, DEMOUNT, EJECT, and AUDIT).
- ▶ Display or change the list of categories currently in use by the host.
- ▶ Validate the connection status of devices in a library to the host.
- ▶ Delete an incorrectly defined library control block in preparation for an input/output definition file (IODF) activation.
- ▶ Issue a diagnostic **state save** to a library when requested by the IBM Support Center.

Important: Do not use this **state save** command for testing purposes. It affects the performance of your IBM Virtual Tape Server (VTS) automated tape library (ATL) because it takes time to get the memory 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 also displays the distributed Library-IDs associated with the composite library. If the Library-ID specified is for a distributed library, the command also displays the composite Library-ID that is associated with the distributed library.

Tip: You can use **DEVSERV QLIB,?** to get the complete syntax of the command:

```
IEE459I 13.16.57 DEVSERV QLIB 040
```

The **DEVSERV 'QLIB'** command has this 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*)
```

QLIB uses the following parameters:

LIST	Indicates that QLIB will display a list of the <i>ACTIVE Library-IDs</i> (the default). You can optionally generate a list of <i>INACTIVE Library-IDs</i> or <i>QUEUED</i> library orders. LIST uses the subparameters ACTIVE , INACTIVE , and QUEUE .
LISTALL	Produces a detailed list of all libraries, including the devices and port-ids within each library. LISTALL uses the subparameters ACTIVE and INACTIVE .
LIBID	Indicates that the request is for a specific library. LIBID uses the subparameters ACTIVE , INACTIVE , VALIDATE , QUEUE , and DELETE .
DDDD	Indicates that the request is either for the library that contains device <i>dddd</i> , or is for the device <i>dddd</i> itself. A subparameter is required when DDDD is specified. DDDD uses the subparameter SS .
DDR	Displays the limit on storage usage for a tape dynamic device reconfiguration (DDR) swap.
SS	Indicates that QLIB will send 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.
CATS CATS(XXX*)	Displays or updates the library partitioning category codes. For a request to change the library partitioning category codes, the first three digits of the category can be modified with the last digit being fixed and representing the media type. If the library partitioning category codes are modified by using the DQ QL,CATS command, the corresponding changes must also be reflected in the DEVSUPxx PARMLIB member. If not, an IPL reverts the category codes to what is specified in DEVSUPxx .

Important: Using the **DS QL,CATS(XXXX)** command on an active running system can be problematic if the categories in the IBM TS7700 are not also changed at the same time. Extreme caution must be used when sending this command. If categories are going to be dynamically changed in the TS7700 and with the **CATS** command, it is important to follow up with the proper **DEVSUPxx** changes before the next boot or initial program load (IPL).

QLIB uses the following subparameters:

ACTIVE	Displays information about the library configuration that is in use by the system.
INACTIVE	Displays information about the library configuration that becomes active following the next IODF activate. The INACTIVE configuration is similar to ACTIVE , but might contain more devices or libraries.
VALIDATE	Displays the same information as the INACTIVE configuration. However, before the configuration is displayed, I/O is sent to each device in the configuration to validate connectivity to the host.
DELETE	Indicates that QLIB must delete the INACTIVE control blocks for library LIBID and not affect the existing ACTIVE library definition.
QUEUE	Lists the library orders that are waiting to be completed, for example, MOUNT , DEMOUNT , EJECT , or AUDIT . When an order completes, the library notifies the host, and the order is removed from the queue. This QLIB display can list orders for all libraries, or be limited to a single library.



Sample job control language

The job control language (JCL) of sample jobs to collect statistical data and run volume reporting is provided. You can use the sample jobs to perform the following tasks:

- ▶ Obtain statistical data from the IBM TS7700 by 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 by using **VEHSTATS** in the IBMTOOLS libraries. To access the IBM Tape Tools, go to the following website:

<ftp://ftp.software.ibm.com/storage/tapetool/>

For the current information about BVIR, see the *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide*, at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101094>

This appendix includes the following sections:

- ▶ BVIR jobs to obtain historical data
- ▶ Extra BVIR reporting
- ▶ VEHSTATS reports
- ▶ Export list volume sample JCL
- ▶ JCL for TS7700 migration scenarios

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 the data to the System Measurement Facility (SMF) log file. These SMF records can then be used as input to **VEHSTATS**. See Example E-1.
- BVIRHSTU** Requests the BVIR historical data for one or more days, and writes the data to a disk data set with RECFM=U. The output of this job can then be used as input to **VEHSTATS**. See Example E-2 on page 874.
- BVIRHSTV** Requests the BVIR historical data for one or more days, and writes the data to a disk data set with RECFM=VB. The output of this job can then be used as input to **VEHSTATS**. See Example E-3 on page 875.

These jobs are also available as members in *userid*.IBMT00LS.JCL after you have installed the IBMT00LS.exe on your host. See 11.15, “Alerts and exception and message handling” on page 713.

After you run one of these jobs, you can create various reports by using **VEHSTATS**. See “VEHSTATS reports” on page 886.

BVIRHSTS

Example E-1 lists the JCL in *userid*.IBMT00LS.JCL member **BVIRHSTS**.

Example E-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 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 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..IBMTOOLS.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)
// PENDING
//*
//RUNPROC EXEC PUTBVIR
//STEP2.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.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
// PENDING
//*
//RUNPROC EXEC COPYBVIR
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.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 E-2 lists the JCL in *userid*.IBMT00LS.JCL member **BVIRHSTU**.

Example E-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 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..IBMT00LS.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..IBMT00LS.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=YMMDD,      YMMDD BEGINNING DATE
//      EDATE=YMMDD      YMMDD 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=YMMDD,EDATE=YMMDD  HERE AS YMMDD
//STEP3.SYSCNTL DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
/**

```

BVIRHSTV

Example E-3 lists the JCL in *userid*.IBMTOOLS.JCL member **BVIRHSTV**.

Example E-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 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=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..IBMTOOLS.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..IBMTOOLS.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=YMMDD,      YMMDD BEGINNING DATE
/**      EDATE=YMMDD      YMMDD ENDING DATE
/**
/**STEP1 EXEC PGM=IEFBR14
/**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..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..HSTV.D&SDATE..D&EDATE.,
//          DCB=(RECFM=VB,BLKSIZE=22000,LRECL=21996),UNIT=SYSDA,
//          DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE)
// PENDING
//*
//RUNPROC EXEC COPYBVIR,SDATE=YMMDD,EDATE=YMMDD HERE AS YMMDD
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//*
//

```

Extra BVIR reporting

The IBM Tape Tools libraries also provide JCL for more 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 grid configuration, separate requests must be sent to each cluster to obtain the map of all physical volumes.

Example: If this is an IBM TS7720 cluster, the following record is returned:

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

Example E-4 shows the JCL to obtain the Volume Map report, which is also contained in the *userid*.IBMTOOLS.JCL member **BVIRVTS**.

Example E-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 ENSURE 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 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
//* 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
/** 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 BVIRMCH 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..IBMTOOLS.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 the cartridges that are in the cache of one specific cluster.

You can use the same JCL as shown in Example E-4 on page 877 for the cache report by replacing the last statement (written in bold) with the statement listed in Example E-5, which creates a report for Cluster 0.

Example E-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=

Clarification: The Cache Contents report refers to the specific cluster to which the request volume was written. In a TS7700 grid configuration, separate requests must be sent to each cluster to obtain the cache contents of all of the clusters.

Copy Audit report

Before removing a cluster from a grid, the Copy Audit request can be used to ensure that all logical volumes are copied in the remaining members. It can also be used when one of the grid members is no longer available, such as in a site disaster or a test procedure, where you must determine which volumes (if any) on the remaining TS7700 tape drives do not have a valid copy.

To obtain the Copy Audit report, use the same JCL shown in Example E-4 on page 877, but replace the last statement (written in bold) with the statement shown in Example E-6, and update the following parameters:

- ▶ VTSID=
- ▶ MC=

Example E-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 E-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 install the IBM Tape Tools.

Example E-7 JCL to obtain Volume Status report

```
//ITS01 JOB CONSOLE,  
// MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,  
// TIME=1440,REGION=2M  
/*JOBPARM SYSAFF=*  
/*  
/*  
/* TO ENSURE 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 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
/** 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..IBMTOOLS.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 can use **BVIRPOOL** to create an unformatted snapshot of the status of physical volumes. The sample JCL is listed in Example E-8.

Example E-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 ENSURE 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 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
//*
//
```

Example: If this is a TS7720 cluster, the following record is returned:

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

Physical Volume Status report

For a formatted report of the physical volume status, use **BVIRPHY** (Example E-9). The output of **BVIRPHY** can then be processed by using **BVIRPRPT** (see “Physical Volume and Pool Status Report Writer” on page 884) to produce a fully formatted report.

Example E-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 ENSURE 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 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
//* 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=IEFBRI4
//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, use **BVIRPLNN** (Example E-10). The output of **BVIRPLNN** can then be processed by using **BVIRPRPT** (see “Physical Volume and Pool Status Report Writer” on page 884) to produce a fully formatted report.

Example E-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 ENSURE 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 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
/* 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 E-11 shows the sample JCL.

Example E-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 window.
*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 can be used to process the history files that are created by **BVIRHSTS**, **BVIRHSTU**, and **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 partitioned data set extended (PDSE) where each report is a separate member.

These three **VEHSTATS** jobs are also in *userid*.IBMT00LS.JCL. Example E-12 lists the sample JCL for **VEHSTPO**.

Example E-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 ACCOMODATE 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
//WEKHSMRY DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSMRY.TXT,
// DCB=(RECFM=FB,BLKSIZE=0)
//DAYHSMRY DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSMRY.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..IBMTOOLS.JCL(EXPIRE)
// DD DISP=SHR,DSN=&USERHLQ..&SITE..IBMTOOLS.JCL(&ORDER)
// DD *
*****
* COMPLETE THE FOLLOWING RECORDS AS APPROPRIATE: *
*****
CUSTOMER= TITENAME 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 RATHER than 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 DAYHSMRY
*IGNOREHEADER; DO NOT 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 AVTIVITY
*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 AVTIVITY
*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 WEKHSMRY 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
//H20VIRT EXEC COPYRPTS,RPT=H20VIRT
//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

This section provides sample JCLs to create an export list volume. Example E-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 Chapter 11, "Performance and monitoring" on page 635.

Example E-13 serves as input for a Copy Export function, which enables you to do offsite vaulting of physical volumes from a TS7700.

Example E-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 migration scenarios

Several JCL and REXX examples are provided that can help you with the TS7700 migration scenarios described in previous sections.

Using EDGUTIL to validate tape configuration database inconsistencies

The JCL in Example E-14 can help you identify inconsistencies in the Removable Media Management (RMM) control data set (CDS) and the tape configuration database (TCDB).

Example E-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 DFSMSrmm Implementation and Customization Guide, SC23-6874*. The job must be run before the migration starts.

IDCAMS example to delete a library definition in the TCDB

The JCL in Example E-15 can help you delete library information in the TCDB.

Example E-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 the TCDB

The JCL in Example E-16 can help you list all entries in the TCDB.

Example E-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 to change the TCDB

The JCL in Example E-17 can help you change the TCDB for scratch or private volumes in the TS7700.

Example E-17 JCL for changing the TCDB to a new TS7700

```
//*****  
//**** 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 E-18 can help you change volumes in RMM in the TS7700.

Example E-18 JCL for changing volumes in DFSMSrmm to a new TS7700

```
//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 IBM Resource Access Control Facility (RACF) Facility class named STGADMIN.EDG.FORCE. Verify that you have the required authorization.

REXX EXEC to update the library name

Example E-19 provides a sample REXX EXEC program for updating the library name for volumes in the TCDB.

Example E-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(source)lib                    */
/*              OUTFILE(ddname)                       */
/*                                                     */
/*      The list generated has the following format:    */
/*      VOLUME-ENTRY----Vvolser                       */
/*      For command specifics, see "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

Even though the IBM TS7700 does not possess its dedicated Library Manager anymore, Table F-1 lists all the default Library Manager volume categories, the platforms on which they are used, and their definitions if they are still valid. As defaults, these categories can be modified to address the operational needs of the system environment.

In Release 3.2 of the TS7700, all categories that are defined as scratch inherit the Fast Ready attribute. There is no longer a need to use the Management Interface (MI) to set the Fast Ready attribute to scratch categories. However, the MI is still needed to indicate which categories are scratch.

Remember: IBM z/OS users can define any category from 0x0001 - 0xFEFF (0x0000 and 0xFFxx cannot be used) with the **DEVSUPxx** member SYS1.PARMLIB. IEASYSxx must point to the appropriate member. If you use the library with other operating systems, or with multiple z/OS sysplexes in a partitioned environment, review your category usage to avoid potential conflicts.

Table F-1 Library Manager volume categories

Category (in hex)	Used by	DFSMSrmm	Definition
0000	Null category	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	CST	Indicates scratch MEDIA1. MEDIA1 is a standard-capacity cartridge system tape.
0002	DFSMS	ECCST	Indicates scratch MEDIA2. MEDIA2 is an enhanced-capacity cartridge system tape.

Category (in hex)	Used by	DFSMSrmm	Definition
0003	DFSMS	HPCT	Indicates scratch MEDIA3. MEDIA3 is the IBM 3590 Extended High-Performance Cartridge Tape.
0004	DFSMS	EHPCT	Indicates scratch MEDIA4. MEDIA4 is the IBM 3590 Extended High-Performance Cartridge Tape.
0005	DFSMS	ETC	Indicates scratch MEDIA5. MEDIA5 is the IBM 3592 tape cartridge.
0006	DFSMS	EWTC	Indicates scratch MEDIA6. MEDIA6 is the IBM 3592 tape cartridge Write Once, Read Many (WORM).
0007	DFSMS	EETC	Indicates scratch MEDIA7. MEDIA7 is the IBM 3592 tape cartridge Economy.
0008	DFSMS	EEWTC	Indicates scratch MEDIA8. MEDIA8 is the IBM 3592 tape cartridge Economy WORM.
0009	DFSMS	EXTC	Indicates scratch MEDIA9. MEDIA9 is the IBM 3592 tape cartridge Extended.
000A	DFSMS	EXWTC	Indicates scratch MEDIA10. MEDIA10 is the IBM 3592 tape cartridge Extended WORM.
000B	DFSMS	EATC	Indicates scratch MEDIA11. MEDIA11 is the IBM 3592 tape cartridge Advanced.
000C	DFSMS	EAWTC	Indicates scratch MEDIA12. MEDIA12 is the IBM 3592 tape cartridge Advanced WORM.
000D	DFSMS	EAETC	Indicates scratch MEDIA13. MEDIA13 is the IBM 3592 tape cartridge Advanced Economy.
000E	DFSMS	N/A	Indicates an error volume. Volumes in this category are scratch volumes for which the software detected an error during processing.
000F	DFSMS	N/A	Indicates a private volume. Volumes in this category contain user data or are assigned to a user.
0010 - 007F	DFSMS	N/A	Reserved. These volume categories can be used for library partitioning.
0080	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH0.
0081	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH1.
0082	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH2.
0083	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH3.
0084	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH4.

Category (in hex)	Used by	DFSMSrmm	Definition
0085	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH5.
0086	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH6.
0087	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH7.
0088	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH8.
0089	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCH9.
008A	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCHA.
008B	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCHB.
008C	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCHC.
008D	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCHD.
008E	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCHE.
008F	DFSMS/VM including VSE guest	N/A	Indicates that the volume belongs to the VM category SCRATCHF.
0090 - 009F	N/A	N/A	Not assigned.
00A0	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH00.
00A1	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH01.
00A2	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH02.
00A3	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH03.
00A4	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH04.

Category (in hex)	Used by	DFSMSrmm	Definition
00A5	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH05.
00A6	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH06.
00A7	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH07.
00A8	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH08.
00A9	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH09.
00AA	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH10.
00AB	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH11.
00AC	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH12.
00AD	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH13.
00AE	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH14.
00AF	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH15.
00B0	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH16.
00B1	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH17.
00B2	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH18.
00B3	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH19.
00B4	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH20.
00B5	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH21.
00B6	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH22.
00B7	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH23.
00B8	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH24.
00B9	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH25.

Category (in hex)	Used by	DFSMSrmm	Definition
00BA	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH26.
00BB	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH27.
00BC	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH28.
00BD	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH29.
00BE	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH30.
00BF	Native z/VSE	N/A	Indicates that the volume belongs to the VSE category SCRATCH31.
00C0 - 00FF	N/A	N/A	Not used.
0100	IBM OS/400® (MLDD)	N/A	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)	N/A	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 - 012B	N/A	N/A	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
012C	Tivoli Storage Manager for AIX	N/A	Indicates a private volume. Volumes in this category are managed by Tivoli Storage Manager.
012D	Tivoli Storage Manager for AIX	N/A	Indicates an IBM 3490 scratch volume. Volumes in this category are managed by Tivoli Storage Manager.
012E	Tivoli Storage Manager for AIX	N/A	Indicates an IBM 3590 scratch volume. Volumes in this category are managed by Tivoli Storage Manager.
012F - 0FF1	N/A	N/A	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
0FF2	Basic Tape Library Support (BTLS)	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH2.
0FF3	BTLS	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH3.
0FF4	BTLS	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH4.

Category (in hex)	Used by	DFSMSrmm	Definition
0FF5	BTLS	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH5.
0FF6	BTLS	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH6.
0FF7	BTLS	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH7.
0FF8	BTLS	N/A	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH8.
0FF9 - 0FFE	N/A	N/A	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	N/A	Indicates a scratch volume. Volumes in this category belong to the default scratch pool used by BTLS. Tip: If you are planning to change to DFSMS, you must use this default scratch category only.
1000 - FOOD	N/A	N/A	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	N/A	Indicates a volume in error. Volumes are assigned to the error category during unmount if the volume serial specified for unmount does not match the external label of the volume being unmounted.
F00F - FEFF	N/A	N/A	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	N/A	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 can be updated by operator interaction through Librarian Workstation Support.
FF01	VTS (VTS) and IBM TS7700	N/A	Stacked Volume Insert category for a VTS and TS7700. A volume is set to this category when its volume serial number is in the range that is specified for stacked volumes for any VTS library partition.
FF02	VTS	N/A	Stacked Volume Scratch category 0 for a VTS and TS7700. This category is reserved for future use for scratch stacked volumes.
FF03	VTS	N/A	Stacked Volume Scratch category 1 for a VTS and TS7700. This category is used by the VTS for its scratch stacked volumes. This category is not used if the Licensed Internal Code (LIC) is 527 or later.

Category (in hex)	Used by	DFSMSrmm	Definition
FF04	VTS and TS7700	N/A	Stacked Volume Private category for a VTS and TS7700. This category includes both scratch and private volumes (since VTS LIC level 527).
FF05	VTS and TS7700	N/A	Stacked Volume disaster recovery category for a VTS and TS7700. A volume is set to this category when its volume serial number is in the range that is specified for stacked volumes for any VTS library partition and the Library Manager is in disaster recovery mode.
FF06	VTS and TS7700	N/A	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 Product Field Engineering (PFE) tool called "movedata" as a temporary category.
FF07	VTS and TS7700	N/A	This category is reserved for future hardware functions.
FF08	VTS	N/A	This category is used by the VTS to indicate that a Read-Only-Recovery Stacked Volume with active data cannot be recovered.
FF09	TS7700	N/A	Stacked Volume Copy Export category for TS7700. This category is used by the TS7700 as a category to represent which physical stacked volumes are being copy exported or have already been copy exported as part of a program initiated copy export operation.
FF0A	TS7700	N/A	Stacked Volume Copy Export Hold category for TS7700. This category is used by the TS7700 as a category to represent which physical stacked volumes have been moved to the copy export hold state as part of a program initiated copy export operation.
FF0B - FF0F	N/A	N/A	Reserved for future hardware functions.
FF10	Library Manager	N/A	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. Tip: Logical volumes cannot be ejected from the library. They can be deleted or exported.

Category (in hex)	Used by	DFSMSrmm	Definition
FF11	Library Manager	N/A	<p>Bulk-Eject category.</p> <p>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.</p> <p>Tip: Logical volumes cannot be ejected from the library. They must be deleted or exported.</p>
FF12	VTS	N/A	<p>Export-Pending category.</p> <p>A logical volume to be exported is assigned to this category at the beginning of a VTS 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.</p>
FF13	VTS	N/A	<p>Exported category.</p> <p>Set when the VTS 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.</p>
FF14	VTS	N/A	<p>Import category.</p> <p>Stacked volumes that contain logical volumes to import into the VTS 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.</p>
FF15	VTS	N/A	<p>Import-Pending category.</p> <p>Logical volumes to be imported from a stacked volume are added to the Library Manager inventory and assigned to this category when the VTS 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 that is assigned to this category fails</p>
FF16	VTS and TS7700	N/A	<p>Unassigned Category.</p> <p>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).</p>
FF17	VTS	N/A	<p>Export-Hold category.</p> <p>Physical volumes are assigned to this category on completion of processing for an export stacked volume.</p>

Category (in hex)	Used by	DFSMSrmm	Definition
FF18	TS7700	N/A	Sunset Media Eject-Hold category for TS7700. Stacked volumes are assigned to this category by the TS7700 when the media is empty and can no longer support writes due to limitations of the current drive configuration. Only empty media, either empty by default or made empty through reclamation, is assigned to this category.
FF19	N/A	N/A	Reserved for library. These categories are reserved for future hardware functions.
FF20	Peer-to-Peer (PTP) VTS and TS7700	N/A	Corrupted-Token Volume Category. In a PtP VTS, volumes are assigned to this category when it is determined that the tokens that are associated with the volume have been corrupted. This is to prevent the volume from being selected by a category mount request.
FF21 - FFF3	N/A	N/A	Reserved for library. These categories are reserved for future hardware functions.
FFF4	Library Manager	N/A	3592 Cleaner Volume Category. Cleaner volumes for 3592-type devices in the library are assigned to this category automatically.
FFF5	Library Manager	N/A	3592 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 3592 device.
FFF6	Library Manager	N/A	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 and FFF8	N/A	N/A	Reserved for library. These categories are reserved for internal library functions.
FFF9	Library Manager	N/A	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	N/A	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	DFSMSrmm	Definition
FFFB	Library Manager	N/A	Purge-Volume Category. 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. No volumes are associated with this category.
FFFC	Library Manager	N/A	Unexpected-Volume Category. This category is reserved for future use.
FFFD	Library Manager	N/A	3590-Cleaner-Volume Category. Cleaner volumes for 3590-type devices in the library are assigned to this category automatically.
FFFE	Library Manager	N/A	3490-Cleaner-Volume Category. Cleaner volumes for 3490-type devices in the library are assigned to this category automatically.
FFFF	Library Manager	N/A	VOLSER-Specific Category. 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.



IBM TS7700 parameter examples

This appendix explains two different parameter scenarios. The first parameter example shows how different definitions and parameters interact. The second example set shows how the tape partitions can be used to influence the configuration and performance.

Important: *These examples are not leading practices or recommended configurations to be adopted.*

The purpose of this appendix is to demonstrate how some operational choices or parameter interactions can affect your TS7700 subsystem by walking you through some options and settings, and evaluating the effects of different settings or choices.

This appendix describes four examples of how consistency policies work and how certain parameters influence the behavior of the configurations.

This appendix explains the usage of the following objects:

- ▶ Different Copy Consistency Policies
- ▶ Scratch allocation assistance (SAA) and device allocation assistance (DAA)
- ▶ Retain Copy Mode
- ▶ Override settings
- ▶ Synchronous deferred on Write Failure Option
- ▶ Cluster family

The examples are meant as a drill to exercise some setting options and evaluate the effects on the grid. They are only hypothetical implementations, for the sake of the settings exercise.

Although the distance between data centers has an influence on latency, the distance has no influence on the function.

These examples show no Time Delay Replication copy policy. A Time Delay Replication copy is only made after a certain amount of time has expired (after creation / after last access). While the timer is not elapsed, this type of copy behaves regarding the dependencies to the parameter mentioned in these examples like a No copy. When the timer is elapsed, the copy behaves like a copy produced in Deferred mode copy. Therefore, no specific examples need to be added.

General example setup

This appendix explains the following examples in detail:

- ▶ Two-cluster grid for high availability (HA) and disaster recovery (DR)
- ▶ Two-cluster grid for HA and DR with selected Copy Override Policies and Retain Copy Mode
- ▶ Three-cluster grid for HA and DR
- ▶ Four-cluster grid for HA and DR
- ▶ Four-cluster grid for HA and DR with cluster families

Every example has four Management Classes (MCs):

- ▶ MC1: Synchronous mode for object access method (OAM) object support and hierarchical storage management (HSM) Migration Level 2 (ML2) workload.
- ▶ MC2: At least two clusters, which are defined with Rewind Unload (RUN) for data that must be immediately copied to the DR site.
- ▶ MC3: Deferred for workload types, where a deferred copy can be considered.
- ▶ MC4: An MC that is limited to a specific cluster (No Copy for all other clusters). This MC is needed for Bulk Volume Information Retrieval (BVIR) and Copy Export runs.

The data in cache statement applies only to the condition when all clusters are available. In outages, the normal rules apply. Synchronous goes to *synchronous deferred* (if the synchronous write failure option is enabled), and RUN copies go to the *Immediate-Deferred* copy queue. As soon as the failing cluster is recovered and is available in the grid again, the copies are made according to their policies.

Each of the examples also shows one example of the specific influence of SAA, DAA, override policies, the synchronous write failure option, and the service preparation mode of a cluster. They also describe the copy policy behavior if a disaster occurs.

Without DAA, there is no pre-selection of the mount point for a non-scratch mount. This is addressed only in the four-cluster grid example.

The Tape Volume Cache (TVC) selection for scratch mounts depends on the Copy Consistency Policy. For non-scratch mounts, there is a general rule that if a cluster has a valid copy of the logical volume in cache, this cluster TVC is selected as the I/O TVC.

Table G-1 shows the influence of the features as explained in the following examples.

Table G-1 Features mapped to examples

Feature	Where to find	Comment
Scratch allocation assist (SAA)	Example 4	
Device allocation assist for private volumes (DAA)	Example 1	
Retain Copy Mode	Example 1	Only used in the MC MCD
Override settings: Prefer local cache for fast ready mounts	Example 2	
Override settings: Prefer local cache for non-fast ready mounts	Example 2	

Override settings: Force volumes to be mounted on this cluster in local cache	Example 2	
Override settings: Copy Count Override	Example 3	
Synchronous Deferred on Write Failure option	Example 1, 3, 4, 5: ON Example 2: OFF	
Cluster family	Example 4	

Example 1: Two-cluster grid for HA and DR

With a two-cluster grid, you can configure the grid for disaster recovery (DR), high availability (HA), or both. Configuration considerations are described for two-cluster grids. The scenario that is presented is a typical configuration. Other configurations are possible and might be better suited for your environment.

Figure G-1 shows a homogeneous TS7740 cluster grid. You can also choose to introduce a TS7720 only, or a hybrid cluster. See 2.4.1, “Homogeneous versus hybrid grid configuration” on page 95 to choose the best configuration to meet your demands.




Figure G-1 Example of two-cluster grid for HA limited distance DR

Setting up the configuration

The environment of a two-cluster grid for HA and DR is described in Table G-2.

Table G-2 Environment for a two-cluster grid for HA and DR

Fact	Number or state	Comment
Number of clusters	Two	Divided in two data centers.
Type of cluster	TS7740 homogeneous	
Host connection		All hosts connected to all clusters.
SAA	Disabled	
DAA for private volumes	Disabled	The default of DAA for private volume is enabled. The disablement is for educational purposes only.
Override settings	None	
Synchronous Deferred on Write Failure option	On	
Cluster family	None	

Results of this parameter setting

MCA has a parameter setting of S/S:

- ▶ Data in cache: Data is written synchronously to both clusters.
- ▶ Mount behavior in normal conditions: Controlled by client (job entry subsystem 3 (JES3)/JES2).
- ▶ Mount behavior in outage conditions: If one of the clusters is unavailable, the mount is still performed on the remaining cluster. This happens because the *Synchronous Deferred on Write Failure Option* is set. The default is that this option is not selected. In this case, the mount will fail (see the next examples).
- ▶ TVC Selection for scratch mounts: Both clusters are treated equally regarding TVC selection.

MCB has a parameter setting of R/R:

- ▶ Data in cache: At RUN time, a valid copy is in cache at both locations.
- ▶ If one cluster is unavailable, the copies are set to immediate copy deferred and run as soon the cluster is available again.
- ▶ Mount behavior in normal conditions: Controlled by client (JES3/JES2).
- ▶ Mount behavior in outage conditions: If one cluster is unavailable, the mount is still run.
- ▶ TVC Selection for scratch mounts: Both clusters are treated equally regarding TVC selection.

MCC has a parameter setting of R/D:

- ▶ Data in cache: At RUN time, a valid copy is in one cache. The other cluster has no valid copy currently. The deferred copy is run later.
- ▶ Mount behavior in normal conditions: Controlled by client (JES3/JES2).

- ▶ Mount behavior in outage conditions: If one cluster is unavailable, the mount is still run.
- ▶ TVC Selection for scratch mounts: Cluster 0 is preferred for TVC selection. That means that if a mount is run in Cluster 1, the TVC from Cluster 0 will probably be selected. However, there are still some cases when the TVC from Cluster 1 will be selected (for non-scratch mounts or a heavy workload on Cluster 0).

MCD has parameters R/N and N/R:

- ▶ Data in cache: At RUN time, only one copy of data in the chosen cluster is available. No copy is run to the other cluster.
- ▶ All data that is stored by using this MC is in only one location. In a disaster, data loss is the result. Also, consider that a media failure can also result in data loss.
- ▶ Mount behavior in normal conditions: Controlled by client (JES3/JES2).

Mount behavior in outage conditions: If one cluster is unavailable, a scratch mount is still run. The MC from the mount cluster will be selected. This is always a RUN in this example. Retain Copy mode is only valid for non-scratch mounts.

Note: Because a copy is located only in one of the two clusters, private mounts might fail in an outage if the targeted volume is not in the cluster that remains available.

- ▶ TVC selection for scratch mounts: Local cache is selected.

These MCs are necessary for BVIR processing, DR volume testing, and Copy Export runs.

MCD with Retain Copy Mode

Assume that you have not specified Retain Copy Mode for this MC. The scratch mount was placed on Cluster 0, and the volume was created on Cluster 0 only (R,N). Now, you use this volume for several runs. Without DAA, the z/OS can select a virtual drive either from Cluster 0 or Cluster 1. If Cluster 1 is selected, the MC is again acknowledged (N,R), and a second copy of the volume is created. To avoid this, you can specify the *Retain Copy Mode*.

In this case, the origin MC (R,N) is selected, and no additional copy by a specific mount is created.

Effect of features on the grid behavior

SAA can be used in this example. However, the benefit is limited because both clusters are TS7740 and might present the same workload characteristics. If your TS7740s have different cache sizes or are differently connected to the host (performance), SAA might be considered to prefer a specific TS7740.

DAA for private volumes might be chosen to ensure that a cluster with a valid copy is selected. In this example, the benefits are limited. For MCs S/S, R/R, and R/D, there must always be a valid copy available in both clusters. Therefore, DAA made no difference. But the MCD example (RN/NR) benefits because a private mount is directed to the cluster that holds the valid copy.

Because this is a two-cluster grid, cluster families do not provide value in this configuration.

Special considerations for a consistency policy with “R,D” and “D,R”

[R,D] and [D,R] are used when the local cluster is meant to always be the I/O TVC. But, it might result in unexpected immediate copies during certain private mount operations, such as when the location of the R swaps. This can happen in a job execution when a volume is created in one cluster (R), and almost immediately, the next step of the execution mounts the same volume to the other cluster (swapping the R location).

It is better to use D/D with the preferred local for Fast Ready mounts, which eliminates any unexpected immediate copies from occurring.

Example 2: Two-cluster grid for HA and DR

Copy policies override settings and Retain Copy Mode. Example 2, as shown in Figure G-2, has the same configuration as shown in Example 1. However, several features are now applied.




Figure G-2 Example of two-cluster grid for HA and limited distance DR

Setting up the configuration

The parameter settings are described in Table G-3. These examples are for this exercise only and are not recommendations or preferred practices.

Table G-3 Environment for a two-cluster grid for HA and DR

Fact	Number/State	Comment
Number of clusters	Two	Divided in two data centers.
Type of cluster	TS7740 homogeneous	
Host connection		All hosts connected to all clusters.
SAA	Disabled	
DAA for private volumes	Disabled	
Retain Copy Mode	Disabled	
Override settings: Prefer local cache for fast ready mounts	On	Is set on both clusters.
Override settings: Prefer local cache for non-fast ready mounts	On	Is set only on Cluster 1.
Override settings: Force volumes to be mounted on this cluster in local cache	Off	This override setting overwrites the previous two, if turned on.
Override settings: Copy Count Override	Off	
Synchronous Deferred on Write Failure option	Off	
Cluster family	None	

Results of this parameter setting

MCA has a parameter setting of S/S:

- ▶ Data in cache: Data is written synchronously to both clusters.
- ▶ Mount behavior in normal conditions: Controlled by customer (JES3/JES2).
- ▶ Mount behavior in outage conditions: The mount fails because the Synchronous Write Failure is set to OFF. This situation might occur in the following situations:
 - Service preparation (microcode update or upgrades)
 - Actual failure of one cluster
 - Actual failure of all grid network links
 - Real disaster situation
- ▶ Synchronous mode spells an absolute need of data protection. It is your choice to set Synchronous Deferred on Write Failure ON or OFF. When OFF, applications that use this MC must have both clusters available always, ruling out otherwise concurrent activities, such as microcode updates or upgrades in the equipment.

- ▶ With this flag ON (one cluster is temporarily unavailable), the application is still able to mount volumes to the remaining cluster. Copies are rolled back to Synchronous-Deferred mode, which is the highest priority of deferred copy queue. The composite library enters the Synchronous Deferred State, exiting from it only when all Synchronous-Deferred copies have been run in all distributed libraries of the grid.
- ▶ For TVC selection for scratch mounts, each cluster has a defined number of virtual drives. If a mount occurs, a virtual drive is selected. The Override policy for Fast Ready Categories ensures that the local TVC is selected. No remote mount occurs.

MCB has a parameter setting of R/R:

- ▶ Data in cache: At RUN time, a valid copy is in each cache.
- ▶ If one cluster is not available, the copies are set to immediate copy deferred and run as soon the cluster is available again.
- ▶ Mount behavior in normal conditions: Controlled by customer (JES3/JES2).
- ▶ Mount behavior in outage conditions: If one cluster is not available, the mount is still run.
- ▶ TVC selection for scratch mounts: As described in MCA, the local TVC is used. No remote mount occurs.

MCC has a parameter setting of R/D:

- ▶ Data in cache: At RUN time, a valid copy is in one cache, and the other cluster has no valid copy. The deferred copy is run later.
- ▶ Mount behavior in normal conditions: Controlled by customer (JES3/JES2).
- ▶ Mount behavior in outage conditions: If one cluster is not available, the mount is still run.
- ▶ TVC selection for scratch mounts: As described in MCA, the local TVC is used. No remote mount occurs.

MCD has parameter settings R/N and N/R:

- ▶ Data in cache: At RUN time, only one copy of data in the chosen cluster is available. No copy is run to the other cluster.
- ▶ Mount behavior: If the cluster is not available, the mount is still run.

Note: Because a copy is located only within one of the two clusters, private mounts might fail in an outage if the targeted volume is not in the cluster that remains available.

- ▶ TVC selection for scratch mounts: As described in MCA, the local TVC is used. No remote mount occurs.

Important: These MCs are necessary for BVIR processing, DR volume testing, and Copy Export runs.

Special considerations for non-Fast Ready mounts and Copy Policy Override

Assume that you have a private non-Fast Ready mount for a logical volume. MC C with R/D is selected. Prefer local cache for non-fast ready mounts is selected on Cluster 1, but not on Cluster 0.

If Cluster 0 is selected as the virtual drive mount point, the TVC might be selected where a valid copy exists. That can result in a remote write to Cluster 1, and the data will be in Cluster 1 after RUN. If the data was modified, a copy is processed to Cluster 0 at RUN time. If the data was not modified, no copy occurs.

If Cluster 1 is selected, it prefers to use the TVC of Cluster 1 because of the Copy Policy Override. If no valid copy of the logical volume exists in the cache of Cluster 1, a recall from a stacked volume occurs. After RUN, a valid copy will be in Cluster 1. If the data was modified, a copy is processed to Cluster 0 at RUN time. If the data was not modified, no copy occurs.

If Cluster 1 has no valid copy of the data (in cache or on a stacked volume), Cluster 0 is selected for TVC.

Special considerations: “Force volumes to be mounted on this cluster in local cache”

This override policy is valid for all mounts. It forces the selected I/O cluster to use the local TVC. If, for any reason, the virtual node (vnode) cluster is unable to act as the I/O Tape Volume Cache (TVC), a mount operation fails even if remote TVC choices are still available when this override is enabled.

Example 3: Three-cluster grid for HA and DR

Assume that two or three TS7700 clusters are in separate locations and are separated by a distance dictated by your company’s requirements for DR. In a three-cluster grid configuration, DR and HA can also be achieved simultaneously by ensuring that two local, high-availability clusters possess volume copies and have shared access to the host, and the third and remote cluster possesses deferred volume copies for DR.

During a stand-alone cluster outage, the three-cluster grid solution maintains no single points of failure that prevent you from accessing your data, assuming that copies exist on other clusters as defined in the Copy Consistency Point.

In this example, Cluster 0 and Cluster 1 are the HA clusters and are local to each other (less than 10 kilometers (6.2 miles) apart). Cluster 2 is at a remote site that is away from the production site or sites. The virtual devices in Cluster 0 and Cluster 1 are online to the host and the virtual devices in Cluster 2 are offline to the hosts on Site A. The optional host is not installed. The host accesses the 512 virtual devices that are provided by Cluster 0 and Cluster 1.

Figure G-3 on page 914 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 extra 256 virtual devices at the host for a total of 768 devices.

In this configuration, each TS7720 replicates to both its local TS7720 peer and to the remote TS7740, depending on their Copy Consistency Points. If a TS7720 reaches the upper threshold of usage, the oldest data that has already been replicated to the TS7740 might be removed from the TS7720 cache, depending on the Copy Consistency Policy.

If you enable the TS7720 to remove data from cache, consider applying the selective dual copy in the TS7740. In this case, the TS7720 can remove the data from its cache, and then, the copy in the TS7740 is the only valid copy. Therefore, consider protecting this last valid copy against a physical media failure.

Copy Export can be used from the TS7740 to have a second copy of the migrated data, if required.

Figure G-3 illustrates a combined HA and DR solution for a three-cluster grid.

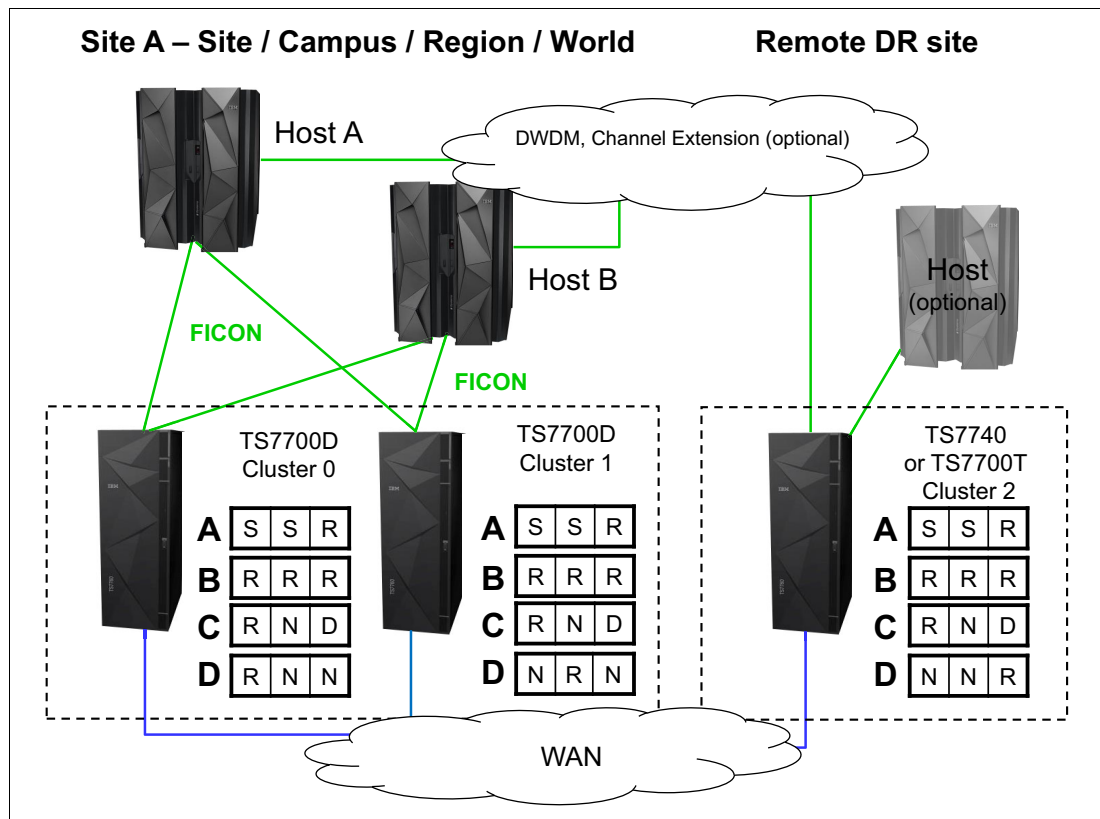


Figure G-3 Three-cluster HA and DR with two TS7700Ds and one TS7700T

Setting up the configuration

The parameter settings are described in Table G-4. All settings are for exercise purposes only.

Table G-4 Environment for a three-cluster grid for HA and DR

Fact	Number/State	Comment
Number of clusters	Three	Divided in two or three data centers.
Type of cluster	Two TS7720s within metro distance and one TS7740 in a DR location as a hybrid cluster	
Host connection		Hosts are connected only to the local Cluster 0 and Cluster 1.
SAA	Disabled	
DAA	Disabled	
Override settings	Copy Count Policy Override is set to 2.	
Synchronous Deferred on Write Failure Option	ON	
Cluster family	None	

Results of this parameter setting

MCA has a parameter setting of S/S/R:

- ▶ Data in cache: Because of the Synchronous mode copy, the data is written synchronously to Cluster 0 and Cluster 1. During RUN, the data is also copied to the DR location.
- ▶ Mount behavior in normal conditions: Controlled by customer (JES3/JES2); Cluster 0 and Cluster 1 have only host attachment.
- ▶ Mount behavior in outage conditions: The same rules for Synchronous mode regarding Write failure ON/OFF apply as in a two-cluster grid, even if a third cluster is available. If Cluster 0 or Cluster 1 is unavailable, the mount is satisfied because the Synchronous Deferred on Write Failure flag is on. The state of Cluster 2 does not have an influence in this S/S mount.
- ▶ TVC selection for scratch mounts: Clusters with S are preferred against a cluster with a Run as the I/O TVC.

MCB has a parameter setting of R/R/R:

- ▶ Data in cache: If you do not use the Copy Count Override Policy, at RUN time, a valid copy is in each cache (there are three copies, one in each cluster).
- ▶ In the example, the Copy Count Override Policy was set to a number of 2. Therefore, by the time that two RUNs complete successfully, which is sufficient for the TS7700 to signal device end, and the job finishes successfully. For MCB, that can result in the situation where two RUNs are processed in the TS7720 in the production environment, and the DR location has no valid copy at RUN time. The remaining RUN copies are produced later.
- ▶ Mount behavior in normal conditions: Controlled by customer (JES3/JES2); Cluster 0 and Cluster 1 have only host attachment.
- ▶ Mount behavior in outage conditions: Scratch mounts can be run if one TS7720 is available. If both TS7720s are not available, scratch mounts cannot be run because the TS7740 is not connected to the hosts.
- ▶ Private mounts can be satisfied if one TS7720 is available.
- ▶ TVC selection for scratch mounts: All three clusters are treated as equal in relation to TVC selection.

MCC has a parameter setting of R/N/D:

- ▶ Data in cache: At RUN time, Cluster 0 has a valid copy of the data. Cluster 1 has no copy, and Cluster 2 has no copy at RUN, but Cluster 2 receives a copy later. If Cluster 2 was selected as TVC (due to special conditions), Cluster 2 can also have a valid copy at RUN time. The copy to Cluster 0 is processed then on RUN time.
- ▶ Mount behavior in normal conditions: Controlled by customer (JES3/JES2); Cluster 0 and Cluster 1 have only host attachment.
- ▶ Mount behavior in outage conditions: For scratch mounts, if Cluster 0 fails, the mount might be run on Cluster 1. Even if Cluster 1 has no valid copy, the TVC of Cluster 2 is used as the remote mount. After RUN, Cluster 2 will have a valid copy in the cache, whereas Cluster 1 has no valid copy.
- ▶ Private mounts can be run if the deferred copy has already created a valid copy of the logical volume. Cluster 1 is selected as the mount point and Cluster 1 uses the TVC of Cluster 2. After RUN, a valid copy will be only in Cluster 2.
- ▶ TVC selection: Cluster 0 is always preferred. Cluster 2 is accepted if Cluster 0 is unavailable. Cluster 1 is not used as the TVC.

Important: This copy policy implies that if you have an outage of one component (either Cluster 0 or Cluster 2), only one valid copy of data is available.

MCD has parameters RNN, NRN, and NNR. These MCs are necessary for BVIR processing, DR volume testing, and Copy Export runs.

Effect of features on the grid behavior

SAA can be introduced in this example if you want to direct a specific type of workload to a certain TS7720. This might be useful if your TS7720s have different configurations.

DAA might be chosen to ensure that a cluster with a valid copy is selected. In this example, the data associated with MCC and MCD benefits.

The usage of override settings for local cache influences only the TVC selection, as described in “Example 2: Two-cluster grid for HA and DR” on page 910.

Cluster families provide no value in this configuration.

Example 4: Four-cluster grid for HA and DR

This example has two production sites (Site A and Site B) within metro distances. Cluster 0, Cluster 1, Cluster 2, and Cluster 3 are HA clusters and are local to each other (less than 10 kilometers (6.2 miles) apart). All virtual drives connect to the host. The host accesses the 1,024 virtual devices that are provided by Cluster 0, Cluster 1, Cluster 2, and Cluster 3. See Figure G-4 on page 917.

In this configuration, there are many different possible consistency point policies available. The section covers three valid configurations and one configuration that is not recommended.

In this example, if a TS7720 reaches the upper threshold of usage, the oldest data that has already been replicated to the TS7740 can be removed from the TS7720 cache. Copy Export can be used from the TS7740 to have a second copy of the migrated data, if required.

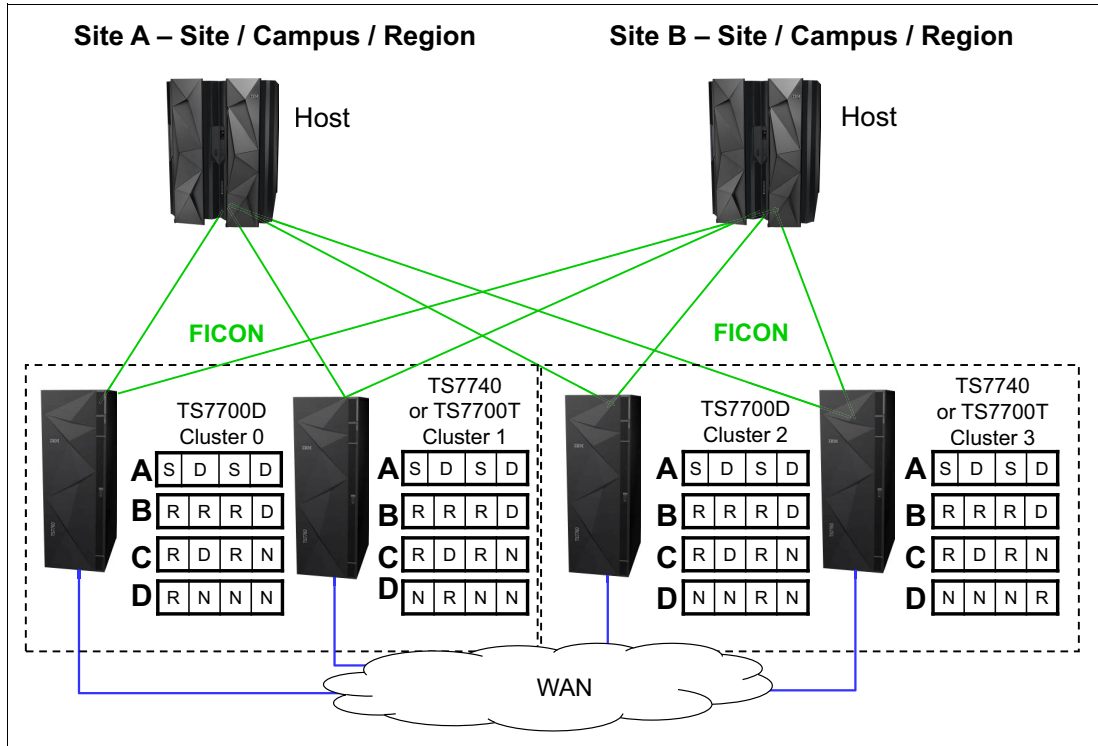


Figure G-4 Examples of a four-cluster grid for HA and DR

Setting up the configuration

The parameter settings are described in Table G-5.

Table G-5 Environment for a four-cluster grid for HA and DR

Fact	Number/State	Comment
Number of clusters	Four	Divided in two data centers.
Type of cluster	One TS7720 and one TS7740 in data center A. One TS7720 and one TS7740 in data center B.	
Host connection		Hosts are connected to all clusters.
SAA	Enabled	TS7720s are selected as scratch mount candidates for MCA and MCB. TS7740 (Cluster 1) is selected as the scratch mount candidate for MCC.
DAA for private volumes	Enabled	
Override settings	None	
Synchronous Deferred on Write Failure option	ON	
Cluster family	None	

Results of this parameter setting

MCA has a parameter setting of S/D/S/D:

- ▶ Data in cache: Because of the Synchronous mode copy, the data is written synchronously to Cluster 0 and Cluster 2 (TS7720). Cluster 1 and Cluster 3 (TS7740) receive the data later.
- ▶ Mount behavior in normal conditions: Due to the use of SAA, Cluster 0 and Cluster 2 are selected for scratch mounts. For private mounts, DAA selects a cluster according to the rules of DAA.
- ▶ Mount behavior in outage conditions: The same rules for Synchronous mode about Write Failure ON/OFF apply as in a two-cluster grid. In the example, Write Failure is set to ON, which enables mounts to be satisfied if one of the synchronous clusters is available.
- ▶ If Cluster 0 and Cluster 2 are not available, the mount is not run because SAA is enabled. In this case, you need to disable SAA or select the TS7740 as the scratch candidate mount. Private mounts are run if DAA can find any valid copy of the data in the grid.
- ▶ TVC selection for scratch mounts: Clusters with S are preferred against clusters with a Deferred as the I/O TVC.

MCB has a parameter setting of R/R/R/D:

- ▶ Data in cache: At RUN, Cluster 0, Cluster 1, and Cluster 2 have a valid copy in cache. Cluster 3 receives a copy later.
- ▶ Mount behavior in normal conditions: Due to the use of SAA, Cluster 0 and Cluster 2 are selected for scratch mounts. For private mounts, DAA selects a cluster according to the rules of DAA.
- ▶ Mount behavior in outage conditions: If both Cluster 0 and Cluster 2 are not available, the mount is not run because SAA is enabled. In this case, you need to disable SAA or select the TS7740 as the scratch candidate mount. Private mounts are run, if DAA can find any valid copy of the data in the grid.
- ▶ TVC selection for scratch mounts: Clusters with R are preferred over clusters with a Deferred as the I/O TVC.

MCC has a parameter setting of R/D/R/N:

- ▶ Data in cache: At RUN time, a valid copy is in Cluster 0 and Cluster 2. Cluster 1 (TS7740) receives the data later. Cluster 3 does not receive a copy.
- ▶ Mount behavior in normal conditions: Due to the use of SAA, Cluster 0 and Cluster 2 are selected for scratch mounts. For private mounts, DAA selects a cluster according to the rules of DAA.
- ▶ Mount behavior in outage conditions: The TS7740 (Cluster 1) is the only cluster that is selected as a scratch candidate in this example. Therefore, scratch mounts can be run only if Cluster 1 is available. Private mounts are run if DAA can find any valid copy of the data in the grid.
- ▶ TVC selection for scratch mounts: Cluster 0 and Cluster 2 are preferred, and Cluster 1 might be selected. Cluster 3 is not selected due to the No Copy setting.

Special consideration for this Management Class

This example of MC setup shows characteristics that you need to avoid when you use scratch candidate selection:

- ▶ Cluster 1 is the only cluster that is selected for a scratch candidate. This is a single point of failure, which is not necessary in such a configuration.
- ▶ Cluster 1 has a definition of Deferred copy. In combination with the scratch candidate selection, the likelihood of remote mount rises.
- ▶ Auto removal is allowed for the TS7720. With Auto Removal allowed, some volumes might have only a consistent copy in one of the TS7720 tape drives. Consider having a copy in both of the TS7740s to protect the data against a site loss and against a physical media failure.

MCD has these parameter settings: R/N/N/N, N/R/N/N, N/N/R/N, and N/N/N/R. These MCs are necessary for BVIR processing, DR volume testing, and Copy Export runs.

Effect of features on the grid behavior

The effect of the Override Policy: Copy Count Override is explained in the three-grid configuration and also applies here.

Cluster families have a major influence on the behavior and are introduced in the next example.

Example 5: Four-cluster grid for HA and DR by using cluster families

This example has a similar environment, but there are two changes:

- ▶ Cluster 0 and Cluster 1 are defined as family A; Cluster 2 and Cluster 3 are defined as family B.
- ▶ SAA is disabled.

See Figure G-5 for the new configuration.

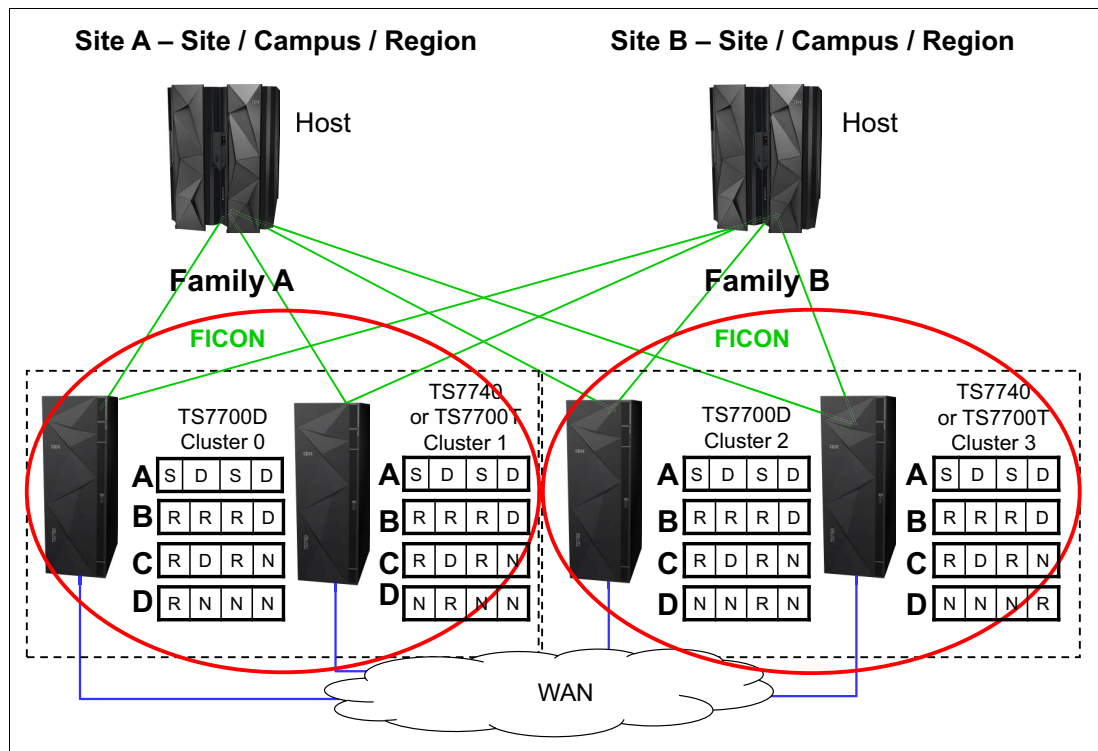


Figure G-5 Examples of a four-cluster grid for HA and DR with cluster families

Setting up the configuration

The parameter settings are described in Table G-6.

Table G-6 Environment - four-cluster grid for HA and DR

Fact	Number/State	Comment
Number of clusters	Four	Divided in two data centers.
Type of clusters	One TS7720 and one TS7740 in data center A (Site A), and one TS7720 and one TS7740 in Site B.	
Host connection		Hosts are connected to all clusters.
SAA	Disabled	
DAA for private volumes	Enabled	
Override settings	None	
Synchronous Deferred on Write Failure option	ON	
Cluster family	Two cluster families, with one cluster family in each site.	

Results of this design

MCA has a parameter setting of S/D/S/D:

- ▶ For this Copy Consistency Point, the influence of a family is small.
- ▶ Data in cache: There is no change to the number of copies available at a certain point in time. Only the copy source is different. Without cluster families defined, all copies were requested from the cluster with the selected I/O cache. With cluster families defined, the copy is requested inside the family from either Cluster 0 (Family A) or Cluster 2 (Family B).
- ▶ Mount behavior: No change.
- ▶ TVC selection: For a remote mount, normally, the cluster with S is selected. However, a cluster inside the family overrules a cluster outside the family. If Cluster 0 needed a remote mount, the cluster prefers Cluster 1 rather than Cluster 2, even if a physical mount needs to be run.

MCB has a parameter setting of R/R/R/D:

- ▶ For this Copy Consistency Point, the introduction of cluster families has the following effect.
- ▶ Data in cache: Assume that the Cluster 0 or Cluster 1 TVC was selected as the I/O cache. At RUN, the data is in the cache of Cluster 0, Cluster 1, and Cluster 2. Cluster 3 receives the copy later, but not from the original TVC cache. Instead, Cluster 3 receives the copy from Cluster 2 because it is a member of the same family.

Remember: All RUN copies are processed as defined.

For deferred copies, only one copy is transferred between the two sites. All other deferred copies are produced inside the defined family.

- ▶ Mount behavior in normal conditions: The mount behavior itself remains the same as family clusters. It is under your control to select the appropriate scratch mount candidates or to disable virtual drives. Due to the use of SAA, Cluster 0 and Cluster 2 are selected for scratch mounts. For private mounts, DAA selects a cluster according to the rules of DAA.
- ▶ TVC selection: Normally, a cluster with R is preferred against a cluster with Deferred. However, if, in a cluster family, a remote mount occurs, the family overrules this behavior. Therefore, if a cluster needs a remote mount, the cluster prefers a cluster inside the family over a cluster with a RUN outside the family. In the example, that can lead to following situation.

Cluster 2 receives a private mount. Assume that Cluster 2 has no valid copy and initiates a remote mount. Without a cluster family, Cluster 0 (TS7720) is selected if Cluster 0 has a valid copy in the cache. Instead, Cluster 2 prefers to select Cluster 3 as the TVC, which might result in a recall from a stacked volume. If the volume is modified, Cluster 3 already has a valid copy and transfers this copy to all other clusters because of the Copy Consistency Point.

MCD has these parameter settings: R/N/N/N, N/R/N/N, N/N/R/N, and N/N/N/R. These MCs are used for BVIR processing, DR volume testing, and Copy Export runs. For this Copy Consistency Point, the definition of cluster families has no influence.

General example setup for Tape partitions

This example shows a customer environment, with six different types of data. The following table shows the type of data, the likelihood that this data is read again, and the data volumes.

All data volumes are calculated without growth for a basic calculation in this example.

Table G-7 shows the type of data, the likelihood of read, and the daily and total volume of TB stored.

Table G-7 Data scenario for Tape Partitions

	Likelihood read / customer requirement	Daily volume compressed data	Total volume in TB / expiration
HSM ML2	High	0.3 TB	150 TB
DB2 archive logs	Customer wants to keep 10 days in cache	1 TB	17 TB
DB2 image copies	Customer wants to keep 3 days in cache	5 TB	150 TB
Archive data - 4 years and longer	Low	0.2 TB	300 TB expiration 4 years and longer,
Other Backups A	Low	15 TB	500 TB, 30 days
Other Backups B	Low		
Other data, like SMF	High during the next week, then low	0.5 TB	730 TB (4 years)
Other data	low	0.1 TB	unpredictable

The following scenarios could be considered:

1. All data in cache: No physical backend.
2. All data will be premigrated ASAP to physical tape: No usage of delay premigration.
3. HSM ML2 will be kept in cache only, all other data will be premigrated. Tape partitions will be used.
4. Delay premigration is used to expire data in cache.

Especially for Example 3 and Example 4, there are several different options. In this appendix, only two options are covered to explain the theory.

The night batch jobs producing the 15 TB from “other backups” are running 5 hours each day. A throughput for 3 TB in an hour is therefore necessary. Assume that this is a steady workload, a compressed host I/O from 875 MBps is expected.

Basic considerations how to find the best configuration and setup

Firstly, the customer has defined some requirements. This is a good starting point for the configuration and setup. Alternatively, there is a total amount of throughput, cache capacity, and cache bandwidth requirements, depending on how the data is treated. The same applies to the amount of premigration queue depth (FC 5274). In the next paragraphs, we calculate to determine whether the configuration and setup would be valid to satisfy the customer needs.

Example 1: All data in cache

To determine, if that is a valid option, it is only necessary to add the total amount of compressed data. In our example, it is obvious that we need more cache than the maximum amount of data we can deliver in one TS7760T (1.3 PB).

So for this example, this is not a feasible solution.

Example 2: All data on physical tape (premigrated ASAP), and with only one tape partition

First, determine what the minimum requirement for the cache is. That would be at least the amount of data that is written on one day. Adding the numbers, a one-drawer TS7760 with 31.42 TB would be sufficient. However, that does not satisfy the requirements of the customer.

To satisfy the customer requirement, we need to calculate the following.

1. DB2 archive log= 1 TB a day = 10 days in cache = 10 TB
2. DB2 image copy = 5 TB a day = 3 days in cache = 15 TB
3. Other backup = 15 TB
4. Other data = 0.8 TB
5. HSM = unpredictable

Adding the numbers $(10 + 15 + 15 + 0.2 + 0.5 + 0.1) = 40.8$ TB

A 2 drawer configuration (approx. 63 TB) would allow also to have some HSM data in cache.

However, this configuration has the following restrictions:

1. It cannot handle the necessary Host I/O sustained throughput.
2. Without Tape partitions, the cache can be controlled only with PG1/PG0, therefore the DB2 data and SMF data might not be in cache as requested.

So this is also not a feasible solution.

Example 3: HSM ML2 will be kept in cache only, all other data will be premigrated, and tape partitions will be used

To find out the configuration, we first calculate the minimum cache capacity:

1. CP0: HSM ML2, 150 TB + 10% free space and contingency = 165 TB
2. CP1: DB2: 25 TB + 10% contingency = 28 TB (rounded)
3. CP2: Other data = 15.8 TB. SMF data will be treated with PG1. All other data will be treated as PG0.

In total 208.8 TB are requested, so seven drawers with approx. 219 TB total capacity needs to be installed to satisfy the minimum request for cache capacity.

Looking to the cache bandwidth (see CHAPTER PERFORMANCE), a seven-drawer configuration could provide the necessary sustained throughput, as long no RUN or SYNC copies shall be produced.

If RUN or SYNC is needed (and we strongly suggest that you use them when HSM ML2 synchronous mode copy is used), then a seven-drawer configuration is not sufficient - or no premigration could run during the night batch. This could result in an issue, because 15 TB will not fit in the premigration queue. So we would not recommend this configuration if RUN or SYNC is needed.

Regarding the premigration queue depth, we can not follow the recommendation to be able to keep a full day of premigration data in the queue, because the customer produces 22 TB a day. In the 5-hour peak, approx. 15 TB are written. That means, that either two TS7700 needs to be installed or the customer needs to accept that premigration during the peak time is essential to not run in throttling. In addition, the customer should consider to review the LI REQ,SETTING2,PRETHDEG / COPYWDEG/TVCWDEG values to be prepared in an unavailable situation of the physical library.

So this could be one feasible option.

Example 4: Delay premigration will be used to expire data in cache

To find out the configuration, we first calculate the minimum cache capacity:

1. CP0: HSM ML2, 150 TB + 10% free space and contingency = 165 TB
2. CP1: DB2: 25 TB + 10% contingency = 28 TB (rounded)
3. CP2: Other data (backups with 30 days will expire in cache)= 500 TB, + 10 TB for other data. SMF treated with PG1, all other data will be treated as PG0.

That means, that we have to provide 703 TB, which results in 23 drawers. This amount of drawers is capable to run the host I/O in sustained mode, and do also some copy activity.

Because all data from the daily backup will expire in cache, the amount of FC 5274 for the premigration queue needs to be recalculated.

in this example. the daily amount of compressed data is 7 TB, which results in 7 * FC 5274. Depending on the workload profile, you might consider to install only 5 TB to cover the DB2 portion only.

In this example also the number of backend cartridges - and maybe even backend drives would be less than in Example 3.

If 7* FC 5274 is installed, also an unavailability of the physical tape library of 24 hours can be allowed without any throttling issue.

This is also a feasible solution, if the customer allows that no physical copy exists for a part of the backup data. Keep in mind that this data has only a short lifecycle anyway.



Extra IODF examples

This appendix covers input/output definition file (IODF) and input/output configuration program (IOCP) examples explained for several different configurations.

Important: These examples are *not* preferred practices or suggested configurations to be adopted. They are for reference only and need to be tailored to your environment and needs.

This appendix includes the following sections:

- ▶ General IODF principles
- ▶ Using switches
- ▶ Directly connecting
- ▶ Upgrading to eight 8-Gb channels
- ▶ Adding devices beyond 256
- ▶ Sharing ports between different hosts
- ▶ LIBPORT-IDs in the MVSCP

General IODF principles

When you set up an IODF you define up to four (or six with IBM z13) channel subsystems (CSSs), which logical partitions (LPARs) are in which CSS, which channel-path identifiers (CHPIDs) are defined to which LPARs, which control units (CUs) use which CHPIDs, and which IODEVICES are defined to which CUs. The IOCP is a text file that can be migrated to an IODF, or all the definitions can be done by using the Hardware Configuration Dialog (HCD).

A CHPID is a two-digit value from 00 to FF depending on how many and which CHPIDs are defined to the processor. Physical channel IDs (PCHIDs) represent a physical location on the central processor complex (CPC) and are hardcoded. CHPIDs are arbitrary and plug into the PCHIDs.

PATH is an IOCP statement that is a set of CHPIDs (path group, up to 8) defined to the processor.

A CHPID statement (Example H-1) specifies the following information:

- ▶ The CSS of the LPARs with access to the CHPID
- ▶ Whether it is shared (i.e. can be used by more than one system in a CSS on the CPC)
- ▶ Type of channel
- ▶ Which Switch the channel cable plugs into (Hex ID of the switch)
- ▶ Which LPARs on the CPC have access to the CHPID
- ▶ The PCHID (physical port) on the CPC that the channel cable is plugged into

Example H-1 Format of a CHPID definition from the IOCP

CHP68	CHPID	PATH=(CSS(0),68),SHARED,TYPE=FC,SWITCH=65,	X
		PART=((ZOS1,ZOS2,ZOS3,ZOS4,ZOS5,ZOS6,ZOS7,ZOS8,ZOS9),	X
		(=)),	X
		PCHID=5B8	

Using switches to connect to the control unit

LINKs are a two-digit or four-digit (four-digit if the switches are cascaded) ports on a blade in the switch. The LINKs are positional, such that the communication that uses the previous CHPID it exits out the two-digit port that is indicated, so the CPC knows how to address the device it wants to communicate with. The cables from the switch outbound port to the cluster can be plugged into any port on the Hankie card. A cluster looking back at the switch can display which switch port it sees on which Hankie port.

If a CU definition specifies LINKs, multiple CPCs can talk to that cluster if the proper LINKs are used in the IOCP/IODF. LINKs are the same on each CPC, even though the CHPIDs are probably different. The CHPIDs go to the switch that the links are from. The CHPID definition specifies which switch that particular CHPID goes to and the CU definition specifies which outbound port (LINK) goes to that device. A single switch or multiple switches (up to 8) can be used for every CHPID in the PATH statement. It just depends on which switch the CHPID runs to.

A CPC can use fewer than eight ports, but the LINKs are still the outbound switch ports, you have a PATH statement with fewer CHPIDs/LINKs. The cables from the outbound switch ports are arbitrary. They can be connected in no particular order, or can be connected so they mean something to whomever is plugging in the cables, such as to aid in troubleshooting.

Example H-2 shows a CU definition example that uses switches on eight channels.

Example H-2 Control unit definition example that uses switches on eight channels

```
*ZORO/0
    CNTLUNIT CUNUMBR=2611,
    PATH=((CSS(0),61,63,65,67,69,6B,6D,6F)),
    LINK=((CSS(0),2B,AC,3A,4F,CD,BB,5F,DD)),
    UNITADD=((00,16)),UNIT=3490,CUADD=0
*$HCDC$    DESC='ZORO BARR39'
TAPED300 IODEVICE ADDRESS=(D300,16),UNIT=3490,CUNUMBR=(2611),UNITADD=00
```

Directly connecting

In a direct connect situation, there is no switch and the CHPID channel cable connects only to that device and no others. If all the CHPIDs are direct, you can forgo the link statement. The link fields in the IOCP definition are all asterisks as shown in Example H-3.

Example H-3 Control unit definition example that uses a direct connection

```
*ZORO/0
    CNTLUNIT CUNUMBR=2611,
    PATH=((CSS(0),61,63,65,67,69,6B,6D,6F)),
    LINK=((CSS(0),**,**,**,**,**,**,**,**)),
    UNITADD=((00,16)),UNIT=3490,CUADD=0
*$HCDC$    DESC='ZORO BARR39'
TAPED300 IODEVICE ADDRESS=(D300,16),UNIT=3490,CUNUMBR=(2611),UNITADD=00
```

Upgrading to 8-Gb channels

There are no changes that are required in the IOCP/IODF to change from 4-Gb channels to 8-Gb channels when the number of devices and number of channels remains the same.

Adding more devices

To add more device addresses for use with the library, add extra CU definitions. In Example H-4, the first 16 CUs (0 - F) are the original 256 devices. The new devices are being added with another 15 CUs (10 - 1E). Except for the CUADD and CUNUMBR addresses specified in each definition, the definitions are identical.

Example H-4 IOCP statements for increasing device count to 496 on 4 channels

```
*ELWOOD/0
    CNTLUNIT CUNUMBR=0C11,
    PATH=(CSS(1),C1,C5,DA,F1),
    LINK=(CSS(1),1D,3D,4F,66),
    UNITADD=((00,16)),UNIT=3490,CUADD=0
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE100 IODEVICE ADDRESS=(E100,16),UNIT=3490,CUNUMBR=(0C11),
    UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/1
    CNTLUNIT CUNUMBR=0C12,
    PATH=(CSS(1),C1,C5,DA,F1),
```

```

LINK=(CSS(1),1D,3D,4F,66), X
UNITADD=((00,16)),UNIT=3490,CUADD=1
*$HCDC$ DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE110 IODEVICE ADDRESS=(E110,16),UNIT=3490,CUNUMBR=(0C12), X
UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/2
CNTLUNIT CUNUMBR=0C13, X
PATH=(CSS(1),C1,C5,DA,F1), X
LINK=(CSS(1),1D,3D,4F,66), X
UNITADD=((00,16)),UNIT=3490,CUADD=2
*$HCDC$ DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE120 IODEVICE ADDRESS=(E120,16),UNIT=3490,CUNUMBR=(0C13), X
UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/3
CNTLUNIT CUNUMBR=0C14, X
PATH=(CSS(1),C1,C5,DA,F1), X
LINK=(CSS(1),1D,3D,4F,66), X
UNITADD=((00,16)),UNIT=3490,CUADD=3
*$HCDC$ DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE130 IODEVICE ADDRESS=(E130,16),UNIT=3490,CUNUMBR=(0C14), X
UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/4
CNTLUNIT CUNUMBR=0C15, X
PATH=(CSS(1),C1,C5,DA,F1), X
LINK=(CSS(1),1D,3D,4F,66), X
UNITADD=((00,16)),UNIT=3490,CUADD=4
*$HCDC$ DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE140 IODEVICE ADDRESS=(E140,16),UNIT=3490,CUNUMBR=(0C15), X
UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/5
CNTLUNIT CUNUMBR=0C16, X
PATH=(CSS(1),C1,C5,DA,F1), X
LINK=(CSS(1),1D,3D,4F,66), X
UNITADD=((00,16)),UNIT=3490,CUADD=5
*$HCDC$ DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE150 IODEVICE ADDRESS=(E150,16),UNIT=3490,CUNUMBR=(0C16), X
UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/6
CNTLUNIT CUNUMBR=0C17, X
PATH=(CSS(1),C1,C5,DA,F1), X
LINK=(CSS(1),1D,3D,4F,66), X
UNITADD=((00,16)),UNIT=3490,CUADD=6
*$HCDC$ DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE160 IODEVICE ADDRESS=(E160,16),UNIT=3490,CUNUMBR=(0C17), X
UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/7
CNTLUNIT CUNUMBR=0C18, X
PATH=(CSS(1),C1,C5,DA,F1), X
LINK=(CSS(1),1D,3D,4F,66), X

```

```

UNITADD=((00,16)),UNIT=3490,CUADD=7
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE170 IODEVICE ADDRESS=(E170,16),UNIT=3490,CUNUMBR=(0C18),      X
            UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/8
CNTLUNIT CUNUMBR=0C19,      X
            PATH=(CSS(1),C1,C5,DA,F1),      X
            LINK=(CSS(1),1D,3D,4F,66),      X
            UNITADD=((00,16)),UNIT=3490,CUADD=8
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE180 IODEVICE ADDRESS=(E180,16),UNIT=3490,CUNUMBR=(0C19),      X
            UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/9
CNTLUNIT CUNUMBR=0C1A,      X
            PATH=(CSS(1),C1,C5,DA,F1),      X
            LINK=(CSS(1),1D,3D,4F,66),      X
            UNITADD=((00,16)),UNIT=3490,CUADD=9
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE190 IODEVICE ADDRESS=(E190,16),UNIT=3490,CUNUMBR=(0C1A),      X
            UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/A
CNTLUNIT CUNUMBR=0C1B,      X
            PATH=(CSS(1),C1,C5,DA,F1),      X
            LINK=(CSS(1),1D,3D,4F,66),      X
            UNITADD=((00,16)),UNIT=3490,CUADD=A
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1A0 IODEVICE ADDRESS=(E1A0,16),UNIT=3490,CUNUMBR=(0C1B),      X
            UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/B
CNTLUNIT CUNUMBR=0C1C,      X
            PATH=(CSS(1),C1,C5,DA,F1),      X
            LINK=(CSS(1),1D,3D,4F,66),      X
            UNITADD=((00,16)),UNIT=3490,CUADD=B
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1B0 IODEVICE ADDRESS=(E1B0,16),UNIT=3490,CUNUMBR=(0C1C),      X
            UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/C
CNTLUNIT CUNUMBR=0C1D,      X
            PATH=(CSS(1),C1,C5,DA,F1),      X
            LINK=(CSS(1),1D,3D,4F,66),      X
            UNITADD=((00,16)),UNIT=3490,CUADD=C
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1C0 IODEVICE ADDRESS=(E1C0,16),UNIT=3490,CUNUMBR=(0C1D),      X
            UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/D
CNTLUNIT CUNUMBR=0C1E,      X
            PATH=(CSS(1),C1,C5,DA,F1),      X
            LINK=(CSS(1),1D,3D,4F,66),      X
            UNITADD=((00,16)),UNIT=3490,CUADD=D

```

```

*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1D0 IODEVICE ADDRESS=(E1D0,16),UNIT=3490,CUNUMBR=(0C1E),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/E
          CNTLUNIT CUNUMBR=0C1F,      X
          PATH=(CSS(1),C1,C5,DA,F1),      X
          LINK=(CSS(1),1D,3D,4F,66),      X
          UNITADD=((00,16)),UNIT=3490,CUADD=E
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1E0 IODEVICE ADDRESS=(E1E0,16),UNIT=3490,CUNUMBR=(0C1F),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/F
          CNTLUNIT CUNUMBR=0C10,      X
          PATH=(CSS(1),C1,C5,DA,F1),      X
          LINK=(CSS(1),1D,3D,4F,66),      X
          UNITADD=((00,16)),UNIT=3490,CUADD=F
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1F0 IODEVICE ADDRESS=(E1F0,16),UNIT=3490,CUNUMBR=(0C10),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/10
          CNTLUNIT CUNUMBR=2C11,      X
          PATH=(CSS(1),C1,C5,DA,F1),      X
          LINK=(CSS(1),1D,3D,4F,66),      X
          UNITADD=((00,16)),UNIT=3490,CUADD=10
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE200 IODEVICE ADDRESS=(E200,16),UNIT=3490,CUNUMBR=(2C11),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/11
          CNTLUNIT CUNUMBR=2C12,      X
          PATH=(CSS(1),C1,C5,DA,F1),      X
          LINK=(CSS(1),1D,3D,4F,66),      X
          UNITADD=((00,16)),UNIT=3490,CUADD=11
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE210 IODEVICE ADDRESS=(E210,16),UNIT=3490,CUNUMBR=(2C12),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/12
          CNTLUNIT CUNUMBR=2C13,      X
          PATH=(CSS(1),C1,C5,DA,F1),      X
          LINK=(CSS(1),1D,3D,4F,66),      X
          UNITADD=((00,16)),UNIT=3490,CUADD=12
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE220 IODEVICE ADDRESS=(E220,16),UNIT=3490,CUNUMBR=(2C13),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/13
          CNTLUNIT CUNUMBR=2C14,      X
          PATH=(CSS(1),C1,C5,DA,F1),      X
          LINK=(CSS(1),1D,3D,4F,66),      X
          UNITADD=((00,16)),UNIT=3490,CUADD=13
*$HCDC$      DESC='ELWOOD CLUSTER 0 BARR74'

```

```

TAPEE230 IODEVICE ADDRESS=(E230,16),UNIT=3490,CUNUMBR=(2C14),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/14
          CNTLUNIT CUNUMBR=2C15,                                     X
          PATH=(CSS(1),C1,C5,DA,F1),                               X
          LINK=(CSS(1),1D,3D,4F,66),                               X
          UNITADD=((00,16)),UNIT=3490,CUADD=14
*$HCDC$   DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE240 IODEVICE ADDRESS=(E240,16),UNIT=3490,CUNUMBR=(2C15),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/15
          CNTLUNIT CUNUMBR=2C16,                                     X
          PATH=(CSS(1),C1,C5,DA,F1),                               X
          LINK=(CSS(1),1D,3D,4F,66),                               X
          UNITADD=((00,16)),UNIT=3490,CUADD=15
*$HCDC$   DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE250 IODEVICE ADDRESS=(E250,16),UNIT=3490,CUNUMBR=(2C16),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/16
          CNTLUNIT CUNUMBR=2C17,                                     X
          PATH=(CSS(1),C1,C5,DA,F1),                               X
          LINK=(CSS(1),1D,3D,4F,66),                               X
          UNITADD=((00,16)),UNIT=3490,CUADD=16
*$HCDC$   DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE260 IODEVICE ADDRESS=(E260,16),UNIT=3490,CUNUMBR=(2C17),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/17
          CNTLUNIT CUNUMBR=2C18,                                     X
          PATH=(CSS(1),C1,C5,DA,F1),                               X
          LINK=(CSS(1),1D,3D,4F,66),                               X
          UNITADD=((00,16)),UNIT=3490,CUADD=17
*$HCDC$   DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE270 IODEVICE ADDRESS=(E270,16),UNIT=3490,CUNUMBR=(2C18),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/18
          CNTLUNIT CUNUMBR=2C19,                                     X
          PATH=(CSS(1),C1,C5,DA,F1),                               X
          LINK=(CSS(1),1D,3D,4F,66),                               X
          UNITADD=((00,16)),UNIT=3490,CUADD=18
*$HCDC$   DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE280 IODEVICE ADDRESS=(E280,16),UNIT=3490,CUNUMBR=(2C19),      X
          UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/19
          CNTLUNIT CUNUMBR=2C1A,                                     X
          PATH=(CSS(1),C1,C5,DA,F1),                               X
          LINK=(CSS(1),1D,3D,4F,66),                               X
          UNITADD=((00,16)),UNIT=3490,CUADD=19
*$HCDC$   DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE290 IODEVICE ADDRESS=(E290,16),UNIT=3490,CUNUMBR=(2C1A),      X

```

```

UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/1A
  CNTLUNIT CUNUMBR=2C1B, X
    PATH=(CSS(1),C1,C5,DA,F1), X
    LINK=(CSS(1),1D,3D,4F,66), X
    UNITADD=((00,16)),UNIT=3490,CUADD=1A
*$HCDC$  DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE2A0 IODEVICE ADDRESS=(E2A0,16),UNIT=3490,CUNUMBR=(2C1B), X
    UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/1B
  CNTLUNIT CUNUMBR=2C1C, X
    PATH=(CSS(1),C1,C5,DA,F1), X
    LINK=(CSS(1),1D,3D,4F,66), X
    UNITADD=((00,16)),UNIT=3490,CUADD=1B
*$HCDC$  DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE2B0 IODEVICE ADDRESS=(E2B0,16),UNIT=3490,CUNUMBR=(2C1C), X
    UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/1C
  CNTLUNIT CUNUMBR=2C1D, X
    PATH=(CSS(1),C1,C5,DA,F1), X
    LINK=(CSS(1),1D,3D,4F,66), X
    UNITADD=((00,16)),UNIT=3490,CUADD=1C
*$HCDC$  DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE2C0 IODEVICE ADDRESS=(E2C0,16),UNIT=3490,CUNUMBR=(2C1D), X
    UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/1D
  CNTLUNIT CUNUMBR=2C1E, X
    PATH=(CSS(1),C1,C5,DA,F1), X
    LINK=(CSS(1),1D,3D,4F,66), X
    UNITADD=((00,16)),UNIT=3490,CUADD=1D
*$HCDC$  DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE2D0 IODEVICE ADDRESS=(E2D0,16),UNIT=3490,CUNUMBR=(2C1E), X
    UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/1E
  CNTLUNIT CUNUMBR=2C1F, X
    PATH=(CSS(1),C1,C5,DA,F1), X
    LINK=(CSS(1),1D,3D,4F,66), X
    UNITADD=((00,16)),UNIT=3490,CUADD=1E
*$HCDC$  DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE2E0 IODEVICE ADDRESS=(E2E0,16),UNIT=3490,CUNUMBR=(2C1F), X
    UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*

```

The MVSCP resulting from these IOCP statements is shown in Example H-5.

Example H-5 MVSCP with 496 devices connected by using the switch

```
IODEVICE ADDRESS=(E100,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,01),(MTL,NO)),CUNUMBR=0C11
IODEVICE ADDRESS=(E110,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,02),(MTL,NO)),CUNUMBR=0C12
IODEVICE ADDRESS=(E120,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,03),(MTL,NO)),CUNUMBR=0C13
IODEVICE ADDRESS=(E130,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,04),(MTL,NO)),CUNUMBR=0C14
IODEVICE ADDRESS=(E140,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,05),(MTL,NO)),CUNUMBR=0C15
IODEVICE ADDRESS=(E150,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,06),(MTL,NO)),CUNUMBR=0C16
IODEVICE ADDRESS=(E160,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,07),(MTL,NO)),CUNUMBR=0C17
IODEVICE ADDRESS=(E170,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,08),(MTL,NO)),CUNUMBR=0C18
IODEVICE ADDRESS=(E180,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,09),(MTL,NO)),CUNUMBR=0C19
IODEVICE ADDRESS=(E190,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,0A),(MTL,NO)),CUNUMBR=0C1A
IODEVICE ADDRESS=(E1A0,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,0B),(MTL,NO)),CUNUMBR=0C1B
IODEVICE ADDRESS=(E1B0,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,0C),(MTL,NO)),CUNUMBR=0C1C
IODEVICE ADDRESS=(E1C0,16),UNIT=3490,FEATURE=COMPACT,      *
        OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                  *
        USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
        4),(LIBPORT-ID,0D),(MTL,NO)),CUNUMBR=0C1D
```

```

IODEVICE ADDRESS=(E1D0,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,0E),(MTL,NO)),CUNUMBR=0C1E
IODEVICE ADDRESS=(E1E0,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,0F),(MTL,NO)),CUNUMBR=0C1F
IODEVICE ADDRESS=(E1F0,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,10),(MTL,NO)),CUNUMBR=0C10
IODEVICE ADDRESS=(E200,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,11),(MTL,NO)),CUNUMBR=2C11
IODEVICE ADDRESS=(E210,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,12),(MTL,NO)),CUNUMBR=2C12
IODEVICE ADDRESS=(E220,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,13),(MTL,NO)),CUNUMBR=2C13
IODEVICE ADDRESS=(E230,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,14),(MTL,NO)),CUNUMBR=2C14
IODEVICE ADDRESS=(E240,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,15),(MTL,NO)),CUNUMBR=2C15
IODEVICE ADDRESS=(E250,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,16),(MTL,NO)),CUNUMBR=2C16
IODEVICE ADDRESS=(E260,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,17),(MTL,NO)),CUNUMBR=2C17
IODEVICE ADDRESS=(E270,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,18),(MTL,NO)),CUNUMBR=2C18
IODEVICE ADDRESS=(E280,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,19),(MTL,NO)),CUNUMBR=2C19
IODEVICE ADDRESS=(E290,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*
  4),(LIBPORT-ID,1A),(MTL,NO)),CUNUMBR=2C1A
IODEVICE ADDRESS=(E2A0,16),UNIT=3490,FEATURE=COMPACT,      *
  OFFLINE=YES,DYNAMIC=YES,LOCANY=YES,                        *
  USERPRM=((LIBRARY,YES),(AUTOSWITCH,YES),(LIBRARY-ID,BA07*

```

```

4), (LIBPORT-ID,1B), (MTL,NO)), CUNUMBR=2C1B
IODEVICE ADDRESS=(E2B0,16), UNIT=3490, FEATURE=COMPACT, *
OFFLINE=YES, DYNAMIC=YES, LOCANY=YES, *
USERPRM=((LIBRARY,YES), (AUTOSWITCH,YES), (LIBRARY-ID,BA07*
4), (LIBPORT-ID,1C), (MTL,NO)), CUNUMBR=2C1C
IODEVICE ADDRESS=(E2C0,16), UNIT=3490, FEATURE=COMPACT, *
OFFLINE=YES, DYNAMIC=YES, LOCANY=YES, *
USERPRM=((LIBRARY,YES), (AUTOSWITCH,YES), (LIBRARY-ID,BA07*
4), (LIBPORT-ID,1D), (MTL,NO)), CUNUMBR=2C1D
IODEVICE ADDRESS=(E2D0,16), UNIT=3490, FEATURE=COMPACT, *
OFFLINE=YES, DYNAMIC=YES, LOCANY=YES, *
USERPRM=((LIBRARY,YES), (AUTOSWITCH,YES), (LIBRARY-ID,BA07*
4), (LIBPORT-ID,1E), (MTL,NO)), CUNUMBR=2C1E
IODEVICE ADDRESS=(E2E0,16), UNIT=3490, FEATURE=COMPACT, *
OFFLINE=YES, DYNAMIC=YES, LOCANY=YES, *
USERPRM=((LIBRARY,YES), (AUTOSWITCH,YES), (LIBRARY-ID,BA07*
4), (LIBPORT-ID,1F), (MTL,NO)), CUNUMBR=2C1F

```

Sharing ports

One CPC can use four ports and another can use the other four ports. You can have up to eight CPCs, each connected to a switched or direct port. Or you can connect all CPCs to all ports (switched). You can also have one CPC that uses all eight ports and another that uses fewer than eight.

Example H-5 on page 933 uses only four CHPIDs/LINKS in each PATH statement. To use the other four ports available on a second CPC, use those values on the first CPC, then change the values as shown in Example H-6 on the second CPC. The only differences are that the PATHs and LINKs are different values on the second CPC from the first CPC.

Example H-6 IOCP statement for using four ports on a second CPC

```

*ELWOOD/0
      CNTLUNIT CUNUMBR=0C11, X
      PATH=(CSS(1),C0,C4,D9,F0), X
      LINK=(CSS(1),1C,3C,4E,65), X
      UNITADD=(00,16), UNIT=3490, CUADD=0
*$HCDC$  DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE100 IODEVICE ADDRESS=(E100,16), UNIT=3490, CUNUMBR=(0C11), X
      UNITADD=00, PART=(CSS(1),MVSC7,VMT07)
*

```

LIBPORT-IDs in the MVSCP

Table H-1 of LIBPORT-IDs can be helpful. For Cluster 0, 256 devices are 01 - 10 and 496 devices are 01 - 1F. LIBPORT-ID is always one more than CUADD.

Table H-1 LIBPORT-IDs

Distributed Library ID	Logical CUs	Libport/Subsystem IDs
0	0-1E	X'01'-X'1F'
1	0-1E	X'41'-X'5F'
2	0-1E	X'81'-X'9F'
3	0-1E	X'C1'-X'DF'
4	0-1E	X'21'-X'3F'
5	0-1E	X'61'-X'7F'



Case study for logical partitioning of a two-cluster grid

Hardware must be used and managed as effectively as possible to ensure your investments. One of the ways to protect the investment is by using the same hardware for more than one sysplex/host. Important points to consider are described, such as common areas that might need other technical competencies in addition to the storage team to be involved within the information technology (IT) structure for the correct sizing and planning.

It also gives you a practical guideline for aspects of the project, such as naming conventions and checklists. The solution is based on standard functions from IBM z/OS, Data Facility Storage Management Subsystem Removable Media Manager (DFSMSrmm), IBM Resource Access Control Facility (RACF), and functions available in the TS7700 2-cluster grid. A similar implementation can be done in any single-cluster or multi-cluster grid configuration.

The TS7700 R2.0 extended the possibilities of manageability and usability of the cluster or grid by introducing the *Selective Device Access Control (SDAC)*. The SDAC enables you to split the grid or cluster into hard partitions that are accessible by independent hosts or applications.

SDAC, also known as *hard partitioning*, can isolate and secure environments with various requirements and objectives, shielding them from unintended or malicious interference between hosts. This is accomplished by granting access to determined ranges of logical volumes by selected groups of devices in a logical control unit (LCU) granularity (also referred to as *LIBPORT-ID*).

An example of a real implementation of this function is provided, going through the necessary steps to separate the environments Production (named PROD) and Test (named TEST) from each other despite sharing the TS7700 2-cluster grid. This appendix includes the following sections:

- ▶ Overview of partitioning
- ▶ Definitions and settings in z/OS
- ▶ Definitions on the TS7700 Management Interface
- ▶ Verification of changes

Overview of partitioning

This description leads you through the steps for the logical partitioning of two hosts from one client: PROD and TEST.

The setup must be as complete as possible and established in a way that generates the best possible protection against unauthorized access to logical volumes dedicated to the other partition. You must protect against unauthorized user access for logical volumes on PROD from TEST and vice versa.

The function SDAC, introduced with R2.0, is used. It can be ordered as Feature Code 5271.

Other requirements must be agreed on before implementation. Depending on whether this is a multiclient or a single client multi-logical partition (LPAR) environment, consider the following requirements:

- ▶ Acceptance of sharing the two-cluster grid
- ▶ Specific security requirements
- ▶ Bandwidth that is needed per host
- ▶ Number of Fibre Channel Connection (FICON) channels per host
- ▶ Acceptance of shared or dedicated FICON channels
- ▶ Number of virtual drives and logical tapes that are needed per host
- ▶ Number of physical drives and physical tapes needed
- ▶ Tape security in RACF: Volume-related, data set-related, or both

Establish defined naming conventions before making your definitions. This makes it easier to logically relate all definitions and structures to one host when updates are needed.

Figure I-1 gives an overview of the setup. Updates are needed on many places. Adapt your current naming standards to your setup.




Figure I-1 Logical partitioning overview

Definitions and settings in z/OS

The needed definitions are covered in the components that run in z/OS. The same names and numbers must be defined and repeated on the clusters:

- ▶ Hardware configuration definitions (HCDs) to define:
 - CTLUNITS for the 3490E devices
 - FICON channels
 - Esoteric definitions
- ▶ PARMLIB definitions to define for these areas:
 - Object access method (OAM) definitions if the OAM-started task is new on that host.
 - Global resource serialization (GRS) parameters to enable the Auto Switchable (AS) function if needed. This is only valid if you have a sysplex with a tape management system (TMS) that is separate from the tape configuration databases (TCDBs).
 - Missing Interrupt Handler (MIH) values for the device addresses.
 - Category updates.
 - Commands member to vary defined devices online.
- ▶ DFSMSrmm definitions:
 - Volume ranges
 - PRTITION and OPENRULE definitions to avoid usage of volumes from the other host (replacement of REJECT and ANYUSE)
- ▶ JES2 JCL Procedure Library updates
 - JCL definitions to enable start of OAM-started task
- ▶ SMS updates for constructs and ACS routines:
 - Definition of libraries and SMS constructs (Storage Class (SC), Management Class (MC), and so on)
 - Updates to ACS routines according to your conventions
- ▶ RACF updates:
 - Define started task OAM and required RACF profiles
 - Decide and define the required security settings regarding access to tape volumes and data sets on tape
- ▶ Automation updates:
 - If OAM is new, several updates are needed.
 - If you choose to vary devices online by using the automation product, updates are needed.
 - New messages must be evaluated and automated.

Figure I-2 shows the z/OS updates needed in the case study that define specific volume ranges, several device addresses, and scratch (Fast Ready) categories.

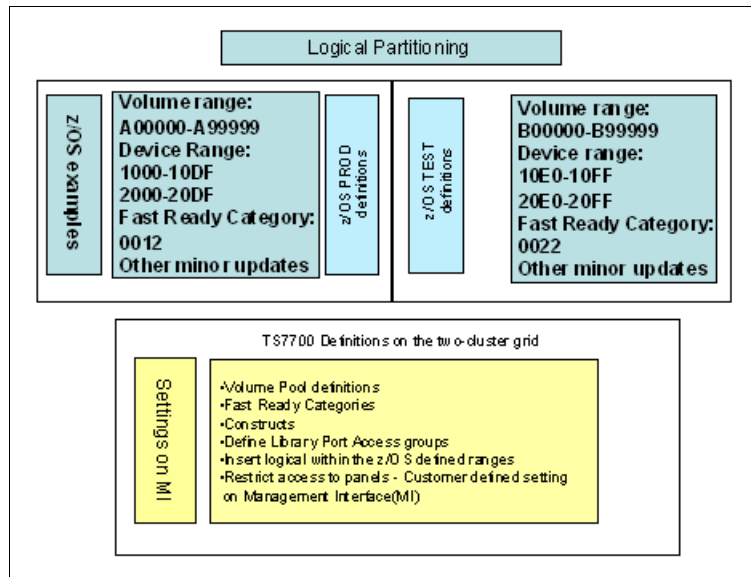


Figure I-2 Software updates with sample values on z/OS

Definitions in HCD

In HCD, you define the devices that are needed for each of the hosts and connect them to the LPARS. This case study defines 28 out of 32 CUs for PROD (2 x 224 logical devices) and 4 out of 32 CUs (2 x 32 logical devices) for TEST. The devices for Cluster 0 have addresses from 1000 - 10FF and for Cluster 1 the values are 2000 - 20FF (Table I-1).

Table I-1 HCD definitions

Host	Cluster	CTLUNIT definition	LIBPORT-ID
PROD	Cluster 0	CUADD 00-0D (1000 - 10DF)	01-0E
	Cluster 1	CUADD 40-4D (2000 - 20DF)	41-4E
TEST	Cluster 0	CUADD 0E-0F (10E0 - 10FF)	0F-10
	Cluster 1	CUADD 4E-4F (20E0 - 20FF)	4F-50

More definitions are needed:

- ▶ The devices must be connected to processors.
- ▶ Devices must be related to esoterics definitions.
- ▶ Definitions regarding the FICON infrastructure and FICON directors are not included.

Normally, you can activate the new definitions dynamically. Details regarding HCD definitions are described in 6.4, “Hardware configuration definition” on page 216.

PARMLIB definitions

Use PARMLIB to define all the essential parameters that are needed to run the z/OS system. Some parameters apply to this case study and definitions can be made with same values on both hosts (TEST and PROD). All the described parameters can be activated dynamically on the current release of z/OS.

See the following resources for complete information regarding options in PARMLIB, definitions, and commands to activate without IPL:

- ▶ *MVS System Commands, SA22-7627*
- ▶ *MVS Initialization and Tuning Reference, SA22-7592*

The updates are within the following members of PARMLIB, where the suffix and the exact name of the PARMLIB data set apply to your naming standards. It is important to make the changes according to normal change rules. If the updates are not implemented correctly, severe problems can occur when the next IPL is planned.

- ▶ IEFSSNxx. These updates apply for TEST and PROD:
 - If OAM is new to the installation, the definitions in Example I-1 are required.

Example I-1 OAM subsystem definition

```
*-----*/
/* OAM - OBJECT ACCESS METHOD - ATL Control */
/*-----*/
SUBSYS SUBNAME(OAM1)
      INITRTN(CBRINIT)
```

- ▶ SCHEDxx. These updates apply for TEST and PROD:
 - If OAM is new to the installation, definitions in Example I-2 are required to start OAM. These definitions require you to start OAM as part of the normal IPL sequence by using your own automation product.

Example I-2 OAM definition in SCHEDxx

PPT	PGMNAME(CBROAM)	/* OAM ADDRESS SPACE	*/
	KEY(5)	/* KEY(5) PROTECTION KEY	*/
	NOSWAP	/* NOSWAP NON SWAPPABLE	*/
	SYST	/* SYST SYSTEM TASK	*/

- ▶ GRSRNLxx. These updates apply for TEST and PROD:
 - If the platform has the prerequisites for use of Auto Switchable (AS) devices, runs in GRS goal mode, or uses Coupling Facilities hardware, AS support can be enabled by the values in Example I-3 on page 942. AS offers the ability to have the devices online on all LPARS in a sysplex and reduces your need for specific products that have similar functions. AS has these requirements:
 - Devices are defined as AS in HCD.
 - Operators or automation products issue **VARY** commands.

Tip: **V 1000,AS,ON** makes the specified address available for AS support. When followed by **V 1000,ONLINE**, it varies the device online. Both commands must be entered on all hosts that need device 1000 online and auto-switchable.

Example I-3 illustrates enabling AS support.

Example I-3 GRS definition to enable AS support

```
/*-----*/  
/* Enable AS support */  
/*-----*/  
RNLDEF RNL(INCL)  
        TYPE(GENERIC)  
        QNAME(SYSZVOLS)
```

Note: A Parallel Sysplex must have a shared TMS - plex (TMSplex), SMSplex, and TCBD. In this case, no partitioning can be defined.

- ▶ IECIOSxx. In this member, you can define specific device ranges, and you must separate TEST from the PROD updates:

- TEST updates are in Example I-4, one line for each range of devices. The MOUNTMSG parameters ensure that the console receives the Mount Pending message (I0S070E) if a mount is not complete within 10 minutes. You can adjust this value. It depends on many factors, such as read/write ratio on the connected host and available capacity in the grid.

Example I-4 IECIOSxx updates for specific TEST device addresses

```
MIH DEV=(10E0-10FF),TIME=45:00  
MIH DEV=(20E0-20FF),TIME=45:00  
MIH MOUNTMSG=YES,MNTS=10:00
```

- PROD updates are in Example I-5, one line for each range of devices.

Example I-5 IECIOSxx updates for specific PROD device addresses

```
MIH DEV=(1000-10DF),TIME=45:00  
MIH DEV=(2000-20DF),TIME=45:00  
MIH MOUNTMSG=YES,MNTS=10:00
```

- ▶ DEVSUPxx. In this member, you can define specific device ranges. You must be specific and separate TEST from the PROD updates:

- DEVSUPxx for TEST is shown in Example I-6 for the categories that apply for TEST.

Example I-6 DEVSUPxx updates for specific TEST category

```
COMPACT=YES,  
MEDIA2=0022,  
ERROR=002E,  
PRIVATE=002F,  
VOLNSNS=YES
```

- DEVSUPxx for PROD is shown in Example I-7 for the categories that apply for PROD.

Example I-7 DEVSUPxx updates for specific PROD category

```
COMPACT=YES,  
MEDIA2=0012,  
ERROR=001E,  
PRIVATE=001F,  
VOLNSNS=YES
```

- ▶ COMMANDxx can be used to vary the range of devices online after IPL:
 - For TEST, apply a specific range of devices as shown in Example I-8.

Example I-8 Vary devices online after IPL for TEST

```
COM='V 10E0-10FF,ONLINE'
COM='V 20E0-20FF,ONLINE'
```

- For PROD, apply a specific range of devices as shown in Example I-9.

Example I-9 Vary devices online after IPL for PROD

```
COM='V 1000-10DF,ONLINE'
COM='V 2000-20DF,ONLINE'
```

DFSMSrmm definitions

In this case study, you have DFSMSrmm as the TMS. Equivalent definitions must be defined if you prefer to use another vendor's TMS. These definitions can be created by using options in DFSMSrmm PRTITION and OPENRULE. PRTITION is the preferred method for partitioning. **REJECT** commands, although still supported, must not be used in new installations. If you use **REJECT** commands, you must convert from the use of **REJECT** commands to use the **PRTITION** and commands.

See *z/OS DFSMSrmm Implementation and Customization Guide, SC23-6874*, for complete information about options for DFSMSrmm. Table I-2 shows the definitions that are needed in this specific case study. You must define the volume range that is connected to the host. Reject use and insertion of volumes that are connected to the other host.

Table I-2 DFSMSrmm options

Host	VLPOOL definitions	PRTITION definitions	OPENRULE definitions
PROD	VLPOOL PREFIX(A*) TYPE(S) DESCRIPTION (PROD DEFAULT') MEDIANAME(*)	PRTITION VOLUME(B*) TYPE(NONRMM) SMT(IGNORE) NOSMT(IGNORE)	OPENRULE VOLUME(B*) TYPE(RMM) ANYUSE(REJECT)
TEST	VLPOOL PREFIX(B*) TYPE(S) DESCRIPTION (TEST DEFAULT') MEDIANAME(*)	PRTITION VOLUME(A*) TYPE(NONRMM) SMT(IGNORE) NOSMT(IGNORE)	OPENRULE VOLUME(A*) TYPE(RMM) ANYUSE(REJECT)

JES2 definitions

If OAM is new to the hosts, OAM JCL must be defined in one of the JES2 procedure libraries. These JCL definitions apply for TEST and for PROD as shown in Example I-10.

Example I-10 JCL for OAM-started task

```
//OAM PROC OSMC=YES,MAXS=2,UNLOAD=9999,RESTART=NO
//IEFPROC EXEC PGM=CBROAM,REGION=64M,
// PARM=('OSMC=&OSMC,APLAN=CBROAM,MAXS=&MAXS,UNLOAD=&UNLOAD',
//       'RESTART=&RESTART')
//*
//SYSABEND DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
```

SMS constructs and definitions

This description does not explain in detail how to create ACS routines or how to create ACS constructs. It states the ACS constructs that are needed and suggests a naming convention that helps you understand the constructs relationship:

- ▶ SC definition:
 - Preference Group 0 assigned to volumes that are unlikely to be accessed
 - Preference Group 1 for volumes likely to be accessed
- ▶ MC definition:

Not relevant for the partitioning, but you must have separate sets for TEST and PROD.
- ▶ DC definitions:

Definition and relation to logical volume sizes, in this example, 800 MB and 6000 MB.
- ▶ SG definitions:

Pointing the SG to the composite library definition
- ▶ ACS and library definitions must also be defined, but they are not described here.

The naming convention in this case defines that all TEST definitions are prefixes with TS and PROD with PR.

ACS constructs and definitions for TEST are in Table I-3. Ensure that the construct names match the names that you define on the Management Interface (MI).

Table I-3 ACS construct names and important values

ACS constructs	ACS construct name for TEST/PROD	Important value
SC for Preference Group 0	TSSCPG0/PRSCPG0	N/A
SC for Preference Group 1	TSSCPG1/PRSCPG1	N/A
MC for one copy in cluster0	TSMCCL0/PRMCCL0	N/A
MC for one copy in cluster1	TSMCCL1/PRMCCL1	N/A
DC for 800 MB volume	TSDC800M/PRDC800M	Media recording must be MEDIA2. Recording tech must be 36TRACKS.
DC for 6000 MB volume	TSDC6GB/PRDC6GB	Media recording must be MEDIA2. Recording tech must be 36TRACKS.
SG to relate to composite library	TSCOMP1/PRCOMP1 ^a	Library name must match the SMS-defined name for composite library.

a. Nomenclature derives from TEST and PROD plus composite library number. Many clients have more than one grid installed.

RACF definitions

General rules for RACF definitions are defined by your policies. Security, in the areas of access to updates of tape information in DFSMSrmm and protection of access to data sets on tape volumes, is improved from IBM z/OS 1.8. However, many installations seem to forget that access to read and write tape volumes and tape data sets is by default not restricted with RACF settings. You need to define your own rules to protect for unauthorized access.

The same rules apply for access to run updates of the content in DFSMSrmm. Various solutions can be implemented.

For more information about options, security options, and access rules, see *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874.

Automation activities

If OAM and TS7700 are new on the host, evaluate these concerns:

- ▶ OAM must start after the IPL.
- ▶ New messages are introduced.
- ▶ Hardware errors and operator interventions occur and must be handled.

See the *IBM Virtualization Engine TS7700 Series Operator Informational Messages*, which is available at the Techdocs library website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101689>

Definitions on the TS7700 Management Interface

Updates and definitions that are needed on the windows of the Management Interface (MI) are covered. The definitions must match the definitions on the z/OS hosts to make it work. Make updates in the areas that are shown in Figure I-3.

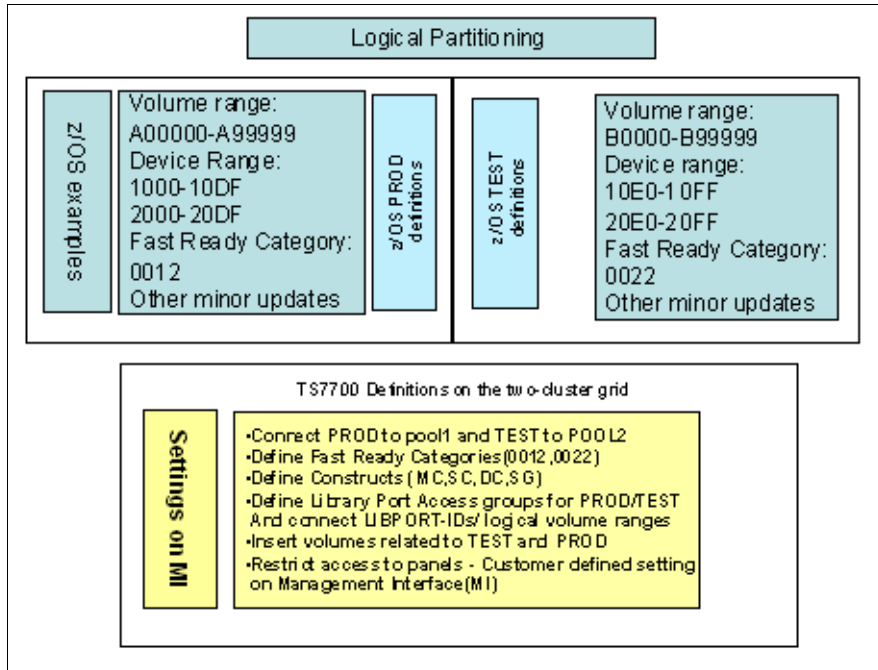


Figure I-3 Management Interface updates for logical partitioning of TEST and PROD

Physical volume pools

You must decide based on your requirements whether TEST and PROD can be on the same physical volumes, share volumes from the same physical volume pools, or must be placed on separate physical volumes:

- ▶ If sharing of same physical volumes is *acceptable*, all logical volumes from TEST and PROD are mixed on the physical VOLSERs, and separate pools for PROD and TEST are not required.
- ▶ If sharing of same physical volume ranges is *unacceptable*, total separation of physical volume ranges can be accomplished by assigning specific physical volume ranges for each pool and by using the *noborrow/keep* options.

The Physical Volume Pool Modify Properties window that is shown in Figure I-4 is used to define reclaim policies for each of the pools. Assign logical volumes to pool 1 for PROD and pool 2 for TEST, as shown in Figure I-9 on page 949.




Figure I-4 Volume pools for PROD and TEST

Scratch (Fast Ready) categories

The TS7700 performs scratch (Fast Ready) mounts directly to the Tape Volume Cache (TVC) with no physical tape involvement. This needs to be enabled on all hosts. *At Release 3.0, all categories that are defined as scratch inherit the Fast Ready attribute.* In prior releases, it was necessary to set this option in the MI.

Defining constructs

Define the constructs on all clusters in the grid with definitions consistent with your policies. Consider the following information:

1. Define an SC for TEST named TSSCPG0. This is an SC for volumes unlikely to be used (level 0 or PG0), as shown in Figure I-5. The other SCs must be defined in a similar way. Ensure that Tape Volume Cache Preference is set according to the defined SC name.




Figure I-5 Define Storage Class for TEST

- Define an MC for TEST named TSMCCL0 with only one copy of data and cache residency in cluster0 as shown in Figure I-6. Set the Copy Consistence Point to RUN on cluster0 and NOCOPY on cluster1. The other MCs are defined in a similar way.

Remember: Without the use of MC from z/OS, the default is a copy in both clusters.

Add/Copy Management Classes

Management class settings

Name:

Secondary pool:

Retain copy mode:

Description:

Management class settings for each cluster

Clusters	Copy Mode	Options
"Yacko[0]" (#BA02A)	<input type="text" value="Rewind Unload (RUN)"/>	Scratch mount candidate: <input checked="" type="checkbox"/>
"Wacko[1]" (#BA02B)	<input type="text" value="No Copy"/>	Scratch mount candidate: <input type="checkbox"/>

Figure I-6 Management Class for TEST

- Define a DC for TEST named TSDC800M as shown in Figure I-7.

"#BA002": Add Data Class

Name:

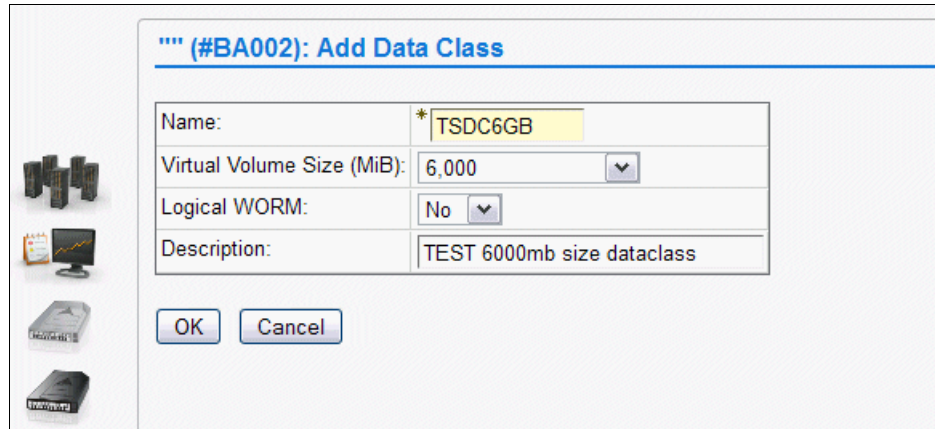
Virtual Volume Size (MiB):

Logical WORM:

Description:

Figure I-7 Data Class for TEST

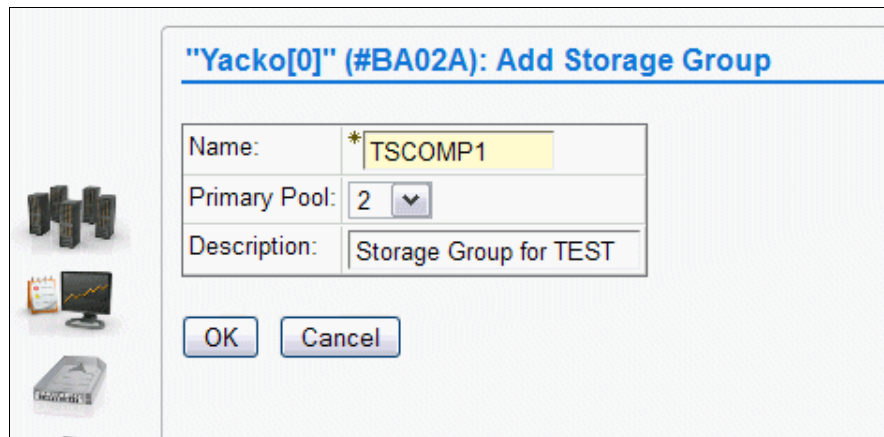
4. Also, define a DC for TEST that is named TSDC6GB (6000 MB), as shown in Figure I-8.



"" (#BA002): Add Data Class	
Name:	*TSDC6GB
Virtual Volume Size (MiB):	6,000
Logical WORM:	No
Description:	TEST 6000mb size dataclass
OK Cancel	

Figure I-8 Data Class for TEST with 6000 MB volume size

5. Define an SG for TEST (named TSCOMP1) as shown in Figure I-9. Remember your requirements for secluded physical volumes as described in "Physical volume pools" on page 946. The definition for PROD is not shown, but it must relate to volume pool 1. Define the SG on both clusters in the grid.



"Yacko[0]" (#BA02A): Add Storage Group	
Name:	*TSCOMP1
Primary Pool:	2
Description:	Storage Group for TEST
OK Cancel	

Figure I-9 Storage Group for TEST

Library Port Access Groups

The library port access group windows are available only when the Selective Device Access Control (SDAC) feature is installed.

Perform the following steps for Library Port Access Groups:

1. Now, define two access groups to relate the CUs (LIBPORT-IDs) to logical volume ranges by selecting **Settings** → **Library Port Access Group** and then clicking **Add** from the menu.

Figure I-10 shows creating the access group TEST and connecting it to Library Ports 0F and 10 on Cluster 0.

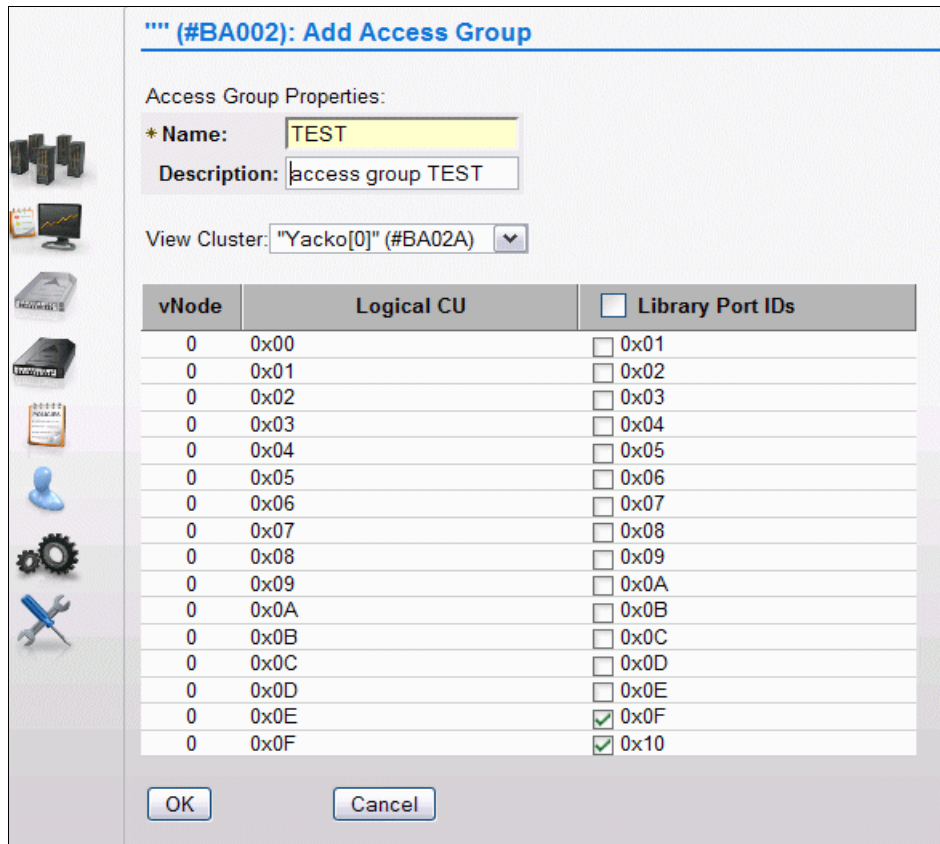


Figure I-10 Add the access group for TEST LPAR

- Now, use the menu to select Cluster 1 to add Lib Ports 4F and 50 to the TEST access group as shown in Figure I-11.

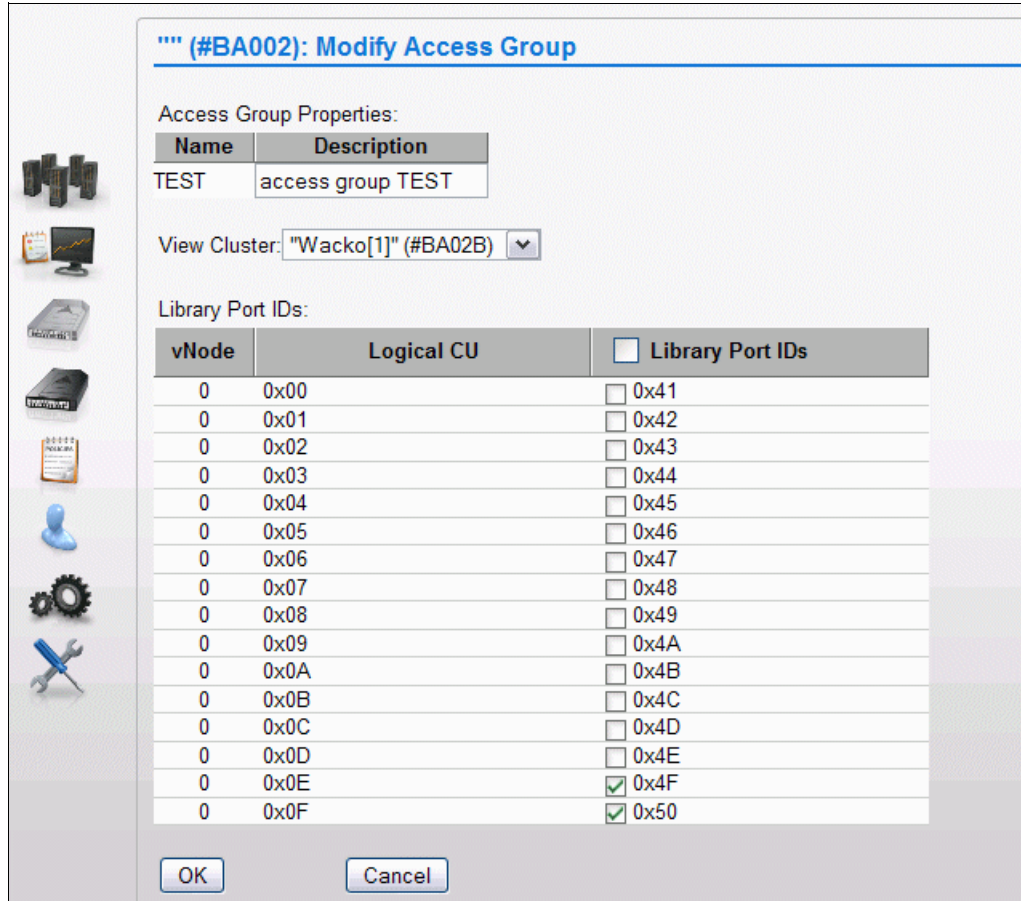


Figure I-11 Modify an access group for TEST LPAR

- PROD is connected to both clusters as well, but by using the opposite Lib Ports, 01-0E on Cluster 0, as shown in Figure I-12.




Figure I-12 Add an access group for PROD LPAR

4. Define 41-4E on Cluster 1 as shown in Figure I-13.




Figure I-13 Modify the access group for PROD LPAR

5. Check to see that each access group's ranges are not overlapping, as shown in Figure I-14.




Figure I-14 Summary display of access groups

- Define logical volume ranges to the access groups as shown in Figure I-15.




Figure I-15 Add VOLSER range for TEST LPAR

- You can see that the ranges are assigned to their correct hosts, as shown in Figure I-16.




Figure I-16 Summary of access groups and VOLSER ranges

Logical volume ranges or insert volumes connected to defined ranges

On the MI window as shown Figure I-17, insert the number of logical volumes that fits your initial need for PROD and TEST. The example shows how to insert 1000 logical volumes for the TEST partition. Normally, you define the inserted volume size to be 800 MB (ECCST). When used on z/OS, the assignment of DC defines the maximum volume size, which is 800 MB or 6000 MB in the case study.

The screenshot displays the 'Virtual Volumes' management interface. On the left is a navigation sidebar with icons for various actions. The main area is titled 'Virtual Volumes' and contains several sections:

- Virtual Volume Details:** A table showing 'Current availability across entire grid':

Currently Inserted:	100,000
Maximum Allowed:	1,000,000
Available Slots:	900,000
- Insert Virtual Volumes:** A section with a 'Show inserted volume ranges between' field set to '000000' and 'ZZZZZZ', with a 'Show' button.
- Insert a new virtual volume range:** A section with the following options:
 - *Starting volser: B00000
 - *Quantity: 1000 (selected)
 - Ending volser: (empty)
 - Initially owned by: "Yacko[0]" (#BA02A)
 - Media type: Enhanced Capacity Cartridge System Tape (800 MiB) (selected)
 - Set Constructs: checked
 - Storage Groups: TSCOMP1
 - Management Classes: TSMCCL0
 - Storage Classes: TSSCPG0
 - Data Classes: TSDC800M

An 'Insert' button is located at the bottom of the configuration section.

Figure I-17 Insert volumes for TEST

User Management on the Management Interface

Depending on your requirements, you can define roles and rules for the users that use the MI to prevent unauthorized access to data and to set user privileges. For more information, see 9.2.9, "The Access icon" on page 446.

Verification of changes

After setting the definitions, evaluate your setup against the one shown in Figure I-18. If you try to read or write to a logical volume belonging to the other host, the job fails and a message presents the reason.

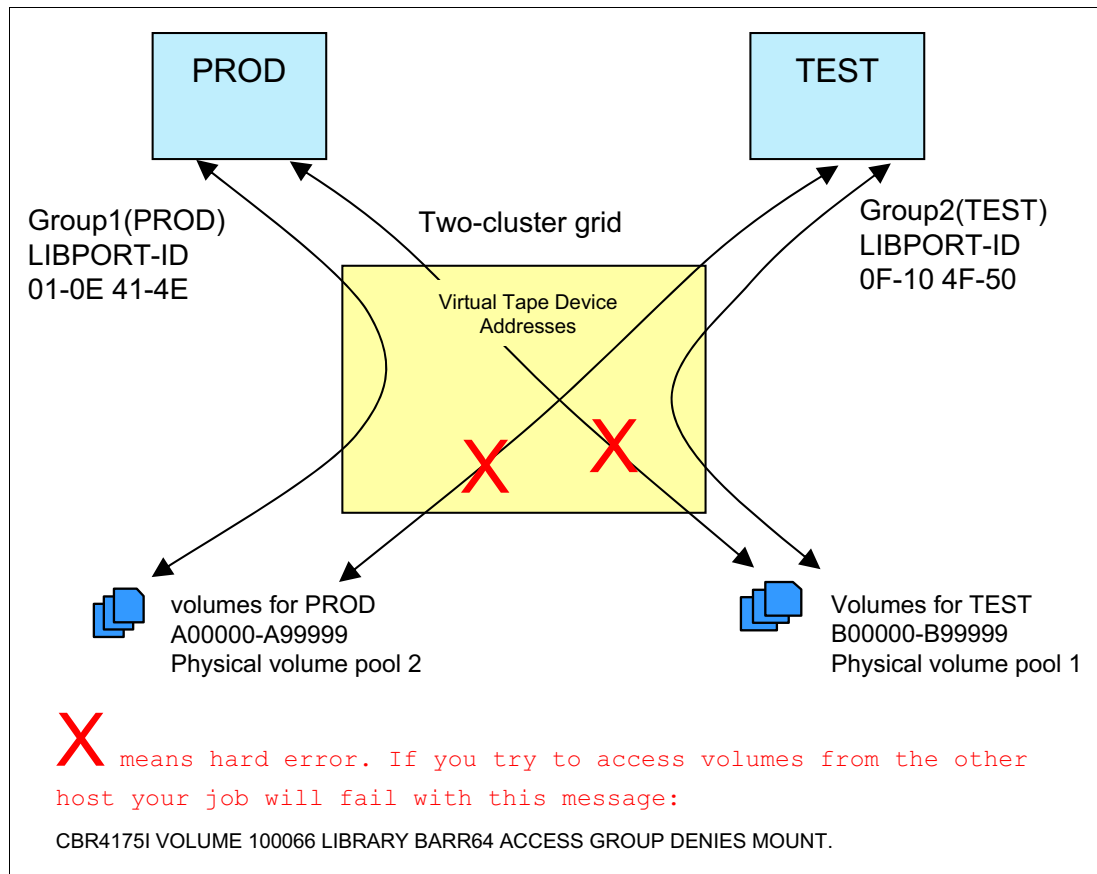


Figure I-18 Result after all definitions are finished

You can run several procedures to ensure that the setup is correct and ready for production. Be sure that you cover the following points:

- ▶ Control that all settings are as expected on z/OS and on the MI.
- ▶ Configure the channels online and vary the devices online on your system.
- ▶ Enter one logical volume and ensure that the volume is entered in the correct partition.
- ▶ Look up the volume by using the started task OAM, or by entering the command:
D SMS,VOL,*volser*,DETAIL
- ▶ Check the values of the scratch tape in DFSMSrmm by using dedicated windows.
- ▶ Create a job that creates a tape and evaluate that the constructs are assigned correctly.
- ▶ Issue D SMS*hi*,VOL,*volser*,DETAIL again to check the assignment of constructs from the grid.
- ▶ Use the Library Request host console commands to evaluate the status on the created private volume and the status of the physical volume to which it was copied.

- ▶ Use the MI to further evaluate the related physical volume pool.
- ▶ Ensure that constructs and definitions on both clusters are tested and evaluated.
- ▶ Make the last and final evaluation after the next IPL of the system to validate that the dynamic commands that are given are reflected in the required data sets, such as PARMLIB.

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics covered in this book.

IBM Redbooks publications

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only:

- ▶ *Continuous Availability - Systems Design Guide*, SG24-2085
- ▶ *Continuous Availability S/390 Technology Guide*, SG24-2086
- ▶ *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 TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *IBM z/OS DFSMSHsm Primer*, SG24-5272
- ▶ *IBM z Systems Connectivity Handbook*, SG24-5444
- ▶ *Introduction to IBM S/390 FICON*, SG24-5176
- ▶ *Introduction to SAN Distance Solutions*, SG24-6408

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- ▶ *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141
- ▶ *FICON Planning and Implementation Guide*, SG24-6497
- ▶ *IBM Encryption Key Manager component for the Java platform Introduction, Planning, and User's Guide*, GA76-0418
- ▶ *IBM Security Key Lifecycle Manager for z/OS IPLA License Info*, GI11-8737
- ▶ *IBM Security Key Lifecycle Manager Installation, Planning, and User's Guide*, SC14-7628
- ▶ *IBM Security Key Lifecycle Manager Version 2.5 Installation and Configuration Guide*, SC27-5335
- ▶ *IBM Security Key Lifecycle Manager Version 2.5 Quick Start Guide*, GI13-2316
- ▶ *IBM System Storage TS1120 and TS1130 Tape Drives and TS1120 Controller Introduction and Planning Guide*, GA32-0555

- ▶ *IBM System Storage TS1120 and TS1130 Tape Drives and TS1120 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 System Storage TS3500 Tape Library with ALMS Operator Guide*, GA32-0594
- ▶ *IBM TotalStorage Enterprise Tape System 3592 Operators Guide*, GA32-0465
- ▶ *IBM TotalStorage UltraScalable Tape Library 3584 Operator Guide*, GA32-0468
- ▶ *IBM TS3500 Tape Library with ALMS Introduction and Planning Guide*, GA32-0593
- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *IBM Virtualization Engine TS7700 Series Introduction and Planning Guide*, GA32-0567
- ▶ *Implementing System Managed Storage*, SC26-3123
- ▶ *VM/ESA DFSMS/VM Removable Media Services User's Guide and Reference*, SC24-6090
- ▶ *z/OS DFSMS Access Method Services Commands*, SC23-6846
- ▶ *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867
- ▶ *z/OS DFSMSdftp Storage Administrator*, SC23-6860
- ▶ *z/OS DFSMSdftp Utilities*, SC23-6864
- ▶ *z/OS DFSMSdss Storage Administration*, SC23-6868
- ▶ *z/OS DFSMSshsm Storage Administration*, SC23-6871
- ▶ *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874
- ▶ *z/OS DFSMSrmm Managing and Using Removable Media*, SC23-6873
- ▶ *z/OS JES2 Initialization and Tuning Reference*, SA32-0992
- ▶ *z/OS JES3 Initialization and Tuning Reference*, SA32-1005
- ▶ *z/OS MVS Initialization and Tuning Reference*, SA23-1380
- ▶ *z/OS MVS Planning: Operation*, SC22-7601
- ▶ *z/OS MVS System Commands*, SA38-0666
- ▶ *z/VM V6R1.0 DFSMS/VM Planning Guide*, SC24-6184
- ▶ *z/VM V6R1.0 DFSMS/VM Storage Administration*, SC24-6186
- ▶ *z/VM V6R2.0 DFSMS/VM Removable Media Services*, SC24-6185
- ▶ *z/VSE System Administration Guide*, SC34-2627
- ▶ *z/VSE System Macros Reference*, SC34-2708

Technical documents on the IBM Techdocs website

IBM publishes many detailed technical documents during the lifetime of a product. IBM makes a great effort to ensure the reliability and accuracy of the content. It is of great benefit to you to use these technical papers.

The documents in IBM Techdocs are active. The content is constantly changing and new documents are being created. To ensure that you reference the newest document, search on the Techdocs website:

<http://www.ibm.com/support/techdocs/atstr.nsf/Web/TechDocs>

From the Search drop-down list, select **All of the Techdocs Library** and enter TS7700 to perform a search for all related documents. The following websites are useful:

- ▶ Common Information Model (CIM)
<http://www.dmtf.org/standards/cim/>
- ▶ Encryption information
<https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>
- ▶ IBM Business Continuity and Recovery Services
<http://www.ibm.com/services/continuity>
- ▶ TS7700 IBM Knowledge Center
http://www.ibm.com/support/knowledgecenter/STFS69_3.3.0/hydra_c_i_chome.html
- ▶ Web-based Enterprise Management (WBEM)
<http://www.dmtf.org/standards/wbem/>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



IBM TS7700 Release 4.0 Guide

SG24-8366-00

ISBN 0738442135



(1.5" spine)
1.5" <-> 1.998"
789 <-> 1051 pages



SG24-8366-00

ISBN 0738442135

Printed in U.S.A.

Get connected

