

IBM TS7700 Release 5.4 Guide

Larry Coyne

Yuki Asakura

Dave Brettell

Ramón A. Minjares Campos

Nielson 'Nino' de Carvalho

Andrew Enriquez

Rin Fujiwara

Nobuhiko Furuya

Lourie Goodall

Joe Hew

Kousei Kawamura

Stefan Neff

Takeshi Nohta

Tomoaki Ogino

Shinya Ohri

Aderson Pacini

Daniel Riggins

Trinidad Armando Rangel Ruiz

Michael Scott

Taisei Takai

Nao Takemura

Erina Tatsumi

Takahiro Tsuda

Shinsuke Ueyama

Chen Zhu



Storage



IBM Redbooks

IBM TS7700 Release 5.4 Guide

November 2024

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

Fifth Edition (November 2024)

This edition applies to Version 5, Release 4, Modification 0 of IBM TS7700 (product number 3957-AGKU).

© Copyright International Business Machines Corporation 2020, 2024. 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	xix
Trademarksxx
Preface	xxi
Summary of contents	xxii
Authors	xxiv
Now you can become a published author, too!	xxvii
Comments welcome	xxviii
Stay connected to IBM Redbooks	xxviii
Summary of changes	xxix
November 2024, Fifth Edition R5.4	xxix
August 2023, Fourth Edition R5.3	xxix
July 2022, Third Edition R5.2	xxx
May 2021, Second Edition R5.1	xxx
September 2020, First Edition R5.0 (minor update)	xxx
March 2020, First Edition R5.0	xxx
April 2019, Third Edition R4.2	xxxii
August 2018, Second Edition R4.1 through R4.1.2 update	xxxii
May 2018, Second Edition R4.1 through R4.1.2 update	xxxii
April 2017, First Edition minor update	xxxiii

Part 1. Architecture and planning	1
Chapter 1. Introducing the IBM TS7700	3
1.1 Overview	4
1.2 New capabilities	5
1.3 Storage virtualization concepts	6
1.4 Benefits of TS7700 Virtualization	14
Chapter 2. Architecture, components, and functional characteristics	15
2.1 TS7700 architecture	16
2.1.1 Monolithic design of an existing IBM Virtual Tape Server	16
2.1.2 Modular design of the TS7700	17
2.1.3 Previous Peer-to-Peer Virtual Tape Server design	19
2.1.4 Principles of grid design	20
2.1.5 TS7700 Models	20
2.1.6 Introduction of the TS7700T	21
2.1.7 Introduction of the TS7700C	26
2.1.8 TS7700O Advanced Object Store for DS8000	27
2.1.9 Management of the TS7700	27
2.2 Stand-alone cluster: Components, functions, and features	32
2.2.1 Views from the Host: Library IDs	32
2.2.2 Tape Volume Cache	34
2.2.3 Virtual volumes and logical volumes	35
2.2.4 Logical volumes and compression	36
2.2.5 Mounting a scratch virtual volume	37
2.2.6 Mounting a specific virtual volume	38
2.2.7 Logical WORM support and characteristics	39

2.2.8	Virtual drives	39
2.2.9	Selective Device Access Control	40
2.2.10	Physical drives	41
2.2.11	Stacked volumes behind the TS7700T	42
2.2.12	Selective Dual Copy function for TS7700T	43
2.3	General TVC management in a stand-alone cluster	44
2.3.1	Basic Rules for TVC Management	44
2.3.2	Scratched virtual volumes and the Delete Expire function	44
2.3.3	Resident-only TVC cache management	46
2.3.4	Premigration of CPx TVC content	46
2.3.5	Removing or Migrating CPx content	49
2.3.6	Recalling volumes into CPx	50
2.3.7	TVC handling in outage situations	50
2.3.8	Copy Consistency Point: Copy policy modes in a stand-alone cluster	51
2.3.9	TVC selection in a stand-alone cluster	51
2.3.10	TVC encryption	51
2.3.11	TS7700T Physical volume pools	53
2.3.12	Logical and stacked volume management	56
2.3.13	Secure Data Erase function	58
2.3.14	Copy Export function	59
2.3.15	Cloud export function	59
2.3.16	Encryption of physical tapes	59
2.3.17	User Management: Roles and profiles	62
2.3.18	Security identification by using Lightweight Directory Access Protocol	63
2.3.19	Service preparation mode	63
2.3.20	Service mode	63
2.3.21	Control Unit Initiated Reconfiguration	63
2.4	Multi-cluster grid configurations: Components, functions, and features	64
2.4.1	Rules in a multi-cluster grid	64
2.4.2	Required grid hardware	65
2.4.3	Data integrity by volume ownership	67
2.4.4	I/O Tape Volume Cache selection	69
2.4.5	Copy consistency points	70
2.4.6	Cluster family concept	71
2.4.7	Override settings concept	73
2.4.8	Host view of a multi-cluster grid and Library IDs	74
2.4.9	Tape Volume Cache	75
2.4.10	Virtual volumes and logical volumes	76
2.4.11	Mounting a scratch volume through specific clusters	76
2.4.12	Mounting a specific virtual volume	76
2.4.13	Logical WORM support and characteristics	77
2.4.14	Virtual drives	77
2.4.15	Device Allocation and Allocation Assistance	78
2.4.16	Selective Device Access Control	81
2.4.17	Physical drives	81
2.4.18	Stacked volumes	81
2.4.19	Selective Dual Copy function	81
2.4.20	General TVC management in multi-cluster grids	81
2.4.21	Expired virtual volumes and the Delete Expired function	82
2.4.22	TVC management for TS7700T/TS7700C CPx in a multi-cluster grid	83
2.4.23	TVC management for disk-only TS7700 clusters in a multi-cluster grid	84
2.4.24	TVC management processes in a multi-cluster grid	87
2.4.25	Copy Consistency Points: Copy policy modes in a multi-cluster grid	88

2.4.26	TVC (I/O) selection in a multi-cluster grid	95
2.4.27	Remote (cross) cluster mounts	96
2.4.28	TVC encryption	96
2.4.29	Logical and stacked volume management	96
2.4.30	Secure Data Erase	97
2.4.31	Copy Export	97
2.4.32	Encryption of physical tapes	97
2.4.33	Autonomic Ownership Takeover Manager	97
2.4.34	Selective Write Protect for DR testing	98
2.4.35	FlashCopy for DR testing	99
2.4.36	Grid resiliency functions	101
2.4.37	Service preparation mode	103
2.4.38	Service mode	104
2.4.39	Control Unit Initiated Reconfiguration	104
2.5	Grid configuration examples	106
2.5.1	Homogeneous versus hybrid grid configuration	106
2.5.2	Planning for high availability or DR in limited distances	106
2.5.3	DR capabilities in a remote data center	107
2.5.4	Configuration examples	108
Chapter 3. IBM TS7700 usage considerations		115
3.1	Introduction	116
3.1.1	History overview	116
3.1.2	Today's business challenges	116
3.1.3	Challenges of technology progress	117
3.2	Gather your business requirements	118
3.2.1	Requirement types	118
3.2.2	Environment: Source of data	119
3.2.3	Backup data, active data, and archive data	120
3.2.4	IBM Db2 archive log handling	122
3.2.5	DFSMSshm Migration Level 2	123
3.2.6	Object access method: Object processing	124
3.2.7	Batch processing: Active data	124
3.2.8	Data type and cache control	125
3.3	Features and functions for all TS7700 models	125
3.3.1	Four TS7700 models: Disk, tape, object, and cloud	125
3.3.2	Stand-alone versus grid environments	127
3.3.3	Sharing a TS7700	128
3.3.4	Tape Volume Cache selection	130
3.3.5	Copy Consistency policy	130
3.3.6	Synchronous mode copy	131
3.3.7	Override policies	131
3.3.8	Cluster family	132
3.3.9	Logical Volume Delete Expire Processing versus previous implementations	132
3.3.10	Logical Write-Once, Read-Many retention function	133
3.3.11	Software compression	134
3.3.12	Encryption	134
3.3.13	DS8000 Object Store and TS7700 Advanced Object Store for DS8000	135
3.3.14	z/OS Allocation with multiple grids that are connected to a single host	136
3.3.15	z/OS Allocation assistance inside a grid	136
3.3.16	25 GB and 65 GB logical volumes	137
3.3.17	Grid resiliency function: "Remote" fence	138
3.3.18	Control Unit Initiated Reconfiguration	138

3.3.19	CUIR grid resiliency improvements	138
3.4	Features and functions available only for the TS7700T and TS7700C	139
3.5	Operation aspects: Monitoring and alerting	139
3.5.1	Message handling	140
3.5.2	Regularly scheduled performance monitoring	140
3.5.3	Optional checks	140
3.6	Choosing a migration method	141
3.6.1	Host-based migration	141
3.6.2	TS7700 Field-Frame-Replacement, push-pull-MES	142
3.6.3	TS7700 internal data migration	142
3.6.4	Tape drive technology behind a TS7700	143
Chapter 4.	Preinstallation planning and sizing	145
4.1	Hardware installation and infrastructure planning	146
4.1.1	System requirements	146
4.1.2	TS7700 specific limitations	155
4.1.3	TCP/IP configuration considerations	157
4.1.4	Factors that affect performance at a distance	165
4.1.5	Host attachments	166
4.1.6	Planning for LDAP	170
4.1.7	Cluster time coordination	171
4.2	Planning for a grid operation	172
4.2.1	Autonomic Ownership Takeover Manager considerations	172
4.2.2	Defining grid copy mode control	173
4.2.3	Defining scratch mount candidates	174
4.2.4	Retain Copy mode	175
4.2.5	Defining cluster families	175
4.2.6	TS7700 cache thresholds and removal policies	175
4.2.7	Data management settings (TS7700T CPx in a multi-cluster grid)	178
4.2.8	High availability considerations	180
4.2.9	Planning for cloud operation	182
4.3	Planning for software implementation	182
4.3.1	Host configuration definition	182
4.3.2	Software requirements	185
4.3.3	System-managed storage tape environments	185
4.3.4	Sharing and partitioning considerations	186
4.3.5	Library Manager Category Usage Considerations	187
4.3.6	Sharing the TS7700 by multiple hosts	187
4.3.7	Partitioning the TS7700 between multiple hosts	188
4.3.8	Logical path considerations	189
4.3.9	Secure Data Transfer	189
4.3.10	MI dual control	190
4.3.11	Planning for logical and physical volumes	190
4.3.12	Volume serial numbering	190
4.3.13	Virtual volumes	191
4.3.14	Logical WORM	194
4.3.15	Physical volumes for TS7700T	194
4.3.16	Data compression	196
4.3.17	Secure Data Erase function	197
4.3.18	Planning for tape encryption in a TS7700T	198
4.3.19	Planning for cache disk encryption in the TS7700	200
4.4	Tape analysis and sizing the TS7700	202
4.4.1	IBM tape tools	202

4.4.2	BatchMagic	206
4.4.3	Workload considerations	206
4.4.4	Education and training	210
4.4.5	Implementation services	211
Chapter 5.	Disaster recovery	213
5.1	TS7700 DR principles	214
5.1.1	Data availability	214
5.1.2	Deferred Copy Queue	215
5.1.3	Volume ownership	216
5.2	Failover scenarios	218
5.3	Planning for DR	219
5.3.1	DR site connectivity IODF considerations	219
5.3.2	Grid configuration	219
5.3.3	Planning guidelines	221
5.4	High availability and DR configurations	222
5.4.1	Example grid configurations	222
5.4.2	Restoring the host and library environments	230
5.5	DR testing	231
5.6	A real disaster	231
5.7	Geographically Dispersed Parallel Sysplex for z/OS	233
5.7.1	Geographically Dispersed Parallel Sysplex considerations in a TS7700 grid configuration	233
5.7.2	Geographically Dispersed Parallel Sysplex functions for the TS7700	235
5.7.3	Geographically Dispersed Parallel Sysplex implementation	235
Part 2.	Implementation and migration	237
Chapter 6.	Implementing IBM TS7700	239
6.1	Hardware configuration definition	240
6.1.1	Defining devices through HCD	241
6.1.2	Activating the I/O configuration	246
6.2	Setting up the TS7700	247
6.2.1	TS7700 definitions	247
6.3	TS7700 software definitions	248
6.3.1	Defining volume catalogs	248
6.3.2	Defining the TS7700 in a z/OS SMStape environment	249
6.3.3	SYS1.Parmlib changes	251
6.3.4	Final steps	253
6.4	Attaching the TS7700 to a physical tape library or cloud	253
6.4.1	Defining cache partitions	253
6.4.2	TS4500/TS3500 tape library definitions	253
6.4.3	Definitions for TS7700T	255
6.4.4	Cloud Tier Settings (TS7700C only)	255
6.4.5	Object Store settings	256
6.4.6	Defining fence actions	257
Chapter 7.	Hardware configurations and upgrade considerations	261
7.1	TS7700 hardware components	262
7.1.1	Common components for the TS7700 models	262
7.1.2	TS7770 components	267
7.1.3	TS7700 tape library attachments, drives, and media	272
7.1.4	TS3000 Total System Storage Console	273
7.1.5	Cables	273

7.2	TS7700 component upgrades	274
7.2.1	TS7700 concurrent system component upgrades	274
7.2.2	TS7700 nonconcurrent system component upgrades	276
7.2.3	TS7770 cache upgrade options	279
7.2.4	Upgrading drive models in a TS7700T	279
7.2.5	Frame replacement of old hardware with new hardware	288
7.3	TS7700 upgrade to Release 5.4	288
7.3.1	Planning for the upgrade.	288
7.4	Adding clusters to a grid	289
7.4.1	TS7700 grid upgrade concept.	289
7.4.2	Considerations when adding a cluster to the configuration.	290
7.4.3	Considerations for merging an existing cluster or grid into a grid	294
7.5	Removing clusters from a grid	299
7.5.1	Reasons to remove a cluster	300
7.5.2	High-level description of the process	300
Chapter 8.	Migration	303
8.1	Migration to a TS7700	304
8.1.1	Host-based migration	304
8.1.2	Tape-based migration	305
8.2	Migration between TS7700s	305
8.2.1	Field frame replacement migration for TS7700T	309
8.2.2	Join and Copy Refresh processing	310
8.2.3	Migration service offering	312
8.2.4	Copy Export and Copy Export Recovery/Merge	313
8.2.5	Grid-to-Grid Migration tool.	313
8.2.6	Cloud-based migration	317
8.3	Moving data for host-based migration	318
8.3.1	Phased method of moving data	319
8.3.2	Quick method of moving data	320
8.3.3	Products to simplify the task	323
8.3.4	Combining methods to move data into the TS7700	324
8.4	Moving data out of the TS7700	324
8.4.1	Host-based copy tools	325
8.4.2	Copy Export and Copy Export Recovery/Merge	325
8.4.3	Cloud Export and Recovery	325
8.4.4	DFSMSHsm aggregate backup and recovery support	326
8.5	Migrating DFSMSHsm-managed data	327
8.5.1	Volume and data set sizes	328
8.5.2	TS7700 implementation considerations	331
8.5.3	DFSMSHsm task-related considerations.	333
8.6	DFSMSrmm and other tape management systems	336
8.7	IBM Spectrum Protect	338
8.7.1	Native or virtual drives.	340
8.7.2	IBM Tivoli Storage Manager parameter settings.	340
8.8	DFSMSdss	341
8.8.1	Full volume dumps	341
8.8.2	Stand-Alone Services	341
8.9	Object access method.	343
8.10	Database backups	344
8.10.1	Db2 data	344
8.10.2	CICS and IMS	346
8.10.3	Batch data.	347

Part 3. Operations	349
Chapter 9. IBM TS7700 Management Interface operations: Part 1	351
9.1 Overview	352
9.2 User interfaces	354
9.3 Tape library management GUI	355
9.4 TS7700 Management Interface	357
9.4.1 Connecting to the Management Interface	357
9.4.2 Using the TS7700 Management Interface	359
9.5 MI Navigation	364
9.6 Systems icon	367
9.6.1 Grid Summary window	367
9.6.2 Actions menu on the Grid Summary page	368
9.6.3 Cluster Families window	371
9.6.4 Grid Identification properties window	373
9.6.5 Lower removal threshold	373
9.6.6 Grid health and details	374
9.6.7 Cluster Summary window	375
9.6.8 Cluster Actions menu	377
9.6.9 Service mode window	379
9.6.10 Cluster Shutdown window	382
9.6.11 Cluster Identification Properties window	383
9.6.12 Cluster health and detail	384
9.7 Monitor icon	387
9.7.1 Events	388
9.7.2 Performance	389
9.7.3 Tasks window	402
9.8 Virtual icon	403
9.8.1 Cache Partitions	404
9.8.2 Incoming Copy Queue window	408
9.8.3 Recall queue	411
9.8.4 Virtual tape drives	412
9.8.5 Virtual volumes	416
9.8.6 Categories	439
9.9 Object Store icon	443
9.9.1 Object Policy	444
9.9.2 Object Store	446
Chapter 10. IBM TS7700 Management Interface operations: Part 2	449
10.1 Physical icon	449
10.1.1 Physical Volume Pools	450
10.1.2 Physical volumes	459
10.1.3 Physical Tape Drives window	471
10.1.4 Physical Media Inventory window	474
10.2 Constructs icon	476
10.2.1 Storage Groups window	476
10.2.2 Management Classes window	478
10.2.3 Storage Classes window	483
10.2.4 Data Classes window	486
10.3 Access icon	492
10.3.1 Dual Control window	493
10.3.2 Security Settings window	495
10.3.3 Roles and Permissions window	505

10.3.4	SSL Certificates window	507
10.3.5	Customer Documentation Settings window	511
10.4	Settings icon	512
10.4.1	Cloud Tier Settings	513
10.4.2	Library Port Access Groups window	516
10.4.3	Cluster Settings page	518
10.4.4	Copy Export Settings window	540
10.4.5	Notification settings	541
10.5	Service icon	543
10.5.1	Ownership Takeover Mode	544
10.5.2	Repair Virtual Volumes window	547
10.5.3	Network Diagnostics window	548
10.5.4	Data Collection window	550
10.5.5	Copy Export Recovery window	551
10.5.6	Copy Export Recovery Status window	553
10.5.7	Cloud Export Recovery	553
10.5.8	Update System	555
Chapter 11. IBM TS7700 common operations and procedures		557
11.1	Call Home and Electronic Customer Care	558
11.1.1	Electronic Customer Care	559
11.1.2	Assist On-site	560
11.2	Common procedures	561
11.2.1	Tape library with the TS7700T cluster	561
11.2.2	TS7700 definitions	578
11.2.3	TS7700 multi-cluster definitions	598
11.3	Basic operations	608
11.3.1	Clock and time setting	608
11.3.2	Physical Tape Library paused or degraded	610
11.3.3	Preparing a TS7700 for service	611
11.3.4	Tape Library inventory	614
11.3.5	Inventory upload	615
11.4	Cluster intervention scenarios	615
11.4.1	Hardware conditions	615
11.4.2	Ownership takeover interventions	620
Chapter 12. IBM z/OS host console operations		625
12.1	System-managed tape	626
12.1.1	DFSMS operator commands	626
12.1.2	MVS system commands	629
12.1.3	Host Console Request function	632
12.1.4	Library LMPOLICY command	637
12.1.5	Useful DEVSERV commands	638
12.1.6	Scratch volume recovery for volumes	640
12.1.7	Ejecting virtual volumes	643
12.2	Messages from the library	645
12.2.1	CBR3750I console message	645
12.2.2	TS7700 Host Console messages	646
12.3	Expire Hold and scratch processing considerations	648
12.3.1	Expire Hold and low on scratch volumes	648
12.3.2	Expire Hold and cache utilization in a TS7700D	649
12.4	Scratch count mismatch	649
12.5	Host cartridge entry processing	650

12.5.1	Removable Media Manager cartridge entry considerations	651
12.6	Effects of changing volume categories	652
12.7	Library messages and automation	653
12.8	Mount retry	653
12.8.1	Enhanced mount retry defaults	654
12.8.2	Enhanced mount retry example	654
12.9	CUIR for tape	655
12.9.1	LIBRARY REQUEST commands to enable or disable CUIR	655
12.9.2	Other commands built to support CUIR functions	655
12.10	Cloud Storage tier considerations (R4.2 enhancement)	658
12.11	Return-to-scratch enhancement (OA48240)	658
12.12	OAM Object SYSZTIOT enhancement	659
12.13	Enhanced SSMHONOR support	659
12.14	DFSMSHsm RECYCLE Enhancement for TS7700C	660
12.15	DFSMSHsm RECYCLE Considerations when using zEDC	660
12.16	LWORM retention changes	661
12.16.1	Adjusting a volume's expiration date	662
12.16.2	Mirroring retention settings	662
12.16.3	Using the WHILECATALOG function	664
12.16.4	Keeping LWORM volumes permanently	665
Chapter 13.	Monitoring	667
13.1	Overview	668
13.2	Base information: Types of statistical records	670
13.2.1	Point-in-time statistics	670
13.2.2	Historical statistics	670
13.3	Web-based Monitoring method	671
13.3.1	TS7700 Management Interface: Performance page	672
13.3.2	TS7700 Management Interface: Other windows	682
13.3.3	TS4500 Management GUI	683
13.3.4	TS3500 Tape Library Specialist	684
13.4	Bulk Volume Information Retrieval	687
13.4.1	BVIR overview	688
13.4.2	BVIR Prerequisites	690
13.4.3	BVIR Request data format	690
13.4.4	BVIR Response data format	695
13.4.5	BVIR Response data	696
13.5	IBM Tape Tools	703
13.5.1	IBM Tape Tools overview	703
13.5.2	IBM Tape Tools installation	708
13.5.3	VEHSTATS tool overview	710
13.5.4	Running the VEHSTATS jobs	711
13.5.5	VEHSTATS reports	712
13.5.6	Performance evaluation tool for VEHSTATS reports	723
13.5.7	VEHGRXCL tool overview	725
13.5.8	VEHAUDIT tool overview	728
13.5.9	The Other IBM Tape Tools	729
13.6	Host Console Request Commands for monitoring	730
13.6.1	LI REQ,distlib,CACHE2	730
13.6.2	LI REQ,complib,STATUS,GRID	731
13.6.3	LI REQ,distlib,STATUS,GRIDLINK	732
13.6.4	LI REQ,distlib,STATUS,GRLNKACT	732
13.6.5	LI REQ,distlib,PDRIVE	733

13.6.6	LI REQ,{complib distlib},DIAGDATA	734
13.7	IBM z/OS commands for monitoring	735
13.7.1	DISPLAY SMS command	735
13.7.2	LIBRARY command	737
13.7.3	DEVSERV command	737
13.8	Alerts and exception and message handling	738
13.8.1	Alerting of specific events	738
13.8.2	Handling Replication Exceptions.	741
Chapter 14.	Performance considerations	747
14.1	Overview	748
14.2	TS7700 performance history.	750
14.3	Basic performance overview.	752
14.3.1	TS7700 components and task distribution	753
14.3.2	Grid considerations and replication modes.	754
14.3.3	Workload profile from your hosts	756
14.3.4	Lifecycle Management of your data	756
14.3.5	Parameters and customization of the TS7700	756
14.3.6	Throughput terminology	757
14.3.7	Throttling in the TS7700	758
14.3.8	Compression methods	761
14.3.9	Internal SSD/HDDs for TS7700 engine.	762
14.3.10	SSD/HDDs for TVC.	762
14.4	TS7700 throughput: Host I/O increments	762
14.4.1	Host Throughput Feature Codes	763
14.4.2	Tuning for Host I/O	764
14.5	Considerations for Virtual Device Allocation	764
14.6	Cache throughput and cache bandwidth.	765
14.6.1	Tuning Cache bandwidth: Premigration	766
14.6.2	Premigration and premigration throttling values	766
14.6.3	Performance consideration for a cache DDM's rebuild.	768
14.7	Grid link and replication performance	769
14.7.1	Mixing different grid link adapters and traffic from Cloud attach or DS8000 object store considerations	769
14.7.2	Bandwidth and quality of the provided network	770
14.7.3	Selected replication mode.	771
14.7.4	Tuning possibilities for copies: COPYCOUNT Control	772
14.7.5	Tuning to avoid the throttling.	773
14.7.6	Tuning possibilities for copies: Deferred copy throttling	773
14.8	Considerations for the backend TS7700T.	776
14.8.1	Number of back-end drives.	776
14.8.2	Tune back-end drive usage	777
14.8.3	Number of back-end cartridges.	780
14.8.4	Tuning of the usage of Back-end cartridges with VEHSTATS	780
14.9	Cloud Tiering.	780
14.9.1	Network bandwidth and premigration queue size	781
14.9.2	Logical volume size.	781
14.9.3	Premigrate and Recall time out.	781
14.10	TS7700 Advanced Object Store for DS8000	781
14.10.1	Network bandwidth	782
14.10.2	Network latency	782
Chapter 15.	Copy Export	783

15.1 Copy Export overview and considerations	784
15.1.1 Control of Copy Export	784
15.1.2 Workflow of a Copy Export process	784
15.1.3 General considerations for Copy Export	785
15.1.4 Copy Export grid considerations	792
15.1.5 Reclaim process for Copy Export physical volumes	794
15.1.6 Copy Export process messages	796
15.1.7 Copy Export and DFSMSrmm	799
15.2 Implementing and running Copy Export	800
15.2.1 Setting up data management definitions	800
15.2.2 Validating before activating the Copy Export function	801
15.2.3 Running the Copy Export operation	804
15.2.4 Canceling a Copy Export operation	808
15.2.5 Host completion message	808
15.3 Using Copy Export Recovery	810
15.3.1 Planning and considerations for testing Copy Export Recovery	810
15.3.2 Lifecycle Management	812
15.3.3 Performing Copy Export Recovery	813
15.3.4 Restoring the host and library environments	817
15.4 Using Copy Exported tape for damaged volume recovery	818
Chapter 16. Disaster recovery testing in a grid configuration	819
16.1 DR testing overview	820
16.2 DR testing methods	820
16.2.1 Method 1: DR Testing by using FlashCopy	820
16.2.2 Method 2: DR Testing by using Write Protect Mode on DR clusters	822
16.2.3 Method 3: DR testing without the use of Write Protect Mode on DR clusters	823
16.2.4 Method 4: Breaking the grid links	823
16.3 DR testing general considerations	824
16.3.1 Setup and restore of the DR Host tape environment	825
16.3.2 Protecting Production Data	825
16.3.3 Cartridge entry considerations	827
16.3.4 Ownership takeover	828
16.3.5 DR Volume Copy policies	829
16.3.6 Clean up phase of a DR test	829
16.3.7 Tier to Cloud considerations	830
16.4 DR for FlashCopy concepts and command examples	831
16.4.1 Basic requirements and concepts	831
16.4.2 FlashCopy and Write Protect enablement/disablement enhancement R4.1.2 and R4.2	832
16.4.3 DR Family	832
16.4.4 LIVECOPY enablement in a DR Family	834
16.4.5 LIVEACC option	835
16.4.6 Write Protect and FlashCopy enablement/disablement	835
16.4.7 Starting FlashCopy and Write Protect Mode for a DR Family	836
16.4.8 Stopping FlashCopy and Write Protect Mode for a DR Family	836
16.4.9 Commands to check volume status during a DR test	837
16.5 DR testing methods examples	842
16.5.1 Method 1: DR Testing by using FlashCopy	842
16.5.2 Method 2: Using Write Protect Mode on DR clusters	845
16.5.3 Method 3: DR Testing without Write Protect Mode	847
16.5.4 Method 4: Breaking the grid links	849
16.6 Expected failures during a DR test	852

Chapter 17. RESTful API	855
17.1 Overview	856
17.1.1 URL structure	857
17.1.2 Data types	857
17.1.3 Endpoints	857
17.2 Access Token	860
17.2.1 Creating a token by using curl	860
17.2.2 Renewing a token by using curl	861
17.2.3 Formatting the RESTful API response by using the jq tool	861
17.3 Query, Filter, and Sort	862
17.3.1 Query	862
17.3.2 Filter	865
17.3.3 Sort function	867
17.4 RESTful API Pagination	867
17.5 Content type	867
17.6 API Description Document	868
17.6.1 Query for the Description Document	868
17.7 Error Handling	870
Chapter 18. IBM TS7700 support for zTape Air-GAP	871
18.1 Overview	872
18.2 Description	872
18.3 TS7700 Attachment to Tape drives and Library	872
18.4 Methods Supported	873
18.5 User Interface	873
18.6 Reference	873

Part 4. Appendixes

Appendix A. Feature codes and requests for price quotations	877
Feature codes	878
3952/3948 F07 features	878
Server features for 3957/3948-VED	879
Cache Controller Drawer features for 3956/3948-CSB	880
Cache Expansion Drawer features 3956/3948-XSB	880
Cache Controller Drawer features for 3956/3948-CFC and 3956/3948-XFC	881
Requests for price quotations	881
Server VED	881
3952/3948 F07 RPQ	881
Appendix B. IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments	883
Software implementation in z/VM	884
Software requirements	884
TS7700 multi-cluster grid environments disaster recovery	884
z/VM native support that uses DFSMS/VM	884
Implementing Outboard Policy Management for z/VM	886
z/VM guest support	887
Software implementation in z/VSE (Virtual Storage Extended)	890
Software requirements	890
Native z/VSE	890
Implementing Outboard Policy Management for z/VSE hosts	892
Software implementation in z/OS Transaction Processing Facility	894
Software requirements	894

Usage considerations for TS7700 with z/TPF	894
Performance considerations for TS7700 multi-cluster grids with z/TPF	896
Implementing Outboard Policy Management for z/TPF	897
Appendix C. JES3 examples and information	899
JES3 support for system-managed tape	900
Library device groups	900
Updating the JES3 INISH deck	902
Example with two separate tape libraries	904
LDG definitions that are necessary for the first example	905
Device statements that are needed for this configuration	905
SETNAME statements that are needed for this configuration	905
HWSNAME statement that is needed for this configuration	906
Example with three Tape Libraries	906
LDG definitions that are needed for the second configuration example	907
Device statement needed	908
SETNAME statements needed	909
High-watermark setup name statements	909
More examples	909
Processing changes	910
JES3/DFSMS processing	911
Selecting UNITNAMEs	911
New or modified data sets	911
Old data sets	911
DFSMS catalog processing	911
DFSMS VOLREF processing	912
Fetch messages	912
JES3 allocation and mounting	912
Multi-cluster grid considerations	913
Scratch allocation assistance and device allocation assistance	914
Appendix D. DEVSERV QLIB command	919
Appendix E. Sample job control language	923
BVIR jobs to obtain historical data	924
BVIRHSTU	924
BVIRHSTV	927
Extra BVIR reporting	929
Volume Map report	929
Cache Contents report	930
Copy Audit report	931
Volume Status report	931
Physical volume status	932
Physical Volume Status report	933
Physical Volume Pool Status report	935
Physical Volume and Pool Status Report Writer	936
VEHSTATS reports	937
Creating Volume Maps for logical volumes on tape or in object stores	948
Using EDGUTIL to validate tape configuration database inconsistencies	948
REXX EXEC to update the library name	949
Appendix F. Library Manager volume categories	951
Appendix G. IBM TS7700 parameter examples	959

General example setup	960
Example 1: Two-cluster grid for HA and DR.	961
Example 2: Two-cluster grid for HA and DR.	964
Example 3: Three-cluster grid for HA and DR.	967
Example 4: Four-cluster grid for HA and DR.	970
Example 5: Four-cluster grid for HA and DR by using cluster families	973
General example setup for tape partitions	975
Basic considerations how to find the best configuration and setup.	976
Example 1: All data in cache.	976
Example 2: All data on physical tape is premigrated now with one tape partition.	976
Example 3: HSM ML2 is kept in cache only, all other data is premigrated, and tape partitions are used.	976
Example 4: Delay premigration is used to expire data in the cache	977
Appendix H. Extra IODF examples	979
General IODF principles.	980
Using switches to connect to the control unit	980
Directly connecting.	981
Upgrading to 8-Gb channels	981
Adding more devices	981
Sharing ports	989
LIBPORT-IDs in the MVSCP	990
Appendix I. Case study for logical partitioning of a two-cluster grid	991
Overview of partitioning	992
Definitions and settings in z/OS	993
Definitions in HCD.	994
Parmlib definitions.	995
DFSMSrmm definitions	998
JES2 definitions	999
SMS constructs and definitions.	999
RACF definitions	1000
Automation activities	1000
Definitions on the TS7700 Management Interface.	1001
Physical volume pools	1001
Scratch categories	1002
Defining constructs	1002
Library Port Access Groups	1005
Logical volume ranges or insert volumes that are connected to defined ranges	1010
User Management on the Management Interface.	1010
Verification of changes.	1010
Appendix J. Configuring externally managed encryption.	1013
Encrypting physical tape cartridges by using external key management.	1014
Disk storage encryption with external key management	1018
Encryption of Data at Rest (EDaR) for TS7770	1018
Encryption by using EKM servers for TS7770	1019
Enabling External Key Management.	1020
Use of digital certificates on TS7700 to EKM connections.	1023
Updating the TS7700 to trust the attached EKMs.	1023
Updating EKM to trust the TS7700	1028
Uploading a certificate into an IBM Security Guardium Key Lifecycle Manager trust.	1030
Managing IBM Security Guardium Key Lifecycle Manager device groups for TS7700	1033
Adding devices to a device group	1035

Creating SPECTRUM_VIRT Device Group for CSB/CFC (TS7770)	1036
More information about IBM Security Guardium Key Lifecycle Manager management. .	1037
Using Thales CipherTrust Manager	1037
General Configuration for External Encryption Key Server (Thales).	1038
Upload an External CA.	1038
Download KMIP Certificate.	1039
KMIP Interface Details	1040
TS7700 Management Interface Configuration	1044
Installing every Certificate into the TS7700 HTTPS	1044
Final TS7700 Configuration, completing the External Encryption Enablements.	1045
More information about Thales CipherTrust Manager (Thales) Configurations	1045
Abbreviations and acronyms	1047
Related publications	1051
IBM Redbooks publications	1051
Other publications	1052
Online resources	1053
Technical documents on the IBM Support website	1053
Help from IBM	1053

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](http://www.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.

Accesser®	IBM Security®	Redbooks®
AIX®	IBM Services®	Redbooks (logo)  ®
CICS®	IBM Spectrum®	S/390®
Db2®	IBM Z®	Storwize®
DS8000®	IBM z13®	Tivoli®
FICON®	NPS®	WebSphere®
FlashCopy®	Open Liberty®	z/OS®
GDPS®	OS/400®	z/VM®
Guardium®	Parallel Sysplex®	z/VSE®
HyperSwap®	Power8®	z13®
IBM®	Power9®	
IBM Cloud®	RACF®	

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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 covers IBM TS7700 R5.4 and R5.4 PGA1. 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.

Building on over 25 years of experience, the R5.4 release includes many features that enable improved performance, usability, and security. Highlights include the IBM TS7700 Advanced Object Store, an all flash TS7770, grid resiliency enhancements, and Logical WORM retention.

Note: In this IBM Redbooks publication, the following terms are used to denote specific types of TS7700:

- ▶ TS7700D: Disk Only
- ▶ TS7700T: Tape Attach (FC 5273)
- ▶ TS7770C: Cloud Attach
- ▶ TS7770O: Advanced Object Store (FC 5282 now 5283, IBM DS8000® offload of objects to a TS7700 cache partition)

For more information, see 3.3.1, “Four TS7700 models: Disk, tape, object, and cloud” on page 125.

By using the same hierarchical storage techniques as the TS7700T, the TS7700C can also offload to object storage. Because object storage is cloud-based and accessible from different regions, the TS7700 Cloud Storage Tier support essentially allows the cloud to be an extension of the grid. As of this writing, the TS7700C supports the ability to offload to IBM Cloud Object Storage, and Amazon S3.

This publication explains features and concepts that are specific to the IBM TS7700 as of release R5.3. The R5.3 microcode level provides IBM TS7700 Cloud Storage Tier enhancements, IBM DS8000 Object Storage enhancements, Management Interface (MI) dual control security, and other smaller enhancements. The R5.3 microcode level can be installed on the IBM TS7770 only.

The TS7760 with R5.2 and the new TS7780 with R6.0 can be placed in a grid for data migration purposes.

Note: Release 5.2 was split into two Phases:

- ▶ R5.2 Phase 1 (also referred to as R5.2.1 and R5.21)
- ▶ R5.2 Phase 2 (R5.2.2 and R5.22)

For more information, see this [TS7700 Code Update Recommendation](#).

TS7700 provides tape virtualization for the IBM Z® environment. Offloading to physical tape behind a TS7700 is used by hundreds of organizations around the world. Tape virtualization can help satisfy the following requirements in a data processing environment:

- ▶ Improved reliability and resiliency
- ▶ Reduction in:

- Time that is needed for the backup and restore process
- Services downtime that is caused by physical tape drive and library outages
- Cost, time, and complexity by moving primary workloads to virtual tape
- ▶ Increased efficient procedures for managing daily batch, backup, recall, and restore processing
- ▶ On-premises and off-premises object store cloud storage support as an alternative to physical tape for archive and disaster recovery

New and existing capabilities of the TS7700 5.4 release includes the following highlights:

- ▶ Data Class MI page that includes LWORM Retention with Dual Control
- ▶ Full Base Frame SSD Cache (R5.3 GA-available 1H23)
- ▶ RESTful API with CIM Server Replacement
- ▶ New TSSC (7063-CR2)
- ▶ Unjoin cluster with Objects
- ▶ Replace Spectrum Scale mcstore
- ▶ LWORM retention customer configurable by using the Management Interface (MI)
- ▶ Support other EKM: Thales
- ▶ Larger Logical Volumes (65000 MiB)
- ▶ Productize current RPQ 8B3670 (Top Exit Cabling) and RPQ 8B3669 (Key Locks on Door)

The TS7700T writes data by policy to physical tape through attachment to high-capacity, high-performance IBM TS1160, IBM TS1150, and IBM TS1140 tape drives that are installed in an IBM TS4500 or TS3500 tape library.

The TS7770 models are based on high-performance and redundant IBM Power9® technology. They provide improved performance for most IBM Z 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 [IBM TS7700 documentation](#).

A series of technical documents and white papers that describe many aspects of the IBM TS7700 are available. Although the basic information about the product is 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 suitable technical document. For these more technical documents, see the IBM Support website and search for the [topic within TS7700](#).

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 115. These chapters provide a functional description of all 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 145 for hardware information. Information about planning for Software begins in 4.3, “Planning for software implementation” on page 182. Chapter 5, “Disaster recovery” on page 213 describes the use of the TS7700 in disaster recovery (DR). Chapter 6, “Implementing IBM TS7700” on page 239 describes the implementation and installation tasks to set up an IBM TS7700.

If you have an IBM TS7700 or even an IBM 3494 Virtual Tape Server (VTS) installed, see Chapter 7, “Hardware configurations and upgrade considerations” on page 261. Chapter 8, “Migration” on page 303 describes migrating to a TS7700 environment.

Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351, Chapter 10, “IBM TS7700 Management Interface operations: Part 2” on page 449, and Chapter 11, “IBM TS7700 common operations and procedures” on page 557 provide 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 12, “IBM z/OS host console operations” on page 625 provides information about 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 13, “Monitoring” on page 667 and Chapter 14, “Performance considerations” on page 747. Although these chapters provide a good overview, more information is available in the technical documents that are available at [IBM Support topics for the TS7700](#).

For availability and disaster recovery specialists, and individuals who are involved in the planning and operation that is related to availability and disaster recovery, see Chapter 15, “Copy Export” on page 783.

For more information about disaster recovery, see Chapter 5, “Disaster recovery” on page 213 and Chapter 16, “Disaster recovery testing in a grid configuration” on page 819.

New to TS7700 R5.3 Guide fourth edition is the introduction chapter for RESTful API. See Chapter 17, “RESTful API” on page 855 and zTape Air-Gap option (Chapter 18, “IBM TS7700 support for zTape Air-GAP” on page 871).

In addition, the following appendixes conclude this book:

- ▶ For more information about feature codes and requests for price quotation (RPQ), see Appendix A, “Feature codes and requests for price quotations” on page 877, which lists all the available features for the IBM TS7700.
- ▶ For more 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 883. This appendix gives a short overview and scheme for the IBM TS7700 implementation.
- ▶ For more information about job entry subsystem 3 (JES3), an operating system component, see Appendix C, “JES3 examples and information” on page 899. This appendix provides more information to assist you if you are running an IBM z/OS® system with JES3.

- ▶ For more 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 919.
- ▶ For more information about job control language, see Appendix E, “Sample job control language” on page 923, which gives you examples of jobs that are needed for installation and operational tasks.
- ▶ For more information about categories, see Appendix F, “Library Manager volume categories” on page 951, which gives you a full list of all category codes that are used in both the IBM TS7700 and the IBM 3494 VTS.
- ▶ For more information about parameters, see Appendix G, “IBM TS7700 parameter examples” on page 959, which provides parameter examples in different grid configurations.
- ▶ For more information about the input/output definition file (IODF) and the input/output configuration program (IOCP), see Appendix H, “Extra IODF examples” on page 979.
- ▶ For more information about a partitioning case study, see Appendix I, “Case study for logical partitioning of a two-cluster grid” on page 991, which provides a scenario about a partitioned IBM TS7700 hardware configuration.
- ▶ Appendix J, “Configuring externally managed encryption” on page 1013 describes the process for configuring the TS7700 Virtual Engine to apply external key management when using data encryption capabilities.

Authors

This book was produced by a team of specialists from around the world working with the IBM Redbooks, Tucson Center.

Larry Coyne is a Project Leader at the IBM® International Technical Support Organization, Tucson, Arizona, center. He has over 35 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 Project Management from George Washington University. His areas of expertise include client relationship management, quality assurance, development management, and support management for IBM storage management software.

Yuki Asakura is a software development engineer in Japan at IBM Tokyo Laboratory. He has been in charge of developing the IBM TS7700 since he joined IBM Japan in 2018. He is responsible for developing and supporting TS7700 software components that are related to hierarchical storage management for disk, tape and cloud storage, data transfer functions, and grid functions.

Dave Brettell is an Advisory Engineer with IBM Storage Infrastructure, based out of the IBM laboratory in Tucson, Arizona, US. He is a Test Engineer for the IBM TS7700 Virtualization Engine. David has worked in mainframe back-end tape and virtual tape storage for over 20 years. He has held various rolls while under the product umbrella, including, Technical Writer, Hardware Microcode Integration Engineer, and Early Verification Test Engineer.

Ramón A. Minjares Campos is a software back-end developer engineer at IBM, where he plays a role as a team leader in the IBM Mexican Software Laboratory group. He has been in the IBM Storage Systems division for 8 years. Ramon is in charge of developing code and providing development support for IBM TS7700 Systems Storage, specifically focusing on file systems, disk, and cache components for the IBM TS7700.

Nielson “Nino” de Carvalho is a Level 2 certified IT specialist at IBM Switzerland with over a decade of experience in IBM Mainframe computing. He specializes in IBM Z, LinuxONE, z/OS, z/VM®, GDPS®, DS8000, TS7700, and SAN solutions, serving IBM clients across various sectors globally.

Andrew Enriquez is the team lead for DFSMS Object Access Method (OAM) support based in Tucson, Arizona, US. His primary focus is providing level 2 support to clients with any OAM tape and object issues. He holds a degree in Electrical and Computer Engineering from the University of Arizona. Andrew joined the IBM team in 2021.

Rin Fujiwara works in the technical support team for DFSMS and high-end tape products in Japan. She has 15 years experience in the IBM Z area as a Subject Matter Expert. Rin has supported DFSMS and TS7700 for five years.

Nobuhiko Furuya retired from IBM Japan at the end of October 2019, started his own company, V-SOL Inc., and continues to work with IBM and clients on TS7700 and DFSMS implementations. He has over 36 years experience at IBM and 3 years in V-SOL for IBM Z and the storage area.

Lourie Goodall is a Senior Technical Staff Member (STSM) with IBM’s Storage Infrastructure, based out of the IBM laboratory in Tucson, Arizona, US. She is the IBM Enterprise TCT Architect for the IBM TS7700 Virtualization Engine and DS8000 storage products. Lourie has worked in mainframe virtual tape and cloud storage for over 21 years and has authored or co-authored hundreds of design documents and numerous technical papers. Lourie is one of the original architects of the DS8000-TS7700 TCT feature.

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 tests 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).

Kousei Kawamura is a software engineer in Tokyo, Japan. He joined IBM Japan in 2012, and worked on software development of the TS7700 hierarchical storage manager (HSM) component for 4 years. In 2013, he started to work as a TS7700 domain manager (HDM) software developer as well. In 2016, Kousei moved from the HSM development team to work as a software developer of the TS7700 Grid component.

Stefan Neff is an IBM Leading Technical Sales Professional for IBM DRI Solutions, especially Mainframe Tape, within the IBM Germany Systems Technical Sales organization. Stefan has more than 20 years of experience in backup and archive solutions and holds a Master’s degree in Electrical Engineering and Computer Science from the University of Mainz-Bingen, Germany. He also focuses on IBM Tape Encryption solutions, IBM Spectrum® Protect backup solutions and IBM Spectrum Archive EE solutions. He is a Level 2 IBM Master Certified IT Specialist and an IBM High-End-Tape Certified Specialist. Stefan is the chairperson of the German Technical Focus Group “DRI and Tape.” He holds storage virtualization patents and is an IBM 2nd plateau inventor.

Takeshi Nohta is a software development engineer with IBM Storage Infrastructure in Tokyo, Japan. He joined IBM in 1999 and started with DASD (7133/2105) products as a Test Engineer. He then moved to the DASD PFE group for DASD (7133/2105/DS8K) products until 2007. From 2007, his primary responsibility was the TS7700 Grid component, Takeshi's role changed in 2022 and he works on RESTful API and general tests with the Hardware Microcode Integration team.

Tomoaki Ogino is a Level 1 certified IT specialist at IBM Japan. He has 30 years of experience in IBM Z® and worked at IBM for 28 years in the field directly with customers. He is a Subject Matter Expert in the areas of DFSMS and high-end tape products and supporting post-sales and pre-sales for two years.

Shinya Ohri is a Subject Matter Expert for IBM Enterprise Tape TS7700s in Japan. He has five years experience supporting tape subsystems in Asia Pacific. He joined IBM in 2005 as a Service Representative and was responsible for servicing several bank customers. In 2018, he moved to the PFE group supporting Hydra in AP and is the leader of Enterprise tape support.

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 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®.

Daniel Riggins is an Enterprise Tape Product Field Engineer/SME residing near Austin, Texas. He has 2 and a half years in PFE and over 9 years as an IBM SSR in Central Texas, focusing on Storage and Power Systems. Currently he is the technical owner and architect of several internal IBM knowledge transfer systems and tools.

Trinidad Armando Rangel Ruiz is a Computer Engineer from the Universidad Tecnologica de Guadalajara, graduated in 2009, the same year he joined IBM. For IBM, he started as a Storage Technical Consultant with the Guadalajara Development Lab (GDL) Executive Briefing Center. Trinidad is a Top Gun for IBM System x and Enterprise Storage and has worked on the TS7700 since 2016. Currently, he is the TS7700 Field Support Test (FST) team lead.

Michael Scott is a Senior Data Facility Storage Management Subsystem (DFSMS) Technical Support Engineer in the IBM Systems, Client Enablement, and Systems Assurance team. He has 23 years of experience in DFSMS technical support. He holds a Master's degree in Business Administration and a Bachelor of Sciences in Mathematics. Michael has 14 patents that are issued in the computer sciences field. He is also a DFSMS Technical Advocate.

Taisei Takai is a Level 2 Certified Technical Specialist in IBM Japan. His area of expertise is storage based on the z/OS platform. He has more than 25 years of experience in technical support for planning and installing mainframe tape products. He is working as a Client Technical Specialist and is involved with most of the TS7700 related projects in Japan.

Nao Takemura is a member of IBM Systems Lab Services in Tokyo, Japan. He joined IBM Japan in 2007 and worked as a technical specialist of compilers and assembler on z/OS, supporting the customers to upgrade their compilers. He works as a technical specialist of TS7700, and supports the customers with TS7700 planning and configuration, implementation, and migration.

Erina Tatsumi is a software engineer in Japan at the IBM Tokyo Laboratory. She joined IBM Japan in 2020, and worked on software development and support of TS7700 HMI for three years. From 2022, she moved to the IBM FICON® development team as a software developer.

Takahiro Tsuda is a software development engineer in Japan at the IBM Tokyo Laboratory. He has worked as an IBM TS7700 Virtualization Engine microcode developer since he joined IBM in 2007. He is responsible for developing and supporting TS7700 software components and functions that are related to Hierarchical Storage Management between disk and tape storage and data transfer between networks.

Shinsuke Ueyama is a Subject Matter Expert with IBM Storage Systems group, supporting the IBM TS7700. He has 15 years of experience in technical support, worked on the IBM Storwize® and various IBM tape products.

Chen Zhu is a Consulting System Service Representative at the IBM Global Technology Services (GTS) in Shanghai, China. He joined IBM in 1998 to support and maintain IBM Z products for clients throughout China. Chen has been working in the Technical Support Group (TSG) providing second-level support to IBM Z clients since 2005. His areas of expertise include IBM Z hardware, IBM Parallel Sysplex®, IBM Tape Library, and IBM FICON connectivity.

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

Felipe Barajas, Ralph Beeston, Erika Dawson, Joe Hayward, Duke Lee, Khanh Ly, Kohichi Masuda, Jeff Pilch, Kerri Shotwell, Sam Smith, Vanessa Sobik, Joe Swingler

IBM Systems

Enete Gomes Dos Santos Filho, Karim Walji

IBM Technology Support Services

Tom Koudstaal

E-Storage B.V.

Thanks to the authors of the previous editions:

Ole Asmussen, Yuki Asakura, Larry Coyne, Katja Deneleh, Derek Erdmann, Monica Falcone, Rin Fujiwara, Nobuhiko Furuya, Lourie Goodall, Joe Hew, Kousei Kawamura, Tony Makepeace, Kohichi Masuda, Sosuke Matsui, Marcelo Lopes de Moraes, Erich Moraga, Stefan Neff, Takeshi Nohta, Tomoaki Ogino, Alberto Barajas Ortiz, Aderson Pacini, Trinidad Armando Rangel Ruiz, Michael Scott, Joe Swingler, Taisei Takai, Nao Takemura, Takahiro Tsuda, and 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 IBM Redbooks® residency project and help write a book in your area of expertise, while honing your experience by using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

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

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

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

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

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

Summary of changes

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

Summary of Changes
for SG24-8464-04
for IBM TS7700 Release 5.4 Guide
as created or updated on January 23, 2025.

November 2024, Fifth Edition R5.4

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

New and existing capabilities of the TS7700 5.4 include the following highlights:

- ▶ Data Class MI page includes LWORM Retention with Dual Control
- ▶ Full Base Frame SSD Cache (R5.3 GA-available 1H23)
- ▶ RESTful API with CIM Server Replacement
- ▶ New TSSC (7063-CR2)
- ▶ Unjoin cluster with Objects
- ▶ Replace Spectrum Scale mcstore
- ▶ LWORM retention customer configurable through by using the MI
- ▶ Support other EKM: Thales (“Using Thales CipherTrust Manager” on page 1037)
- ▶ Larger Logical Volumes (65000 MiB)
- ▶ Productize current RPQ 8B3670 (Top Exit Cabling) and RPQ 8B3669 (Key Locks on Door)

August 2023, Fourth Edition R5.3

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

New and existing capabilities of the TS7700 5.3 include the following highlights:

- ▶ A single machine type of 3498
- ▶ Introduction to FC 5283 TS7700 Advanced Object Store
- ▶ Storage Expert Care services for TS7700
- ▶ Increased maximum cache capacity of SSD-based TS7700s
- ▶ Limited REST(ful) API support
- ▶ IBM TS7700 zTape Air-Gap function (FC 5995)
- ▶ IBM Lab Services for migration (for VEDs) from FC 5282 to FC 5283 Advanced Object Store for DS8000
- ▶ CRC Check End-to-End of virtual volumes
- ▶ Support for IBM TS1160 Tape drives and JE/JM media (introduced 5.2 PGA1)

New to TS7700 R5.3 Guide fourth edition is Chapter 17, “RESTful API” on page 855 and Chapter 18, “IBM TS7700 support for zTape Air-GAP” on page 871.

Updates August 15:

- ▶ Added update to 2.4.15, “Device Allocation and Allocation Assistance” on page 78
- ▶ Added “Other considerations” on page 292 in 7.4.2, “Considerations when adding a cluster to the configuration” on page 290
- ▶ Swapped TOC position of Chapter 3, “IBM TS7700 usage considerations” on page 115 and Chapter 4, “Preinstallation planning and sizing” on page 145

TS7700 Redpapers updated for R5.3:

- ▶ *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.
- ▶ *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583.

July 2022, Third Edition R5.2

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

New and existing capabilities of the TS7700 5.2.2 include the following highlights:

- ▶ Eight-way Grid Cloud, which consists of up to three generations of TS7700
- ▶ Synchronous and asynchronous replication of both virtual tape and TCT Objects
- ▶ Grid access to all logical volume and object data independent of where it exists
- ▶ An all-flash TS7700 option for improved performance
- ▶ Full Advanced Object Store Grid Cloud support of DS8000 Transparent Cloud Tier
- ▶ Full AES256 encryption for data that is in-flight and at-rest
- ▶ Tight integration with IBM Z and DFSMS policy management
- ▶ DS8000 Object Store AES256 in-flight encryption and compression
- ▶ Regulatory compliance through Logical WORM and LWORM Retention support
- ▶ Cloud Storage Tier support for archive, logical volume version, and disaster recovery
- ▶ Optional integration with physical tape
- ▶ 16 Gb FICON throughput that exceeds 5 GBps per TS7700 cluster
- ▶ Grid Resiliency Support with Control Unit Initiated Reconfiguration (CUIR) support
- ▶ IBM Z hosts view up to 3,968 common devices per TS7700 grid
- ▶ TS7700 Cache On Demand feature that is based on capacity licensing
- ▶ TS7700 support of SSD within the VED server

Note: The latest Release 5.2 was split into two phases:

- ▶ R5.2 Phase 1 (also referred to as *R5.2.1* and *R5.21*)
- ▶ R5.2 Phase 2 (*R5.2.2* and *R5.22*)

For more information, see this IBM Support [web page](#).

May 2021, Second Edition R5.1

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

New and existing capabilities of the TS7700 5.1 include the following highlights:

- ▶ Eight-way Grid Cloud, which consists of up to three generations of TS7700
- ▶ Synchronous and asynchronous replication
- ▶ Full AES256 encryption for replication data that is in-flight and at-rest
- ▶ Tight integration with IBM Z and DFSMS policy management
- ▶ Optional target for DS8000 Transparent Cloud Tier by using DFSMS

- ▶ DS8000 Object Store AES256 in-flight encryption and compression
- ▶ Optional Cloud Storage Tier support for archive and disaster recovery
- ▶ 16 Gb IBM FICON throughput up to 5 GBps per TS7700 cluster
- ▶ IBM Z hosts view up to 3,968 common devices per TS7700 grid
- ▶ Grid access to all data independent of where it exists
- ▶ TS7770 Cache On Demand feature that is based on capacity licensing
- ▶ TS7770 support of SSD within the VED server

September 2020, First Edition R5.0 (minor update)

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized in the following list:

- ▶ Updated Tape Tools link and description about the content of the link in 4.4.1, “IBM tape tools” on page 202.
- ▶ Updated links, content, Table 13-5 on page 705 in 13.5, “IBM Tape Tools” on page 703. Change bars are turned on.

March 2020, First Edition R5.0

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

New and existing capabilities of the TS7700 5.0 include the following highlights:

- ▶ Eight-way Grid Cloud, which consists of up to three generations of TS7700
- ▶ Synchronous and asynchronous replication
- ▶ Full AES256 encryption for replication data that is in-flight and at-rest
- ▶ Tight integration with IBM Z and DFSMS policy management
- ▶ Optional target for DS8000 Transparent Cloud Tier by using DFSMS
- ▶ Optional Cloud Storage Tier support
- ▶ Optional integration with physical tape
- ▶ 16 Gb FICON throughput up to 5 GBps per TS7700 cluster
- ▶ IBM Z hosts view up to 3,968 common devices per TS7700 grid
- ▶ Grid access to all data independent of where it exists
- ▶ TS7770 Cache On Demand feature that is based on capacity licensing
- ▶ TS7770 support of SSD within the VED server

Book organization updates:

- ▶ Operations chapter is divided into three chapters:
 - Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351
 - Chapter 10, “IBM TS7700 Management Interface operations: Part 2” on page 449
 - Chapter 11, “IBM TS7700 common operations and procedures” on page 557
- ▶ The Performance and Monitoring chapter has been divided into two chapters:
 - Chapter 13, “Monitoring” on page 667
 - Chapter 14, “Performance considerations” on page 747

Note: For more information about the TS7770 Cloud Object Storage solution and how to implement and integrate this solution into your enterprise, see *IBM TS7770 R5.3 Cloud Storage Tier Guide*, REDP-5573.

For more information about planning and implementing the function of using the TS7700 as an optional target for DS8000 Transparent Cloud Tier by using DFSMS, see the following publications:

- ▶ *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583
- ▶ *IBM DS8000 Transparent Cloud Tiering: DS8000 Release 9.3*, SG24-8381

April 2019, Third Edition R4.2

Note: *IBM TS7700 Release 5.0 Guide* is now *IBM TS7700 Release 5.2.2 Guide*, SG24-8464. *IBM TS7700 Release 4.2 Guide*, SG24-8366 remains the same.

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

TS7700 R4.2 delivers the following capabilities:

- ▶ TS7760C supports the ability to offload to IBM Cloud® Object Storage and Amazon S3

Note: For more information about the TS7760 Cloud Object Storage solution and how to implement and integrate this solution into your enterprise, see *IBM TS7760 R4.2 Cloud Storage Tier Guide*, REDP-5514.

- ▶ 8-way Grid Cloud consisting of any generation of TS7700
- ▶ Synchronous and asynchronous replication
- ▶ Tight integration with IBM Z and DFSMS policy management
- ▶ Optional Transparent Cloud Tiering
- ▶ Optional integration with physical tape
- ▶ Cumulative 16 Gb FICON throughput up to 4.8GBps * 8
- ▶ IBM Z hosts view up to 496 * 8 equivalent devices
- ▶ Grid access to all data independent of where it exists

August 2018, Second Edition R4.1 through R4.1.2 update

Minor update to 7.2.1, "TS7700 concurrent system component upgrades" on page 274. Added the following sentence:

The 3957-VEC with FC 3402/3403 (16 Gbps FICON) allows a max of 40 performance increments.

May 2018, Second Edition R4.1 through R4.1.2 update

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

TS7700 R4.1 and R4.1.2 delivers the following capabilities:

- ▶ 7 and 8 way Grid support through approved request for price quotation
- ▶ 16 Gb FICON adapter support for TS7760 (R4.1.2)
- ▶ Optimized host data compression that is based on software (not FICON adapter hardware) compression algorithm (R4.1.2)
- ▶ Control Unit Initiated Reconfiguration (CUIR) for code load improvement (R4.1.2)
- ▶ Grid Resiliency Improvements (R4.1.2)
- ▶ System Events Redesign (R4.1.2)
- ▶ Remote System Log Processing Support in (R4.1.2)

April 2017, First Edition minor update

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized in the following list:

- ▶ Clarified that the IBM Security® Key Lifecycle Manager for z/OS external key manager supports TS7700 physical tape, but does not support TS7700 disk encryption.
- ▶ IBM Security Key Lifecycle Manager (formerly IBM Tivoli® Key Lifecycle Manager) supports TS7700 tape and disk encryption.



Part 1

Architecture and planning

This part introduces the IBM TS7700 R5.4 (IBM TS7700) family. The family consists of the IBM TS7770 disk-only virtualization solution, the IBM TS7770T (Tape Attached with TS4500 or TS3500 support), and the IBM TS7770C (Cloud Attached).

This part also 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 TS7770, TS7770C, TS7770T, TS7760, TS7760C, and TS7760T.

This part includes the following chapters:

- ▶ Chapter 1, “Introducing the IBM TS7700” on page 3
- ▶ Chapter 2, “Architecture, components, and functional characteristics” on page 15
- ▶ Chapter 3, “IBM TS7700 usage considerations” on page 115
- ▶ Chapter 4, “Preinstallation planning and sizing” on page 145
- ▶ Chapter 5, “Disaster recovery” on page 213



Introducing the IBM TS7700

The IBM TS7700 is now in its seventh generation, enabling over 25 years of IBM Z virtual tape support. The IBM TS7700 also supports object store concepts when integrated with a DS8000. The IBM TS7770 for IBM Z is the latest model of the TS7700 family.

This chapter provides an introduction to the IBM TS7700 for IBM Z.

This chapter includes the following topics:

- ▶ 1.1, “Overview” on page 4
- ▶ 1.2, “New capabilities” on page 5
- ▶ 1.3, “Storage virtualization concepts” on page 6
- ▶ 1.4, “Benefits of TS7700 Virtualization” on page 14

1.1 Overview

With exponential storage growth, IBM Z organizations need a cost-effective way to manage 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. The storage solution must also fit within today's budget constraints.

Storing infrequently accessed data on costly disk storage does not make sense. At the same time, physical tape libraries or cloud object stores can introduce longer access times, which make the use unacceptable in transactional storage infrastructures. It is in this situation where virtualized tape and object storage comes into play.

This publication explains features and concepts that are specific to the IBM TS7700 product family. IBM TS7770 Release 5.4 includes and introduces:

- ▶ LWORM Retention, which includes a redesign of the data class GUI panel
- ▶ Increases the maximum LVOL capacity to 65 GB
- ▶ Extended REST(ful) API support by adding new queries
- ▶ Support of other EKM: Thales - for TS7770D "disk-only"
- ▶ Productized F07-Frame RPQs to Feature Codes (for door locks and Top-Exit cabling)
- ▶ Support for "Unjoin" cluster with objects
- ▶ Increased maximum cache capacity of SSD-based TS7700s (full populated base frame), introduced with R5.3 PGA1
- ▶ New TSSC/TS3000 Service Console as part of the 3948-VED ship-group, introduced with R5.3PGA1
- ▶ TLS version 1.3 Support for RSYSLOG (starting with R5.4 pga1)
- ▶ Dual Control for adding users (starting with R5.4 pga1)

R5.4 is supported on VED TS7770 only and requires same as R5.3 before 128 GB of VED Server Memory (FC3479).

Note: For more information about the TS7700 Cloud Object Storage solution and how to implement and integrate this solution into your environment, see *IBM TS7700 R5.4 Cloud Storage Tier Guide*, [REDP-5573](#).

The last two generations of the IBM TS7700 consist of the following models:

- ▶ Current VED model configurations:
 - TS7770 (universal reference to any TS7770 configuration)
 - TS7770T (tape-attached, uses a TS3500 or TS4500 library and TS1100 drives)
 - TS7770C (cloud-attached, uses IBM TS7700 Cloud Storage Tier)
 - TS7770D (disk only, which is upgradeable to TS7770T or TS7770C)
 - TS7700O (object offload, DS8000 offload of objects to a TS7700 cache partition)
- ▶ Previous VEC model configurations:
 - TS7760 (universal reference to any TS7760 configuration)
 - TS7760T (tape-attached, uses a TS3500 or TS4500 library and TS1100 drives)
 - TS7760C (cloud-attached, use IBM TS7700 Cloud Storage Tier)
 - TS7760D (disk only, which is upgradeable to TS7760T or TS7760C)

The IBM TS7770 is the seventh generation of IBM virtual tape system technology. The IBM TS7770 delivers next-level cybersecurity and maintains the reliability and availability of the mainframe enterprise storage platform to support business across hybrid multicloud deployments. TS7770 optimizes data protection and business continuance for logical volumes and object store data for IBM Z data and improves storage economics, system resiliency, and data stability for mission-critical workloads. It also includes support for object data and acts as a fully resilient, enterprise object store repository for IBM Z unstructured object store data as a direct target for DS8000 Transparent Cloud Tier (TCT) offloaded objects.

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

Note: For more information about the DS8000 Transparent Cloud Tiering (TCT) solution, an introduction to the TS7700 Object Store, and TS7700 Advanced Object Store solutions, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).

These capabilities can improve performance and capacity, which reduces the total cost of ownership and simplifies overall IBM Z processing. A TS7700 can improve IBM Z operations by reducing tier one primary storage through emulated tape on disk and optional automatic tiering to cost-effective tape or cloud storage. In addition, the TS7700 supports the ability to accept and manage objects received directly from an attached DS8000 as part of the DFSMS DS8000 TCT support.

With regard to IBM Z virtual tape, the TS7770 tape virtualization can help satisfy the following requirements in a data processing environment:

- ▶ Improved reliability and multi-regional resiliency
- ▶ Reduction in:
 - Time that is needed for critical batch windows
 - Services downtime that is caused by physical tape drive and library outages
 - Cost, time, and complexity by moving primary workloads to virtual tape
- ▶ Greater efficiency managing daily batch, backup, recall, and restore processing
- ▶ Utilization of on-premises and off-premises object store cloud storage support as an alternative to physical tape for archive and disaster recovery

For DS8000 TCT objects, the TS7700 Advanced Object Store can help satisfy the following requirements in a data processing environment:

- ▶ Improved reliability and multi-regional resiliency with full grid support
- ▶ Reduce IBM Z CPU and FICON SAN fabric use by having DS8000 stored content move directly to a TS7700 through IP versus the traditional FICON attached tape path
- ▶ Eliminate the need of traditional migration level 2 (ML2) recycle processing
- ▶ Reduced time that is needed for critical batch windows

1.2 New capabilities

Building on over 25 years of experience, the R5.3 release included many features that enable improved performance, usability, and security. Highlights included increased maximum cache

capacity of SSD-based TS7700s, zTape Air-Gap option (Chapter 18, “IBM TS7700 support for zTape Air-GAP” on page 871), CRC Check End-to-End of virtual volumes, and TS7700 RESTful API (Chapter 17, “RESTful API” on page 855). The R5.4 release follows on from this. The R5.4 release includes again increased maximum cache capacity of SSD-based TS7700s, increased maximum LVOL size of 65 GB, adds LWORM Retention to the data class GUI panel, and Dual Control functionality to augment security for the business.

New and existing capabilities of the TS7700 5.4 release includes the following highlights:

- ▶ Support for IBM TS1160 Tape Drives and JE/JM media
- ▶ Eight-way Grid Cloud, which consists of up to three generations of TS7700
- ▶ Synchronous and asynchronous replication of virtual tape and TCT objects
- ▶ Grid access to all logical volume and object data independent of where it resides
- ▶ An all flash TS7770 option for improved performance
- ▶ Full Advanced Object Store Grid Cloud support of DS8000 Transparent Cloud Tier
- ▶ Full AES256 encryption for data that is in-flight and at-rest
- ▶ Tight integration with IBM Z and DFSMS policy management
- ▶ DS8000 Object Store with AES256 in-flight encryption and compression
- ▶ Regulatory compliance through Logical WORM and LWORM Retention support
- ▶ Cloud Storage Tier support for archive, logical volume versioning, and disaster recovery
- ▶ Optional integration with physical tape
- ▶ 16 Gb FICON throughput that exceeds 4 Gbps per TS7700 cluster
- ▶ Grid Resiliency Support with Control Unit Initiated Reconfiguration (CUIR) support
- ▶ IBM Z hosts view up to 3,968 3490 devices per TS7700 grid
- ▶ TS7770 Cache On Demand feature that capacity-based capacity based licensing
- ▶ TS7770 support of SSD within the VED server

Note: For more information about planning and implementing the TS7700 as a target for DS8000 Transparent Cloud Tier by using DFSMS, see *IBM DS8000 Transparent Cloud Tiering: DS8000 Release 9.3*, [SG24-8381](#).

1.3 Storage virtualization concepts

A virtual tape subsystem presents emulated tape drives to the host, and stores tape data on emulated tape volumes. These volumes are stored 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. Dynamic Disk Pools within the TS7770 use Distributed RAID (DRAID) to safely store volumes that are written by the host. The disk space that is provided is called a *Tape Volume Cache (TVC)*.

Any DS8900 present can also use TCT support to allow DFSMS to request the DS8900 offload content directly to the TS7770 through IP connectivity. This content is stored as objects within the same TVC.

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

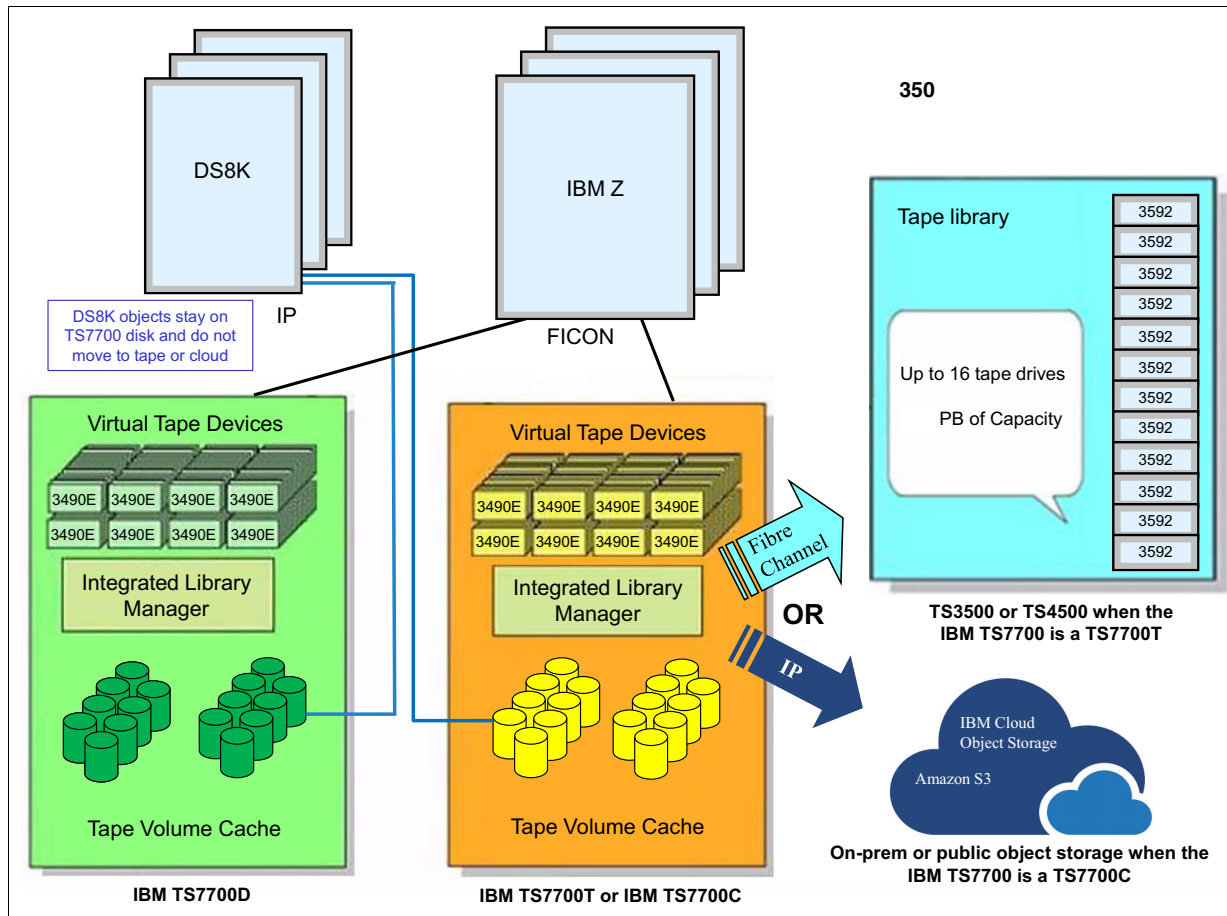


Figure 1-1 Main components of the TS7700

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*) that are optionally attached in TS7700T configurations.

The IBM Z host does not know whether there are physical tape drives or attached cloud object stores. For TS7700D (disk-only) configurations, there are no physical tape or cloud attachments.

An application that supports only 3490E tape technology can use the TS7700 without any changes. Therefore, applications see improved performance while also benefiting from business continuance concepts and the storage capacities of physical tape and cloud.

The IBM TS7700C can also emulate an object store for DS8000 Object Store or Transparent Cloud Tier (TCT) support. That is, the TS7700 can be a direct target of DS8000 TCT offload. The IBM TS7700C can emulate both 3490E drives and an object store at the same time, though data created by one method cannot be directly accessed through the other method.

For IBM Z tape workloads, all virtual tape data must be written to or read from emulated volumes in the disk-based TVC over the FICON channel interface. These emulated tape volumes in the TVC are called *virtual volumes* or *logical volumes* interchangeably.

With the Advanced Object Store support, IBM Z content can be written in object form into the same disk-based TVC over 1 Gb or 10 Gb Ethernet (highly recommend 10 Gb) connections. The object data can then be recalled through the DS8000 object interface.

IBM Z virtual tape content and object data can coexist within a TS7700, but each type must always be accessed by using the same method in which it was created.

As of R5.1, the in-flight data can be encrypted by using AES256 encryption and previously compressed before leaving the DS8000. For more information regarding TS7700 Advanced Object Store for DS8000 (FC 5283) and DS8000 Object Store (FC 5282), see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).

When IBM Z requests a private virtual tape volume for read or append (DISP=MOD), the TS7700 targets a copy in the TVC. If the only available copy exists on back-end physical tape or an attached object store, it is automatically brought back into the TVC by a process called recall and only then can the host mount complete. Therefore, the host does not require knowledge of the location of the logical volume. All scratch mounts are always satisfied directly from the TVC as no previous volume content is accessed for scratch mounts.

Although you can define maximum sizes for your volumes, a virtual volume takes up only the space in the cache that the compressed data on the volume requires. For this reason, tape virtualization makes efficient use of disk capacity. In IBM TS7700T (tape-attached) configurations, the virtual volumes are copied from disk to tape. They also need only the amount of tape physical capacity that is occupied by the logical volumes that are stacked end-to-end, which makes efficient use of disk and tape capacity.

This is the same for cloud attached IBM TS7700C (cloud storage tier) configurations.

Another benefit of tape virtualization is the large number of drives that are 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 this issue by providing many virtual tape drives. For tape-attached configurations, the TS7700T manages all physical tape drives and physical volumes in the tape library partition that is assigned to the TS7700T. It also controls the movement of logical volume data between TVC and physical tape.

In the TS7700T and TS7700C, data that is written from the host into the TVC is optionally scheduled for copying to tape or cloud later. The process of first copying data to tape or cloud is called *premigration* and the volume exists on both disk and tape or cloud. Depending on policies, the copy in the TVC can later be removed from disk cache. After it is removed, the logical volume is viewed as *migrated*, which means it exists on only physical tape or cloud behind that particular TS7700 cluster.

For a TS7700T, 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 stacking concept does not apply to any non-tape attached TS7700 because no physical tape devices are attached to it.

Many applications cannot fill the high-capacity media of modern physical tape technology. This issue can lead to under-used capacity and excessive tape cartridge counts and slots. Virtualization and stacking resolve this problem without changes to your applications or job control language (JCL).

The TS7700T and TS7700C support TVC partitioning that allows the cache to be divided into smaller partitions. Through policy management, IBM Z workloads can then target a specific partition allowing the disk cache footprint to be unique to that set of workloads.

In addition, DS8900 TCT Objects can target a specific partition that is dedicated to TS7700 Advanced Object Store for DS8000 objects. As of R5.4, DS8900 TCT Objects targeting the TS7700 do not support tiering to physical tape or cloud.

When space is required in a particular logical tape partition, volumes that were premigrated to tape or cloud for that partition can be removed from the TVC. By default, removal is based on a *least recently used* (LRU) algorithm, which allows the most accessed content to stay disk cache resident.

For a TS7700D or for data in the resident-only partition of a TS7700T or TS7700C, movement to physical tape or cloud is not supported. Only logical tape volumes that target a TS7700T or TS7700C migration-supported partition can move to physical tape or cloud. DFSMS policy management can be used to determine which partition a workload targets. Command line requests can be used to move logical tape data between partitions.

When a TS7700 is a member of a multi-cluster grid, virtual volumes in a TS7700 cache can be automatically removed if a peer contains a copy of the same logical volume. This function is referred to as the *Automatic Removal policy* and is most applicable to configurations in which a large capacity difference exists between peers. For example, if one peer is disk only (TS7700D) and another peer contains back-end tape (TS7700T). The total capacity of the tape-attached TS7700T easily exceeds the TS7700D. Automatic removal then allows the smaller capacity TS7700D to contain only the most important data.

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. When two or more different cluster types exist in the same grid, such as a TS7700D and TS7700T, it is referred to as a *Hybrid Grid*.

Figure 1-2 shows IBM TS7700T and IBM TS7700C TVC processing.

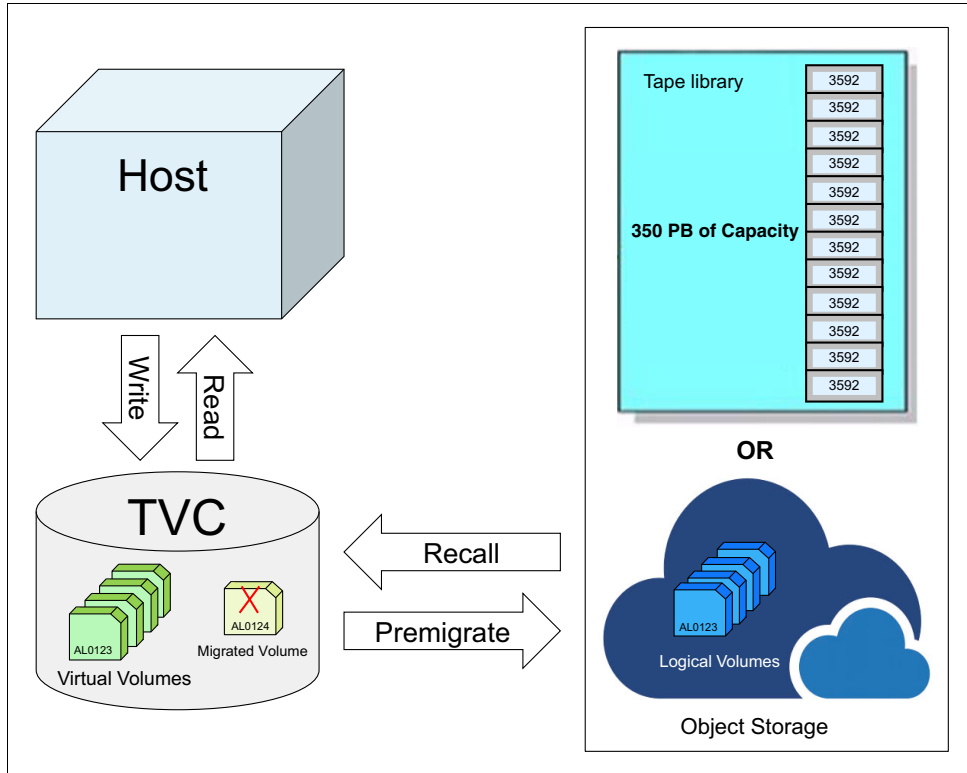


Figure 1-2 IBM TS7700T and IBM TS7700C Tape Volume Cache processing

With a TS7700D, all logical tape volumes are always contained within the TVC. With a TS7700 supporting DS8000 Advanced Object Storage, all objects that are contained within the object partition always remain in the TVC.

Figure 1-3 shows the IBM TS7700D and TS7700 with DS8000 Object Storage TVC processing.




Figure 1-3 TS7700D and TS7700 with DS8000 Object Storage Tape Volume Cache processing

Data replication functions are another benefit of tape virtualization. A maximum of eight IBM TS7700 devices can be interconnected through TCP/IP to form a grid. The connections can be through different combinations of the following sets of links:

- ▶ Two or four 1 Gigabit per second (Gbps) Ethernet links (copper)
- ▶ Two or four 10 Gbps Ethernet links (longwave (LW) fiber)

These sets of links form a *multi-cluster grid configuration*. Link types cannot be mixed in a single cluster, but they *can* vary within a grid, which requires switches that can support 1 Gb and 10 Gb links. The number of links per cluster can be different. Each TS7700 has a one-for-one relationship with the IP ports in other clusters. For example, if one cluster has four links and the other clusters have only two, then the IP traffic uses two grid links.

Logical volume attributes, object attributes, and data are replicated across the clusters in a grid. Any data that resides in the grid is accessible by any other cluster in the grid configuration. Every grid cluster needs to be able to see all the other clusters of the grid via the Internet Protocol network.

Through remote volume access, you can reach any virtual volume through any virtual device. Similarly, the TS7700 Advanced Object Store for DS8000 objects can also be accessed remotely through any DS8K TCT entry point into the grid.

Setting replication policies on the TS7700 determines how many copies are made, where copies are kept, and how copies occur. For logical tape volumes, each policy can be set at logical volume granularity, through DFSMS Management Class. For objects, each policy can be unique per configured DFSMS Cloudname. The logical volume and object cloud-name granular approach allows flexibility for workloads. For example, some workloads can use synchronous replication while others can use asynchronous replication. Also, some workloads might require no copies while others require many copies.

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, TS7700 devices in a grid provide the following benefits:

- ▶ High availability (HA)
- ▶ Simplified disaster recovery (DR)
- ▶ Metro and global business continuance
- ▶ Multiple data center solution sharing
- ▶ Data center and workload migration

Each IBM TS7700 in a grid can support up to 496 devices. An eight-way grid can support up to 3,968 total devices that appear to IBM Z as one global composite tape library with access to all logical volumes contained within it. Device ranges can also be partitioned, which allows different system plexes or tenants to access dedicated devices and volumes.

An Advanced Object Store enabled grid can support up to 256 configured cloud names, which allows different workloads, plexes, and tenants data to be managed differently within a grid.

The synchronous and asynchronous replication of logical volumes and objects is handled entirely by the TS7700s within a grid. Replication occurs at logical volume and object granularity and has no dependency on disk-based replication techniques. You configure policies that inform the TS7700 how many copies you want, where you want copies, and how those copies are to be made.

Throughout this document, you might see the following copy policies denoted by their abbreviations:

- ▶ Sync (S)
- ▶ RUN (R)
- ▶ Deferred (D)
- ▶ Time Delayed (T)
- ▶ No Copy (N)

For more information, see 2.4.5, “Copy consistency points” on page 70.

Figure 1-4 shows multiple IBM TS7700 devices that are configured with multiple host configurations.

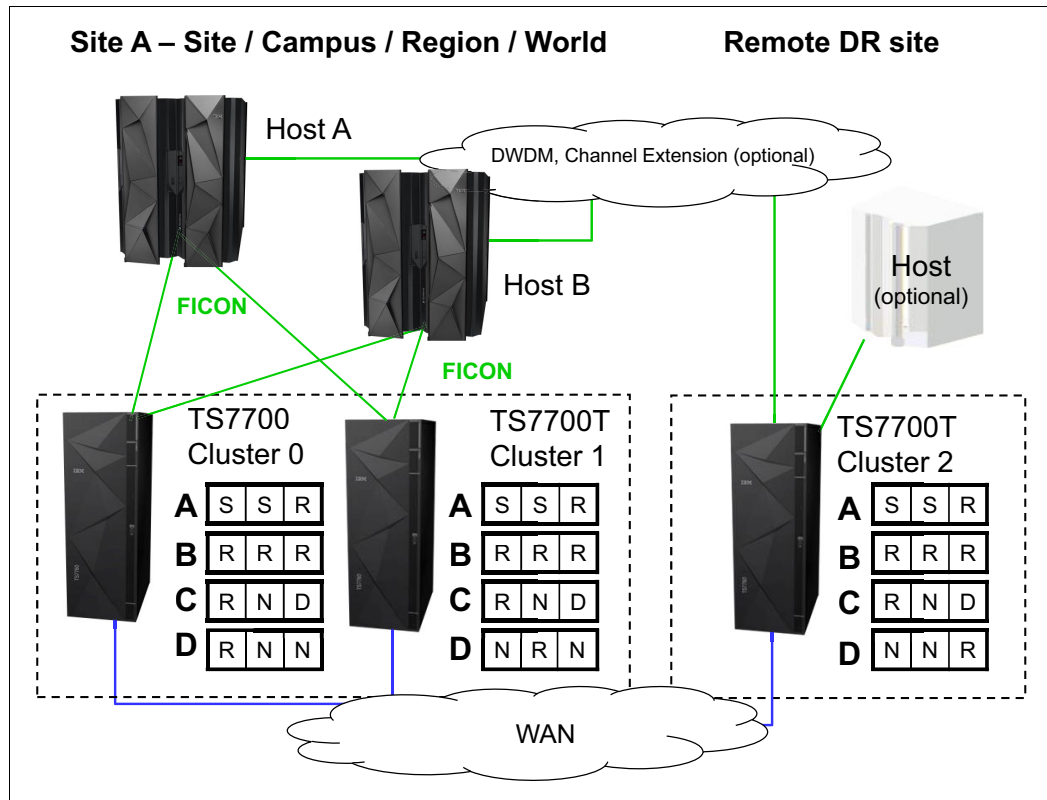


Figure 1-4 Multiple TS7700 tape products that depict possible host and grid connections

For grid configurations in which one or more TS7700T clusters exist, each TS7700T manages its own set of physical volumes. Each TS7700T maintains the relationship between logical volumes and the physical volumes on which they are located.

For grid configurations in which one or more TS7700C clusters exist, all clusters are synchronized as to what and where volumes exist in the cloud. When one TS7700C completes a premigration to a cloud, all TS7700C clusters in the same grid have immediate access to that data in the cloud. For more information, see *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.

The clusters in a TS7700 grid might be geographically dispersed. They can exist next to each other or across oceans. Many clients configure a grid with metro and global business continuance in mind; for example, two or more TS7700 clusters that are within 100 km (62 miles) of each other, whereas the remaining clusters can be located more than 1000 km (621.37 miles) away.

Cloud storage tier attachment can further increase the redundancy of your configuration by putting logical volume copies within an object store that can also be redundant across global regions of the world. This cloud level of redundancy can be used for full or partial disaster recovery, additional copy redundancy, and multiple point-in-time recovery backups for Safeguarded recovery.

A multi-cluster grid supports the concurrent growth and reduction of cluster counts, which allows new locations to be easily introduced or removed.

1.4 Benefits of TS7700 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 a wide portfolio, including Storage Systems, software, and services that are built on decades of innovation. These offerings are designed to provide industry-leading solutions that meet the most demanding business requirements. IBM's offerings continually evolve to answer the always changing challenges that are associated with rapid data growth, applications, dynamic workloads, and regulations. IBM can give your business a competitive advantage.

Information services must provide maximum performance and be resilient, secure, and simple to manage. In addition, they must provide these benefits while also meeting business continuance requirements, which usually mean spanning two or more locations.

The IBM TS7700 is a critical building block for IBM Z operations that is designed exclusively for IBM Z users. With over 25 years in the virtual tape market and the recent joining of the object store market, IBM continues to lead with innovation for IBM Z tape processing.



Architecture, components, and functional characteristics

This chapter describes the architecture of the IBM TS7700. The description includes general virtualization concepts, new concepts, and functions that were introduced with the TS7700 product family up to the current release R5.4. In addition, configuration examples are addressed.

The chapter begins with a description of features that apply to stand-alone and multi-cluster configurations. features that apply to multi-cluster grid configurations only.

This chapter includes the following topics:

- ▶ 2.1, “TS7700 architecture” on page 16
- ▶ 2.2, “Stand-alone cluster: Components, functions, and features” on page 32
- ▶ 2.3, “General TVC management in a stand-alone cluster” on page 44
- ▶ 2.4, “Multi-cluster grid configurations: Components, functions, and features” on page 64
- ▶ 2.5, “Grid configuration examples” on page 106

2.1 TS7700 architecture

The architectural design of the IBM TS7700 and many of its capabilities are described in this section. 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 the following different models:

- ▶ IBM TS7770, including the IBM TS7770D, IBM TS7770T, IBM TS7700O, and IBM TS7770C
- ▶ IBM TS7760, including the IBM TS7760D, IBM TS7760T, and IBM TS7760C

Although some differences exist between these models, the underlying architecture is the same. Any function or feature that is unique or behaves differently for a specific model is clearly stated. If this distinction is not made, assume that it is common across all models.

Note: Consider the following points:

- ▶ When the TS7700 is referenced, it implies all models and types, including the TS7770, TS7760 and all D, T and C submittals.
- ▶ When the function is applicable to models that are disk-only, TS7700D is used.
- ▶ If the functions are applicable to all tape-attached models only, TS7700T is used.
- ▶ If the functions are applicable to all cloud storage tier-supported TS7700s, TS7700C is used.
- ▶ If the functions are applicable to all Tape Object Offload models only, TS7700O is used.
- ▶ If the function is applicable to only a specific version of the TS7700, such as the TS7770, TS7770D, TS7770T, TS7770C, TS7760, TS7760D, TS7760T, or TS7760C, the product name or names are referenced.

2.1.1 Monolithic design of an existing IBM Virtual Tape Server

IBM virtual tape has existed since 1997. The first three generations of the IBM 3494 VTS (B16, B18, and B10/20) performed all functions within a single IBM System p server, including handling all RAID functions.

Peer-to-peer (PTP) functions were created for the B18 and B10/20 models by using more hardware components. The complex design reached an architectural limit that made it difficult to add enhancements, especially around three or more site-based business continuances. 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, which was released in 2006.

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 feature 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 feature 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), which runs on a single System p server.

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

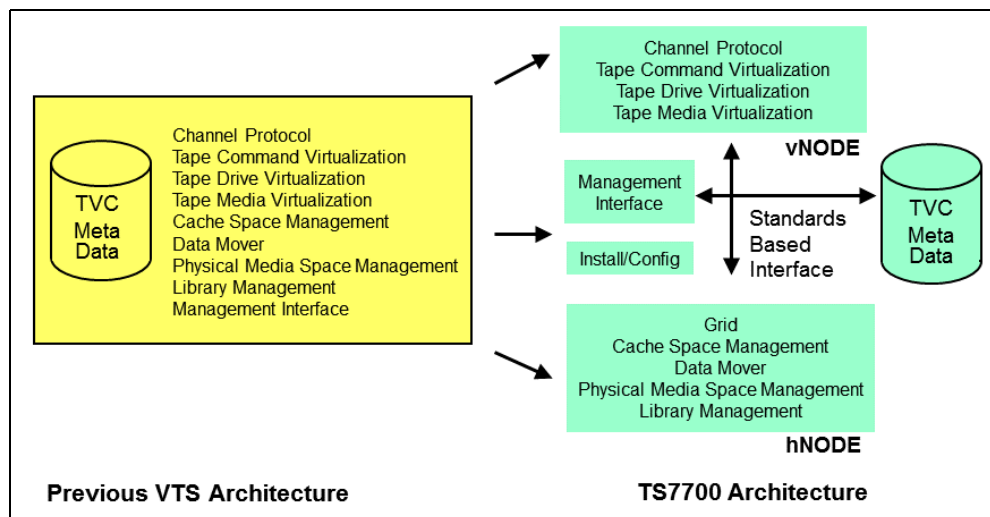


Figure 2-1 TS7700 virtualization design that is compared to a VTS design

Nodes

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

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

Virtualization node

A vnode is a code stack that presents the virtual image of a library and virtual tape drives to an IBM Z 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 disk cache subsystem to represent the virtual tape image. After the logical volume is created or altered by the host system through a vnode, it exists initially in the disk cache only.

Hierarchical data storage management node

A *hierarchical data storage management node (hnode)* is a code stack that manages all logical volumes that are in disk cache, physical tape, or cloud object store. 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.

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 features a vnode and hnode that are combined in an IBM Power8® (TS7760) or IBM Power9 (TS7770) processor-based server.

Figure 2-2 shows a relationship between nodes.

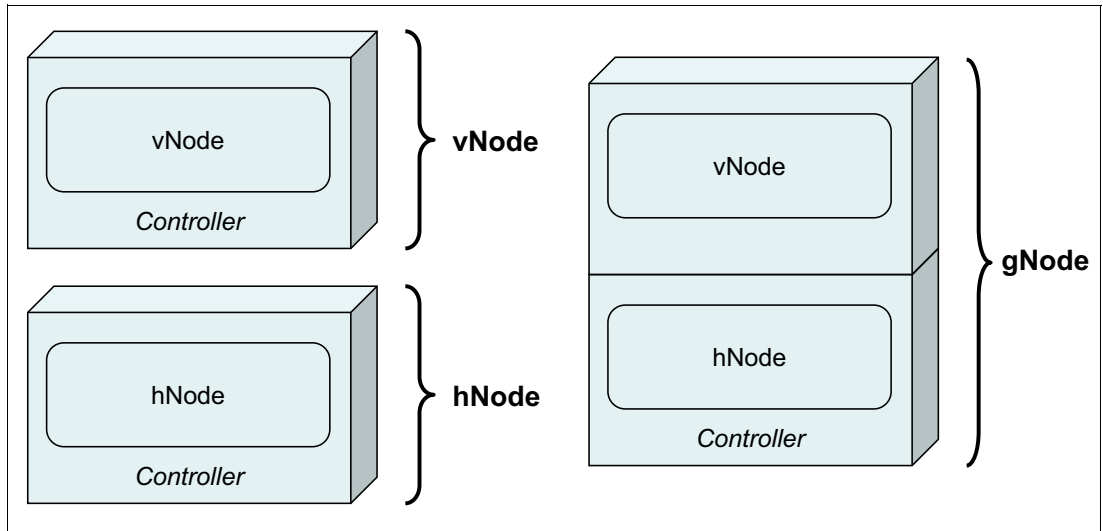


Figure 2-2 Node architecture

Cluster

The TS7700 cluster combines a vnode, hnode (or gnode), and an external disk cache system into one operational server.

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




Figure 2-3 TS7700 cluster

A TS7700 cluster provides a 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 TS7700T cluster attaches to a defined TS3500 or TS4500 Tape Library partition and physical tape drives by using fiber switches.

A single TS7700 frame contains one disk cache controller and several expansion drawers. One or more expansion frames can be added that enables adding disk cache controllers and expansion drawers.

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

One or more controllers and their expansion drawers are collectively referred to as the TS7700 Tape Volume Cache (TVC). The amount of cache available per TS7700 TVC depends on your physical configuration and installed feature codes.

The TS7760 and TS7770 use dynamic RAID6 technology that accelerates rebuild times.

2.1.3 Previous Peer-to-Peer Virtual Tape Server design

In the IBM 3494 PtP VTS, you needed external 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 concept with an industry-leading new technology that is referred to as a *grid*.

2.1.4 Principles of grid design

The TS7700 R5.4 *grid configuration* is a series of two or more TS7700 clusters, with a maximum of eight clusters in a single grid. These clusters are all connected to each other through Ethernet. This network is referred to as the *grid network*, which enables high availability and disaster recovery solutions.

Note: Seven- and eight-cluster grid configurations are available with an individual request for price quotation (RPQ).

A grid configuration and all virtual tape drives that are emulated in all configured clusters appear as one large library to the attached IBM Z 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 within any cluster to be accessed remotely, which enables ease of access to content anywhere in the grid.

In general, any data that is 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 model types, including the following types:

- ▶ TS7770C
- ▶ TS7770D
- ▶ TS7770T
- ▶ TS7700O
- ▶ TS7760C
- ▶ TS7760D
- ▶ TS7760T

When a mixture of models is present within the same grid, it is referred to as a mixed configuration. When one or more models support back-end tape or cloud and others do not, it is referred to as a *hybrid grid*.

The term *multi-cluster grid* is used for a grid with two or more clusters. For more information, see 2.4, "Multi-cluster grid configurations: Components, functions, and features" on page 64.

2.1.5 TS7700 Models

R5.2 introduced a flash-only or solid-state version of the TS7770. It is referred to as the *TS7770 Performance Model*. One or two drawer configurations were initially supported, which provided 60 TB or 120 TB of usable capacity. The tape volume cache throughput of the new performance model is up to five times faster than the SAS-based TS7770 tape volume cache. In fact, a single drawer of SSD has equal or better tape volume cache performance than a 10 drawer SAS-based TS7770.

R5.3 introduced an increased maximum SSD cache capacity of 260 TB (240 TiB) before compression. This resulted in a total of maximum 4 cache drawers (1x 3948-CFC plus 3x 3948-XFC). The SSD-based cache modules are not required to scale in number of pairs as

the SAS-based HDD do. There can be any odd number of SSD drawers configured as well, although choosing an even number of cache modules would be recommended.

In the second half of 2023, R5.3 pga1 introduced the support for a full F07 base frame of SSD cache modules. This translates to 10 cache modules in total with a usable capacity of 640 TB (590 TiB) before compression: 1x 3948-CFC plus 9x 3948-XFC.

All described maximum configurations that are mentioned previously are also valid and supported if the TS7700 is installed in a client-supplied rack.

Note: For the client supplied rack support with the SSD cache option, an individual request for price quotation (RPQ) must be submitted prior to the order (RPQ 8B3749).

R5.0, introduced the TS7770. A disk-only model is available (referenced as TS7770D). The TS7770T provides a tape attachment to a physical library (IBM TS3500 or IBM TS4500). The TS7770C provides for attachment to Object Storage in the cloud. The TS7770 can provide up to 3.9 petabytes (PB) of usable disk cache capacity in the R5.0 pga1 release and later.

R4.0, introduced the TS7760. A disk-only model is available (referenced as TS7760D). The TS7760T provided tape attachment to a physical library (IBM TS3500 or IBM TS4500). The TS7760 provided up to 2.5 petabytes (PB) of usable disk cache capacity.

TS7760 and TS7770 models can attach to an IBM TS3500 or the IBM TS4500 Tape Library, which is the next generation of the TS3500 library. In an IBM TS3500, all types of tape drives of the TS1100 family are supported. In an IBM TS4500, the TS1160, TS1150, and TS1140 can be used.

Previous generation TS7740 and TS7720 models can coexist within R5.0 and later mixed grids for migration purposes, but are not described in this IBM Redbooks publication. For more information about the TS7740 and TS7720 models, see *IBM TS7700 Release 4.2 Guide*, [SG24-8366](#).

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 that are 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 tape industry evolved, a need emerged for large-capacity, disk-only tape solutions that enabled 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 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 possible. The evolution of the TS7700 combined the benefits of the TS7720 and TS7740 models to create the TS7760T. Now, the latest tape attach TS7700 is the TS7770T.

Through the combination of the technologies, Hybrid grid benefits can now be achieved with a single product that is configured in different ways. All features and functions of the older TS7720 and TS7740 were maintained, although more features and functions were introduced to help with the industry's evolving use of IBM Z virtual tape. These features and functions include the following examples:

- ▶ 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 multiple petabyte disk cache storage, the traditional TS7740 disk cache management style is not adequate for many workload types. Different workloads can include different disk cache residency requirements and treating all types with one recently used algorithm is not always sufficient. A method to manage disk cache usage at workload granularity might be needed.

The TS7700T supports the ability to create 1–7 disk partitions. Each partition is user defined in 1 TB increments. Workloads that are directed to a disk partition are managed independently concerning disk cache residency. After you create 1–7 disk 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 publication, the disk cache partitions are referred to as CP1–CP7, or generically as cache partitions (*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 partitions and which pools are used for a specific workload are independent.

Storage Class (SC) is used to direct workloads to a specific partition. Content cannot be automatically moved between partitions. However, content can be moved by using mount/demount sequences, or the **LIBRARY REQUEST** command.

TS7700T disk cache partitions (CP1-CP7 and CPx)

At least one partition must exist in a TS7700T configuration that is used for offloading to physical tape. The default is CP1.

However, you can configure new partitions and delete partitions if one *CPx* partition remains. 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 specific *CPx* partition are managed against data in 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 configuration enables each workload to include 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 for that target cluster. 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.

If IBM DS8000 Object Store support is enabled, one of the CP1 through CP7 partitions is set aside for DS8000 Objects. This special partition can also be dynamically resized, but it does not support offloading to physical tape. It resembles the *CP0* or resident only partition.

For more information, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).

TS7700T CPx tape 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 TS7700T feature codes (FCs) 5274 and 5279. Each FC 5274 feature provides 1 TB of premigration queue; FC 5279 provides 5 TB of premigration queue. The minimum is one feature for 1 TB and the maximum is of 10 1 TB features and 10 5 TB features. When 60 TB of total premigration queue is licensed, the premigration queue limit is unbound and limited only by the amount of licensed disk cache capacity.

Content queued for premigration is compressed, so the premigration queue size is based on post-compressed capacities. For example, if a host workload 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, premigration activity is minimal until the PMPRIOR threshold is crossed. When crossed, the premigration activity increases based on the defined premigration drive count for the target pool.

If the amount of content in the premigration queue continues to increase, the PMTHLVL threshold is crossed, and the TS7700T intentionally throttles inbound host and copy activity into all CPx partitions to maintain the premigration queue size. The TS7700T enters the sustained state of operation. The PMPRIOR and PMTHLVL thresholds can be no larger than the licensed premigration queue size. For example, if three FC 5274 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 disk cache and physical tape until the migration policies determine whether and when the volume is deleted from disk cache.

How many FC 5274 and FC 5279 features are installed is based on many factors. 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 is not necessary unless the data aged a certain amount of time. Delay premigration provides a method to retain data that is only in the disk cache until a delay criteria is met, and only then queuing it for premigration.

The data is never moved to tape if the logical volume expires before this delay period. This feature reduces physical tape activity to only the workload that is viewed as archive content. It can also greatly reduce the back-end physical tape reclamation processing that can result from data that expires quickly. Any volume that is in the grace period and did not meet its delay criteria is not counted against the premigration queue size limit.

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

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 independently on each CPx partition. Migration is the process of removing a logical volume in a disk cache after first moving one or more copies on physical tape and meeting the following criteria:

- ▶ A copy of the logical volume must be premigrated to a primary physical pool and an optional secondary pool if configured.
- ▶ Peer clusters in a TS7700 grid configuration completed copies of the logical volume, or they are unavailable due to an hours long service outage and space is needed.
- ▶ The preference group criteria also were met.
 - PG0: Volumes are removed from disk cache immediately after this criteria is met.
 - PG1: Volumes are removed from disk cache based on a *least recently used* algorithm that is relative to other volumes in the same CPx partition. 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 that is 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 automatically disabled. Otherwise, automatic removal is enabled by default and can be disabled by using the **LIBRARY REQUEST** command.

The CP0 usable size is determined by the remaining configured or licensed capacity after defining one or more CPx and DS8000 Object Store partitions. The CP0 partition must be at least 2 TB and can be as large as the licensed 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 IBM FlashCopy® processing, and for overcommit or overspill, as described next.

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. A partition can enter the overcommitted state by using one of the following methods:

- ▶ A user decreases the size of a partition to a value that is smaller than the amount of data that is contained within the partition.
- ▶ An excess of volume content is moved from one partition to another as part of a policy change, followed by mount activity or **LIBRARY REQUEST PARTRFSH** activity.
- ▶ CPx receives more content during a workload than can be premigrated. This issue might be because of the CPx partition being smaller than the premigration queue size. Or, it can be a result of an inability to perform premigration. When either of these conditions occur, space is taken from the CP0 partition, which 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 temporary capacity for these expected use cases.

If CP0 has no remaining free space, further overspill is prevented. In addition, a LI REQUEST tunable is provided to limit how much overspill can occur into CP0.

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 enables a partition assignment change to occur without a mount/demount sequence. When **PARTRFSH** is used, 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. The command must be issued to a TS7700T distributed library for each logical volume that must be moved.

In either case, logical volumes can be moved from CP0 to CPx, from CPx to CP0, and from CPx to a different CPx partition.

The following rules apply for CPx to CPx movements:

- ▶ Virtual volumes that are still in disk cache are immediately reassigned to the new partition and the active content value for the source and target partition are automatically adjusted.
- ▶ If **PARTRFSH** is used, the last access or creation time is not modified. Therefore, any delay in premigration remains acknowledged if the target partition can accommodate the delay.
- ▶ If a mount/demount is used, the last access time is modified. Therefore, any delay in premigration based on access time is delayed further.
- ▶ Any other changes in constructs, such as preference group, premigration delay rules, and pool attributes, are kept only 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 and not yet on back-end physical tape 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 invalidated.
- ▶ Virtual volumes that are in CPx and only on physical tape have the partition reassigned, but a copy is not automatically recalled into CP0 disk cache. If a recall occurs later, the instance that is recalled into the CP0 disk cache becomes the only copy on that cluster. Any previous physical tape instances are removed from physical tape.
- ▶ Any other changes in constructs, such as removal properties, are kept only 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 the disk cache is immediately reassigned.
- ▶ If no delay premigration is active for the assigned SC, the volume is immediately queued for premigration for physical tape.
- ▶ If a delay premigration is active for the assigned SC, the delay criteria based on last access or creation time are retroactive. The movement does not alter the last access or creation time reference point.

- ▶ If the logical volume was 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 kept only if the movement is the result of a mount/demount sequence.

2.1.7 Introduction of the TS7700C

A Cloud Storage Tier enabled cluster, or TS7700C, can be a member of any grid if the peer cluster code levels are compatible. Peers can be more TS7700C clusters, TS7700T clusters, or disk-only TS7700D clusters

By using management class policies, virtual volumes are replicated among peers in the grid. Those peers with Cloud Storage Tier support then can offload to an object store. After one TS7700C cluster offloads a logical volume, other peers that attempt to offload the same logical volume detect that the volume exists in the object store and can skip the premigration phase. Either cluster can then recall the logical volume from the cloud, if needed, including newly joined clusters.

Logical volumes offload in their entirety to the object store. If the only available logical volume copy in the grid is within an object store, every TS7700C cluster can recall it into its disk cache. After the entire volume is present in the disk cache, the content is accessible by using the grid techniques. Because a volume must be recalled entirely, choosing logical volume sizes that are smaller can provide faster time to access when the entire content of the logical tape is not being accessed.

With R5.1 the TS7700C supports two new features:

- ▶ Volume version retention retains old versions of logical volumes in the cloud for later recovery.
- ▶ Cloud Export, Restore, and Recovery exports one or more cloud pools with a point in time database backup into the cloud. This backup can be restored to any empty TS7700C.

As of this writing, the TS7700C supports the ability to offload to IBM Cloud Private/Public and Amazon Web Services. Multiple cloud pools are supported, which means the TS7700C can direct workload into different buckets (even on different Cloud providers).

Similar to TS7700 Tape Attach, the TS7700C features disk cache partitions (CPx) that are available to manage the disk cache footprint of the TS7700C. Those workloads that benefit from larger disk cache footprints can use a large disk cache partition. Those workloads that do not require as much disk cache residency can target smaller partitions.

Logical volumes that target any partition other than the residency-only partition premigrate to an object store immediately. The policy that is assigned to the logical volume and how much space is available in the partition determines whether and when the logical volume is removed from disk cache after the premigration. That is, as with the TS7700T, the copy in the disk cache is removed, which makes the cloud instance the only available copy in that cluster. A recall of the logical volume from the object store into the disk cache is required if that cluster was chosen for the tape volume cache (TVC) or a copy source.

With the TS7700C, the same techniques as with the TS7700T apply in terms of moving logical volumes between CP partitions.

One other option is available with TS7700C, called **MMOUNT**. It is an option within the **LI REQ PARTRFSH** command. This option can be used to update the cloud attributes of a logical volume. For example, if a logical volume must be premigrated into the cloud without issuing a mount/demount.

For more information about this feature, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

Similar to hybrid grids that are made up of disk-only TS7700 clusters and TS7700T clusters, the TS7700C can be mixed within the same grid. At the time of this writing, a specific cluster cannot be Tape Attach *and* support Cloud Storage Tier at the same time; therefore, the two features are mutually exclusive within the same cluster.

For more information about the TS7700C solution and how to implement and integrate this solution into your enterprise, see *IBM TS7700 R5.4 Cloud Storage Tier Guide*, [REDP-5573](#).

2.1.8 TS7700O Advanced Object Store for DS8000

DS8000 object store allows for the TS7770 to act as an object store for transparent cloud tiering with IBM DS8000 (DS8K), DFSMSHsm (HSM), and native DFSMSdss (DSS). Through DFSMS policy management, applications (such as DFSMSHsm) can request the DS8K offload data sets in the form of objects directly to the TS7770 instead of a traditional object store device (MLC instead of ML2). Offload occurs through the 1 Gb or 10 Gb Ethernet (highly recommend 10 Gb). The DS8K can target up to two TS7770s per grid.

FICON created content and DS8K offload data can co-exist by using Logical Disk Cache Partitions within the TS7770 because logical cache partitioning of objects versus virtual volumes exists. Existing cache storage and Ethernet are used.

Data that is in-flight between the DS8K and TS7770 is provided with full AES256 bit encryption through TLS1.2. Intelligent compression exists within the DS8K before offloading objects to the TS7770. Data that was compressed or encrypted is excluded from compression.

Through Advanced Object Management, all TS7770 object-enabled clusters are aware of all objects in the grid. Access to all objects is available from any cluster in the grid whether the cluster has a local copy. Changes during a cluster outage are automatically reconciled.

Object policy management allows for selecting how many copies to create, where to create them, and whether synchronous or asynchronous. Multi-cloud support allows for configuration of up to 8 cloud targets per DS8900F whether public cloud, private cloud, and TS7770 as an object store. Combinations can be public, private, and TS7770s and multiple TS7700 grids or the same grid.

For more information, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).

2.1.9 Management of the TS7700

Management of the TS7700 consists of the following key components:

- ▶ TS7700 Management Interface (MI)
- ▶ TS4500 or TS3500 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

These components are described next.

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.

TS4500/TS3500 web interface

The TS4500 and TS3500 web interface is used to configure and operate the optional attached physical 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 by using 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 as part of every logical tape mount.

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, which enables you and the TS7700 to better manage your volumes. For the other IBM Z platforms, constructs can be associated with the volumes when the volume ranges are defined through the library's MI.

DFSMS constructs in the IBM Z platform and their equivalents in TS7700

In the IBM Z platform, the following DFSMS constructs exist:

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

Each of these constructs is used by the TS7700 to determine how to manage your data. Ideally, each construct features a matching definition within the TS7700, although default behaviors are inherited if no match is found. You can predefine these constructs on the TS7700 MI. For more information, see "Defining TS7700 constructs" on page 593. If constructs are sent to the TS7700 without having predefined constructs on the TS7700, the TS7700 creates the construct by using the currently configured default parameters for that construct.

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 decide the following factors:

- ▶ Whether data is SMS-managed
- ▶ The level of performance of a data set
- ▶ 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 into storage pools, so you allocate by storage pool. Storage pools can also consist of tape volumes. This feature 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 the library and SG levels. If an SG is connected to specific 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 primary physical tape or cloud pool. 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 are otherwise 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.

Data Class in SMS

The DATACLAS construct defines the look of a file. The Data Class ACS routine is always started, even if a file is not SMS-managed. A Data Class is only ever assigned when a file is created and cannot be changed.

A file is described by how many volumes it can span and its:

- ▶ Data set organization
- ▶ Record format
- ▶ Record length
- ▶ Space allocation
- ▶ Data compaction
- ▶ Media type
- ▶ Recording information

Data Class TS7700

The DATACLAS 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 to all other clusters in the grid. In R4.1.2, new compression optimization options are supported along with the 3490 block handling counters.

Important: The DATACLAS 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. For more information, see the following chapters in this publication:

- ▶ Chapter 6, “Implementing IBM TS7700” on page 239
- ▶ Chapter 8, “Migration” on page 303
- ▶ Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351
- ▶ Appendix F, “Library Manager volume categories” on page 951

The following 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 started 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. When settings are changed, the TS7700 behavior can change for all the hosts that use the TS7700, which must be considered when settings are changed by using the LI REQ command. For more information, see IBM TS7700 Series z/OS Host Command Line Request User's Guide Version 5.4.</p>
DS QLIB	Use the DEVICE SERVICES QUERY LIBRARY command to display library and device-related information for the composite and distributed libraries.

A subtle difference exists between the **DS QLIB** and **LI** commands, 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

The TS7700 provides a method within the MI to customize if and how alerts are surfaced by the TS7700. Each can be independently modified in the following way:

- ▶ Customized severity
- ▶ Whether the notification is surfaced
- ▶ To which notification channels (Host, SNMP, RSYSLOG server) the message is sent
- ▶ Sending more custom text information to the notification channels to help with automation

The changes to the event notification settings are grid wide and persistent if new microcode levels are installed.

In addition, you can back up these settings and restore them independently on other grids to improve the ease of managing maintenance of a multi-grid environment.

For information about content-based retrieval (CBRxxxx) messages, see [TS7700 Series Operator Informational Messages](#).

Tools

Many helpful tools are available for the TS7700. For more information, see Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351.

Remote System Log Processing Support

Because security is becoming a critical aspect for most customers, all events and messages can be sent to one or two different RSYSLOG servers. Remote System Log Processing Support (RSYSLOG) is an open standard utility that uses TCP for the message transport.

An MI interface is provided to configure a TCP port and an IP address to an RSYSLOG server to send the events from TS7700.

Starting with R5.4 pga1, RSYSLOG also supports adding SSL certificates for RSYSLOG with TLS version 1.3. For more information how to configure it, refer to 10.3, “Access icon” on page 492. Also, see Chapter 10, “IBM TS7700 Management Interface operations: Part 2” on page 449. You can also see the following documentation:

- ▶ [Adding SSL certificates for RSyslog with TLS](#)
- ▶ [R Syslog](#)

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
- ▶ In the associated TS3500 or TS4500 library and tape drives (if present)
- ▶ 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.

Call Home events send data that is related to the problem 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. Following analysis, an appropriate course of action is defined to resolve the problem. For example, 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) 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.

Starting with R5.3 pga1 (late summer 2023) the TSSC/TS3000 Service Console is part of the 3948-VED server ship-group. This means that the TSSC/TS3000 specific feature codes within the F07 are discontinued and are no longer listed in a configuration.

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

In general, any cluster can be used as a stand-alone cluster. The TS7700 features several internal characteristics for High Availability (RAID6 protection, dual power supplies, and others). However, a grid configuration can be configured for both more high availability (HA) and Disaster Recovery (DR) functions with different levels of business continuance. For more information, see Chapter 3, “IBM TS7700 usage considerations” on page 115.

In this section, general information is provided about the components, functions, and features that are used in a TS7700 environment.

For more information about deviations and the multi-cluster grid, see 2.4, “Multi-cluster grid configurations: Components, functions, and features” on page 64.

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.

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 are available: 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 2-4 shows the host view of a stand-alone cluster configuration.

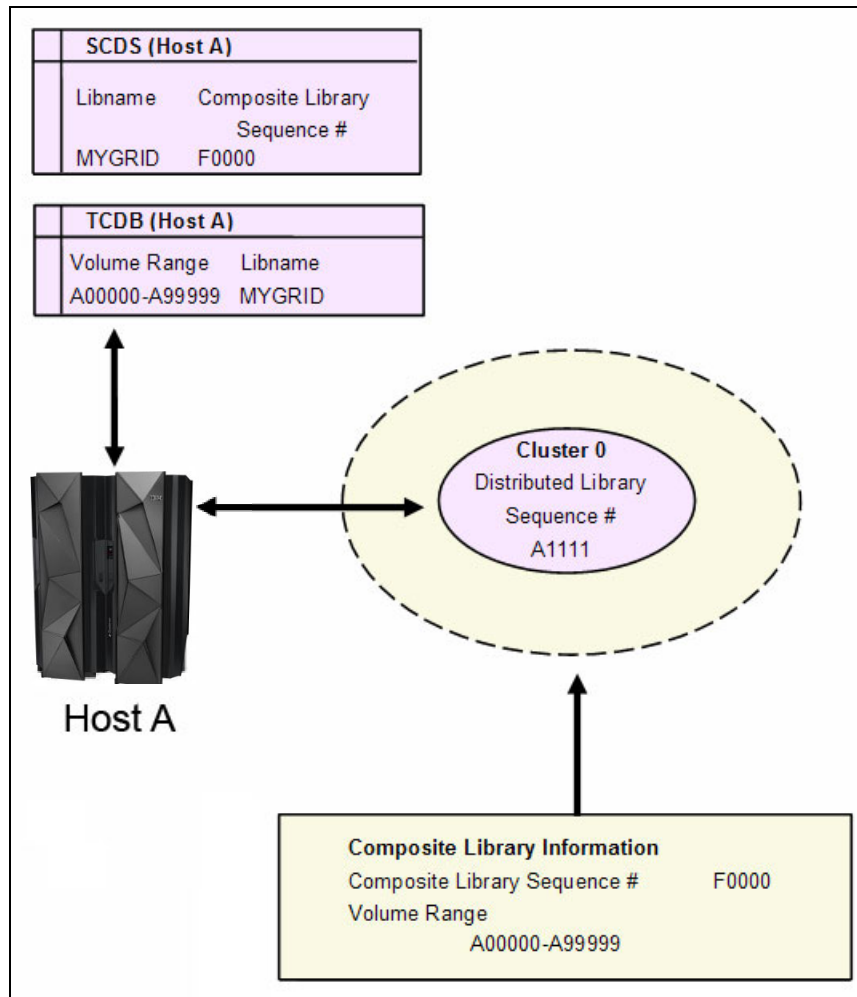


Figure 2-4 TS7700 stand-alone cluster configuration

Distributed library

Each cluster in a grid is a distributed library, which consists of any type of TS7700. At the host, the distributed library is also defined to SMS. It is defined by using the existing ISMF windows and does not include any defined tape devices. The virtual tape devices are defined to the composite library only. It is important to define the distributed libraries so that library-specific attentions and alerts are properly handled by SMS.

A distributed library consists of the following cluster hardware components:

- ▶ A virtualization engine
- ▶ A TS7700 TVC
- ▶ A 3952-F06 frame
- ▶ Attachment to a physical library (TS7700T)
- ▶ Several physical tape drives (TS7700T)
- ▶ Connectivity to an object store in the cloud (TS7700C)

Important: A composite library ID must be defined 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 multi-cluster configuration, the composite library ID must differ from any of the unique distributed library IDs. 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 hardware.

2.2.2 Tape Volume Cache

The TS7700 TVC is a disk buffer that receives all emulated tape write data. It also 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 any optional physical tape drives or cloud object stores that are attached to the TS7700.

Originally, the TS7760 was delivered with 4 TB disk drive support. Since August 2017, only 8 TB disk drives are available. You can mix 4 TB or 8 TB drives within a frame, but you cannot mix them within a drawer or enclosure.

The TS7760 CSA/XSA uses Dynamic Disk Pooling (DDP) while the TS7770 CSB/XSB and TS7770 CFC/XFC use Distributed RAID (DRAID) to protect all logical volume data. These distributed or dynamic RAID6 configurations provide the resiliency of RAID6 protection, but with greatly accelerated rebuild speeds when a disk failure occurs.

The DDP on the TS7760 is built up from 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 are installed, the last drawer in the frame is configured as a single drawer DDP.

DRAID on the TS7770 is at single drawer granularity. Up to two disks per physical drawer can be lost without losing access to data. TS7770 CSB/XSB drawers must always be added in physical pairs while the TS7770 CFC drawer configuration minimum is one. In addition, the physical available capacity is then licensed by using the TS7700 Cache On Demand features. All physically installed spindles are used independent of whether all available capacity is licensed.

The distributed RAID technologies do not use a global spare concept; instead, they reserve free space in each distributed disk pool. If a DDM failure occurs, the parity-generated data is read from all remaining DDMs and written to the reserved free space of the remaining DDMs. Because all DDMs receive rebuilt data, the speed of the rebuild is not governed by a single spare. This procedure is called *reconstruction*.

If a second DDM fails during the reconstruction, the drawer pauses the first reconstruction and starts a "Critical reconstruction" of any RAID6 stripes that depend on both failed disks. This process allows faster rebuild times to return the disk cache to at least RAID5 protection. After the Critical reconstruction is finished, the paused reconstruction is restarted.

After the reconstruction completes and new DDMs are inserted into the disk cache drawers, the content in the reserved space is redistributed back to the non-reserved space of all DDMs. Though similar to traditional RAID6 copy back, it is faster because a single replaced DDM is not the only source or target of the copy back content.

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 file-like objects that are in the TVC. They can optionally replicate to peer locations and optionally offload to back-end physical tape or cloud.

As with a real volume, each logical volume features the following characteristics:

- ▶ Has a unique volume serial number (VOLSER) that is known to the host and 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.
- ▶ Newly inserted logical volumes, scratch-mounted volumes and expired-data volumes all allow the first write to occur at any position between BOT and just after the first tape mark. These volumes always present an auto-created 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. The exception is if Tape Write Immediate mode is used which 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 that are 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, 25000, or 65000 MiB when using the larger volume size options that are assigned by DC.
 - 4,194,303 blocks were written, since 3490E devices are limited to a 22-bit block ID.

The default logical volume sizes of 400 MiB or 800 MiB are defined by the media type that is chosen at the logical volume insert time. By using DC constructs, these established volume sizes can be overridden during any scratch mount or any private mount where a write from BOT or just past the label occurs.

You can direct your SMS-managed tape operations to a TS7700 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 that are defined through the ACS routines.

A maximum number of 4,000,000 virtual volumes per composite library or grid can be supported. The default maximum number of supported logical volumes is 1,000,000 per grid. Support for extra logical volumes can be added in increments of 200,000 volumes by using FC 5270.

VOLSERs for logical volumes are defined through the MI when they are newly inserted. Virtual volumes go through similar cartridge entry processing as though they were native cartridges that were inserted into a tape library that is attached directly to an IBM Z 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, each host operating system interacts by using the object access method (OAM) with the Library to do a category inventory of all volumes in the insert category.

Depending on the definitions in the DEVSUPxx and EDGRMMxx parmlib members, and whether the host has a VOLSER matching Tape Management System (TMS) entry, 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.

2.2.4 Logical volumes and compression

The TS7700 supports the following compression algorithms:

- ▶ FICON: IBMLZ1 ASIC-accelerated compression
- ▶ LZ4: LZ4 FAST software-based compression
- ▶ ZSTD: Z Standard software-based compression

The compression algorithm that is used for a specific tape workload is determined through the Data Class construct. This definition allows you to select the most suitable compression algorithm for different workloads.

Three primary factors differ between the algorithms:

- ▶ How well they compress, starting with FICON, then LZ4, and then ZSTD having the strongest compression.
- ▶ How much of the TS7700 CPU is used to perform the compression, starting with FICON, then LZ4, and then ZSTD by using the most TS7700 CPU.
- ▶ The performance, which can vary depending on how many concurrent jobs are running.

With fewer jobs, FICON and LZ4 provide the best performance, while ZSTD can provide the best performance when many jobs run in parallel.

Consider an example in which 64 or more parallel devices are being used per TS7700. If data is compressed or encrypted before it reaches the TS7700, FICON compression can expand the data by 12.5%, and LZ4 and ZSTD do not expand the data.

It is best to disable TS7700 compression for such workloads by using `COMPACTION=NO` within your IBM Z Data Class parameters. In this way, the TS7700 avoids attempting to compress the workloads.

To use LZ4 and ZSTD compression algorithms, all clusters in a grid must have R4.1.2 or later microcode. If older VEB/V07 clusters exist, the CPU usage might be unacceptable for those clusters. Consider the use of only FICON or LZ4 compression when production workloads write to VEB or V07 clusters.

Higher compression has a positive effect in the following areas:

- ▶ Cache resources (cache bandwidth and cache space) required
- ▶ Grid link bandwidth
- ▶ Physical tape resources
- ▶ Object stores capacity

- ▶ Premigration queue length (FC 5274)
- ▶ Recovery point objective, depending on the configuration

The host data that is stored on a virtual CST or ECCST volume is displayed by the **LI REQ** commands and in the **MI**.

The virtual tape volumes have a range of 400 MiB–65000 MiB. With a compression ratio of 3:1, the uncompressed size can vary from 1200 MiB–195,000 MiB.

2.2.5 Mounting a scratch virtual volume

When a request for a scratch is sent to the TS7700, the request specifies one or more source categories. The TS7700 selects a virtual VOLSER from the candidate list of scratch volumes in the primary or alternative category. Whether one or two category options are provided depends on if the ACS routine's chosen Data Class is assigned a specific media type (for example, MEDIA2 or MEDIA1). If a media type is not explicitly assigned to the ACS routine-assigned Data Class, both MEDIA2 and MEDIA1 are provided. If both are provided, the TS7700 first chooses from MEDIA2 volumes and then MEDIA1 (if present).

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 were in the source category the longest. Volume serials are also toggled between odd and even serials for each volume selection.

Note: Step 1 of how scratch volumes are chosen can be skipped by configuring the **LI REQ PFRL0C0** option. Skipping step 1 can help with scratch usage rotation by viewing all volumes in the entire grid by age first (step 2) instead of only those locally owned, which might include only recently returned to scratch.

For all scratch mounts, the volume is temporarily initialized as though the volume was initialized by using the **EDGINERS** or **IEHINITT** program. The volume includes an IBM-standard label that consists of a VOL1 record, HDR1 record, and 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 2.4.4, "I/O Tape Volume Cache selection" on page 69.

When the Fast Ready attribute is set or implied, no recall of content from physical tape or Cloud Object Storage is required in a TS7700T or TS7760C. 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 and always emulates a simple VOL1/HDR1 label.

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. The Management Class construct determines which clusters should be used for a workload when SAA is enabled by using the **LI REQ DEVALLOC** command.

2.2.6 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. For more information, see 2.4.12, "Mounting a specific virtual volume" on page 76.

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

- ▶ A valid copy exists in the TVC. In this case, the mount completes quickly and the host can access the data immediately.
- ▶ A valid copy does *not* exist in the TVC. In this case, the following options are available:
 - If it is a TS7700T, the volume exists in CP1-CP7 and was copied to physical tape. The virtual volume is also recalled from a stacked volume. Mount completion is signaled to the host system only after the entire volume is available in the TVC.
 - If it is a TS7700C, the volume exists in CP1-CP7 and was copied to an object store. The virtual volume is also recalled from the cloud tier. Mount completion is signaled to the host system only after the entire volume is available in the TVC.
 - Assuming a stand-alone cluster, if no consistent copy exists in the disk cache or within the attached tape or cloud, the mount fails. In a grid, a peer's TVC can be chosen as well.

Any 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 a recalled volume is modified, the new instance of the volume is premigrated to tape or cloud and the previous instance is invalidated. If the recalled volume is not modified, no premigration occurs unless a pool configuration change occurred as part of the mount.

For example, if the primary or secondary pool changes as part of a mount operation, a premigration occurs to one or both pools. Any instance in a previous pool is invalidated. Furthermore, copies to remote TS7700 clusters in a grid configuration are not required if modifications were not made, unless the Management Class also changed and retain copy modes is not enabled.

Any specific or private mount must target a logical volume that is assigned to a non-scratch category. 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. If this process occurs, the temporary tape header is created, and no recalls occur.

The DFSMS Removable Media Manager (DFSMSrmm) or other TMS likely fail 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 within the MI.

2.2.7 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
- ▶ During the assignment of LWORM to a volume's characteristics, it generates 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 appends only 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 was bound as an LWORM volume.

TS7700 reporting volumes (BVIR) cannot be written in LWORM format. For more information, see 13.4, "Bulk Volume Information Retrieval" on page 687.

Clarification: Cohasset Associates, Inc. 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.

2.2.8 Virtual drives

From a host perspective, each TS7700 emulates 16–31 logical IBM 3490E tape CUs. Each CU has 16 unique drives that are attached through any of the up to 8 FICON channels.

With 31 CUs per TS7700, a total of 496 devices can be emulated per TS7700. In an eight-way grid, up to 3,968 devices can be defined for the same composite library.

As with physical IBM 3490 systems, the IBM Z host uses the hardware configuration definition (HCD) to configure the devices. Defining a preferred path for the virtual drives gives you no benefit.

Each virtual drive features 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 it is started 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 include an ICL. For scratch mounts, the use of 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 to a maximum of 496 logical drives per cluster.

Each physical TS7700 FICON channel supports 512 logical paths. A single path is reserved for each LPAR chpid that communicates through that FICON channel to a particular LCU. For example, if a single LPAR wants to see all 496 devices or 31 LCUs down all 8 FICON channels, each of the 8 FICON channels reserves 31 logical paths for that LPAR, which leaves 481 logical paths per FICON channel for other LPARs.

If only a subset of the physical FICON channels is used for a specific LPAR, only the FICON channel used reserve paths, which allows more LPAR combinations to exist by partitioning which FICON channels are used by different LPARs. In addition, paths are reserved only for LCUs that are configured for that LPAR. That is, if an LPAR needs to see only 32 devices (2 LCUs), it reserves two paths per FICON channel.

2.2.9 Selective Device Access Control

Having protection in place to prevent system plexes from accessing other plexes' volumes is important. Hard partitioning was used for many years as a way to assign specific device ranges to different systems.

Selective Device Access Control (SDAC) provides another layer of security when *hard partitioning* is used. The primary intent of SDAC is that in a multi-tenant environment, SDAC can be used to prevent one IBM Z logical partition (LPAR) or sysplex from modifying, accessing, or removing data that is owned by another system or tenant.

SDAC can be valuable in a setup where isolated production systems are used, which includes test systems and development systems all sharing a common grid or TS7700 cluster. It can also be used in a multi-tenant configuration to prevent tenants from accessing each other's data. Lastly, it is often used when different IBM Z operating systems 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. By using SDAC you can to create a corresponding SDAC group for the same LCU definitions and associate that group with one or more dedicated volume ranges. This process creates a relationship between LCU devices and the volumes that those devices can access.

SDAC is a useful function when multiple partitions or plexes feature the following characteristics:

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

SDAC allows you to assign one or more LCUs of the entire grid to an access group (CUs in ranges of 16 devices based on the LIBPORT definitions in HCD). Those groups are then associated with one or more logical volume range definitions. After it is enabled, it controls access to host-initiated mounts, ejects, attribute or category changes, and optionally volume inventory requests. For more information about the implementation of SDAC, see Appendix I, “Case study for logical partitioning of a two-cluster grid” on page 991.

Implementing SDAC requires planning and orchestration with other system areas in defining how devices are partitioned among systems and how the input/output definition file (IODF) or HCD is defined and managed. From the TS7700 subsystem perspective, SDAC definitions are set up by using the TS7700 MI.

Important: SDAC is based on the availability of LIBPORT definitions, which are fixed values for all possible LCU instances in a theoretical maximum configured eight-way grid. It is assumed that the IBM Z IODF or similar is managed by a common administrator who oversees the entire configuration. SDAC protected partitions must be defined on full LCU boundaries (in groups of 16 devices).

2.2.10 Physical drives

The physical tape drives that are used by a TS7700T are installed in an IBM TS4500 or IBM TS3500 tape library. The physical tape drives are not addressable by any attached host system and are controlled entirely by the TS7700T. The TS7700T supports TS1160, TS1150, and TS1140 when the IBM TS4500 is used. When an IBM TS3500 is used, the following physical tape drives are supported:

- ▶ TS1160
- ▶ TS1150
- ▶ TS1140
- ▶ TS1130
- ▶ TS1120
- ▶ IBM 3592-J1A

Remember: Do *not* change the assignment of physical tape drives that are attached to a TS7700T in the IBM TS4500 IBM or IBM TS3500 Tape Library web interface. This procedure must always be done by your IBM SSR.

The TS7700T primarily supports homogeneous drive type configurations. The exception is when a drive upgrade occurs. A mixed drive configuration is supported to help with any JA and JB media compatibility issues.

A mixed configuration supports a maximum of two drive types: one newer type (a TS1150 or TS1160) and a second type, which is a previous generation TS1140 or older. The newer type is used for read and write activity; the older type is used for read-only access to JA and JB media types. At least four of the newer type and at least two of the previous generation type is required. A maximum total of 16 drives can be installed.

The read-only media in such a mixed configuration is referred to as *sunset media*. Sunset media can be present indefinitely. Sunset media can be reclaimed naturally or more aggressively based on the different settings. For more information, see 7.1.3, “TS7700 tape library attachments, drives, and media” on page 272.

2.2.11 Stacked volumes behind the TS7700T

Physical cartridges that are used by the TS7700T to store logical volumes are under the control of the TS7700T. 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 initialized or used previously. If the internal label does not match the external label, the TS7700T rejects the volume. The exception is if a **LI REQ RELABEL** command was issued against the physical tape before use.

For more information, see “Relabeling cartridges with the TS7700T” on page 459.

After an IBM Z host creates a volume in a TS7700T, or a volume is replicated to a TS7700T, the storage management software that is inside the TS7700T schedules the virtual volume to be premigrated onto one or more physical tape cartridges. This use case assumes that the logical volume targets a CPx partition and any delay premigration criteria was met.

The TS7700T attempts to maintain a minimal number of stacked volumes per pool. The TS7700T mounts a defined number of drives in parallel for a specific pool and attempts to favor unfilled stacked volumes before empty volumes are used.

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. Although all logical volume compression that occurred within the host or within the TS7700 is retained, compression is still enabled within the physical tape drive, which can provide small improvements for certain workloads.

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.

Because of 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 53.

Through the TS3500/TS4500 web interface, physical cartridge ranges are assigned to the appropriate library partition that is associated with your TS7700. This feature enables them to become visible to the correct TS7700.

The TS7700 MI must also be used to define which pool physical tapes are assigned to when they are 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.

2.2.12 Selective Dual Copy function for TS7700T

In a 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, that TS7700T can lose access to one or more logical volumes that are contained on the damaged physical tape.

If a grid configuration is present, an alternative copy in the grid is used. But, it can still be preferred to also have a second copy behind the TS7700T for stand-alone use cases or when the grid peer might also be unavailable.

The TS7700 provides a method to create redundant copies on independent physical tapes to help reduce the risk of such a loss. It can also be used for Copy Export or the ability to eject an incremental point in time backup of the TS7700T.

With the Selective Dual Copy function, storage administrators can selectively create two copies of logical volumes within two pools on the same TS7700T. The Selective Dual Copy function can be used for added resiliency or with the Copy Export function to provide a secondary offsite physical copy for DR purposes. For more information about Copy Export, see 2.3.14, “Copy Export function” on page 59.

The second copy of the logical volume is created in a separate physical pool to help ensure physical cartridge separation. Control of Dual Copy is through the MC construct (see 10.2.2, “Management Classes window” on page 478). The second copy is created after 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 the 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 available only when the primary volume cannot be recalled or is inaccessible. It cannot be accessed separately, such as when the primary volume is being used by another operation. The second copy provides an error path recovery backup if the primary volume is damaged or inaccessible.

Selective Dual Copy is defined in the TS7700T and includes 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 or inbound copy complete, the secondary pool assignment is determined and premigration of the logical volume is scheduled. The primary must first be successfully premigrated before the secondary is queued for premigration.
- ▶ The logical volume in the disk cache is retained until both copies to back-end tape are completed.

2.3 General TVC management in a stand-alone cluster

The TS7700 cluster manages logical volume data within the TVC cache. By using policy settings and LI REQ settings, you can influence certain behaviors that are related to managing data in the TVC cache. For example, defining how data moves to back-end tape or cloud, or defining how long to retain data that is associated with expired scratched volumes.

2.3.1 Basic Rules for TVC Management

Cache Management can be summarized by the following rules:

- ▶ 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. If a virtual volume is written to the physical tape or the cloud, it uses only the space that is occupied by the compressed data and resulting metadata.
- ▶ All virtual volume data within any TS7700 can be auto-expired from TVC cache by using scratch category delete expire processing. After a volume is moved to a configured scratch category, any data that is associated with it can be auto-deleted after a defined grace period passes or the volume is reused during a scratch mount.
- ▶ TVC data's ability to offload to tape or cloud is managed by definitions in the SC.
- ▶ Active data in a stand-alone disk-only TS7700D, the CP0 partition of a TS7700T or TS7700C, and a DS8000 Object Store partition always remains in the TVC cache.
- ▶ In a TS7700T or TS7700C, data in a CPx partition is scheduled to be premigrated to tape or cloud by using SC policies. Delay premigration policies can be used to defer movement to tape so that only aged or archived data premigrates to tape.
- ▶ If a TS7700T or TS7700C CPx partition runs out of space, the cache management function removes or migrates previously premigrated volumes from TVC cache. The volumes that are candidates for removal are chosen based on PG0/PG1 preference groups. 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.
- ▶ In a TS7700T or TS7700C, volumes that are not in TVC cache during a tape volume mount request are scheduled to be brought back into the disk cache from a physical tape device or cloud object store by using a recall. The entire volume is recalled into the TVC cache before it is accessible.

2.3.2 Scratched virtual volumes and the Delete Expire function

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

- ▶ TVC might fill up with large amounts of expired data
- ▶ Stacked volumes might retain an excessive amount of expired data
- ▶ Stacked volumes fill up with expired data
- ▶ Object stores continue to use excessive capacity

To help manage this TMS expired content, the TS7700 supports a function that is referred to as *delete expire*. When enabling delete expire processing against a configured scratch category, you can set a grace period for expired volumes 1 hour–2000 years.

If the volume was not reused when the delay period passes, the volume is marked as a candidate for auto deletion or delete expire. A background process then periodically deletes candidate volumes. 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 to a maximum Value of 5000.

After the background handler deletes the volume, the volume's active space in TVC is freed. If it was also stacked to one or more physical tapes, that region of the physical tape is marked inactive. If it is contained in an object store, it is then marked for pending deletion where another background task periodically deletes objects in the cloud.

The start timer for delete expire processing is set when the volume is moved to a designated scratch category that includes an assigned delete expire value. If the scratch category has no delete expire value that is assigned in the MI, the timer is not set. Setting an initial expire time in the MI categories window is not retroactive for existing content in scratch. Only volumes that are returned to scratch from that point forward are candidates for delete expire. Also, a volume is not a candidate for expire processing until at least 12 hours passes since its last host access.

If the logical volume is reused during a scratch mount before the expire grace period passes and the background process deletes the data, the existing content is immediately deleted at the time of first write.

For more information about expired volume management, see "Defining the logical volume expiration time" on page 592. The explicit movement of a volume that is 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 or CP0 partition of a TS7700T and TS7700C.

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 TS7700T and more cloud object space in a TS7700C. The time that it takes a physical volume to fall below the reclamation threshold is also increased on a TS7700T 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 TS7700T 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 passes. 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 passes. The only host command that is acknowledged during this hold period is a category change back to a private category.

This extra 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.

Note: Important: Consider the following points:

- ▶ 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 stacked onto physical tape remains recoverable through an IBM services salvage process if the physical tape is not yet reused or if the secure erase process was not performed against it. Contact your IBM SSR if these services are required. Also, disabling reclamation when any return to scratch mistake is made can help retain any content still present on physical tape.
- ▶ 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.
After it is enabled, any other changes to the Delete Expired duration and toggling Expire Hold are retroactive for all scratched content.

2.3.3 Resident-only TVC cache management

In a stand-alone disk-only TS7700D, virtual volumes always remain in the TVC because no physical tape or cloud is attached. In a TS7700T and TS7700C configuration, contents in the CP0 partition are also not candidates for moving to physical tape or cloud. Lastly, a DS8000 object partition retains all objects in disk cache and does not yet support movement to physical tape or cloud. For virtual volumes in this state, only Delete Expire processing, a logical volume EJECT, or a reuse can free the logical volume TVC space.

For a DS8000 Object partition, SMS-initiated object deletes free the space for the offloaded object by requesting an attached DS8000 to delete the object and any metadata objects that are associated with it. In a TS7700T or TS7700C, **LI REQ PARTRFSH** can be used to move resident-only volumes from CP0 to a CPx partition as another method to free CP0 capacity.

If a disk-only TS7700D or resident-only partition runs out of space, warning messages are surfaced to all attached hosts and the library enters a degraded state. If capacity continues to grow and the free space is exhausted, a critical message is surfaced and the TS7700D or resident-only partitions enter a read-only state. The TS7700D enters the Out of Cache condition and moves to a read-only state. If a TS7720T/TS7760C 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.

2.3.4 Premigration of CPx TVC content

The TS7700T and TS7700C support the ability to automatically offload logical volumes from TVC cache to back-end attached tape or cloud. *Premigration* is the term that is used for putting logical volumes onto a physical tape or object store. The next few sections describe how this premigration process works.

Premigration with multiple disk cache partitions

You can specify 1–7 independent disk cache partitions (CPx). Each partition features its own independent cache management that allows different workloads to have different disk cache footprint characteristics. For example, workloads that benefit from a high disk cache hit ratios can use a large disk cache partition. Workloads that are rarely accessed and benefit little from disk cache residency can use a smaller disk cache partition.

By providing each workload type with its own partition, the hierarchical storage management concepts can apply to other volumes with similar workload characteristics only.

The Storage Group (SG) construct is used to determine which partition a workload targets. After it is in a CPx partition, it is then a candidate for premigration or the copying to back-end tape or cloud. By default, the queuing of the premigration occurs within minutes of volume close. But, it can be delayed by using Time-delayed premigration.

The total amount of logical volume content that can be queued for premigration to tape or cloud is limited by the number of FC 5274 and FC 5279 features. The **LI REQ PMTHLVL** command is used to determine when the TS7700T slows down inbound host activity to prevent the premigration queue from growing much further.

For the TS7700C, **LI REQ PMTHLVL** command is not applicable because it has an equivalent value that is fixed. This concept is referred to as *Host Throttling* and is a normal behavior for the TS7700T and TS7700C.

When throttling occurs to limit the growth of the premigration queue, the TS7700T or TS7700C is said to be within the sustained state of operation. Basically, the rate of which data is copied to tape or cloud limits the rate to which new compressed data can be accepted by host operations.

Premigration to back-end tape or cloud also occurs only if the TS7700C/T cluster is fairly idle or the **LI REQ PMPRIOR** threshold was crossed. This limiting of premigration allows more resources to be given to host activity as the premigration queue increases. This period when the queue grows without offloading to tape or cloud is referred to as the *peak state of operation*.

All logical volumes in all CPx partitions share the premigration queues. One queue exists for tape offloading and a separate queue exists for cloud or object store offloading.

Time-delayed premigration

Traditionally, queuing of a volume for copying to back-end tape occurs within minutes of the completion of the unmount process. Optionally, you can use the delay premigration to delay the queuing of the logical volume for premigration.

The delay premigration is managed by the following major variables:

- ▶ As part of the SC construct, an optional delay premigration setting can be enabled. It is based on units of hours since either Volume creation or when the Volume was last accessed. You choose whether the creation time or last access is used. The delay time can be 0 - 65535 hours. 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 disk partition. You can also have multiple SCs with different delay premigration definitions that point to different disk partitions

- ▶ Each CPx partition also has a maximum amount of delay premigration content that can exist at any time. This value is used to limit excessive delay of content for a specific partition. The minimum is 500 GB; the maximum is the CPx partition size minus 500 GB. You can adjust the premigration delay limit within the resize partition window. The default limit is 500 GB.

If the amount of delay premigration content in the specific CPx partition exceeds its configured limit, the TS7700T “fast forwards” a theoretical clock and finds the next logical volumes that can meet the delay criteria if more time passed. These volumes are then queued for premigration ahead of their configured criteria. Therefore, the oldest volumes are not necessarily are queued early; instead, it is those volumes that were next if time passed.

One important aspect of the use of delay premigration is that the content that is delayed for premigration is not added to the premigration queue until its delay criteria is met or it is selected as an early premigration candidate. That is, 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.

The rate at which the premigration content increases matches the rate at which the volumes were originally created or last accessed. The longer the delay, the less content is likely to still exist and thus less data is queued. Therefore, this rapid increase is more applicable to short delays.

Ensure that your FC 5274 and FC 5279 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 disk partitions and a unique SC for each one. The following SC definitions are available:

- ▶ 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 less host activity occurs. 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 at 10:00 in the morning.

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

Therefore, all 9 TB of the workload meets its delay criteria at roughly the same time, which produces a large increase in premigration activity. If the premigration queue size is not large enough, real-time workloads into the TS7700T might be throttled until the premigration process can reduce the premigration queue size. Ensure that the number of FC 5274 and FC 5279 features are sufficient or plan the delay times so that they do not all expire at the same time or when zero to minimal host activity occurs.

2.3.5 Removing or Migrating CPx content

After a volume that is contained within a CPx partition is premigrated to tape or cloud, you can determine when the CPx copy that remains is flushed/removed/migrated from TVC cache. This process is accomplished by assigning a preference group within the SC construct that is defined in the MI.

The following SC preference choices are available:

- 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 when they are copied to tape or cloud. This concept is called Preference Group 0 (PG0). This control is suitable for data that is unlikely to be read again because a recall is always required when accessed.
- Level 1** Copied volumes remain in the TVC until more space is required for a specific CPx partition and then volumes in that partition are removed from disk cache in a least recently used order. Before any Level 1 volumes are removed from the CPx partition, volumes in scratch that are also in the same CPx partition are removed first. Next, any premigrated Level 0 volumes in the same CPx partition are removed. Only if space is still needed are Level 1 volumes then removed based on LRU against other volumes in the same CPx partition.

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. If the TS7700 has no matching SC, the SC is auto-created by using the currently defined default SC.

For environments that are not z/OS (SMS) environments, the MI can be used to predefine an SC to a range of volumes when they were initially inserted. The MI can also be used to assign an SC to a range of volumes that were inserted.

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 assigned 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 that is based on the Preference Level 0 or 1 of the SC that is assigned to the volume.

When a preference level is assigned to a volume during mount/demount, 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/demount occurs, the new policy also takes effect.

The **LI REQ PARTRFSH** operation also acknowledges any preference group changes that occurred against the volume before the **LI REQ** operation.

2.3.6 Recalling volumes into CPx

When a logical volume is present on back-end tape or cloud only, a recall or staging back into TVC disk cache is required for any host-started mount. Whether the host plans to modify the volume or only read it, the entire contents of the logical volume must be staged back into the TVC within the assigned CPx partition before the mount completion status is surfaced to the host.

After a recalled logical volume is unmounted, the default behavior is to have its assigned preference group acknowledged. Preference Group 0 or 1 is used to determine whether the volume is retained in disk cache, or flushed or migrated from disk cache immediately.

An option exists that forces volumes that were recalled for read only purposes to be migrated or flushed from cache immediately as described next.

Prefer immediate migration for recalled volumes

Normally, a volume recalled into the TVC 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 the cache and moves to the end of the list of PG1 candidates to migrate because of how the LRU algorithm functions. The default behavior assumes any recall of a volume into the TVC might follow with more host access.

However, use cases might exist in which volumes that are recalled into cache are known to be accessed only once and should be removed from disk cache when they are read (for example, during a multi-volume data set restore). In this case, you do not want the volumes to be kept in cache because they use disk cache space that is more valuable to other resident friendly workloads.

Each TS7700T features 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 they retain their previously assigned preference group. 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 RECLPG0** option:

- ▶ When **DISABLED**, which is the default, logical volumes that are recalled into cache and read-only 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 and read-only 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.

2.3.7 TVC handling in outage situations

If a TVC disk cache becomes unavailable, the TS7700 quickly fences itself and moves into a non-operational state. This process is done to prevent any risk of data corruption or loss.

If an attached tape library or object store becomes unavailable, the TS7700T and TS7700C can remain operational, but function that is related to premigration and recall is paused for that particular cluster. A **PHYSLIB** configuration parameter is used to determine whether new inbound host operations are allowed on a TS7700T with a degraded TS3500 or TS4500 library. You can also determine whether the TS7700T ignores premigration throttling in such a situation.

The default TS7700T behavior is to stop inbound data acceptance when a TS3500 or TS4500 library becomes degraded. Therefore, review these settings if you want operations to continue. The LI REQ settings that govern this behavior are under the PHYSLIB keyword: TVCWDEG, PRETHDEG, and COPYWDEG.

For DS8000 Object Store support, inbound objects are not affected by the state of any back-end tape library or object store.

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

In a stand-alone cluster, you cannot define any Copy Consistency Points. The local and only cluster always contains a consistent copy of a volume, whether in its TVC disk cache, on back-end physical tape, or within an attached object store.

2.3.9 TVC selection in a stand-alone cluster

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

2.3.10 TVC encryption

The TS7760 and TS7770 support full AES256 disk encryption for all data at rest.

When TVC encryption is enabled, the entire TVC is enabled for encryption by using the same key-managed process. A supported method is not available to partially encryption portions of the TVC or to use different key management processes for different portions or content within the TVC.

When a TVC is encryption-enabled, any local or external key exchange occurs during the power on phase of each disk cache string that is attached to a TS7700.

Only the distributed version of IBM Security Key Lifecycle Manager is supported for TS7700 TVC disk encryption. As of this writing, the z/OS version of IBM Security Key Lifecycle Manager does *not* support TS7700 disk encryption.

If you use TS7700T with physical tape encryption, you must use the distributed version of IBM Security Key Lifecycle Manager if you also want to use external key management for TVC disk encryption. If you use only internal or local key management for TS7700 TVC disk encryption, the z/OS version of IBM Security Key Lifecycle Manager can be used for TS7700T physical tape encryption.

After encryption is enabled, any future FC 4017 Manufacturing Cleanup process that is run against the TS7700 provides the SSR an option to run a cryptographic erase. As part of the cleanup process, all data encryption keys are destroyed if this cryptographic erase option is enabled. A repurposing of the TS7700 at that point requires a complete file system rebuild.

TS7760 TVC Encryption

All TS7760 configurations contain self-encrypting disk drives or DDMS. Encryption can be internally or externally key managed. Encryption can be enabled retroactively and concurrently for any TS7760 in the field. Any TS7760 with internal key management enabled can also be concurrently upgrade to external key management.

Internal key management manages lock keys for each CSA disk string within the CSA controllers and the TS7760 server. These internally managed keys can be rotated through the MI or through an SSR-initiated process. These keys can also be optionally exported to DVD.

External key management is supported by using the IBM Security Key Lifecycle Manager (formerly IBM Tivoli Key Lifecycle Manager). Up to two IBM Security Key Lifecycle Manager servers are supported per TS7700. The TS7760 uses an IBM proprietary protocol for communications between the TS7760 and the IBM Security Key Lifecycle Manager server, which is optionally protected by using TLS 1.2.

When external key management is enabled, the lock keys that are used for each CSA disk string are managed entirely outside of the TS7760. No alternative method is available to unlock the CSA encryption keys without an exchange with an IBM Security Key Lifecycle Manager server. The externally managed keys can be rotated through the MI or an SSR-initiated process.

TS7770 TVC Encryption

All TS7770 configurations support encryption for data at rest. The TS7770 does not rely on self-encrypting disk drives or DDMs. Encryption instead occurs within the CSB/CFC disk cache controller by using hardware acceleration. Because of this approach differs from the TS7760 approach, encryption support on the TS7770 must be enabled or disabled at manufacturing time.

No method exists to enable encryption retroactively on a TS7770 after it is installed in the field. When enabled at manufacturing time, local key management is initially enabled. External key management can then be enabled after it is installed in your lab or later concurrently.

Local key management manages lock keys for each CSB/CFC disk string within USB flash drives that are connected to each redundant CSB/CFC subcontroller. Two backup USB flash drives also exist. That is, a total of four USB flash drives are used per CSB/CFC disk string (two USB flash drives plugged in and two for backup).

The lock key that is stored within the USB sticks is used to unlock the main encryption key, which is scattered among the DDMs of the string. Subsets of the DDMs contain only a random piece of the locked encryption key, which makes it non-valuable without the remaining DDMs. The portability of the USB flash drives allows for their removal if the TS7770 is at risk. For example, if your data center is being evacuated because of approaching weather, the TS7700s and CSB/CFC controllers can be powered down and their USB sticks removed.

External key management is supported by using the IBM Security Key Lifecycle Manager. Up to two IBM Security Key Lifecycle Manager servers are supported per TS7770. The TS7770 uses the KMIP protocol for communications between the TS7770 and the IBM Security Key Lifecycle Manager server, which is protected by using TLS 1.2.

When external key management is enabled, the lock keys that are used for each CSB/CFC disk string are managed entirely outside of the TS7770. No alternative method is available to unlock the CSB/CFC encryption keys without an exchange with an IBM Security Key Lifecycle Manager server. The externally managed keys can be rotated only through an SSR-initiated process. This limited rotation approach is used because of the need to swap out USB flash drives during the rekey process.

2.3.11 TS7700T Physical volume pools

When the TS7700T is used, you can group physical volumes into pools. By using policy management, you can then direct logical volumes to one or more pools.

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

- ▶ A method to separate customer data so that two or more customer's data do not exist on the same physical tapes.
- ▶ Customers want to “see, feel, and touch” their data by having only their data on dedicated media, which can be removed if needed.
- ▶ Customers need separate pools for different environments, such as test, user acceptance test (UAT), and production.
- ▶ Provides a method to track how many physical tapes are being used by a specific customer or workload.
- ▶ 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, so you can 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. This cascading allowed the most archived data to end up on media that is less likely to require reclamation.
- ▶ A second dedicated pool for key workloads to be Copy Exported.

Because the use of physical volume pools includes many benefits, plan for the number of physical pools. For more information, see “Relationship between reclamation and the number of physical pools” on page 57.

Using physical volume pools

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

Each 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. For pools 01–32, you set it up to borrow scratch stacked cartridges from the common scratch pool (Pool 00) if a scratch shortage occurs. This setup can be done on a temporary or permanent basis.

Each pool can be defined to borrow a single media type (for example, JA, JB, JC, or 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.

Note: The common scratch pool must have at least three scratch cartridges available to prevent low scratch warnings.

General-purpose pools (Pools 01–32)

A total of 32 general-purpose pools are available for each TS7700T cluster. These pools can contain empty, full, and filling stacked volumes. All physical volumes in a TS7700T cluster are distributed among available pools according to the physical volume range definitions within the MI.

The distribution is also based on the pools' borrow and return attribute settings. For example, you can have pools 01–32 be the target of a specific physical volume range to have a specific volume that is used for that pool. In this use case, the targeted pool should feature borrow/return disabled so that scratch volumes are not returned to the common scratch pool 00.

A pool's properties can be 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 56.

By default, one pool (Pool 01) is available and the TS7700T stores all newly created virtual volumes on any stacked volume that is available to it. This configuration creates an intermix of logical volumes from different sources.

The user cannot influence the physical location of the logical volumes within a physical pool, Defining more than one pool can be valuable for the following reasons:

- ▶ Need to separate different clients or LPAR data from each other
- ▶ Need to separate media types
- ▶ Set up specific pools for Copy Export
- ▶ Set up pool or pools for encryption with unique encryption key labels
- ▶ Set a reclamation threshold at the pool level to match the reclamation needs of the data that is there
- ▶ Set up reclamation cascading from one pool to another so that aged data ends up in pools with less reclamation activity
- ▶ Set different numbers of physical back-end devices for premigration for different workloads
- ▶ Assign specific physical volser ranges to specific workloads

Physical pooling of stacked volumes is identified through pool numbers, as shown in Figure 2-5.

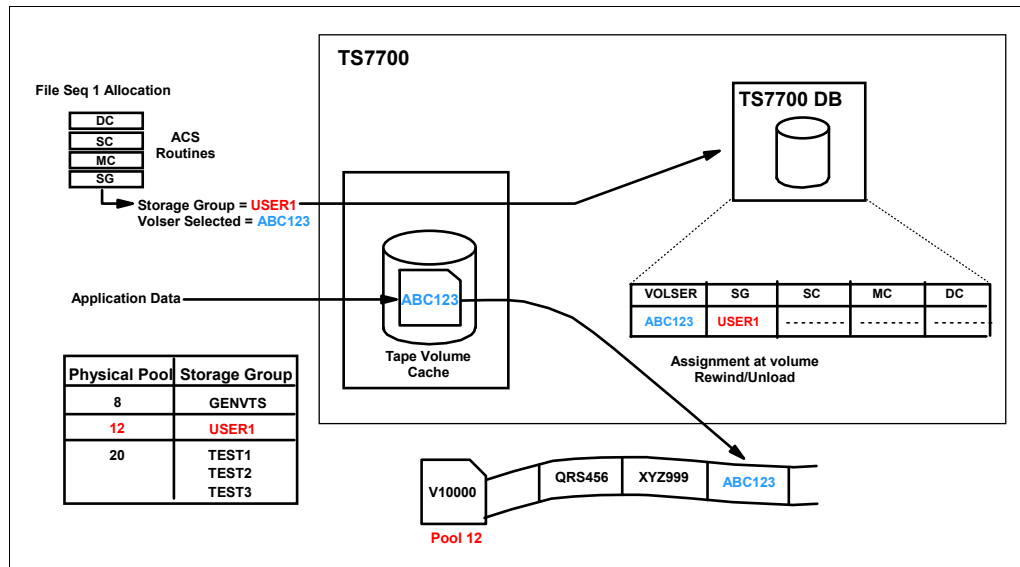


Figure 2-5 TS7700T Logical volume allocation to specific physical volume pool flow

Through the MI, you assign which primary pool is used for a workload within the SG attributes. During a scratch mount, a logical volume is assigned to a selected SG and the assigned pool for that SG is targeted. 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 to target a specific home pool at insert time. Changing the home pool of a range has no effect on the volumes that are in the library. For pools with borrow/return enabled, target the common pool 00. For pools where borrow/return is disabled, set up the range table to include volumes that specifically target that pool. Each pool that contains any amount of active data must always have at least two scratch-stacked physical volumes to prevent low scratch warnings.

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

Borrow and return: Physical Scratch Management

By using the concept of borrowing and returning, physical scratch distribution can be automatically addressed.

With borrowing, stacked volumes can move from pool to pool based on each pool's assigned media type and media needs. When all borrow/return enabled pools with active data contain at least two scratch-stacked volumes, any excess physical scratch is returned to the common pool 00. You need at least two empty stacked volumes in pool 00 and any pool that contains active data. If borrow/return is not enabled for a specific pool, you must manage its scratch by using the physical VOLSER ranges and inserting more tapes when counts go below three.

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

Note: Pools that feature borrow/return enabled and contain no active data eventually return all 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 selection can be used to create separate pools of 3592 tape cartridges with various capacities of 128 GB–10 TB, depending upon the type of media and tape drive technology that is 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.

2.3.12 Logical and stacked volume management

Whenever a logical volume is created or modified, any data from a previous use of this logical volume that is on a stacked volume becomes obsolete. The new or modified virtual volume is placed in the cache and written to a new location on a stacked volume in the appropriate pool. The previous copy on a stacked volume is invalidated but it still uses up space on the physical tape. Auto deleting expired scratched volumes or ejected volumes can also result in the creation of inactive space on physical tapes. This inactive space can be reclaimed by using reclamation.

Physical volume reclamation

The reclamation process periodically checks for active volume usage percentages for each fully stacked cartridge within each pool. As data is removed or expired, the active data value decreases. When it drops below an optional threshold, reclamation occurs, which allows all active data that remains to be copied to new physical stacked tapes. The now emptied tape is returned to a scratch status.

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

- ▶ A host 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 assigned the logical volume to a scratch category. The category features a nonzero delete-expired data parameter value. The parameter value was exceeded and the TS7700T deleted the logical volume.
- ▶ A host modified the contents of the volume. This modification 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 TS7700T tracks the amount of active data on a physical volume in increments of 0.1% and always rounds up. 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 TS7700T tracks is 1/10 of 1%, it rounds down, so even 1 byte of inactive data drops the percentage to 99.9%.

Physical volume reclamation details

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 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 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 through the disk cache. Instead, it is transferred directly from the source to the target tape cartridge by using only memory that is contained within the TS7700 server. 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 pool and is expressed in a percentage (%) of tape usage.

Volume reclamation can be followed 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 2.3.13, “Secure Data Erase function” on page 58.

Because reclamation requires the use of physical drives and TS7700T resource, you can schedule when it cannot run. Consider *not* running reclamation during peak workload hours of the TS7700T, which ensures that recalls and migrations are not delayed because of physical drive shortages. Choose the best period for reclamation by considering the workload profile for that 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.

It is 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 must be available. For more information, see 10.1.1, “Physical Volume Pools” on page 450.

Relationship between reclamation and the number of physical pools

The reclaim process is done on a per pool basis and each reclamation process needs two drives. Defining too many pools can lead to a situation in which the TS7700T cannot process the reclamation for all pools in an appropriate manner. Eventually, pools can run out of space (depending on the borrow definitions), or more stacked volumes are needed 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 that run concurrently with the **LI REQ** setting. That way, you can prevent too many drives from being used at any time for reclamation purposes.

2.3.13 Secure Data Erase function

When physical tape is present, one concern is the security of stale data that was on the cartridges. The TS7700T provides a pool level method to run a secure erase against physical tapes before they were moved back into a scratch state.

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, which makes all data that was present now inaccessible through traditional read operations.

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. This process provides a method to have volumes run through reclamation based on the age of expired data versus only the amount of expired data.

When a physical volume is encrypted, the TS7700T can run a fast erase of the data by erasing the EKs on the cartridge. It erases only the portion of the tape where the encryption 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. 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, which ejects them before completing the erase.

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

Volumes that are returned to the TS7700T that were offset reclaimed as part of a copy export operation also have the erase run on them before returning them to a scratch status.

2.3.14 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 seconds. If you do not require such low recovery times for all or a mixture of your workload or you require another PIT copy for air gap purposes, a function that is called *Copy Export* is available for the TS7700T.

The Copy Export function enables a copy of selected logical volumes that are written to secondary physical tape pools within the 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 volumes remain accessible to the TS7700T.

The following logical volumes are excluded from the export:

- ▶ Volumes that are mounted during any portion of the export process
- ▶ Volumes that cannot create a valid primary and secondary pool copy
- ▶ Volumes that did not complete replication into the exporting TS7700T at the start of the export process

These excluded volumes are 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.

An offsite reclamation process runs against copy-exported stacked volumes. This process does not require the movement of physical cartridges. Rather, non-expired logical volumes are written to a new secondary copy-exported volume and the original copy-exported stacked volume is marked invalid. For more information, see 15.1.5, “Reclaim process for Copy Export physical volumes” on page 794.

2.3.15 Cloud export function

With R5.1 and later, a point-in-time backup of the internal database can be stored in the cloud. An empty TS7700C can then be used as a restore point-in-case of disaster recovery. The export is started by using the Library Export IBM Z operation (such as Copy Export).

The backup is not an “extra copy” of the data in the cloud. Instead, it is a snapshot of the database that points to all the data that existed or was recently put into the cloud during export operation.

For more information, see Chapter 17, “Cloud Storage Tier export, recovery, and testing”, in *IBM TS7700 R5.4 Cloud Storage Tier Guide*, [REDP-5573](#).

2.3.16 Encryption of physical tapes

The importance of data protection became 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:

- ▶ An external key manager
- ▶ The TS1160, TS1150, TS1140, TS1130, and TS1120 encryption-enabled tape drives
- ▶ The TS7700T

Encryption key manager

For physical tape, the TS7700 can use one of the following encryption key managers:

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

In this publication, the general term *key manager* is used for all EK managers.

Important: Consider the following points:

- ▶ IBM Security Key Lifecycle Manager for z/OS does not support external key management for TS7700 disk cache storage. Therefore, if you plan to enable external key management for the TS7700 disk cache, you must use IBM Security Key Lifecycle Manager for distributed systems for disk cache and physical tape.
- ▶ *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 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 4.3.18, “Planning for tape encryption in a TS7700T” on page 198.

The TS1160, TS1150, TS1140, TS1130, and TS1120 encryption-enabled tape drives

The IBM TS1160, TS1150, TS1140, TS1130, and TS1120 tape drives provide hardware that performs the encryption without reducing the data transfer rate.

The TS7700T Physical Tape Encryption

The TS7700T provides the means to enable encryption at physical tape pool granularity. 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.

Physical tape encryption on the 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 TS7700T. Storage pool encryption parameters are configured through the TS7700T MI under Physical Volume Pools.

For encryption support, all drives that are attached to the TS7700T must be Encryption Capable and encryption must be enabled. If 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* that are defined at the key manager server:

► Specific EKs

Each pool that is defined in the TS7700T can have its own unique EK. As part of enabling a pool for encryption, you must enter two key labels for the pool and an associated key mode. The two key labels might not be the same. Two key labels 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, which allows different pools to use different underlying encryption keys.

The second key is used for data exporting or sharing where an external party is granted access to the secondary key without the need to share the primary key because either of the keys can be used to decrypt the data later.

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 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 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

All logical volumes that are directed to a pool that is enabled for encryption are 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 was last mounted.

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 TS7700T determines whether a new physical volume must be mounted.

If a new cartridge is required, the TS7700T directs the drive to use encryption during the mount process. The TS7700T also provides the drive with the key labels that are specified for that pool. When the first write command is received by the drive, a connection is made to a key manager through the TS7700T and the key that is needed to perform the encryption is obtained. After the first write occurs, the same physical tape requires its drive to communicate with the key server for every mount.

Any partially filled physical volumes continue to use the encryption settings that are 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 2-6 shows that the method for communicating with a key manager is through the same Ethernet interface that is used to connect the TS7700T to your network for access to the MI.

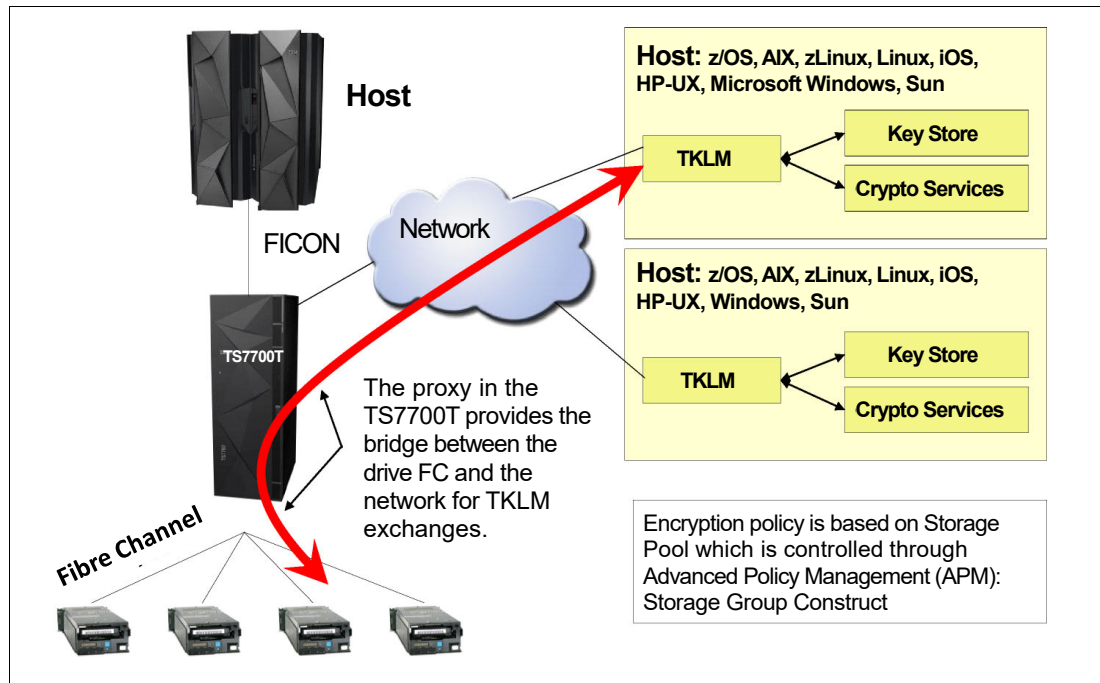


Figure 2-6 TS7700T physical tape encryption

The request for an EK is directed to the IP address of the primary key manager. Responses are passed through the 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 must be read from a physical volume in a pool that is enabled for encryption (as part of a recall or reclamation operation), the TS7700T uses the key manager to obtain the necessary information to decrypt the data.

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

Remember: If you want to use external key management for cache and physical tapes, you must use the same external key manager instance that must be the distributed systems version of IBM Security Key Lifecycle Manager.

2.3.17 User Management: Roles and profiles

The TS7700 offers you internal user management and external user management through LDAP support. You can use this user management to define user IDs and assigned roles. The role identifies the access rights for each user. You can use this method to restrict the access to specific tasks for each user or group of users. The most minimal role that is supported is a read-only mode that allows access to panels without the ability to make changes.

Consider restricting access to specific panels so that only specific users are allowed to change critical settings.

2.3.18 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), the TS7700 clusters and TS3000 System Console (TS3000 TSSC) include native support for an LDAP.

After 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 appropriate roles that are defined to log in to the TS7700.

If your LDAP server is not available, you cannot 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.

IBM RACF® can also be used to control the access when LDAP is contained within IBM Z, which means that users are defined to RACF and verified through a RACF database. Roles and profiles must still be maintained because the RACF database provides password authentication only.

When enabling LDAP, two optional settings in the MI enablement page for LDAP allow a local SSR or IBM second-level remote support to bypass LDAP.

For more information about the implementation of LDAP/RACF, see [IBM TS7700 LDAP Security - Introduction to Access Management](#).

2.3.19 Service preparation mode

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

2.3.20 Service mode

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

2.3.21 Control Unit Initiated Reconfiguration

The CUIR function is available in a multi-cluster grid only.

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

Multi-cluster grids are combinations of two to eight clusters that work together as one logical entity. All clusters in a grid can be of the same type or they can be mixed. For example, a TS7700D, TS7700T, and TS7760C can be combined as a *hybrid grid*. 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 2-7 shows a four-cluster hybrid grid. The configuration consists of two TS7720s and two TS7740s. For more information about other examples, see 2.5, “Grid configuration examples” on page 106.

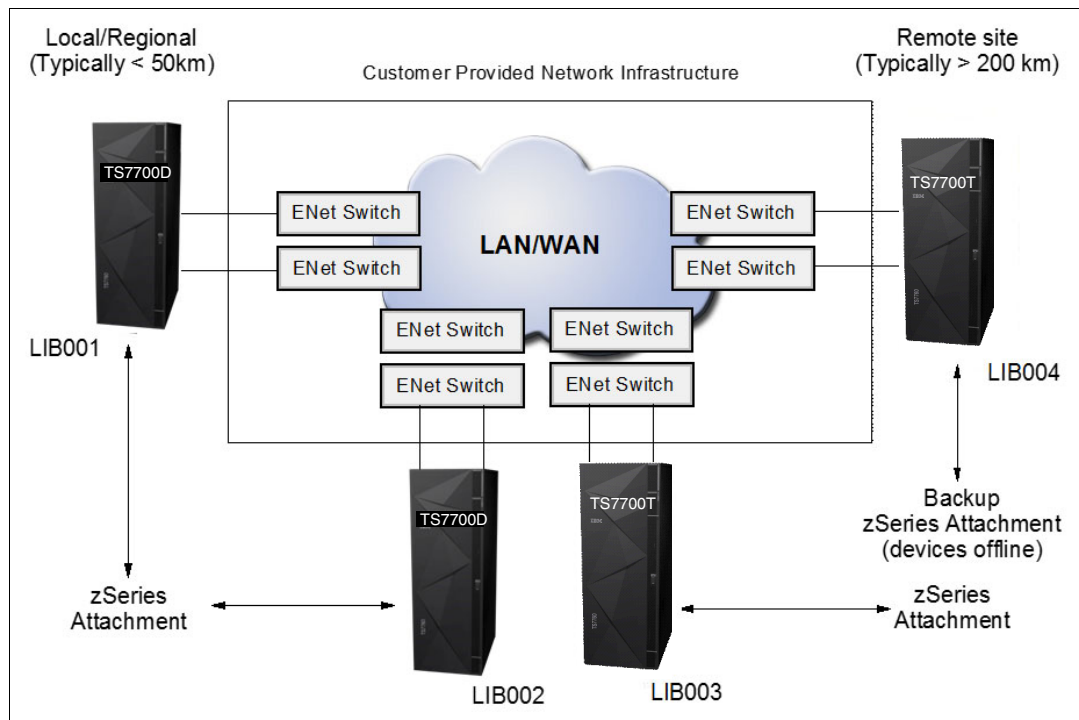


Figure 2-7 TS7700D 4-cluster grid

2.4.1 Rules in a multi-cluster grid

In a multi-cluster grid, the following general rules apply:

- ▶ A grid configuration resembles a single tape library and tape drives to the attached hosts.
- ▶ The grid is a composite library with underlying distributed libraries.
- ▶ Up to eight clusters can be contained in a single grid.

- ▶ Each TS7700 emulates up to 496 unique logical devices that provide access to all logical volumes that are contained within any cluster in the grid.
- ▶ All TS7700 models can coexist in a grid. If a mixture of cluster types exists (for example, TS7700T, TS7700D or TS7700C), it is referred to as a *hybrid grid*.
- ▶ If one cluster is not available, the grid still continues to operate.
- ▶ Clusters can be grouped into cluster families.
- ▶ Scratch and private mounts can be satisfied from any cluster in the grid.
- ▶ Synchronous and asynchronous replication can occur to one or more peers at volume granularity.

In a multi-cluster grid, the following 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 in different clusters in the grid up to the number of clusters in the grid.
 - Each TS7700T cluster in the grid can store dual copies on physical tape.
 - Each TS7700C cluster can contain a copy in an on-premises or off-premise object store.
- ▶ 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 if at least one cluster is available that contains a valid copy of the logical volume.
- ▶ Any logical volume (replicated or not) is accessible from any other cluster in the grid.
- ▶ Each distributed library can access any logical volumes within the composite library.
- ▶ Access to specific clusters is determined by which clusters have devices varied online to IBM Z or by how Scratch Allocation Assist is configured.
- ▶ A virtual Volume that is copied into the Cloud can be accessed by any other cluster in the Grid, even recently added clusters.

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

2.4.2 Required grid hardware

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

- ▶ Each of the TS7700 clusters must have the Grid Enablement feature installed.
- ▶ Each of the TS7700 clusters 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. Not all clusters require the same number of links.

Grid enablement

FC 4015 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 features two Ethernet adapters that are connected to the TCP/IP infrastructure. Depending on which features are installed, the two adapters provide four total 1 Gb copper links or four 10 Gb longwave fiber links. The speed of this network has a major performance effect on the entire grid.

Previous TS7700 VEC models might also contain four 1 Gb shortwave fiber links, although they are no longer available for new TS7700 models. For more information, see 7.1.1, “Common components for the TS7700 models” on page 262.

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 or synchronous mode 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, which results in a host job failure or a Synchronous mode copy break. When it is a synchronous mode copy break, the job continues in the asynchronous deferred state or fails, depending on the configured MC settings.

Secure Data Transfer

The TS7700 grid supports the ability to encrypt logical volume data, which is transferred across the grid network. When the Secure Data Transfer feature is installed on two or more clusters, the clusters use TLS 1.2 to create an AES128 or AES256 encrypted connection. Only clusters that are running release level 5.0 or later include the secure data transfer feature and have it enabled can communicate by using encrypted links. The AES128 or AES256 encryption uses the IBM Power8 and IBM Power9 encryption instruction set that provides hardware acceleration with minimal to no overhead.

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 by using the configured time of the lowest two cluster indexes in the grid. An external time source (such as an NTP server) is not needed, even in a grid with large distances between the clusters. However, an NTP server is supported. The lowest two configured cluster indexes can optionally synchronize with the NTP server and the remaining peers (if any) synchronize with the lowest two clusters.

Because all clusters are configured to Coordinated Universal Time, the time that is presented from the grid (VEHSTATS, LIBRARY REQUEST, and others) might not show the same time as your LPARs, which can lead to some confusion during problem determination or for reporting because the different timestamps do not match. The MI adjusts timestamps based on the time zone of the web client.

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 (see Figure 2-8).

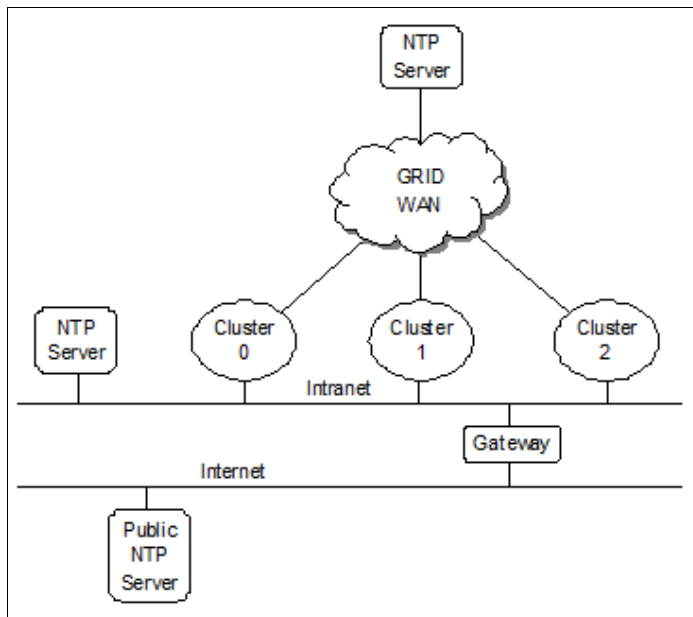


Figure 2-8 Time coordination with NTP servers

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 should be able to communicate with the same NTP server that is defined in VPD. At a minimum, the lowest two clusters by index must communicate with the NTP server. In the absence of an NTP server, all clusters coordinate time with Cluster 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.

2.4.3 Data integrity by volume ownership

In a multi-cluster grid, only one cluster at a time can modify volume data or attributes, which is managed through the concept called “Volume Ownership.

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 that is 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 must mount a logical volume on one of its virtual devices and it is not the current owner of that volume, it must obtain ownership first. When requested, the current owning TS7700 cluster 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 are in a grid and if 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.

If a cluster is not reachable, clusters do not automatically assume or take ownership of a logical volume without being directed. This process can be done manually or be automated by using Autonomic Ownership Takeover Manager (AOTM). Service outages include 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 not enabled or the failure cannot be managed by AOTM.

Note: Ownership can also be transferred manually by using the `LI REQ,OTCNTL` command for special purposes. For more information, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

For more information, see 5.1.3, "Volume ownership" on page 216.

To support the concept of ownership, a token exists per logical volume.

Tokens

Tokens are used to track changes to the ownership, data, or properties of a logical volume. Key token attributes are mirrored to each cluster that participates in a grid and represent the current state of the logical volume. Tokens feature the following characteristics:

- ▶ Every logical volume includes a corresponding token.
- ▶ The grid component manages the tokens.
- ▶ Tokens are maintained in an IBM Db2® database.
- ▶ Each cluster's Db2 database includes 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 maintaining its 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 in 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 a cluster outage occurs. When AOTM is not enabled or cannot proceed, you must enable takeover manually by using the MI or `LI REQ` command.

The following options are available through the MI:

► **Read-only Ownership Takeover**

When Read-only Ownership Takeover is enabled for an unavailable cluster, ownership of a volume is taken from the unavailable TS7700 cluster when the volume is accessed. Only read access to the volume is allowed through the remaining available TS7700 clusters 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 takeover mode for the unavailable cluster remains in place until another mode is selected or the unavailable cluster is restored. Only one available cluster must enable the takeover mode because all peers share the enablement after it is enabled.

► **Write Ownership Takeover (WOT)**

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

The mode for the unavailable cluster remains in place until another mode is selected or the failed cluster is restored. If volumes were taken over with the read-only takeover mode, enabling write-ownership takeover allows any previous takeovers and future takeovers to enable updates.

Scratch mounts continue to prefer volumes that are owned by the available clusters. Only after all available candidates are 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 or LI REQ.

More than one cluster can have ownership takeover that is enabled. For example, if two clusters both become unavailable, ownership takeover must be enabled against both clusters individually so that the remaining cluster or clusters can access volumes that are owned by the two unavailable clusters.

Important: You cannot change the takeover mode of a cluster in service.

For more information about an automatic takeover, see 2.4.33, “Autonomic Ownership Takeover Manager” on page 97.

2.4.4 I/O Tape Volume Cache selection

All vnodes in a grid feature 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 mount is called a *remote mount*. The TVC is then accessed by the grid network. Selecting the TVC can be influenced by using several methods.

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
- ▶ Dynamic Allocation Assist (DAA) for specific mounts
- ▶ Override settings
- ▶ Cluster family definitions
- ▶ Disk cache residency and recall times
- ▶ Network latency times

2.4.5 Copy consistency points

In a multi-cluster grid configuration, several policies and settings can be used to influence the location of logical volume copies and when the copies occur.

Consistency point management is controlled through the MC storage construct. By 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?
- ▶ Do you want some copies to occur immediately while others are deferred?

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.

The following consistency point settings are available:

Sync As data is written to the volume, it is compressed (if enabled) 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 process provides a zero RPO at tape sync point granularity.

Tape workloads in IBM Z environments assume sync point hardening through explicit sync requests or during close processing, which allows 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.

If one of the two sync points cannot be maintained, options exist that allow the failed copy to be made later asynchronously or to have the job fail.

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 older generation 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 process is also called <i>Asynchronous</i> 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 made to those sites with a Time Delayed consistency point. For Time Delayed, you can specify whether the specified time begins after creation or after each access in the MC
No Copy	No copy is made.
Exists	This policy is not a configurable policy but one that might result after an MC change. When a cluster’s consistency point changes from any consistent copy type to “No Copy”, the current copy if it exists by default is not automatically deleted. The copy remains and the consistency point is referred to as “Exists”.

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. When “Retain Copy Mode” is enabled, only the first write to a new output tape causes the copy policy to be updated.

For more information, see 2.4.5, “Copy consistency points” on page 70.

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

2.4.6 Cluster family concept

When two or more clusters are near each other, certain operations can be made more efficient. For example, TVC selection can choose a much more local cluster over a far more remote cluster, or copies to remote clusters can pass data across distant links after allowing far remote adjacent clusters to share later.

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 another 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. Then, the two remote clusters can share the data by way of local replication.

Also, when a cluster determines where to source a volume, it gives higher priority to a cluster in its own 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 mount I/O and copy source TVC selection.

Figure 2-9 shows how *cooperative replication* occurs with cluster families. Cooperative replication is used for Deferred copies only. When a cluster must 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. In this way, the volume travels only once across the long grid distance.

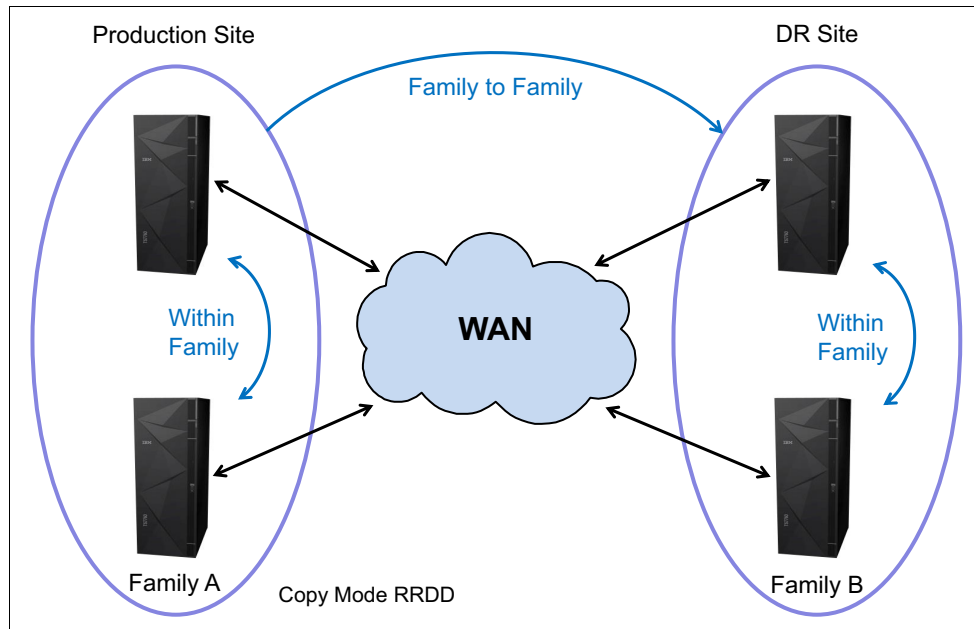


Figure 2-9 Cluster families

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 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. After copies for volumes from outside the family into the family are finished, the family members replicate among themselves.

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 received their copy. With families, a cluster's responsibility to keep the volume in the cache is released after all families that need a copy have it. This process enables PGO volumes in the source cluster to be removed from cache sooner. One of the target family members then retain a copy in the disk cache until all peers in the family that require a copy completed their copy process.

Another benefit is the improved TVC selection in cluster families. For cluster families that use cooperative replication, the TVC algorithm favors the use of a family member as a copy source. Lastly, 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](#).

2.4.7 Override settings concept

With the TS7700, you can define and set the optional override settings that influence the selection of the I/O TVC and replication responses by using the MI.

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

TS7700 overrides I/O TVC selection and replication response

The settings are unique per cluster, which means that each cluster can have separate settings, if wanted. The settings take effect for any mount requests that are received after the settings are changed. Independent of which MC is used, all mounts use the same override settings. Mounts that are in progress are not affected by a change in these 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 or in the cloud, a recall occurs versus the use of 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 that is based on the definitions of the MC that is associated with the volume that is 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 includes 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 made to the local TVC before mount completion. The copy source can be any participating TS7700 in the grid.

In a TS7700T or TS7700C, the logical volume might have to be recalled from a stacked cartridge or cloud. If for any reason the vnode cluster cannot act as the I/O TVC, the 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 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 process 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*.

► 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 source volumes are preferred in copy priority at target clusters so that the source cluster can migrate the volumes sooner. When this option is checked, preference group 0 and one volume are treated 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. Other requirements might exist, such as DR testing, where all I/O must go only to the local TVC to ensure that the correct copy policies are used and that data is available where required.

In certain GDPS use cases, you can 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.

2.4.8 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. However, each of the multiple TS7700s must have a unique distributed library that is defined. The host must be configured 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 were upgraded to support 496 drives.

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

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

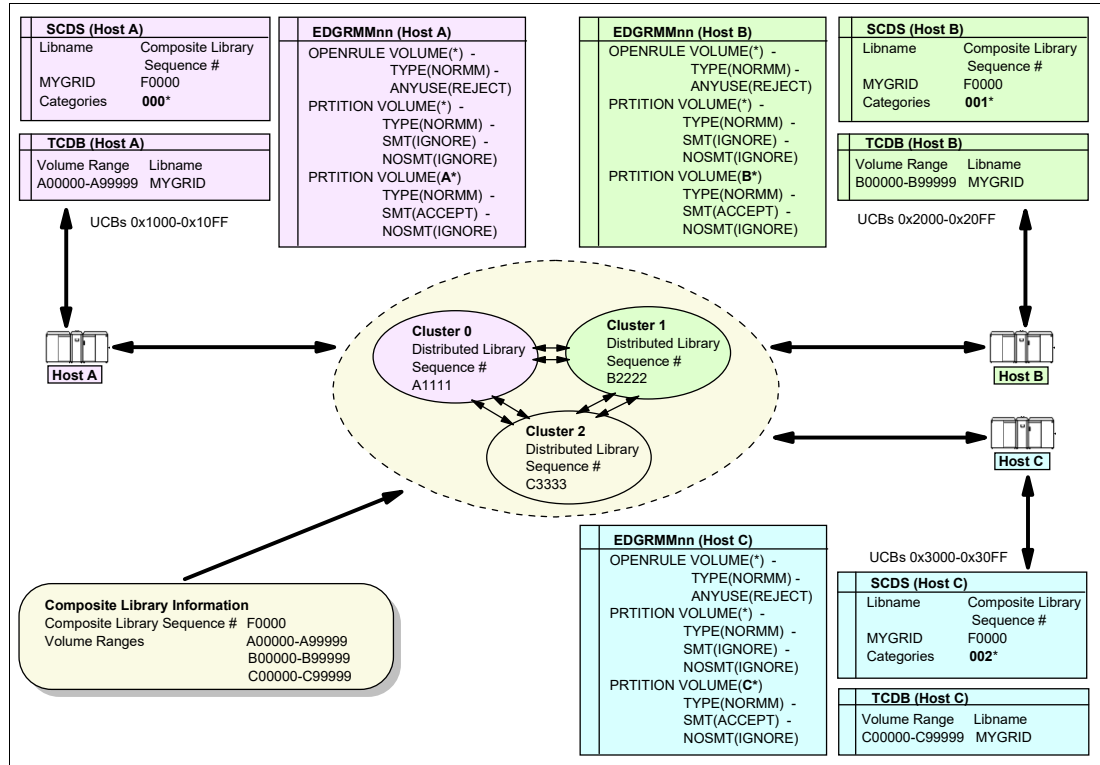


Figure 2-10 TS7700 3-cluster grid configuration

Distributed library

Each cluster in a grid is a distributed library, which consists of any model TS7700. Each distributed library can have up to 31 3490E tape control units per cluster. Each control unit emulates 16 IBM 3490E tape drives and is attached through up to eight FICON channel attachments per cluster. However, the virtual drives and the virtual volumes are associated with the composite library.

No difference exists for a stand-alone definition, except that the composite library consists of only one distributed library.

2.4.9 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 2.4.20, “General TVC management in multi-cluster grids” on page 81, and 2.4.25, “Copy Consistency Points: Copy policy modes in a multi-cluster grid” on page 88.

2.4.10 Virtual volumes and logical volumes

No difference exists between multi-cluster grids and stand-alone clusters. The composite library, whether one distributed library or up to eight distributed libraries, is limited to 4,000,000 logical volumes.

Note: Consider the following points:

- ▶ A cluster is initially configured to support 1,000,000 logical volumes. Support for extra logical volumes can be added in increments of 200,000 volumes by using FC 5270.
- ▶ All clusters in a grid must have the same quantity of installed instances of FC 5270 configured. If you configured a different number of FC 5270s in clusters that are combined to a grid, the cluster with the lowest number of virtual volumes constrains all the other clusters. Only this number of virtual volumes is then available in the grid.

2.4.11 Mounting a scratch volume through specific clusters

Through MC, scratch allocation assist can be used to choose which clusters are candidates for scratch allocations, which use that MC. For more information, see “Scratch allocation assistance” on page 80.

2.4.12 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 other assistance, the mount uses the TVC I/O selection process to locate a valid version of the volume.

The following scenarios are possible:

- ▶ A valid copy exists 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.
- ▶ No valid copy exists in the TVC of the cluster where the mount is placed. In this case, the following options are available:
 - Another cluster has a valid copy in the 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 the disk cache. In this case, a TS7700T/TS7760C CP1–CP7 is chosen to recall the volume from physical tape or cloud object store 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 a cache, or on a stacked volume or cloud. The mount fails. Clusters in service preparation mode or in service mode are considered inactive.

To optimize your environment, DAA is used in IBM z/OS configurations automatically to choose the TS7700 that can provide the most ideal TVC selection. For more information, see “Device allocation assistance” on page 79.

If you do not specify the Retain Copy Mode, the copy policies from the mount cluster are refreshed at mount time. When clusters do not agree on MC settings, this issue can result in more copies being made that were not intended because the consistency points can change with each mount.

DAA often prevents such conditions, but it is still recommended to enable Retain Copy Mode when peers do not have identical MC settings. Existing copies are not removed if they are present. Remove any non-required copies by using the **LIBRARY REQUEST REMOVE** command.

If a virtual volume is modified during a specific mount operation, it is premigrated to back-end tape or cloud (if present) and has all copy policies acknowledged. The virtual volume is transferred to all defined consistency points.

If no modification occurs while the volume is mounted, the TS7700 does not schedule another copy operation and any current copy of the logical volume on a stacked volume or in the cloud remains active. Furthermore, copies to remote TS7700 clusters are not required if modifications were not made

2.4.13 Logical WORM support and characteristics

No differences exist between LWORM in multicluster and stand-alone cluster environments. The LWORM behavior for any logical volume is consistent across all clusters in the grid.

2.4.14 Virtual drives

From a technical perspective, no differences exist 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 Default number of maximum virtual drives in a multi-cluster grid

Cluster type	Number of maximum virtual drives
Stand-alone cluster	256
Dual-cluster grid or two-cluster grid	512
Three-cluster grid	768
Four-cluster grid	1024
Five-cluster grid	1280
Six-cluster grid	1536
Seven-cluster grid	1792
Eight-cluster grid	2048

With FC 5275, you can add one LCU with 16 drives up to the maximum of 496 logical drives per cluster. This addition results in the maximum number of virtual drives that are listed in 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 and two-cluster grid	992
Three-cluster grid	1488
Four-cluster grid	1984
Five-cluster grid	2480
Six-cluster grid	2976
Seven-cluster grid	3472
Eight-cluster grid	3968

To support this number of virtual drives, specific authorized program analysis reports (APARs) are needed to install the appropriate program temporary fixes (PTFs). For more information about identifying required PTFs see 8.2.1, “Field frame replacement migration for TS7700T” on page 309.

2.4.15 Device Allocation and Allocation Assistance

The load-balancing algorithm in the z/OS host has two options EQUAL (the default before z/OS 3.1) and BYDEVICES (the default starting with z/OS 3.1). With the TS7700 potentially having hundreds of devices available for an allocation request, the recommendation is to use BYDEVICES to obtain a better distribution of the mount requests across clusters within a grid and across grids. BYDEVICES can be specified in the ALLOCxx parmlib member that uses the TAPELIB_PREF setting TAPELIB_PREF(BYDEVICES). BYDEVICES takes the number of online and available devices into account as it randomizes across the eligible devices resulting in a distribution of mounts that better reflects the devices that are available to each system.

So, if cluster 0 has 128 devices online to SYSTEM1 and cluster 1 has 256 devices online to SYSTEM1 (a total of 384 eligible devices), over time, cluster 1 should receive approximately 2/3 of the scratch allocations. The BYDEVICES algorithm also better distributes the workload for specific requests across the eligible clusters. For more information, see the z/OS DFSMS Planning, Installation, and Storage Administration Guide for Tape Libraries [here](#).

With the ALLOCxx TAPELIB_PREF(EQUAL) setting (the default before z/OS 3.1) in z/OS parmlib, devices up to the first 7 LCUs only per cluster in the control block that is created internally in z/OS during IPL are considered for device selection in DAA, but devices exceeding 7 LCUs are not considered for DAA and are selected randomly. Therefore, if devices exceeding 7 LCUs are defined and devices within 7 LCUs are not available at the time of allocation and other devices are used, device selection in the DAA is not considered. This control block is determined at IPL (or during dynamic IO configuration changes), but the order of its LCUs is undefined and not user-controllable.

Scratch and private allocations in a z/OS environment can be more efficient or more selective by using the allocation assistance functions that are 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 in all TS7700 clusters by default. 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 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.

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 mount is requested. DAA returns to the host a ranked list of clusters (the preferred cluster is listed first) where the mount should be run.

The selection algorithm orders the clusters in the following sequence:

1. Clusters with the highest Copy Consistency Point.
2. Clusters that have the volume already in the cache.
3. Clusters in the same cluster family.
4. Clusters that have a valid copy on tape.
5. 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, remote mounts and recalls can still occur (even if DAA is enabled).

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 the 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, the use of the copy policy overrides to *Prefer local TVC for Fast Ready mounts* provides the best overall performance.

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 for TS7700T/TS7700C configurations. However, this performance gain includes a reduction in the effective grid cache size because multiple clusters are maintaining a copy of a logical volume in the disk cache. To regain the same level of effective grid cache size, an increase in physical cache capacity might be required.

DAA (JES2) 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, different clusters might have a benefit over others for certain workload types. For example, OAM or DFSMSHsm Migration Level 2 (ML2) migration might favor a large TS7700D disk-only cluster over a smaller configured TS7700T/C cluster. Or, an archive workload might favor a TS7700T or TS7700C cluster so that the data can be offloaded quickly to back-end tape or cloud.

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 2-11 shows the process of scratch allocation.

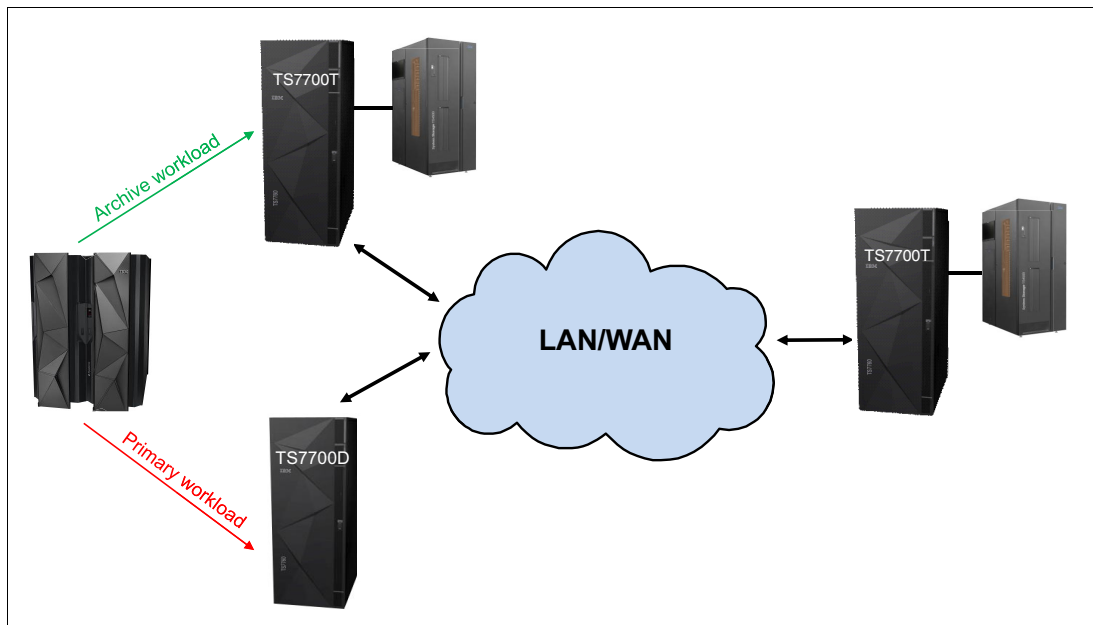


Figure 2-11 Scratch allocation direction to preferred cluster

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 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.

If in the Management Class defines only unavailable clusters are provided, the mount is processed. However, the job is still unable to run because the selected TVC is unavailable. You see CBR4000I and CBR4171I messages and receive a CBR4196D for a reply. This condition can occur if clusters become available from a grid perspective, but their devices are not yet varied online to the host. It is best to disable SAA *before* a scheduled outage where the last candidate cluster in a SAA MC configuration is expected to enter a service state.

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 extra routines that are needed to obtain the candidate list of mount clusters from a specific 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.

2.4.16 Selective Device Access Control

No difference exists 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.

2.4.17 Physical drives

In a multi-cluster grid, each TS7700T can have different physical back-end drives, media types, and Licensed Internal Code levels.

2.4.18 Stacked volumes

No difference exists between stacked volumes in multicluster and stand-alone environments.

2.4.19 Selective Dual Copy function

The Selective Dual Copy function is often used in stand-alone clusters. However, you can also use it in a multi-cluster grid. No difference exists in its usage in a multicluster and a stand-alone environment. For example, you can still use copy export from one member of the grid where selective dual copy is configured for one or more physical volume pools.

2.4.20 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 a 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 must manage which clusters receive copies. This process is done by using copy policies (how many copies of the logical volume must 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 larger effective capacity without losing availability or redundancy. For example, in a four-way configuration, one local and one remote cluster can contain a copy while the other two do not. Approximately 50% of your data is in one local or remote pair and the other 50% is in the adjacent pair. In addition, any logical volume that is stacked in physical tape or in the cloud can be recalled into TVC, which makes them available to any cluster in the grid.

Replication order

Copies are queued at target clusters in the following order:

- ▶ RUN copies
- ▶ Synchronous Deferred copies
- ▶ RUN deferred copies
- ▶ Deferred Copies
- ▶ Family Deferred Copies
- ▶ COPYRFSH Copies

TS7700T and TS7700C Replication Handling

Logical volumes that must be replicated to one or more peer clusters are retained in the disk cache, regardless of their preferred group assignments. This feature 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 groups that are 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, which enables data that is not yet replicated to be migrated to tape. These volumes likely must be recalled into the disk cache later for replication to complete. The migration of not yet replicated data might be expected when replication is not completing because of an extended outage within the grid.

Volumes that are written to an I/O TVC that is configured for PG0 have priority by default for replication to peers. This priority is to help the source cluster flush this data from its cache as quickly as possible. This behavior overrides a pure FIFO-ordered queue. A new setting is available 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.

2.4.21 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. Each cluster is limited to a maximum number of deletions per hour. The default is 1,000 per hour, which can be configured in the **LI REQ** command.

2.4.22 TVC management for TS7700T/TS7700C 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 be 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. Copies that are made to peers by default do not acknowledge the preference group that is defined at the target location; instead, they treat it as PG0 so that they are flushed from disk cache immediately.

For example, in a two-cluster grid, consider that you set up a Copy Consistency Point policy of RUN, RUN, and that the host can access all virtual devices in the grid. After that process, 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 can be removed from cache. The result is that each TVC is filled with unique newly created or modified volumes, which roughly double 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 be 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 in cache. The use of the default setting results 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. Consider the following points:

- ▶ 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 that is 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

No difference exists in a stand-alone cluster environment.

2.4.23 TVC management for disk-only TS7700 clusters in a multi-cluster grid

Disk-only TS7700 clusters in a grid feature several unique options for cache management.

Enhanced Removal Policies

The Enhanced Volume Removal Policy provides tuning capabilities in grid configurations where one or more disk only TS7700D clusters are present. Traditionally, other TS7700T or TS7700C clusters also exist which have a larger available capacity. Enhanced removal policy provides a method to retain only certain workloads within the smaller capacity TS7700 solutions as data ages.

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, which is also true for a TS7700T/TS7700C CP0.

In addition, when the automatic removal is run, it implies an override to the current Copy Consistency Policy in place, which results in a reduced number of consistency points compared with the original configuration that is 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 that are described next are ignored for scratch (Fast Ready) volumes.

To help ensure that data always is in a TS7700D or CP0 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 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, the following policies are available for each volume in a TS7700 CP0:

► Pinned

The copy of the volume is never removed from this cluster. No volume retention time is applicable and it 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 or creation.
- The 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 to be 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) elapsed since the last access or creation.
- The free space on the cluster fell below the removal threshold.
- Volumes with the Prefer Remove (LRU Group 0) policy were 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 the use of 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. Regardless of their removal policies, scratch category volumes 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 do PG0 and PG1 volumes become candidates 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 2-12 shows the TS7720 cache removal priority.




Figure 2-12 TS7700 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 or TS7700C. For more information, see [TS7700 Library Request Command V5.4](#).

The following host console requests are related:

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

Time Delayed Replication can change the behavior of the auto-removal algorithm so that removal of volumes where one or more delayed replication consistency points exist can occur only after those delayed replications are completed. If families are defined, only delayed consistency points within the same family must complete.

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, which enables candidates for removal to be created. If no candidates for removal exist and Time Delayed copies are still pending to one or more peers, the TS7700C cluster requests the peers to begin their copies early so that removal can find candidates.

TS7700 CP0 cache full mount redirection

If the Enhanced Volume Removal Policies were not defined correctly or are disabled, a TS7700D TVC can become full. This issue is also true for a TS7700T/TS7700C 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 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 directed only to a TVC in an Out-of-Cache Resources state if no other eligible (TVC) candidate exists. When the only private mount TVC candidate is a full TS7700 cluster, the private mounts are mounted with read-only access.

When all TVC candidates are 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 demount 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 scratch 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 TS7700 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 plans are made to take 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 space before a DR test with IBM FlashCopy®. During the DR test, no autoremoval or delete expire process is allowed. Therefore, use the temporary removal threshold to ensure that enough free space for the usual productions and the extra flash copies is available in the clusters in the DR family.

2.4.24 TVC management processes in a multi-cluster grid

The TVC management processes are the same as for stand-alone clusters. In addition to premigration management and free space management, the following processes are available:

- ▶ Copy Data Residency Retention (TS7700T/TS7700C CP1–CP7)

This process applies to only a multi-cluster grid configuration in which one or more TS7700T/TS7700C clusters exist. TS7700T and TS7700C clusters retain data in the disk cache longer than needed if it believes that no other peers require a copy. The goal is to reduce recalls for copy.

After the volume is copied to the required peers, the volume is no longer held in disk cache for this purpose alone and acknowledges its PG0 or PG1 LRU algorithm. If the copies target a remote cluster family, the target family retains the data on behalf of its family peers, and the original source cluster view it as completed.

This process is done to avoid a possible recall operation from being started 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*.

- ▶ RUN Copy Throttling

This process applies to multi-cluster grid configurations where the RUN Copy Consistency Point is used. When enabled, it limits the host input rate if it believes that the RUN copies cannot keep up otherwise. 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 does not help, the job changes to Immediate/RUN Deferred mode and an alert is posted to the host console that the TS7700 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. Sometimes, the RUN Copy throttling can be disabled if high bandwidth and minimal latency between peers exists. Because Synchronous mode copy is treated as Host I/O to the remote cluster, RUN throttling is not applicable to Synchronous copies.

2.4.25 Copy Consistency Points: Copy policy modes in a multi-cluster grid

In a TS7700 Grid, you likely 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 wrote the virtual volumes.

Copy management is controlled through the MC storage construct. By 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 the following 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 must exist on a cluster's TVC, you indicate a Copy Consistency Point for that cluster. If you do not want one or more clusters to have a copy of the data, you can specify the No Copy option for those clusters.

You can define Sync, Run, Deferred, Time Delayed, or No Copy. For more information, see 2.4.5, "Copy consistency points" on page 70.

Note: The default MC is *deferred* at all configured clusters, including the local cluster. The default settings are applied whenever a new construct is defined through the MI or to a **mount** command where MC was not defined. The default MC settings can be changed in the MI to change what settings are viewed as "default".

Synchronous mode copy

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

Data is written into one TVC and simultaneously written to as secondary TVC. This process differs from all other copy modes that start at RUN time or later and require a read from the source.

Because synchronous mode copies are written at the same time, no reading of source data is required. One or both synchronous locations can be remote. Synchronous mode copy works for new allocations and private mounts with MOD and volume size is not limited.

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 the persistent disk cache, which provides a zero RPO of all data up to that point on tape. Mainframe tape operations do not require that each tape block is synchronized, which enables improved performance by hardening data at critical sync points only.

Applications that use data set-style stacking and migrations are the expected use cases for SMC. However, 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 is provided if the zero RPO cannot be achieved. For more information, see [IBM TS7700 Series Best Practices - Synchronous Mode Copy](#).

The options that are available with the synchronous mode copy are described next.

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 a write failure occurs, 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, the host operations fail.

During allocation, if both S locations are unavailable, an R and then D site is chosen as the primary consistency point. If a write occurs or is anticipated and this option is not selected, the mount or write operation fail.

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

Open Both Copies On Private Mount option

Enable this option to open both written S locations when a private mount occurs. If one or both S locations are on back-end tape or cloud, the tape copies are first recalled into the 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 not used.

Consider the following 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 started.
- ▶ If a write operation occurs, both TVC locations receive write data and must be synchronized during each implicit or explicit synchronization command.
- ▶ If a write operation occurs and either location fails to synchronize or only one copy was opened, the host job 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 an append is anticipated. Consider the following points:

- ▶ If `DISP=OLD` is specified, the TS7700 assumes that only a read occurs 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 as described with `DISP=OLD`.
- ▶ If `DISP=MOD` is specified, the TS7700 assumes that an append occurs and opens both TVC copies in anticipation of a synchronized append.

Some applications open the virtual volume with a **DISP=OLD** parameter and still append to the volume. In this case, the append succeeds or fails, depending on whether the Synchronous Deferred On Write Failure option is enabled.

Tip: Use this option for DFSMSHsm 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. Other than the original, two or more other RUN copies are produced in parallel. Options are available to override this requirement for performance tuning purposes, as explained in 2.4.7, "Override settings concept" on page 73.

Note: Logical volume sizes of 25 GB or 65 GB are not compatible with RUN copies. If an MC uses RUN copies and the DC specifies 25 GB or 65 GB logical volume sizes, the RUN consistency points fall back to Deferred.

Deferred

If a Copy Consistency Point of Deferred is defined, the copy to that cluster's TVC can occur anytime after the **Rewind/Unload** command is processed for the I/O TVC.

Time Delayed Replication Policy

In the TS7700, all types of data can be stored. Some of this data often has a short lifetime, and is replaced with other (more current) data. This issue 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 logical volumes must be treated as archive data (for example, GDGs).

Time delayed replication provides a method to which volumes are replicated to only one or more peers after a delay criteria is met and the volume did not yet expire, which treats the copy target as an archive only target. A deferred copy is made to all T sites after X hours 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 under the same MC. You can specify 1–65,535 hours.

Data that is expired is still copied to the target clusters in the following circumstances:

- ▶ The TMS did not yet return the volume to scratch.
- ▶ The logical volume is scratch but not reused and the scratch category does not include an expire delete definition.
- ▶ The logical volume is scratch but not reused and the scratch category has an expire delete setting, which was not yet reached for this specific logical volume. Ensure that the expire delete setting for the scratch category and the Time Replication time combination fit together. That is, add a day or two after the expected expiration time to allow housekeeping and delete expire processing to complete before the T time passes.

By using the Time Delayed policy, the automatic removal in the TS7700 CP0 can be influenced. The following rules apply:

- ▶ In a grid without cluster families, all T copies must be processed before an automatic removal can occur on any TS7700 CP0 in the grid.
- ▶ If cluster families are defined, all T copies in the family must be processed before auto removal to any TS7700 CP0 in the cluster family can occur. However, a logical volume in a TS7700 CP0 can be removed, even if all T copies in a different family were not processed.

A TS7700 CP0 might run out of removal candidates and the only candidates are those time-delayed replication volumes that have not yet had their time expire. In this case, the TS7700 CP0 detects this condition and triggers a subset of peer time-delayed copies to replicate early to create removal candidates. The TS7700 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 more information about examples of how Copy Consistency Policies work in different configurations, see 2.5, “Grid configuration examples” on page 106.

A mixture of Copy Consistency Points can be defined for an MC, which enables each cluster to have a unique consistency point. In addition, each cluster in the grid can define a unique MC so that the copy consistency behavior is different, depending on where the logical volume mount operation was issued.

Tip: The Copy Consistency Point is considered for scratch and specific mounts. Specific mounts do not update the consistency points when the “Retain Copy Mode” setting of the MC is enabled.

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 valid reasons exist for doing so. For example, you might have a three- or four-way configuration in which only one local cluster and one remote cluster receives a copy of the logical volume. Which local cluster receives the scratch mount operation determines a different consistency point outcome that is based on that cluster’s MC definition. In these scenarios, use the Retain Copy mode option to prevent consistency mode changes for private mounts. When the Retain Copy mode is enabled against the defined MC, the assigned copy modes are retained independently of the current MC definition during a private mount.

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. It acknowledges the latest MC for each mount, except for private mounts where the Retain Copy Mode setting is enabled.

Retain Copy mode across grid mount points

Retain Copy mode is an optional setting in the Management Class where a volume’s Copy Consistency Points are used rather than applying the Copy Consistency Points that are defined at the mounting cluster during private mounts. This setting applies to private volume mounts for read/write appends. It is used to prevent more copies of a volume in the grid than wanted because previous copies are not automatically removed.

Figure 2-13 shows a four-cluster grid where Cluster 0 replicates to Cluster 2, and Cluster 1 replicates to Cluster 3. The goal is that only two copies of data remain in the grid after the volume is accessed. One of the copies is local and one of the copies is remote. Later, when a 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).

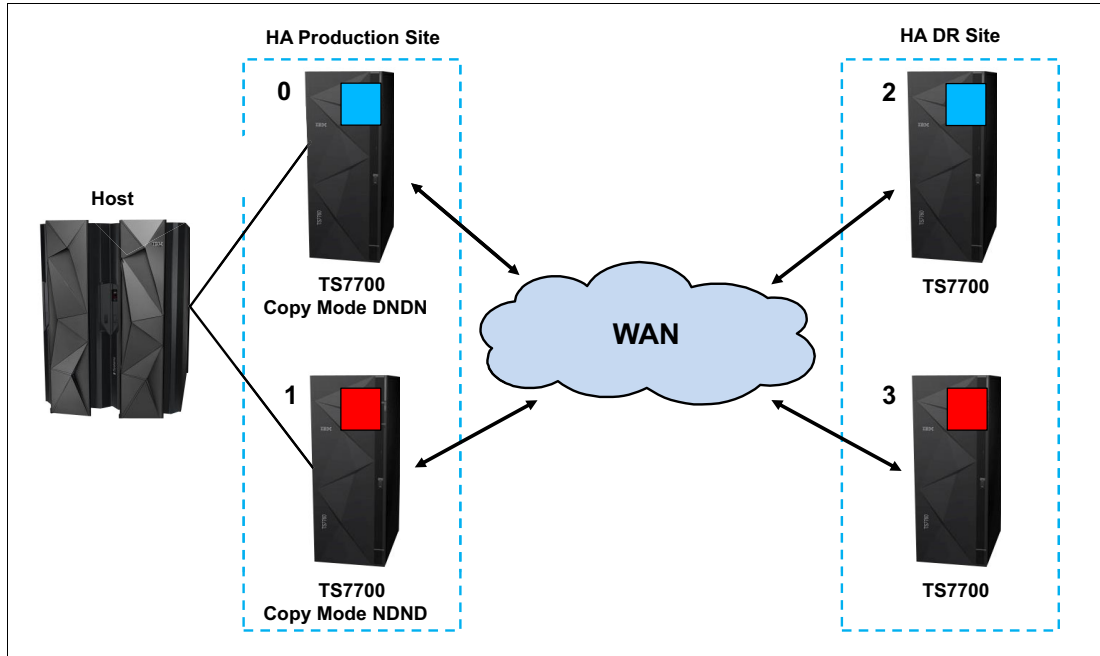


Figure 2-13 Four-cluster grid with DAA

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

On systems where DAA is not supported, the host allocates 50% of the time to the cluster that does not have a copy in its cache. When the alternative cluster is chosen, the copies remain present, and more copies are made to the new Copy Consistency Points that are defined in the Management Class, which results 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 must be made to Cluster 1 and Cluster 3.

Four copies remain for a read operation. Three copies are created for a write append. This process is shown in Figure 2-14.




Figure 2-14 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, as shown in Figure 2-15.




Figure 2-15 Four-cluster grid without DAA with 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. The creation of more copies applies to JES2 and JES3 without Retain Copy mode enabled (see Figure 2-16).




Figure 2-16 Four-cluster grid, one production cluster down, Retain Copy mode disabled

Figure 2-17 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 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 receives a copy of the altered volume when it becomes available.




Figure 2-17 Four-cluster grid, one production cluster down, Retain Copy mode enabled

For more information, see [IBM Virtualization Engine TS7700 Series Best Practices - TS7700 Hybrid Grid Usage](#).

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

The TVC that is 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 clusters or cluster in service prep are excluded).
- ▶ Mount type:
 - Scratch: Exclude TS7700 clusters where its disk cache cannot contain the newly created volume, such as a full CP0 partition. Also, exclude all No Copy clusters.
 - Private: Exclude TS7700 clusters without a valid copy.
- ▶ Preferences regarding:
 - Consistency Points such as Sync, followed by RUN, followed by Deferred, followed by Time Delayed, followed by Exists.
 - Copy Override Policies.
 - Cluster Families.

With these three elements, an obvious ordered list can be considered. If two or more choices are viewed as equal, more filtering occurs in which choices are ranked by certain performance criteria, such as the following examples:

- ▶ Cache residency
- ▶ Recall times
- ▶ Network latency
- ▶ Host workload

The list is ordered favoring the clusters that are thought to provide the optimal performance while still acknowledging copy consistency point preferences.

Unavailable TS3500 or TS4500 Library

Two LI REQ parameter settings can influence the TVC selection process. You can use the SETTING2,PHYSLIB parameter to determine how a degraded TS3500 or TS4500 library is handled in a TS7700T.

In addition, you can use the LI REQ parameter LOWRANK to give a cluster a lower ranking in the TVC selection process. This parameter influences the TVC selection for Host I/O and the copy and mount behavior. This parameter can also be used under special conditions before you enter Service Mode to help reduce TVC usage or when you want to avoid the use of a specific cluster for mounting TVC or copy sourcing.

In addition, it is a persistent setting and can be set on every cluster independently. Because it is persistent, its use might be temporary and should be moved back to a default setting after it served its purpose; for example, after completing Service Prep.

For more information, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

2.4.27 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 virtualizes 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 that is provided by the TS7700. Meaning, if the TS7700 uses FICON, LZ4, or ZSTD compression, the compression occurs before the data is transferred across the grid network. Read-ahead and write buffering is also used to achieve maximum performance for remote mounts.

2.4.28 TVC encryption

From a technical perspective, TVC encryption is a cluster feature. Each cluster can be treated differently from the others in the multi-cluster grid. No difference exists with a stand-alone cluster.

2.4.29 Logical and stacked volume management

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

2.4.30 Secure Data Erase

No difference exists with stand-alone cluster.

2.4.31 Copy Export

In general, the Copy Export feature has the same functions as a stand-alone cluster. However, consider the following points:

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

Only the physical volumes that are exported from a source TS7700T can be used for the recovery of a source TS7700T. Physical volumes from more than one source 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 start 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 merged later.

- ▶ Recovery that is run by the client is to a stand-alone cluster configuration only. After recovery, the Grid MES offering can be applied to re-create a grid configuration.
- ▶ When a Copy Export operation is started, only the following logical volumes are considered for export:
 - They are assigned to the secondary pool that is 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 runs the Copy Export operation, but that copy has not yet completed when the export is started, it is not included in the current export operation. Ensure that all logical volumes that must be included completed replication to the cluster where the export process is run.

- ▶ A service from IBM is available to merge a Copy Export set in a grid. For more information, contact your IBM SSR.

2.4.32 Encryption of physical tapes

No difference exists with a stand-alone cluster.

2.4.33 Autonomic Ownership Takeover Manager

AOTM is an optional function by which after an unexpected 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 that is 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 configuration 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 failed or is otherwise not operational. For more information, see 10.5, “Service icon” on page 543.

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 or by using the **LI REQ** command.

Without AOTM, an operator must determine whether one of the TS7700 clusters 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 is enabled only when a cluster failed and not when a problem exists only 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 issue 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 extra TS3000 TSSC path is used to determine whether an unavailable cluster is still operational. 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.

2.4.34 Selective Write Protect for DR testing

This function enables clients to emulate DR 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 128 categories can be identified and set to be 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 attributes changed. Write protection is cluster scope and configured through the MI. Settings are persistent.

Also, the function allows full read/write access to volumes that are assigned to a category that is included in the exclusion list. 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 configuration enables the DR test to mount volumes as private that the production environment since returned to scratch (they are accessed as read only).

One exception to the write protection 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 implied to be 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 Chapter 16, "Disaster recovery testing in a grid configuration" on page 819.

2.4.35 FlashCopy for DR testing

This 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 content to segregated volume ranges. A total of 128 write exclusion categories are supported, versus the previous 32 in R4.2 and prior. Write-protected media categories cannot be changed (by the DR host) while the Write Protection mode is enabled. This fact is true for the data and 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 only TS7700Ds, the flash that is started during the DR test is across all TS7700Ds in the configured 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 instance in each cluster's TVC. Therefore, return to scratch processing can now run at the production side during a test, and it does not need to deferred or expire hold used.

The TS7700T or TS7700C can be a part of a DR Family, but they have some limitations for those volumes that are not resident in the disk cache at the time of the Flash.

For more information about FlashCopy setup, see Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351. For DR testing examples, see Chapter 16, “Disaster recovery testing in a grid configuration” on page 819.

Consider the following points regarding FlashCopy for DR Testing:

- ▶ A disk cache snapshot occurs to one or more TS7700 clusters in a DR family at the same time within seconds.
- ▶ All logical volumes in a TS7700T/TS7700C CP0 partition and all logical volumes from CPx kept in cache, are part of the DR-Flash.
- ▶ Volumes in the TS7700T or TS7700C that are stored in CPx partitions that are migrated to physical tape or cloud 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.
- ▶ Volumes in the TS7720T that are stored in CPx partitions and are 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.
- ▶ TS7700 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 clusters.
 - Any volume that is reused in production results in two DR copies (old at time zero and new).
 - Any new content created by the DR hosts, such as a dedicated volume range used by the DR host.
- ▶ Automatic removal is disabled within all clusters during a DR test requiring a pre-removal to be completed before testing.
- ▶ **LI REQ** DR Family settings can be completed in advance, which enables a single **LI REQ** command to be run to start the flash and start DR testing.
- ▶ DR access introduces its own independent ownership and allows 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 instance is the live instance of a volume in a cluster that is the most current true version of the volume. It is live production data or the latest instance that is created by a DR host.</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 be an obsolete or incomplete replication at time zero. Synchronous mode volumes are flashed up to the last successful sync point before time zero. Data after the last sync point can exist, although its validity is based on how much data was written at time zero.</p>
DR family	<p>A set of TS7700 clusters (most likely those clusters at the DR site) that serve the purpose of DR. A total of 1–7 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 flash accessible. Virtual tapes that are assigned to the excluded categories surface live data and live attributes versus the time zero instance.
Time zero	The time when the FlashCopy is taken within a DR family. The time zero mimics the time when a real disaster occurs. Customers can establish the time zero by using LI REQ (Library Request command).

2.4.36 Grid resiliency functions

A TS7700 Grid is made up of two or more TS7700 clusters that are interconnected through Ethernet connections. It is designed and implemented as a business continuance solution with implied enterprise resiliency. When a cluster in the grid encounters a problem, the multi-cluster grid should accommodate the outage and continue the operation, even if the state of the grid is degraded.

Grid-wide problems might occur because of a single cluster in the grid experiences a problem. When problems exist in one cluster that causes it to be sick or unhealthy but not dead (Sick But Not Dead [SBND]), the peer clusters might be greatly affected. Customer jobs end up being affected, such as long mount time, failed sync mode writes, and volume ownership might be hung.

Grid Resiliency Improvements are the functions to identify the symptoms and make the grid more resilient when a single cluster experiences a problem by removing the sick or unhealthy cluster from the grid explicitly or implicitly through different methods. By removing it, the rest of peer clusters can then treat it as "dead" and avoid further communication with it until it can be recovered.

The grid resiliency function is designed to detect permanent effects. It is not designed to react to the following situations:

- ▶ Temporarily impacts (like small network issues)
- ▶ Performance issues due to high workload

Note: Because of the nature of a TS7700 grid, this isolation is not comparable to a mechanism in the disks, such as IBM HyperSwap® or similar techniques. Such techniques are based on local installed devices, whereas TS7700 grids can span thousands of miles. Therefore, the detection can take longer than in a disk world, and also the actions might take longer.

The customer can specify different thresholds (for example, mount timing, handshake and token timings, and error counters) and other parameters to influence the level of sensitivity to events that affect performance. To avoid a false fence condition, use the defaults for the thresholds at the beginning and adjust the parameters only if necessary.

Building upon functions that were introduced in R4.1.2, R5.2 now provides notification to the attached z/OS hosts of unhealthy clusters. When a cluster is fenced as part of Grid Resiliency, any cluster that is connected to the same hosts as the unhealthy cluster notifies the attached hosts about the unhealthy cluster through control unit-initiated recovery (CUIR). Then, the hosts auto-vary offline the devices to the unhealthy cluster.

A single sick cluster can also notify its own attached hosts that it is sick. The unhealthy cluster then notifies the host to auto vary the devices back online after it recovers.

The following mechanisms exist in the grid:

- ▶ Local Fence
- ▶ Remote Fence
- ▶ Manual Fence

Local Fence

As in a stand-alone environment, the cluster decides (based on hardware information) that it suffered an SBND condition and started a local fence action (that is, restart of the local cluster by itself). The function is automatically enabled after R4.1.2 and later is installed on the cluster, even if other clusters in the grid are not yet running on R4.1.2 or higher level. The local fence has no parameters or options and cannot be disabled.

Remote Fence

Depending on the parameters, one of the clusters in a grid might detect an unhealthy state of a peer cluster. In a grid with three or more clusters, all the clusters must concur that a specific cluster is unhealthy for it to be fenced.

In a two-cluster grid, both clusters must agree that the same cluster is SBND. Otherwise, no remote fence occurs.

The remote fence is disabled per default. If the customer enables the remote fence action, the following parameters must be defined:

- ▶ Primary Action:
 - ALERT: An Alert message is sent to the attached hosts and the cluster is fenced. However, the cluster remains online and is still be part of the grid. Therefore, the unhealthy situation is not solved automatically. You might consider this option if you want to be notified about that SBND condition occurred, but want to run the necessary actions manually.
 - REBOOT: The SBND cluster is restarted. If the restart is successful, the cluster automatically is varied back online to the grid. If the restart is not successful, the restart action is repeated twice before the cluster remains offline. Consider this option if availability of the complete grid is the main target, such as when the remaining grid resources cannot handle the workload during peak times.
 - REBOFF: The SBND cluster is restarted, but stays in an offline mode. Consider this option if an analysis of the situation is always requested before the cluster returns to the grid.
 - OFFLINE,FORCE: The SBND cluster is set to offline immediately with no restart. This option provides the quickest shutdown, but the restart action might take longer.
- ▶ Secondary Action

The customer can enable a secondary option. If enabled and the primary option fails (for example, the primary action cannot be run), the cluster is isolated from the grid. Therefore, only the gridlink ports are disabled, and no communication exists between the SBND cluster and all other clusters in the grid.
- ▶ IBM AIX® System Dump Action

The action to take a AIX dump can be enabled when the parameters of REBOOT/REBOFF are defined through the **LI REQ** command.

Manual Fence

TS7700 detects SBND symptoms on the local or against peer cluster. Then, the local/remote cluster fence action is applied based on the defined or configured fence action.

However, a cluster can also be locally and remotely fenced from Management Interface (MI) manually. A manual fence is a “force” mode, and no healthy cluster agreement is required to kickstart the remote cluster fence action (that is, the remote cluster fence action is forcibly applied, even though the other healthy clusters do not think the target cluster is SBND).

For more information, see [IBM TS7700 Series Grid Resiliency Improvements User's Guide](#).

2.4.37 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 cluster is in service. If a volume that is owned by the service cluster is not accessed during the outage, ownership is retained by the original service 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 new copies from other clusters to the cluster that is entering the service mode are started.
- ▶ All logical volumes that are consistent only on the service prep cluster and require at least one peer copy are copied to at least one other cluster in the grid. This situation 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.
- ▶ Any pending updates (hot tokens) to peer clusters where the service-prep cluster is the current owner must be reconciled or the volume ownership must be transferred to an available peer.

When a service prep completes, the cluster enters service mode and is online unless taken offline. The other clusters in the grid remain online and available for processing work load. This service mode stops communication to this cluster from the remaining clusters, which results in queued actions to this cluster. The additional task of taking the cluster offline enables service personnel to run maintenance tasks and hardware diagnostic tests without affecting other clusters.

Only one service prep can occur within a composite library at a time. If a second service prep is attempted while another cluster is still in service prep, the second request fails. You should put only one cluster in the service prep mode at the same time. After service prep completes, another cluster can then be put into service prep.

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.

If you use SAA, you should consider disabling SAA during the maintenance. This step is necessary if you often bring 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 issue causes the job to go into device allocation recovery, if this device was the only SAA candidate.

If you have multiple SAA candidates that are defined, you might still consider disabling SAA. This action would be necessary if otherwise the number of SAA selectable devices are not sufficient to run all the jobs concurrently.

After SAA is enabled again, restart all attached OAM address spaces to ensure that the changed SAA state is recognized by the attached z/OS or vary online another drive. If you do not restart the OAM address spaces or vary online another drive, the system might react as though SAA is still disabled.

2.4.38 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 configuration 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 ideal 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. By using the MI, you must 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. If they cannot be brought back up at the same time, it is ideal to have them brought up in the opposite order of which they were serviced. The clusters must be forced online if one or more peers are still in service mode or unavailable.

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.

2.4.39 Control Unit Initiated Reconfiguration

CUIR is available in R4.1.2 and later to reduce the manual intervention during microcode upgrade processes. Currently, it is your obligation to ensure that all devices are set to offline in all attached z/OS LPARs. With the CUIR function, the TS7700 notifies attached IBM Z LPARs by using an unsolicited interrupt when the cluster enters or exits service-prep state. If customer enabled, the devices are brought offline automatically in the attached z/OS LPARs.

The TS7700 tracks the grouped devices to all path groups and does not enter service until they are all varied offline.

The customer can decide when the cluster returns to an operational state from service if an automatic online (AONLINE) to supported z/OS LPARs is run. Then, the logical paths that are established from the IBM Z LPAR that the cluster surfaced the unsolicited attention receive an unsolicited attention to request the zLPAR to vary the devices back online. Only z/OS LPARs support the CUIR function. Other IBM Z operating systems can exist and hold up service-prep until you manually vary off their devices.

If the customer decides not to use AONLINE, the devices must be varied online again by using the MI. The original z/OS command (**Vary Onl ine**) cannot be used to online the devices in the z/LPAR while in CUIR. The manual function (through the MI) must be used to online the drives that CUIR put offline. At the time of writing, no option is available to bring online drives only partially (for example, to test a system) for quality assurance tests before the cluster is varied online to all systems again.

Note: If a device is varied offline for CUIR reasons and is unintentionally left in this state, the **MVS VARY XXXX, RESET** command can be used to reset the device for CUIR reasons. This command should be used only if devices exist that are left in this state and should no longer be in this state.

CUIR features the following limitations:

- ▶ Can be enabled only when all clusters in the grid are at R4.1.2 or later.
- ▶ Only native LPARS with z/OS 2.2 or later and JES2 can use CUIR. Other setups can exist, but their devices must be manually varied offline and back online.
- ▶ Only Service preparation/Service mode is supported. The CUIR function might be extended in the future.

The following APARs must be installed:

- ▶ OA52398
- ▶ OA52390
- ▶ OA52376
- ▶ OA52379
- ▶ OA52381

New **LI REQ** commands are provided to enable or disable CUIR and AONLINE and to get an overview of the current logical drive or path group information. The default is disabled.

For more information, see [IBM TS7700 Series Control Unit Initiated Reconfiguration \(CUIR\) User's Guide](#).

2.5 Grid configuration examples

Several grid configuration examples are provided in this section. These examples describe the requirements for high availability (HA) and DR planning.

2.5.1 Homogeneous versus hybrid grid configuration

Homogeneous configurations contain only TS7700Ds, TS7700Ts, or TS7700Cs. If you use an intermix of disk-only and tape-attached models, it is a *hybrid configuration*. Consider the information that is provided in this section when you choose whether a mixture of the types is best for you.

Requirement: No physical tape or dark site

Some clients want to eliminate physical tape from one or more data center locations. The TS7760C, TS7700D, or a hybrid configuration supports these requirements. The complete elimination of physical tape might not be the ideal configuration; therefore, TS7700T clusters can exist at alternative locations.

Requirement: Big data

The 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 up to 100 PB of compressed data.

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 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 TS7700T Copy Export with merge (available as a service offering) provide this capability.

2.5.2 Planning for high availability or DR 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 alternative 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 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 points:

- ▶ Plan for the virtual device addresses in both clusters to be configured to the local hosts.
- ▶ Plan a redundant FICON attachment of both sites; that is, an extender that is longer than 10 km (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 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. For more information, see 2.4.33, “Autonomic Ownership Takeover Manager” on page 97.

2.5.3 DR 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. No host connectivity exists 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 must introduce dense wavelength division multiplexing (DWDM) or channel extension equipment. Depending on the distance (latency), a difference might exist in read or write performance compared to the virtual devices on the local TS7700 Cluster.

Consider the following points:

- ▶ 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 only the use of the virtual device addresses in a remote TS7700 Cluster when the local TS7700 is unavailable. If these differences are an important consideration, you must 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 points:

- ▶ Plan for the necessary WAN infrastructure and bandwidth to meet the copy requirements that you need. You generally need more bandwidth if you primarily use 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.

2.5.4 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 replace the TS7740 with a TS7720T or TS7760T, depending on the customer's 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 TS7700Ts 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 (see Figure 2-18).

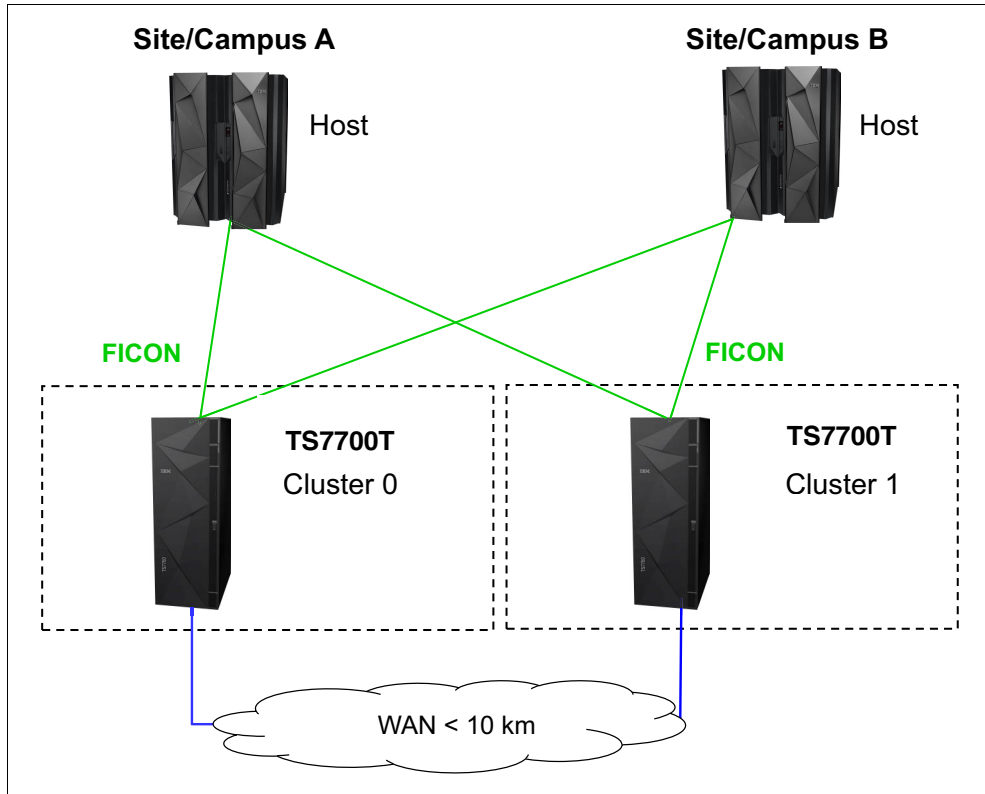


Figure 2-18 Two-cluster grid

Example 2: Three-cluster grid in two locations

In this example (see Figure 2-19), 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 TS7700T is installed.

The client runs many OAM and HSM workloads, so the large cache of the TS7700 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 TS7700T in the remote location.

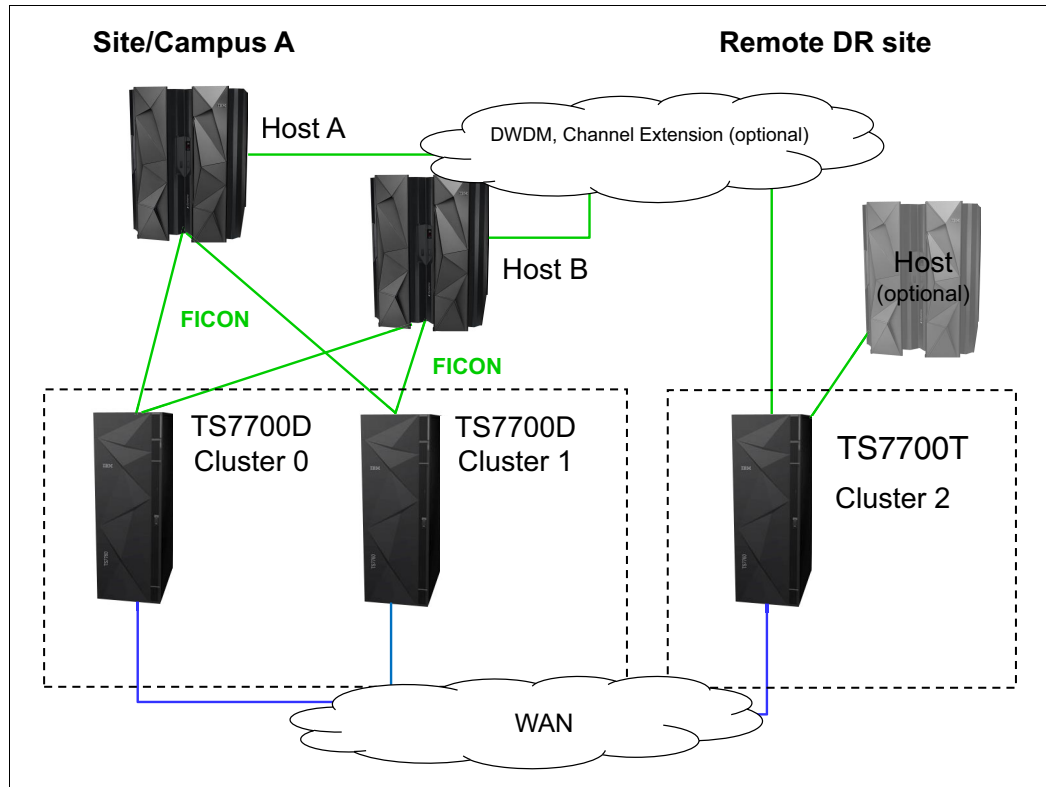


Figure 2-19 Three-cluster grid in two locations

Example 3: Three-cluster grid in three locations

This example is the same as configuration example 2. However, in this case, the two TS7700s 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 a TS7700T was used (see Figure 2-20).

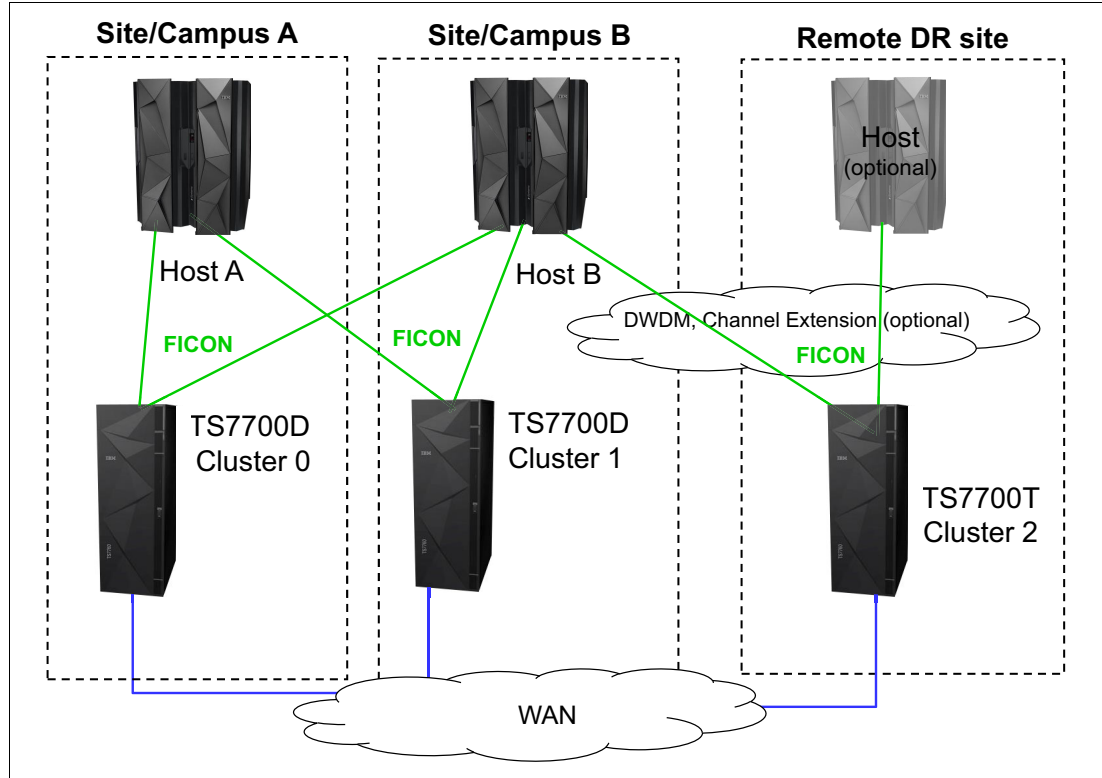


Figure 2-20 Three-cluster grid in three locations

Example 4: Four-cluster grid in three locations

The setup that is shown in Figure 2-21 shows the configuration after a merge of 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.

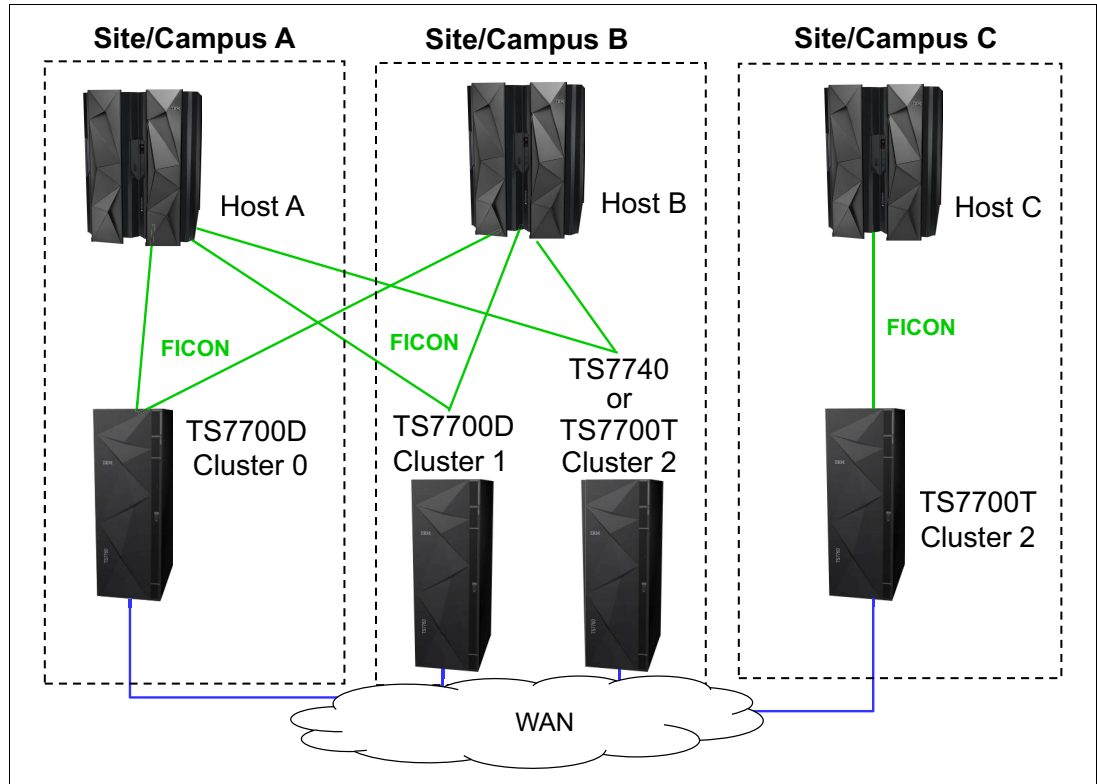


Figure 2-21 Four-cluster grid in three locations

Example 5: Two-cluster grid, all connected to the cloud

Figure 2-22 shows a 2 way grid where both clusters are TS7700C with R5.1 or higher.

Cluster 0 puts a logical Volume into the Cloud and cluster1 can access it, even when the copy modes are set to have no copy on cluster 1 (D,N). This process is possible because of the grid awareness feature. Also, a newly joined cluster that is configured to access the cloud can immediately read the volume in the cloud.

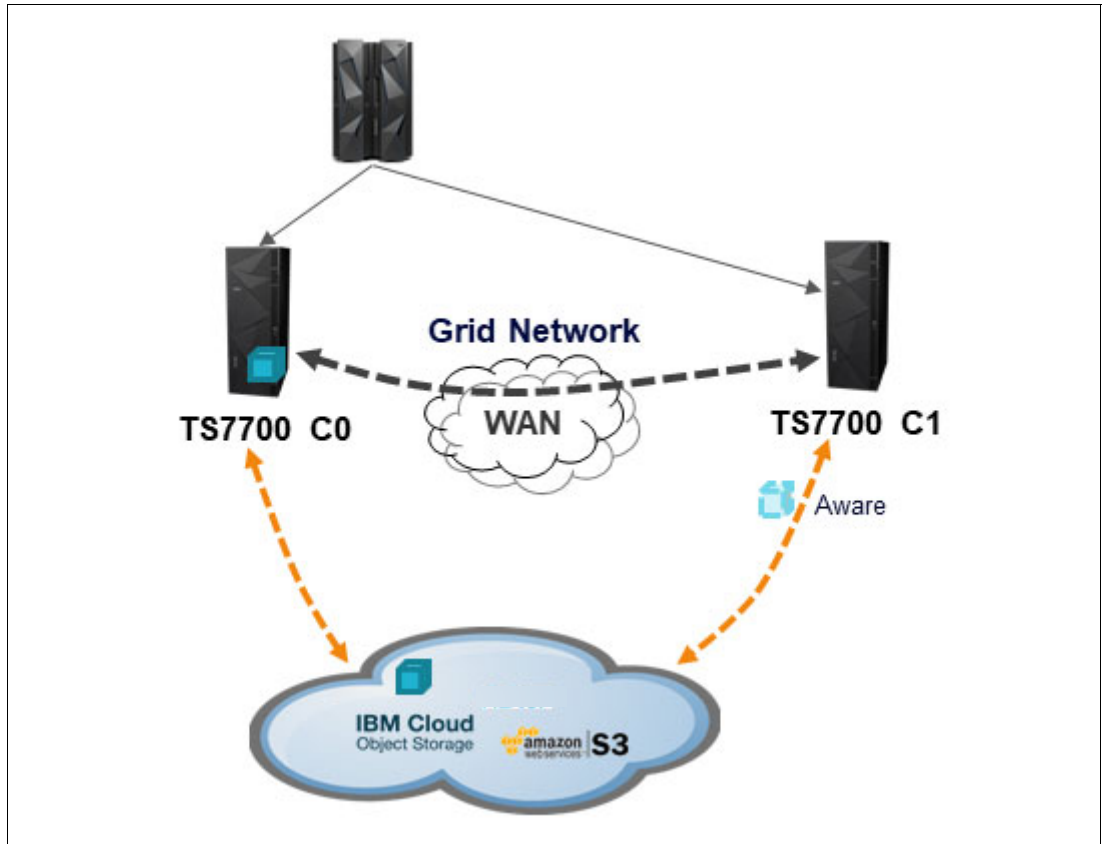


Figure 2-22 Two-cluster grid, all connected to the cloud



IBM TS7700 usage considerations

This chapter provides a general overview of the necessary information and considerations to choose the best configuration for your individual 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 raise only 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 topics:

- ▶ 3.1, “Introduction” on page 116
- ▶ 3.2, “Gather your business requirements” on page 118
- ▶ 3.3, “Features and functions for all TS7700 models” on page 125
- ▶ 3.4, “Features and functions available only for the TS7700T and TS7700C” on page 139
- ▶ 3.5, “Operation aspects: Monitoring and alerting” on page 139
- ▶ 3.6, “Choosing a migration method” on page 141

3.1 Introduction

Since the first days of tape usage, the world has changed dramatically. The amount of data that is stored has increased a lot, 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) product family, IBM 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, but 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 History overview

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 were 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 multi-file 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 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 Today's business challenges

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.

Because of 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 the use of such applications, the tier concept needs to be planned carefully.

3.1.3 Challenges of technology progress

Challenges exist with advanced technology, such as having many new options to meet many client needs. For example, the IBM 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 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 configuration likely increases the configuration capacity and bandwidth needs, which can introduce negative results. Planning and working with your IBM technical 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 as a starting point the lists that are described in this section.

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 is the data accessed, and what retrieval times are expected? Consider sizing and performance.
- ▶ How does 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 future growth.
- ▶ Which type of storage for your backup and archival data is needed, mandatory, and acceptable: Disk, physical tape or Object Storage (Cloud, on-premises, off-premises), or even a combination.
- ▶ What are the performance and recovery point expectations during an outage or disaster event?
- ▶ Plan for protection against cybercrime (gold copies).

It can be difficult to get all 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 tape environment helps 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 (power, cooling, and space)
- ▶ Financial requirements
- ▶ Multiple platforms required (IBM Z operating systems)
- ▶ Monitoring and automation capabilities to identify issues and degradations

- ▶ Maintenance (microcode) and defect repair strategy
- ▶ Capacity forecast
- ▶ Network infrastructure

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. For a single data center, make sure that the clusters are physically separated to minimize risks from fire, technical disruptions, attacks, natural disasters, and others.

If your approach is that each data center can host the total workload, plan your environment. 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, a tradeoff always exists 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 that you store in your TS7700 environment. Depending on your type of data, you can 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 (for example, a new installation).
- ▶ 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 somewhere else.
- ▶ If physical tape is present, migrate UAT-created content 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 nonreplicated 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. It 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. Because support no longer exists for a direct-attached physical tape to your Z System, your tape processing environment can be all virtual only. Nevertheless, the backend storage of an IBM TS7700 Virtual Tape System (VTS) can be disk, physical tape, object storage or cloud, or a combination of these formats.

Note: At the time of this writing, backend tape or backend object storage or cloud is mutually exclusive for a specific cluster.

Backup data

The data on virtual tape is a backup only. Under normal conditions, it is not accessed again. It might be accessed again only if problems exist, such as direct access storage device (DASD) hardware problems, logical database failures, and site outages.

Expiration

The data expiration depends on several aspects, such as legal requirements, the need to keep multiple older generations of data, and the time window for how long the data is useful and beneficial if being restored. It can be a short, medium, or long period.

Availability requirements

If a virtual tape environment or a part of it 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

Backend physical tape recall can normally be tolerated, or at least for previous generations of the backup. The same is true for recalls from an on-premises or off-premises object storage or cloud in the backend.

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 or object storage or cloud. However, physical media might fail, network access to an object storage or cloud might not be possible 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, one copy consistency point of the Deferred character might be sufficient, depending on your requirements.

Active data on tape

The data is stored on virtual tape as a second tier storage only. This data is not also somewhere in DASD such as DS8000. If the data must be accessed, it is read or staged from the virtual 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

Backend physical tape or object store 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 and recalls from an object storage or cloud (on-premises or off-premises).

Multiple copies

Although (virtual) 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 a suitable 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.

The use of the synchronous mode copy is highly recommended, especially for DFSMSHsm ML2, HSM backups, OAM objects, or similar applications.

Archive data on virtual tape

Archive data on virtual tape also is active data. However, archive data is stored for a long period. Expiration dates for up to 30 years to satisfy regulatory requirements are common. In addition, Logical Write Once Read Many (LWORM) virtualization might be required for such regulatory data.

Expiration

The expiration depends on your application, but it is often 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 or object storage or cloud recalls might be tolerated.

Multiple copies

Although the virtual 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 stand-alone environment, dual copies on physical tape are suggested.

Depending on the criticality of the data, choose a suitable Copy Consistency Point Policy.

Archive data must sometimes be kept for up to 30 years. During such long 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 TS7700T, you must also consider the life span of physical tape cartridges. Some vendors suggest that you replace their cartridges every five years, but other vendors, such as IBM, offer tape cartridges that have longer lifetimes.

If you are using a 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 increases. 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, where this data is stored on the stacked volume cannot be directly controlled, and the same condition might also occur in other clusters. Setting up a reclamation schedule that is based on the time since last access or write can help self-audit the contents of the physical tapes. This process can be optimized further by steering long active data into another physical tape pool during reclamation.

If you use a TS7700C and your archive data is in an object storage or cloud (on-premises or off-premises), you must also consider some important aspects. For example, if your cloud provider is no longer available, which can require you to migrate your data from an object storage or cloud storage to a newer or different object storage or cloud.

Therefore, you might consider storing data with such long expiration dates on a specific stacked physical tape 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 datapool.

3.2.4 IBM Db2 archive log handling

With IBM Db2, many choices are available for how to handle your Db2 archive logs. You can put both of them to DASD and might rely on a later migration to tape through DFSMSHsm or an equivalent application. You can also write one archive log to DASD and another one to a virtual tape environment. Alternatively, you can put them both directly to virtual tape.

Depending on your choice, the virtual tape environment is more or less critical to your Db2 application. This choice also depends on the number of active Db2 logs that you define in your Db2 environment. In some environments, because of peak workload, logs are switched every 2 minutes. Db2 stops processing if all Db2 active logs are used and if they cannot be archived successfully to virtual tape.

Example scenario

You have a four-cluster TS7700 grid that is spread over two data center sites. A TS7700D (disk only) and a TS7700T (with backend physical tape) are present at each site. You store one Db2 archive log directly on virtual tape and the other archive log on native mainframe disk (DASD). Your requirement is to have two copies on virtual tape.

Consider the following points:

- ▶ The use of the TS7700D 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 TS7700Ds. If one TS7700D is available, Db2 archive logs can be stored to virtual tape. However, if one TS7700D is unavailable, you have only one copy of the data. In a DR situation where one of the sites is unusable for a long time, you might want to change your policies to replicate this workload to the local TS7700T as well.
- ▶ If the TS7700D enters the Out of cache resources state, new data and replications to that cluster are suspended. To avoid this situation, consider having this workload also target the TS7700T and enable the Automatic Removal policy to available space in the TS7700D. Until the Out of cache resources state is resolved, you might have fewer copies than expected within the grid.
- ▶ If one TS7700D is not available, all mounts must be run on the other TS7700D cluster.
- ▶ In the unlikely event that both TS7700Ds are not reachable, Db2 stops working when all Db2 logs on the disk (mainframe DASD) are used.
- ▶ Having a consistency point of R, N, R, D provides you with three copies, which are stored in both TS7700Ds and in the TS7700T 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 TS7700D can remove the data from cache because there is still an available copy in the TS7700T.

Note: Any application with the same behavior can be treated similarly.

3.2.5 DFSMSHsm Migration Level 2

Several products are available on the IBM Z 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 TS7700D or a TS7700T CP0 at least for a certain period. With time-delay copies to extra tape-attached TS7700T clusters in a grid, you can ensure that the data is kept first in the TS7700D, and later copied to a cluster with tape attachment. Auto removal processing (if enabled) can then remove the content from the TS7700 as it ages.

Ideally, DFSMSHsm ML2 workloads are 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.

The use of 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 reduce the complexity of bringing up operations at a secondary location. In addition, you can use Synchronous mode copy along with DFSMSHsm to partially restore partially filled ML2 virtual tapes that might occur during a data center outage or disaster.

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 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, it is active data.

Users who are accessing the data on virtual tape (in particular, the TS7700T) might have to wait for their document until it is read from physical media back into the TS7700 TVC. The TS7700D or the TS7700T CP0 are traditionally better options for such a workload because the disk cache residency can be longer, even indefinite.

For OAM primary objects, use Synchronous mode copy on two clusters and depending on your requirements, more 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 is 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 is placed on a TS7700D or TS7700T/C CP0, a TS7700T/C, CP1–CP7, or a mix of them. For example, active data that needs quick access is ideal for a TS7700D or the resident partition CP0 of a TS7700T/C. Or, target a CPx partition with enough capacity to contain all or most of the active data likely to be accessed.

Depending on your environment, you can also place this data on a tape partition of a TS7700T and ensure that the data is kept in cache. This process can be done by delaying premigration or defining the size of the tape partition to keep all of this data in the cache.

Rarely, accessed data that does not demand quick access times can be put on a TS7700T/C tape partition with PG0 and a not-delayed migration, or on the TS7700T/C in a second cluster.

Data that becomes less important with time can also use the TS7700D or TS7700T/C CP0 auto-removal policies to benefit from both technologies.

Assume that you have the same configuration as our Db2 archive log example. Consider the following points:

- ▶ With a Consistency Copy Point policy of [N,R,N,R], your data is stored on only the TS7700T/C CPx (fast access is not critical).
- ▶ With a Consistency Copy Point policy of [R,N,R,N], your data is stored on only 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 TS7700T/Cs. This configuration enables 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 that is held on a TS7700, as listed in Table 3-1.

Table 3-1 Type of data

Data type	Application examples	Fits best on	Suitable cache control
Data needs a 100% cache hit	<ul style="list-style-type: none"> ▶ OAM objects (primary data) ▶ HSM ML2 	TS7700 Disk-Only TS7700T/C CP0	Pinned/PG1 CP0/Pinned/PG1
Data that benefits from a longer period in disk cache	<ul style="list-style-type: none"> ▶ Depends on the user requirements ▶ OAM objects (primary data), HSM ML2 	TS7700D with autoremoval TS7700T/C CPx	PG1
Data that is needed for a specific time in the cache, but then should be kept on tape	<ul style="list-style-type: none"> ▶ Db2 log files (depending on your requirements) ▶ Active Batch data 	TS7700T	CPx/PG1 delay premigration
Data with limited likelihood to be read, only cache pass through	<ul style="list-style-type: none"> ▶ Backups ▶ Memory dumps 	TS7700T/C	CPx/PG0
Object data with 100% cache hit	<ul style="list-style-type: none"> ▶ DS8000 TCT object data 	TS7700O with TS7700D/T/C ^a	Pinned only

a. TS7700O Advanced Object store can be enabled on disk-only, tape-attached, or cloud-attached TS7700s. However, object data cannot be tiered to tape or cloud currently.

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 Four TS7700 models: Disk, tape, object, and cloud

A TS7700 system can be ordered and configured in the following versions (from plant or as an MES):

▶ TS7700 (TS7700D)

All data is stored on pre-configured internal disk space only. This configuration is also known as a *disk-based* or *disk-only* TS7700.

Starting with R5.2, you can choose between two internal disk storage technologies: Nearline SAS spinning hard-disk drives (HDDs) or solid-state disk drives (SSDs).

The TS7700 supports usable capacities from 157 TB–3.94 PB before compression with Nearline SAS HDDs, and it scales in steps of 157 TB, which represents two cache enclosures.

The SSD-based TS7700 supports only two configurations as of this writing: One SSD cache enclosure that provides 60 TB of usable capacity, or two SDD cache enclosures that provide a total of 120 TB of usable cache capacity. (A specific 40 TB configuration can be upgraded to a 70 TB configuration at any time.)

► TS7700T

In this configuration, a TS7700 features a physical tape library and physical tape drives (minimum 4, up to 16) attached. In this setup, the internal disk cache is partitioned (up to seven tape partitions for different workloads) and acts as a cache. All data that arrives in one of the tape partitions is a candidate to be copied out of the cache onto a physical tape cartridge and then optionally migrated or deleted from disk cache.

The maximum capacity in this configuration is limited by the amount of physical tape media that is inside the assigned tape library partition or the number of configured logical volumes and their configured sizes. The maximum is up to 100 PB of virtual tape content that is stored within the attached physical tape library.

FC 5273 “TS7760/TS7770 Tape Attach” must be ordered with the initial order or as an MES from a TS7700D. If the feature is present and active a physical tape library with a supported number of drives is mandatory. A TS7700T also provides a disk-only partition (CP0) for data not to be migrated out to a physical tape.

► TS7700C

In this setup, a TS7700 can be attached to an external object storage or cloud. Logical volumes can be copied to the attached object store and then optionally migrated or deleted from disk cache. This type of configuration is similar to the TS7700T usage, but uses an attached object store instead of physical tape.

Also, the cloud can be shared among two or more TS7700C clusters, which allows data that is copied to the cloud to be accessible by other TS7700C clusters in the grid. The internal disk cache is partitioned in this configuration and supports up to 256 cloud pools. Each cloud pool can reference different object stores or clouds, or to different containers within the object stores.

The maximum capacity in this configuration is limited by the maximum available storage capacity of the object storage or cloud. FC 5278 Cloud Enablement must be ordered with the initial order or as an MES to a TS7700D to support that use case. Even if that feature is present and active, an attachment and use of an object store or cloud is *not* mandatory; a TS7700C still can also be operated with its internal disk cache only (partition CP0).

Note: FC 5273 TS7760/TS7770 Tape Attach and FC 5278 Cloud Enablement are mutually exclusive at the time of this writing.

► TS7700O

In this setup, an IBM DS8000 can be attached to a TS7700 for object offload. The DS8000 can target one or two TS7700 clusters to migrate object data to if they support FC 5283 (Advanced Object Store introduced with 5.2.220.x code). Copies of object data can be distributed among any combination of grid members with FC 5283 installed. If two TS7700s are targeted by the DS8000, data offload is load-balanced between the two target clusters, meaning that a single copy is sent to one of the clusters and that cluster's copy policy determines how the object copies are handled. Both targeted clusters must be in the same TS7700 grid.

Object copy policies are created at the TS7700 MI. There are three options for the copy handling for each FC 5283 enabled cluster in the grid. These options are selected while creating or editing a specific object policy and include synchronous (must select at least 2 clusters), deferred, and no copy. For more information, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).

In any of these configuration types of a TS7700, the native TS7700(D) builds the basis with its scalable cache size 157 TB–3.94 PB, or 40 TB/70 TB for the other two use cases, having a physical tape library attached or an object store or cloud.

The type of TS7700 configuration that you choose is influenced by the following factors:

▶ TS7700D

A “disk only” TS7700 provides the fastest way to access your data. Because no recall from physical tape nor external object storage or cloud back into its disk cache is required, you can always access your data immediately. This option can be an organizational requirement.

In cases where you have only a small amount of host data capacity to store (for example, a few hundred TB), this configuration provides a better price point because no other external storage, such as a tape library, drives, tape media, or an object storage, is required.

If you are considering larger capacity configurations (for example, greater than 1 PB), this option might not provide the best price point with its internal disk-based storage rather than a physical tape attachment or object storage. The exact numbers depend on your individual use case and your IBM sales representative can help you determine the best setup for you.

▶ TS7700T

A “backend tape” driven TS7700 focuses on migrating its internal data out to a physical tape media. Therefore, the size of the internal disk cache is typically configured as small as possible but as large as needed to satisfy greater 90% of cache hits for read access.

Physical tape provides some unique advantages over the other options. First, data on a physical tape media has a real air gap because no electrical connection exists. This natural air gap provides the highest available security against attacks and willful manipulation.

Also, because physical tape still provides the best TCO of all available storage products, it can be a cost-effective solution, depending on your environment. Also, physical tape media is portable and can be taken out or moved somewhere else, if needed.

▶ TS7700C

Similar to a TS7700T, the TS7700 focuses on migrating its internal data in this use case out to an external object storage or cloud. This storage can be on-premises, off-premises, or both.

Similar to physical tape, this type of operation does make sense above a specific amount or capacity of managed data. It still can provide a better price point compared to a large TS7700 internal disk cache only, but still does not provide the best TCO compared to a physical tape.

However, use cases might exist where your organization operates a large object storage and the TS7700C can participate, or storing the tape data in a cloud to exchange it easily with another location in a different geographical area. TS7700 R5.1 added several functions and enhancements for the backend cloud attachment of a TS7700.

For more information, see *IBM TS7700 R5.4 Cloud Storage Tier Guide*, [REDP-5573](#).

3.3.2 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 if cluster outage occurs.
- ▶ If a data center is lost, data loss is tolerable or a recovery from Copy Export tapes or an object store or cloud environment is feasible (time and DR site).
- ▶ You can plan and manage outages for Licensed Internal Code loads or upgrade reasons.

If you cannot tolerate any of these items, consider implementing a grid environment with at least a 2-way grid configuration.

3.3.3 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 IBM Z 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, even every individual virtual tape, 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 or cloud pools within a TS7700T/C. 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 or backend object storage and cloud.

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, consider the following points:

- ▶ The TVC is used in common in a TS7700D or a TS7700T/C CP0. You cannot define a physical limit to the amount of space that a client is using in the TVC, unless that client is assigned an exclusive CPx partition in a TS7700T/C. However, through policy management, you can use preference groups differently in these models or the removal policies can be configured differently, which gives more TVC priority to some workloads over others.
- ▶ 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 situation is typically handled through the tape management system (TMS). For DFSMSrmm, this situation 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 situation 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 situation 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, Selective Device Access Control, FC 5271.

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 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 information, 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 situation is valid for the LPAR where the command was run and or 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. For more information, see 3.4, "Features and functions available only for the TS7700T and TS7700C" on page 139.

Note: With APAR OA49373 (z/OS V2R1 and later), the individual IBM MVS LIBRARY command functions (EJECT, REQUEST, DISPDRV, and others) can be protected by using a security product such as RACF. This APAR adds security product resource-names for each of the LIBRARY functions.

3.3.4 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, it 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 [TS7700 Library Request Command](#).

TVC selection might also influence the Copy Export behavior. For more information, see 15.1, “Copy Export overview and considerations” on page 784.

3.3.5 Copy Consistency policy

Define the consistency policy for each Management Class (MC). For more information, see 2.4.5, “Copy consistency points” on page 70.

Consider the following points:

- ▶ 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** is selected, the local mount is an R if it was an N.
- ▶ A consistency point in example [D, D, D, D] means that the selected TVC is treated as RUN, and the extra copies are created asynchronously (Deferred).

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 in example [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 task can be done intentionally to create a remote version as the initial instance of a volume.

If you are not concerned about which TVC is chosen and you prefer a balanced grid with deferred copies, use [D, D, D, D].

By using Time Delayed Replication policy, you can 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 time. Time Delayed Replication policy is targeted especially for multi-cluster grids (3 or more).

You can specify that this type of data is replicated only if the data is still valid and the specified time (after creation or last access) expired. This specification might reduce replication traffic, and the backend activities in TS7700s. In such scenarios, a more far away remote DR site in a vault is most likely targeted.

However, plan to have at least two copies for redundancy purposes, such as on a local TS7700D/TS7700T/C and a remote TS7700D/TS7700T/C.

3.3.6 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 situation 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 DFSMSHsm 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 a private mount.
- ▶ Open only one instance on a private mount.
- ▶ Open both instances on z/OS implied update.

Plan carefully the usage of this option. Dual open is necessary for workloads that can append to tapes. When only reads are occurring, the dual open can introduce unnecessary resource use, especially when one or more of the instances requires a recall from a physical tape. The use of the dual open z/OS implied update helps reduce any resource use to only those mounts where an update of the data 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 sync-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 2.4.5, “Copy consistency points” on page 70.

3.3.7 Override policies

Override policies overrule the explicit definitions of Copy policies.

Note: Synchronous mode is not subject to the 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 TS7700 cluster chooses a TVC selection during the mount operation, and whether a copy must 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 situation if you use this override.

3.3.8 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 group 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.

3.3.9 Logical Volume Delete Expire Processing versus previous implementations

When a host 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 way for the TS7700 to automatically delete the contents after a period passes.

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 (recently increased up to almost 2.000 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. The same is true if an object store or cloud is used as a repository behind a TS7700.

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 situation 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.

3.3.10 Logical Write-Once, Read-Many retention function

TS7700 Logical Write-Once, Read-Many (LWORM) retention function was first introduced in microcode release R1.6. It provides the capability of the software-emulated Write-Once, Read-Many (WORM) that is available on physical tape media for virtual tape volumes.

The TS7700 LWORM retention function was introduced in microcode release 8.51.1.x (R5.1 PGA1). It provides the capability to include TS7700 enforced retention as part of the LWORM function. Initially, an IBM RPQ/SCORE request was required to request help to activate and configure this function because there was no implementation in the TS7700 Management Interface to let the user do that task. The release of R5.4 includes a redesign of the Management Interface panel for the DATA CLASS configuration. The release of R5.4 also includes a new wizard based configuration of the DATA CLASS parameters and the panels and options so that users can define the specific LWORM retention settings. No help from IBM support is required for configuration.

The retention duration is based on the expiration dates that are encoded in the HDR1 data set labels of the logical volume, fixed durations, or a combination of both.

The virtual tape volumes and the data sets that are written on the tape volumes are managed by a tape management system, such as DFSMSrmm. The tape management system manages which tape volumes contain data and specific retention rules for those volumes.

One of the major functions of the tape management system is to manage how long the data sets on the tape volumes must be retained. IBM Z tape management systems and IBM Z applications work together to determine when a volume expires and can be returned to scratch. Any return to scratch is acknowledged by the TS7700 independent of the LWORM function alone.

With the introduction of TS7700 LWORM retention, an unexpected or early return to scratch can be denied or allowed to continue with hold. The TS7700 uses dates that are encoded in the IBM standard label HDR1 of each data set on tape, configured durations, or a combination of both when determining whether any attempt to return a volume to scratch is too early.

Note: The LWORM Retention function is enabled only when all clusters in the same Grid have a code level of 8.51.1.x or later.

The following choices are available for retention duration at volume creation:

- ▶ Fixed duration with option to extend at volume Modify time (MOD): Each append to this LWORM extends the retention period by the period that is defined (x days or years).
- ▶ HDR1 Tape Management System provided expiration date.
- ▶ Fixed /Added Duration at return to scratch time (dynamic expire-hold).

The following settings allow customized behaviors:

- ▶ Fixed durations on create and modify volume time, including an option of “forever”.
- ▶ How to handle TMS dates that imply “Application Managed”.
- ▶ How to handle cases where no HDR1 date is provided.
- ▶ Whether to allow return to scratch before retention period expires.

If allowed, the volume enters an expire-hold state until the retention period expires, which supports returning to private, if needed.

An option is available to extend or introduce a retention period when returned to scratch, which provides a data class granular expire-hold capability.

As described at the beginning of this section, initially, an RPQ/SCORE was required to set up and enable the LWORM retention function by the IBM support team. If you are still running a release below R5.4 and if you want this function, you must work with your IBM representative. For more information, see [TS7700 LWORM Retention Function User's Guide](#). If you are running R5.4 or later, you can enable and configure the LWORM retention by using the redesigned Data Management class window of the TS7700 Management Interface. For further details, refer to 10.2.4, "Data Classes window" on page 486.

Note: The LWORM retention function combines both functions; that is, the delete expire hold retention with the capabilities of granular data class-driven LWORM usage.

For more information about the setup and use of this function, see [TS7700 LWORM Retention Function User's Guide](#).

3.3.11 Software compression

Starting with R4.1.2, two new software compression options can be selected by using DATACLAS. These methods are in addition to the still available classic compression LZ1 by way of the hardware of the FICON channel adapters:

- ▶ FICON compression: The same compression method that was available for previous releases and is embedded on the FICON adapter (LZ1).
- ▶ LZ4 compression: Compression algorithm that prioritizes processing speed.
- ▶ ZSTD compression: As with LZ4, it is a lossless algorithm, but unlike the latter, it aims for the best possible compression ratio, which might come with a performance tradeoff.

No special feature code is needed for these options to be available, but all clusters in the Grid must run R4.1.2 or later (the presence of lower levels of code in the grid prevents the use of this feature). For more information about compression, see 4.3.16, "Data compression" on page 196.

Plan ahead if you also use Copy Export or plan to use the Grid to Grid migration tool in the future. The receiving clusters/grids need to be capable of reading the compressed logical volumes, so they need to have the same capability (R4.1.2 microcode or later).

Using the software compression at least during the migration period can affect capacity planning. The reduction is not only for GB in cache or on physical tape. It also reduces the amount of data in the premigration queue (FC 5274), the amount of data needs to be copied (might result in a better RPO) and the necessary bandwidth for the physical tapes.

3.3.12 Encryption

Depending on your legal requirements and your type of business, data encryption might be mandatory. The TS7700 supports various types of encryption:

- ▶ Internal disk cache encryption:
 - Internal key managed: No external key manager is required
 - External key managed: The use of an external key manager is required
- ▶ Physical tape encryption (always requires an external key manager)

- ▶ Secure data transfer (for encrypted data replication), which requires the enablement of each TS7700 cluster with FC 5281

Consider the following points:

- ▶ 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 must 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 the disk cache can be enabled against only the entire cache repository. Each individual virtual tape volume that is in the TS7700 disk cache is not encrypted.
- ▶ Physical tape encryption can be enabled at any time. You need the no additional charge FC 9900 for enablement within the TS7700T. If you use an external key manager for physical tape and TVC encryption, the same external key manager instance must be used and only the IBM provided key manager IBM Security Guardium® Key Lifecycle Manager is supported at the time of this writing.
- ▶ TS7760 TVC Encryption

TVC encryption can be enabled at any time. After TVC encryption is enabled, it cannot be disabled. It is based on self-encrypting disk drives (DDMs).

Disk-based encryption can be enabled in the field retroactively on all Encryption Capable hardware. Therefore, enabling encryption can occur after the hardware is configured and used.
- ▶ TS7770 TVC Encryption

All TS7770 configurations also support encryption for data at rest. The TS7700 does not rely on self-encrypting disk drives or DDMs. such as the TS7760. Instead, encryption occurs within the CSB disk cache controller by using hardware acceleration. Because this approach differs from the TS7760 approach, encryption must be enabled at manufacturing time and ordered this way if encryption support is wanted on the TS7770 model.

3.3.13 DS8000 Object Store and TS7700 Advanced Object Store for DS8000

TS7700, with DS8000 and DFSMSHsm, delivers DS8000 Object Store. DS8000 transparent cloud tiering (TCT) provides business efficiencies, flexibility, and a reduction in capital expenses by reducing CPU utilization by more than 50% when archiving large data sets.

DS8000 Object Store (FC 5282) introduced the initial redundancy of offloaded data, which was achieved by using the method and functions of forking from the DS8000 to offload data to two independent TS7700 Clusters in a grid.

Starting with R5.2, the redundancy of offloaded DS8000 data within a TS7700 grid can also be achieved by setting up replication policies with the TS7700 grid, which is similar to the management classes that are used for virtual tape volume replication. This function is called TS7700 Advanced Object Store for DS8000 (FC 5283). It provides a simple and granular method to distribute offloaded data within a TS7700 grid and exceeds the limit of a maximum of two instances from the use of the forking functions.

For more information, see *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-55833.

This feature (FC 5283, which is a successor of FC 5282) enables DS8000 TCT object data and traditional FICON logical volume data to coexist within the same physical TS7700 cluster.

Note: FC 5282 and FC 5283 feature the following requirements:

- ▶ FC 5282: TS7770 VED with microcode level 8.50.x or higher with DS888x 8.5.x code or higher or DS8900 with 9.0.x and higher
- ▶ FC 5283: TS7770 VED with microcode level 8.52.200.x or higher with DS8900 at 9.2.x and higher

No other hardware is required because data movement is done by using GRID interfaces between the DS8000 and the TS7700, and data is logically partitioned out of the resident cache.

The DS8000 can target up to two TS7700 models within a grid for synchronous replication. The use of TS7700 as object storage is up to two times faster than conventional cloud storage systems. DS8000 Object Store is available on TS7770 Model VED servers that are running R5.0 code or higher. The feature was withdrawn from marketing on TS7760 Model VEC.

3.3.14 z/OS Allocation with multiple grids that are connected to a single host

You can connect multiple grids to the same sysplex environment, and define them all in the same storage group. In this case, the mounts are distributed depending on the JES2 selected method (EQUAL or BYDEVICE) to the grids.

This allocation routine is aware if a grid crossed the virtual tape scratch threshold and considers this issue for the mount distribution. All other information (such as TVC usage, premigration queue length, and TVC LOWRANK) is not available and are not used for grid selection.

3.3.15 z/OS Allocation assistance inside a grid

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 2.4.15, “Device Allocation and Allocation Assistance” on page 78.

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.

3.3.16 25 GB and 65 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 and as of R5.4, 65 GB logical volumes are also supported. Using 25 GB or 65 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 simpler.
- ▶ Large multi-volume workloads, such as a large database backup, can be stored with fewer logical tapes.

Consider the following points if you choose to use 25 GB or 65 GB logical volumes:

- ▶ 25 GB or 65 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 the 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 or 65 GB logical volumes can create a large increase in disk cache content, which can result in nonoptimal performance.
- ▶ Depending on the grid network performance, and the number of concurrently running copy tasks involving 25 GB or 65 GB volumes, consider increasing the Volume Copy Timeout value from the default of 180 minutes to 240 minutes if copies are timing out. For 65 GB volumes, the timeout value is linearly scaled. For the scaling, the timeout value is multiplied by 65/25. For example, when 180 is set, then $180 \times (65/25) = 468$ is automatically used for 65 GB volumes. This process can be done by using the library request command `Li Req SETTING COPY TIMEOUT <value>` either from MI or the Host Console.
- ▶ There is a limit on the number of concurrent mounts in Sync mode copy with 25 GB or 65 GB logical volumes. If the number exceeds the limit, then mounts result in an error (CBR4171I RSN=42). The default is 128 for each cluster, and it includes the local and remote mount. If it is a mount with Sync Copy mode and both logical volumes on local and remote clusters are opened, the mount is counted on both clusters. You can change the limit by using the `LI REQ SETTING2,CACHE,MAXLGMC` command.
- ▶ When writing approximately 68 GB of data before compression, a 3-byte counter that counts every 4 KB increments wraps at approximately 68 GB (4,096*x"FFFFFF" = 68,719,472,640). If you want to write larger amounts of data to a single volume, you must set Wrap Support instead of Surface-EOT at 3490 Counters Handling in the Data Class to exceed the limit.
- ▶ The 3490 architecture that the TS7700 emulates, has a maximum number of blocks that can be written to a single volume which is approximately 4.1 million (22 bits = 4,194,304).

Writing beyond this limit results in an End of Volume (EOV) condition. The higher the compression ratio, the larger the required block size.

To avoid any performance effects, review your installation before you use the 25 GB or 65 GB volumes.

3.3.17 Grid resiliency function: “Remote” fence

Grid resiliency gives you the ability to automatically identify a “sick but not dead” condition and take a predefined action.

Unlike other IBM Z availability functions such as System failure management for z/OS LPARs or HyperSwap for disks, this feature does not react in seconds. The grid technology is designed for local implementations and remote data placement, which is often thousands of miles away. Therefore, timeouts and retries must be longer to cover temporary network issues.

Although the function supports defining small threshold parameters, change the default only after you analyze the TS7700 grid environment to prevent a false fence.

The secondary action (isolate the cluster from the network) should be considered only if a clear automated action plan is defined and the effect on your production is fully understood.

3.3.18 Control Unit Initiated Reconfiguration

In R4.1.2, CUIR supports only the planned action for service preparation and can be used only in a grid environment. Using this function can reduce the manual effort that is needed for a service maintenance. This feature is beneficial in these circumstances:

- ▶ The TS7700 clusters are shared between many different z/OS systems

Note: The z/OS host must include APAR 0A52376 with code level V2R2 and later.

- ▶ No checkout tests must be run after microcode loads

In case software products like Automatic Tape Allocation Manager (ATAM) or other third-party vendor products are used, review whether CUIR is beneficial for your environment. After CUIR is used to offline the drives, the usual z/OS online command cannot be used with online the devices after the service is complete.

Therefore, your automation, operating procedures, or both need to be reviewed.

3.3.19 CUIR grid resiliency improvements

The Control Unit Initiated Reconfiguration (CUIR) grid resiliency improvements build upon grid resiliency and CUIR functions that were introduced in R4.1.2

Grid resiliency uses thresholds and other methods to determine a TS7700 cluster in a grid is not healthy and fences the cluster. CUIR provides a way for a cluster to notify attached z/OS hosts to automatically vary off devices ahead of a planned outage.

New with this function is the automated notification of z/OS of unhealthy clusters. When a cluster is fenced as part of grid resiliency, any cluster that is connected to the same hosts as the unhealthy cluster notifies the attached hosts about the unhealthy cluster through CUIR. The hosts then auto-vary offline the devices to the unhealthy cluster.

After the fenced cluster returns to an online operational state, the unhealthy cluster then notifies the host to auto vary the devices back online after it recovers.

Note: CUIR requires code level 8.52.100.XX or later on all clusters in the grid and z/OS APAR OA60929

For more information about CUIR configuration and use, see the [TS7700 Series CUIR User Guide V1.2](#) or later.

3.4 Features and functions available only for the TS7700T and TS7700C

With the introduction of tape and object storage or cloud support behind the TS7700T/C, the following features that are unique to the TS7700T/C are now available:

- ▶ Multiple tape-managed partitions or cloud pools (starting with R5.1)
- ▶ Delay premigration to physical tape or object storage, cloud

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 or cloud. 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 or cloud. This situation 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 the 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 or cloud. This configuration reduces your back-end activities (migrate and reclaim).

Note: Tape-attach feature code 5273 is mutually exclusive with cloud enablement FC 5278 on the same cluster.

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 [IBM TS7700 Notification Messages](#).

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.

Starting with R4.1.2, you can now modify the priority or severity for messages that are presented to the attached LPAR by the TS7700. These messages result from conditions or events requiring some level of interaction with an operator (known as “intervention messages”). The z/OS systems identify the intervention messages with the CBR3750I ID. This modification is used only in your environment. The Call home capability will still use the original priority/severity provided by IBM.

In addition, you can enhance the messages and can extend the message text with user-defined content.

If you modify your environment, back up this modification to make sure you can upload your changes if a microcode issue occurs.

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, consider running extra checks, as described in this section.

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, the following methodologies are available:

- ▶ Host-based migration
- ▶ TS7700 Field-Frame-Replacement, push-pull-MES
- ▶ TS7700 internal data migration

TS7700 Release 3.3 introduced 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 virtual 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 others. 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 Field-Frame-Replacement, push-pull-MES

A tape-attached TS7700 (TS7700T) offers a simple and smooth migration method. For example, to newer hardware generation or to another, empty TS7700 system if there is a data center reallocation. This process benefits from the fact that all managed active data is on physical tapes.

The field-frame-replacement method is done by your IBM SSR. It requires the TS7700 to be in Service Mode (like for any other maintenance activity). This process starts to offload (pre-migrate) not yet pre-migrated virtual volumes to physical tape.

Then, a backup is created of the TS7700 internal file systems, configuration settings, and database. Then, these assets can be used to restore another, empty TS7700T. After the restore, the new TS7700T comes back online and immediately provides access to all your data. All settings and configurations are also retained.

3.6.3 TS7700 internal data migration

With the introduction of Release 3.3, there are two different data-migration methods that are provided by the TS7700 technology:

- ▶ Cluster Join and Copy Refresh Processing
- ▶ The GGM tool

Still other possibilities exist; for example, Host Tape Copy.

Cluster 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 TS7700 cluster.
2. Change the MC contents to allow copies to the new cluster.
3. Adjust your ACS routines so that new data is steered to the new cluster.
4. Keep the old TS7700 cluster for read-only purpose online to the host.
5. 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. No Host I/O moves through the FICON adapters, and all data in the TCDB and tape management remains 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 a simple 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 whether the following situations are true:

- ▶ Eight clusters are installed in the grid, which is the maximum number of clusters in a 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, the 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 more BVIR reports. For more information about this command, see [TS7700 Grid To Grid Migration User's Guide](#).

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](#).

For more information, see Chapter 8, "Migration" on page 303, or ask your local IBM SSR.

3.6.4 Tape drive technology behind a TS7700

Before Release 3.3, all tape drives that were attached to a TS7700 had to be homogeneous. Different tape drive generations could not be intermixed.

With Release 3.3, you can mix the TS1150 with *only one* older drive technology. With Release 5.2.1 PGA1 and starting TS1160 support, you can also mix the TS1160 with *only one* older drive technology. In most use cases it will be an intermix of a TS1150 and TS1140, a TS1160 and TS1140, or a TS1160 and TS1150.

This intermix with TS1140 can be for migration purposes because a TS1150 or TS1160 cannot read content from JA and JB cartridges. Or an intermix with a TS1140 can be because the TS1150 is withdrawn from marketing and a client wants to extend the number of backend physical tape drives (initially started with TS1150) and is required to add TS1160s.

Consider the following points:

- ▶ The earlier technology is used only for reads. You cannot write data on the older 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 for reclaim. However, have enough TS1150 or TS1160 tape drives to allow premigration, recalls, and reclaim for newly written data.
- ▶ There must be at least 4 or more newer generation tape drives.
- ▶ At least 2 older generation tape drives must be still available.
- ▶ Use the **LI REQ** to define the values for the alerts for missing physical drives for both technologies and the TS1150 and TS1160.
- ▶ Run **VEHSTATS** to understand the physical drive behavior.

If running an intermix of TS1160 and TS1150 tape drives consider the following points:

- ▶ The JD media are then written with the new TS1160 tape drives with 15 TB capacity and are also only read by the TS1160 tape drives.

- ▶ The existing JD media with active data, which were written with the TS1150 tape drives, remain in 10 TB format and are only read by the TS1150 tape drives.
- ▶ When one of the existing JD media, which was written with the TS1150 tape drives, goes to *scratch*, it is written by the new TS1160 tape drives with the 15 TB capacity.

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 topics:

- ▶ 4.1, “Hardware installation and infrastructure planning” on page 146
- ▶ 4.2, “Planning for a grid operation” on page 172
- ▶ 4.3, “Planning for software implementation” on page 182
- ▶ 4.4, “Tape analysis and sizing the TS7700” on page 202

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 include system requirements and infrastructure requirements.

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 more information, see [System requirements](#).

IBM 3948 and 3952 Tape Frame specifications

A TS7700 is enclosed by a 3948 or 3952 Tape Frame.

A 3948 or 3952 Tape Frame contains the TS7700 Server, TS7700 Cache, and the TS3000 TSSC, and additional components such as switches and power distribution units for connection to independent power sources. The 3952 Tape Frame model F06 or F07 can house either a TS7760 or a TS7770. The 3948 Tape Frame model F07 can house a TS7770.

Also, a 3948 or 3952 Tape Frame can be designated as a TS7770 Storage Expansion Frame to expand the capacity of a TS7700 disk-only Cluster.

Encryption-capable TS7770 physical frame F07 characteristics (FC 7339) are listed in Table 4-1.

Table 4-1 Encryption-capable TS7770 physical frame F07 characteristics (FC 7339)

Measurement	Base Frame	Expansion Frame each (up to 2)	Client Rack Mount (minimum/maximum configuration)
Width	616 mm (24.25 in.)	616 mm (24.25 in.)	483 mm (19.0 in.) of each component
Depth	1274 mm (50.16 in.)	1274 mm (50.16 in.)	Depth of the deepest component 3948-VED or 3757-VED (p9 pSeries) from front to back of the cable management arm and cables in the arm 38.9 inches (98.8 am)
Height	1930 mm (76 in.)	1930 mm (76 in.)	<ul style="list-style-type: none"> ▶ 2-disk module configuration 18U = Minimum configuration ▶ 4 disk module configuration 22U ▶ 6 disk module configuration 26U ▶ 8 disk module configuration 30U ▶ 10 disk module configuration 34U = Maximum configuration
Weight	Maximum 706.7 Kg (1558 lbs)	Maximum 900.8 Kg (1986 lbs)	<ul style="list-style-type: none"> ▶ The minimum weight of components for configuration with 2 disk cache drawers is 243 kg (536.4 lbs) ▶ The maximum weight of components for configuration with 10 disk cache drawers is 456 kg (1,006 lbs)
Power consumption	Maximum 3480 watts	Maximum 4400 watts	max 3480 watts

Measurement	Base Frame	Expansion Frame each (up to 2)	Client Rack Mount (minimum/maximum configuration)
Power phase options	Single phase (240 V AC) or three phase (400 V AC)		Single phase only (240 V AC)
Power frequency	50 Hz–60 Hz (+/- 3 Hz)		
Dry bulb temperature (recommended operating range)	20°C–25°C (68°F–77°F)		
Relative humidity (recommended operating range)	40%–55%		

The 3952 Tape Frame F06 houses the components of older TS7700 models; that is, TS7760. The dimensions of the frame that encloses the TS7700 are listed in Table 4-2.

Table 4-2 Physical characteristics of a maximally configured 3952 F06 Tape Frame

Characteristic	3952 F06
Height	1930.4 mm (76 in.)
Width	616 mm (24.25 in.)
Depth	<ul style="list-style-type: none"> ▶ Closed doors: 1425 mm (56.1 in.) ▶ Open doors (front/rear): 2515 mm (99 in.)
Weight	746 kg (1645 lb.) maximum configuration
Power	240 Vac, 15 amp (single phase)
Unit height	40 U

A 3948 or 3952 Tape Frame that is configured as a 3948 or 3952 Base Frame contains two 1 Gb intranet Ethernet switches that are used for private communication between components within a TS7700 Cluster. A customer network connects directly to the TS7700 Server and uses virtual IP address technology to one or more customer-provided IP addresses across two IBM AIX configured Ethernet connections.

Rack mount

An RPQ is released for the clients to install a TS7770 into their own client-supplied 19-inch rack. RPQ 8B3721 enables clients to order one TS7770 Server 3948 or 3957 Model VED, one Cache Controller 3956 Model CSB, and up to nine Cache Module Model XSBs. RPQ 8B3749 enables clients to order one TS7770 Server 3948 Model VED, one Cache Controller 3948 Model CFC, and up to nine Cache Module Model XFCs (all flash cache) and all other associated features to be installed in their own client-supplied rack.

Environmental operating requirements

Your facility must meet specified temperature and humidity requirements before you install the TS7700. The preferred environmental conditions for the TS7700 are listed in Table 4-3 on page 148.

Table 4-3 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 3948 or 3952 Tape Frame includes one internal power distribution unit. However, FC 1903 for VEC and FC 1913 (before FC 1912 which was discontinued from marketing) for VED, Dual AC power, is required to provide two power distribution units to support the high availability (HA) characteristics of the TS7700. The 3952 Storage Expansion Frame has two power distribution units and requires two power feeds.

FC 1913 is also the new foundation PDU for 3-phase power. To move to 3-phase power, you add the iRPQ8B3723 (3 phase PDU power cords).

TS7770 Base Frame power requirements

Your facility must ensure an available power supply to meet the input voltage requirements for the TS7770 Base Frame. Table 4-4 displays the maximum input power for a fully configured TS7770.

Table 4-4 TS7770 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 (watt)	3480 watts
Input power required	4.8 kVA (single phase)
Thermal units	12.2 kBTU/hr, 3.06 kcal/hr

TS7770 Storage Expansion Frame power requirements

Your facility must ensure an available power supply to meet the input voltage requirements for the TS7770 Storage Expansion Frame. Table 4-5 displays the maximum input power for a fully configured TS7770.

Table 4-5 TS7770 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 (watt)	4400 watts
Input power required	4.8 kVA (single phase)
Thermal units	15.4 kBtu/hr, 4.0 kcal/hr
Exhaust capacity	1030 m3/hr
Noise level	59 db

TS7760 Base Frame power requirements

Your facility must have ample power to meet the input voltage requirements for the TS7760 Base Frame. The maximum input power for a fully configured TS7760 is listed in Table 4-6.

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. The maximum input power for a fully configured TS7760 is listed in Table 4-7.

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)

Power requirement	Value
Current	24 amp
Leakage current	13.5 ma
Inrush current	250 amp
Power (W)	3200 watts
Input power required	4.8 kVa (single phase)
Thermal units	11.2 kBtu/hr, 2.9 kcal/hr
Exhaust capacity	750 m ³ /hr
Noise level	59 db

TS7770 specifications and requirements

The component specifications are listed in Table 4-8.

Table 4-8 TS7770 component specifications

Specification	TS7770 Server		TS7770 Cache Subsystem			
	3948-VED or 3957-VED	I/O expansion drawer (each of 2)	TS7770 cache controller 3948-CFC or 3956-CFC	TS7770 cache expansion drawer 3948-XFC or 3956-XFC	TS7770 cache controller 3948-CSB or 3956-CSB	TS7770 cache expansion drawer 3948-XSB or 3956-XSB
Width	482 mm (18.97 in)	440 mm (17.32 in.) Total width of two joined drawers	483 mm (19.0 in)			
Depth	766.5 mm (30.2 in)	800 mm (31.5 in.)	556 mm (21.9 in)			
Height	86.7 mm (3.4 in)	220 mm (8.66 in.)	87 mm (3.4 in)			
Weight	30.4 kg (67 lbs)	42.2 kg (92.8 lb.) weight of two combined. Reflects weight of each, without cassettes	27.7 kg (61.0 lb.) fully configured	26.7 kg (58.7 lb.) fully configured	27.7 kg (61.0 lb.) fully configured	26.7 kg (58.7 lb.) fully configured
Power	180–240 V AC 50–60 Hz (+/- 3 Hz)	180–240 V AC 50/60 Hz	198–264 RMS V AC 50–60 Hz			
Unit height	2 U	5 U	2 U	2 U	2 U	2 U
Temperature (non-operating)	5°C–45°C (41°F–113°F)	Included in 3952 Tape Frame maximum	-10°C–65°C (14°F–149°F)	5°C–45°C (41°F–113°F)	-10°C–65°C (14°F–149°F)	5°C–45°C (41°F–113°F)

Specification	TS7770 Server		TS7770 Cache Subsystem			
Type and Model	3948-VED or 3957-VED	I/O expansion drawer (each of 2)	TS7770 cache controller 3948-CFC or 3956-CFC	TS7770 cache expansion drawer 3948-XFC or 3956-XFC	TS7770 cache controller 3948-CSB or 3956-CSB	TS7770 cache expansion drawer 3948-XSB or 3956-XSB
Temperature (operating)	5°C–40°C (41°F–104°F) Depends on altitude. For more information, see Table 4-9 on page 152 Suggested: 18°C–27°C (64°F–80°F)	Included in 3952 Tape Frame maximum	10°C–32 °C (50°F–89.6°F)			
Relative humidity	Non-operating: 8%–85% Operating: 20–60%	Included in 3952 Tape Frame maximum	Non-operating: 8%–80% Operating: 20%–80%			
Maximum wet bulb (power off)	Non-operating: 28°C (82°F)	Included in 3952 Tape Frame maximum	23°C (73.4°F)			
Power consumption	1400 watts (2X) maximum	Included in 3952 Tape Frame maximum	810 watts	540 watts	406 watts	213 watts
Thermal output	6655 Btu/hour (maximum)	Included in 3952 Tape Frame maximum	2730 Btu/hr	1820 Btu/hr	1329 Btu/hr	730 Btu/hr
Power source loading	1.6 kVA (maximum config.) 3,048 m (10,000 ft.) maximum altitude	Included in 3952 Tape Frame maximum	0.9 kVa	0.6 kVa	0.409 kVa	0.240 kVa

The Temperature Depending on Altitude requirements are shown in Table 4-9.

Table 4-9 Temperature Depending on Altitude

Temperature (operating)	Altitude
5–40°C (41–104°F) Note: The minimum operating temperature cannot go below 5°C (41°F).	Up to 3050 meters @ 5°C–28°C (41°F–82°F)
	Up to 2875 meters @ 29°C (84°F)
	Up to 2700 meters @ 30°C (86°F)
	Up to 2525 meters @ 31°C (88°F)
	Up to 2350 meters @ 32°C (89°F)
	Up to 2175 meters @ 33°C (91°F)
	Up to 2000 meters @ 34°C (93°F)
	Up to 1825 meters @ 35°C (95°F)
	Up to 1650 meters @ 36°C (96°F)
	Up to 1475 meters @ 37°C (98°F)
	Up to 1300 meters @ 38°C (100°F)
	Up to 1125 meters @ 39°C (102°F)
Up to 950 meters @ 40°C (104°F)	

TS4500/TS3500 Tape Library attachment

A TS7770 Tape Attach or TS7760 Tape Attach attached to a TS4500 or TS3500 Tape Library interfaces directly with tape drives in the library.

When attached to a TS4500 or TS3500 Tape Library, the TS7700 can attach to only 3592 Tape Drives. Up to 16 3592 Tape Drives can be attached. The tape drives must be in either Model L22, L23, D22, or D23 frames in the tape library. Each frame can contain up to 12 3592 Tape Drives.

Communication, control, and data signals travel along Fibre Channel connections between the TS7700 and tape drives that are contained in the TS4500 or TS3500 Tape Library. A pair of Fibre Channel switches routes the data to and from the correct tape drive.

To ensure correct operation of 3592 Tape Drives in a TS3500 or TS4500 Tape Library that is attached to the TS7700, you must:

- ▶ Designate four 3592 Tape Drives as control paths by using the tape library management GUI.
- ▶ If you use an 8 Gb Fibre Channel switch, verify that the ports for the 3592 Tape Drives are set up as Auto (L), not Auto (N).
- ▶ If you use a 16 Gb Fibre Channel switch, verify that the ports for the 3592 Tape Drives are set up as Auto (N), not Auto (L).

Simple Tape Attach

Simple Tape Attach (STA) provides a new method for writing and reading to physical tape media from a TS7700 disk-only solution.

The STA feature can be added only to a TS7700 disk-only solution and provides new methods for writing and reading to physical (LTO) tape media residing in a TS4300 library. This is an alternative solution to the TS7700 Tape Attach feature and copy export and fits within a single rack. This feature is designed for limited and specific use cases and uses new LIBRARY REQUEST commands (STAxxxx) to initiate and to track status.

For more information, see Chapter 18, “IBM TS7700 support for zTape Air-GAP” on page 871 and the white paper [TS7700 zTape Air-Gap FC5995 Users Guide](#).

Tape drives and media support (TS7700T)

The TS7700T supports the following cartridges:

- ▶ 3592 Tape Cartridge (JA)
- ▶ 3592 Expanded Capacity Cartridge (JB)
- ▶ 3592 Advanced Tape Cartridge (JC)
- ▶ 3532 Advanced Data Tape Cartridge (JD)
- ▶ 3532 Advanced Data Tape Cartridge (JE)
- ▶ 3592 Economy Tape Cartridge (JJ)
- ▶ 3592 Economy Advanced Tape Cartridge (JK)
- ▶ 3592 Economy Tape Cartridge (JL)
- ▶ 3592 Economy Tape Cartridge (JM)

The TS7700T supports the 3592 Extended Tape Cartridge (JB) media and requires 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, TS1150, and TS1160 with the corresponding media are supported.

The TS7700 tape encryption (FC 9900) requires that all the backend drives be encryption capable. TS1130 Model E06/EU6, TS1140 Model E07 or EH7, and TS1150 Model E08 or EH8 drives, and TS1160 Model 60F/60G are encryption-capable. TS1120 Model E05 Tape Drives in E05 mode are encryption-capable, with FC 9592 from the factory, or FC 5592 as a field upgrade.

For more information about restrictions for use with TS1140 and TS1150 Tape Drives, see Chapter 7, “Hardware configurations and upgrade considerations” on page 261.

For the media, format, and drive model compatibility to see which tape drive model is required for a specific capability, see Table 4-10 on page 154.

Table 4-10 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	EFMT6
Model J1A	Read/write	Not supported	Not supported	Not supported	Not supported	Not supported
Model E05	Read/write ^a	Read/write	Not supported	Not supported	Not supported	Not supported
Model E06/EU6	Read	Read/write	Read/write	Not supported	Not supported	Not supported
Model E07/EH7	Read ^b	Read ^b	Read/write ^c	Read/write	Not supported	Not supported
Model E08/EH8	Not supported	Not supported	Not supported	Read/write	Read/write	Read/write
Model 60F/60G	Not supported	Not supported	Not supported	Read	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.

The tape drive models, capabilities, and supported media are listed by tape drive model in Table 4-11.

Table 4-11 3592 Tape Drive models and characteristics versus supported media and capacity

3592 drive type	Supported media type	Encryption support	Capacity	Data rate
TS1160 Tape Drive (3592-60F/60G Tape Drive)	<ul style="list-style-type: none"> ▶ JC ▶ JD ▶ JE ▶ JK ▶ JL ▶ JM 	Yes	<ul style="list-style-type: none"> ▶ 7 TB (JC native) ▶ 15 TB (JD native) ▶ 20 TB (JE native) ▶ 900 GB (JK native) ▶ 2 TB (JL native) ▶ 5 TB (JM native) 	400 MBps
TS1150 Tape Drive (3592-E08/EH8 Tape Drive)	<ul style="list-style-type: none"> ▶ JC ▶ JD ▶ JK ▶ JL 	Yes	<ul style="list-style-type: none"> ▶ 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)	<ul style="list-style-type: none"> ▶ JB ▶ JC ▶ JK <p>Media read only:</p> <ul style="list-style-type: none"> ▶ JA ▶ JJ 	Yes	<ul style="list-style-type: none"> ▶ 1.6 TB (JB native) ▶ 4.0 TB (JC native) ▶ 500 GB (JK native) ▶ 4.0 TB (maximum all) 	250 MBps

3592 drive type	Supported media type	Encryption support	Capacity	Data rate
TS1130 Tape Drive (3592-EU6 or 3592-E06 Tape Drive)	<ul style="list-style-type: none"> ▶ JA ▶ JB ▶ JJ 	Yes	<ul style="list-style-type: none"> ▶ 640 GB (JA native) ▶ 1.0 TB (JB native) ▶ 128 GB (JJ native) ▶ 1.0 TB (maximum all) 	160 MBps
TS1120 Tape Drive (3592-E05 Tape Drive)	<ul style="list-style-type: none"> ▶ JA ▶ JB ▶ JJ 	Yes	<ul style="list-style-type: none"> ▶ 500 GB (JA native) ▶ 700 GB (JB native) ▶ 100 GB (JJ native) ▶ 700 GB (maximum all) 	100 MBps
3592-J1A	<ul style="list-style-type: none"> ▶ JA ▶ JJ 	No	<ul style="list-style-type: none"> ▶ 300 GB (JA native) ▶ 60 GB (JJ native) ▶ 300 GB (maximum all) 	40 MBps
<p>Notes: Consider the following points:</p> <ul style="list-style-type: none"> ▶ To use tape encryption, all drives that are associated with the TS7700T must be Encryption Capable and encryption-enabled. ▶ Encryption is not supported on 3592 J1A tape drives. 				

The following media identifiers are used for diagnostic and cleaning cartridges:

- ▶ CE: IBM service representative (formerly 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 TS7700T 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 TS7700T.

When you change the model of the 3592 tape drives of a TS7700T, 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 TS7700 tape drives with older tape drives to TS1160 tape drives.

For more information, see 7.2.4, “Upgrading drive models in a TS7700T” on page 279.

4.1.2 TS7700 specific limitations

Consider the following restrictions when you perform your TS7700 preinstallation and planning:

- ▶ Cloud Storage Tier capabilities can be enabled only in a TS7700 that is *not* attached to a physical tape library.
- ▶ The TS7760 must be at a microcode level of 8.42.x.x or later for cloud support. If the TS7760 is in a grid, then all clusters need to be at 8.42.x.x or later.
- ▶ The TS7760 memory must be upgraded to 64 GB to support Cloud Storage Tier. This is done by using the FC 3466: 32 GB Memory Upgrade.

- ▶ Release 5.2 Phase 1 is supported on TS7760 and TS7770, whereas Release 5.2 Phase 2 or later is supported on TS7770 only (see code update recommendation).

For more information, see [TS7700 Code Update Recommendation](#).

- ▶ TS1120 Tape Drives set in static emulation mode are not supported by the TS7700T. 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 that uses short wavelength attachments at the 4 Gbps speed is up to 150 meters by using 50-micron fiber cable, and up to 55 meters by using 62.5-micron fiber.
- ▶ The distance between a TS7770 Storage Expansion Frame and the TS7770 Base Frame cannot exceed 10 meters. This distance permits connection of the frames by using a 30-meter cable.
- ▶ At 8 Gbps speed, the short wave total cable length cannot exceed the following measurements:
 - 150 meters by using 50-micron OM3 (2000 MHz*km) aqua blue-colored fiber
 - 50 meters by using 50-micron OM2 (500 MHz*km) orange-colored fiber
 - 21 meters by using 62.5-micron OM1 (200 MHz*km) orange-colored fiber
- ▶ At 16 Gbps speed, the short wave total cable length cannot exceed the following measurements:
 - 130 meters by using 50-micron OM4 (4700 MHz*km) aqua blue-colored fiber
 - 100 meters by using 50-micron OM3 (2000 MHz*km) aqua blue-colored fiber
 - 35 meters by using 50-micron OM2 (500 MHz*km) orange-colored fiber
- ▶ Long wavelength attachments (4 Gb, 8 Gb, or 16 Gb) provide a direct link of up to 10 km (6.21 miles) between the TS7700 and the host processor on 9-micron fiber.
- ▶ Short and long wavelength attachments provide for up to 100 km (62.1 miles) between the TS7700 and the host processor by using appropriate fiber switches, and up to 250 km (155.3 miles) with DWDMs. Support is not provided through more than one dynamic switch.
- ▶ 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 (328 feet).
- ▶ The TS7700 does not support capacity scaling of 3592 tape media.
- ▶ The TS7700 does not support physical WORM tape media.
- ▶ The TS3500/TS4500 and TS7700 must be within 30 meters (100 feet) of the TSSC.
- ▶ The 3592 back-end tape drives for a TS7700T cluster must be installed in a TS3500 or TS4500 Tape Library. TS7700T can be connected to only backend tape drives through 16 Gb fiber switches.
- ▶ 3948-VED or 3957-VED cannot be joined into a grid containing a 3957-VEA.

For all models of TS7700, 8.54.x.x microcode supports the ability to have a VED 8.54.x.x cluster (new from manufacturing or empty through a manufacturing cleanup process) join into an existing grid with a restricted mixture of 8.5x.x.x (VEC/VED) clusters.

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. Because one new cluster can be joined in a 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.
- ▶ A cluster cannot be joined or merged to a grid having a write-protected cluster or FlashCopy is active on.

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 also covered.

TS7700 grid and cloud 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 when they are online. Two or four 1-GbE or 10-GbE connections must be in place before grid installation and activation, which includes the following equipment:

- ▶ Intranet Ethernet switches

Intranet Ethernet switches are used primarily for private communication between components within a cluster.

- ▶ ATM switches
- ▶ Ethernet extenders and adapters

An Ethernet extender or other extending equipment is used to complete extended distance Ethernet connections. The following extended grid Ethernet connections can be used directly from the TS7700 cluster:

- 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 (328 feet) by using four pairs of CAT5, CAT6, or CAT7 copper cabling.

Note: 1 Gb Optical SW Grid Ethernet adapter is *not* supported anymore on the TS7700 3948-VED or 3957-VED.

- 10 Gbps optical LW

This 10 Gbps grid optical LW connection provides single or dual port options and a 10-Gbps Ethernet LW adapter for grid communication between TS7700 tape systems. This adapter includes an LC Duplex connector for attaching a 9- μ , single-mode fiber cable. This standard LW (1310 nm) adapter conforms to the IEEE 802.3ae standards. It supports distances up to 10 km (6.21 miles).

The default configuration for a TS7700 server from manufacturing is two dual-ported PCIe 1-GbE adapters having one port of each adapter set active. The second port on each adapter can be activated by using another chargeable feature code (FC 1034 Enable Dual Port Grid Connection) to increase redundancy and grid link replication performance.

For a 3948-VED, 3957-VED, or 3957-VEC, you can use FC 1041 (successor of FC 1038), which is a 10 Gb, dual-port optical LW connection to add support for two 10 Gb optical longwave Ethernet adapters. This feature improves data copy replication with minimum bandwidth redundancy. Configured clusters that use two 10 Gb, four 1 Gb, or two 1 Gb clusters can be interconnected within the same TS7700 Grid. Adapter types can be mixed within a grid, but not within a single cluster.

Tip: For a general overview and summary of the current TS7700 options and specification of inter-connectivity (FICON and Grid), see [TS7700 FICON, Backend-Firre, and GRID specifications 2024 v1.10](#)

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. 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 clusters in a grid must include 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 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 inter-cluster 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
- ▶ PCF (Priority Control Flow) is not supported
- ▶ 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, which results in fewer dropped packets.

Remember: Packets are lost when the grid network capabilities are below TS7700 capabilities. This loss causes TCP to stop, resync, and resend data, which results in a less efficient use of the network. Flow control helps to reduce this behavior. Clusters of 1 Gb and 10 Gb 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.

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 the following requirements:

- ▶ Sufficient bandwidth exists to account for all network traffic that is expected to be driven through the system to eliminate network contention.
- ▶ The flow control between the TS7700 clusters and the switches is supported, 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 handle the data rates that are expected from all 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 element is the most important element for an extended distance topology.

In short, latency between the sites is the primary factor. However, packet loss because of 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 virtual volumes, access virtual 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 copper (RJ-45) or software fiber (single-ported or dual-ported) links. For copper networks, CAT5E, CAT6, or CAT7 Ethernet cabling can be used, but CAT7 cabling is preferable to achieve the highest throughput. Alternatively, two or four 10-Gb LW fiber Ethernet links can be provided.

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
- ▶ 10.10.10.1
- ▶ 10.10.10.2

For TS7700 clusters that are 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

Note: 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.

DNS must be configured in the Cluster Network Settings (under the Management Interface corresponding window) if the selected Cloud Object Store is Amazon S3.

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: For new deployments, it is a preferred practice that each of the up-to-four grid interfaces exists on discrete subnets. Plan different subnets for each grid interface. If the grid interfaces are directly connected (without the use of Ethernet switches), you must use separate subnets.

Local IP addresses for Management Interface access

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. All three provided IP addresses are assigned to one TS7700 cluster for MI access.

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.

Provide one gateway IP address and one subnet mask address.

Note: If FC 9900, Tape Encryption configuration, is installed, this same connection is used for communications between the TS7700 and the Encryption Key Server. Because encryption occurs on attached physical tape drives, encryption support does not exist for non-tape attached instances of the TS7700, and the virtual connection is used exclusively to create redundant paths.

Each cluster in the grid must be configured in the same manner, with three TCP/IP addresses providing redundant paths between the local intranet and cluster.

Connecting to the Management Interface

This section describes how to connect to the IBM TS7700 MI. The following browsers are supported:

- ▶ Microsoft Edge
- ▶ Microsoft Internet Explorer 11
- ▶ Mozilla Firefox
- ▶ Google Chrome

Note: Make sure that your browser is dated so that it includes the latest fixes and security enhancements.

Complete the following steps to connect to the interface:

1. In the address bar of a supported web browser, enter “https://” followed by the *virtual IP* that was entered during the installation process. The virtual IP is one of three IP addresses that are provided during installation. The complete URL uses the following form:

https://<virtual IP address>

2. Press **Enter** or click **Go** in your web browser.

The web browser then automatically 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>`, which does not require an update after a cluster ID change.

The login page for the MI loads. The default login name is admin and the default password is admin.

For the list of required TSSC TCP/IP port assignments, see Table 4-12 on page 163.

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 automatically goes 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 4-1 shows a sample setup for a two-cluster grid.

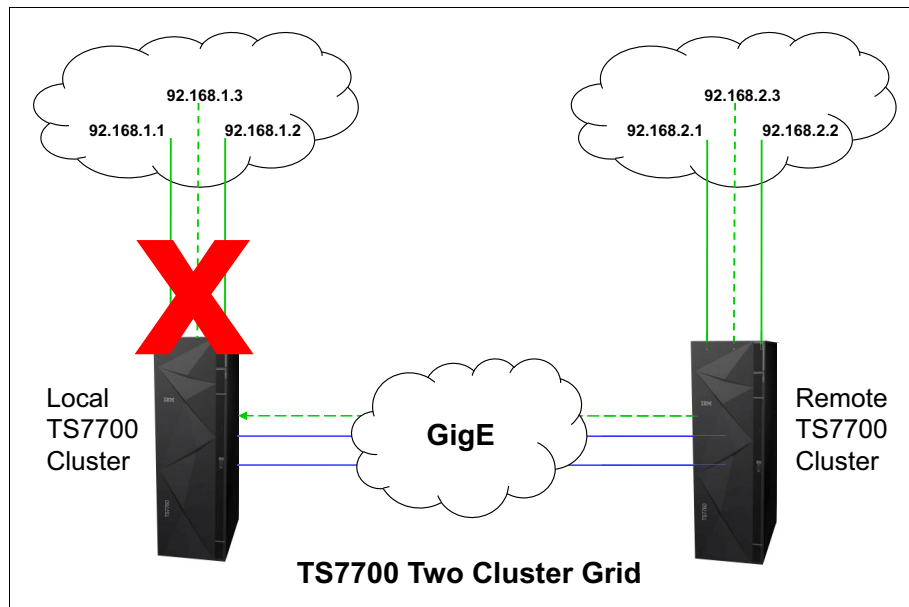


Figure 4-1 TS7700 Management Interface access from a remote cluster

IPv6 support

Internet Protocol Version 6 (IPv6) is the next evolution in Internet Protocol beyond the IPv4 standard currently in use in most networks today. The key IPv6 enhancement is the expansion of the IP address space from 32 to 128 bits, which enables almost unlimited IP addresses. This addressing capability, along with new functions that enable end-to-end security, improve mobility support, and simplify address configuration and management, make IPv6 a critical component in the evolution of e-business and the next generation internet.

The following network interfaces that support monitoring and management functions can now support IPv4 or IPv6:

- ▶ Management Interface (MI)
- ▶ Key manager server: IBM Security Key Lifecycle Manager
- ▶ Simple Network Management Protocol (SNMP) servers
- ▶ Rocket-fast System for Log processing (RSYSLOG) servers
- ▶ Lightweight Directory Access Protocol (LDAP) server
- ▶ Network Time Protocol (NTP) server

Setting up IPv6 in a TS7700 environment requires the following special set of instructions and requirements:

- ▶ Supported TS7700 Servers include 3957-V07, 3957-VEB, 3957-VEC, 3957-VED, and 3948 VED.
- ▶ Grid cluster link function is not supported for use with IPv6.
- ▶ The dynamic host configuration protocol (DHCP), which allows a server to assign network IP addresses, is not supported for use with IPv6.
- ▶ You cannot simultaneously enable IPv4 and IPv6.
- ▶ For grid configurations, *each* cluster can be 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 more information about implementation, see “Enabling IPv6” on page 597.

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.

Note: Starting with R5.3 pga1 in the 2nd half of 2023, the TSSC/TS3000 Service Console is part of the 3948-VED server ship-group. This means the TSSC/TS3000 specific feature codes within the F07 are discontinued and are not listed anymore in a configuration.

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 the other from the alternative switch). If your cluster is a TS7700T, 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 this [IBM Documentation web page](#).

Generally, at least one TSSC should be 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

These two Ethernet adapters are used by advanced functions, such as AOTM, LDAP, Assist On-site (AOS), and Call Home. If you plan to use them, provide one or two Ethernet connections and the corresponding IP addresses for the TSSC. The ports in must be opened in the firewall for the interface links to work correctly. The network port requirements for the TSSC are listed in Table 4-12.

Table 4-12 TSSC TCP/IP port requirements

TSSC interface link	TCP/IP port	Role
TSSC External	80	Call Home
	443	
	22	SSH
	123	NTP
	162	SNMP trap
	389	LDA
	415	RSYSLOG
	636	LDAP with TLS
	16311	LDAP with SAS
	53	Advise to remain open for the domain name server
	Internet Control Message Protocol (ICMP)	
TSSC Grid	443	Autonomic Ownership Takeover Mode (AOTM)
	80	HTTP
	22	Recommended to remain open
	9666	
	ICMP	

For more information about Call Home/AOS configuration, see [IBM TS7700, TS4500, Diamondback and DS8000 System Connectivity and Security v5.04](#).

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 listed in Table 4-13.

Clarification: These requirements apply only to the LAN/WAN infrastructure. The TS7700 internal network is managed and controlled by internal code.

Table 4-13 Infrastructure grid WAN TCP/IP port assignments

Link	TCP/IP port	Role
TS7700 MI	ICMP	Ping - Dead gateway detection
	53	DNS
	80	HTTP
	123 ^a	NTP uses the User Datagram Protocol (UDP) time server
	162	SNMP Trap
	389	LDAP
	415	RSYSLOG
	441	IBM Security Guardium Key Lifecycle Manager TLS (R5.0)
	443 ^b	Access the TS7700 MI (HTTPS)
	636	LDAP with TLS
	3801	IBM Security Guardium Key Lifecycle Manager Server
	5696	KMIP TLS (R5.0)
	16311 ^c	LDAP using SAS
TS7700 GRID	ICMP	Check cluster health
	9	Discard port for speed measurement between grid clusters
	22 ^{bd}	SSH
	80	HTTP
	123 ^a	NTP time server
	1415/1416	IBM WebSphere® message queues
	443 ^b	Access the TS7700 MI (HTTPS)
	350	TS7700 file replication, Remote Mount, and Sync Mode Copy (distributed library file transfer)

- a. Port 123 is used for grid link-time synchronization within clusters, not for an external time server.
- b. The use of ports 20/21 and 80 were deprecated; ports 22 and 443 must be used instead.
- c. System Storage Productivity Center (SSPC) support was deprecated.
- d. Used by PFE/Development during remote access (recommended).

For more information, see [IBM TS7700, TS4500, Diamondback, and DS8000 System Connectivity and Security v5.04](#).

4.1.4 Factors that affect performance at a distance

Fibre Channel distances depend on the following 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 feature a maximum distance of 150 meters (492 feet) with 50-micron diameter, multi-mode, optical fiber (at 4 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 55 meters (180 feet). 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. For more information, see Table 4-14 on page 167.

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 the 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 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. A match must exist 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 and a switch port with only eight buffer credits can read at full performance over a greater distance than it can write because during 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 existed to allocate one buffer credit for every 2 km (1.24 miles) to maintain full performance.

Buffer credits within the switches and directors play a large part 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 that is covered. For more information, see *IBM SAN Survival Guide*, SG24-6143.

4.1.5 Host attachments

The TS7700 attaches to IBM Z hosts through the FICON adapters on the host (FICON LW or SW) at speeds of 4 Gbps, 8 Gbps, or 16 Gbps. Connection speeds of 1 Gbps and 2 Gbps are no longer supported by the newest 16 Gb FICON adapters.

Consider the following points:

- ▶ ESCON channel attachment is not supported.
- ▶ FICON channel extension and DWDM connections are supported.
- ▶ FICON directors and director cascading are supported.

Note: Considerations for host FICON connections:

- ▶ IBM Z servers that use 8 Gbps FICON ports do support only TS7700 connections that are running 4 Gbps and 8 Gbps speeds for direct attachments. However, 2 Gbps connections to TS7700 are also supported if FICON Director provides proper speed conversion.
- ▶ IBM Z 16 Gbps FICON supports only TS7700 16 Gbps and 8 Gbps FICON direct-attached.
- ▶ IBM Z 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 4-2 on page 167 shows a sample diagram that includes the DWDM and FICON Directors specifications. For more information, see “FICON Director support” on page 168.

The maximum distances vary depending on the cable type and on the speed and type of optical transducer. The following basic types of optical cable fiber are available:

- ▶ The orange-colored cables are SW, multimode OM2 type cables.
- ▶ The aqua-colored multimode cables are OM3 and 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.

The relationship between connection speed and distance by cable type is listed in Table 4-14 on page 167.

Figure 4-2 on page 167 shows the supported distances by 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 IBM Z servers by using IBM FICON at distances up to 250 km (155 miles) by using dense wavelength division multiplexing (DWDM) with switches, or more extended distances by using supported channel extension products.

Distances greater than 30 km (18.6 miles) require DWDM 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 (62–155 miles).

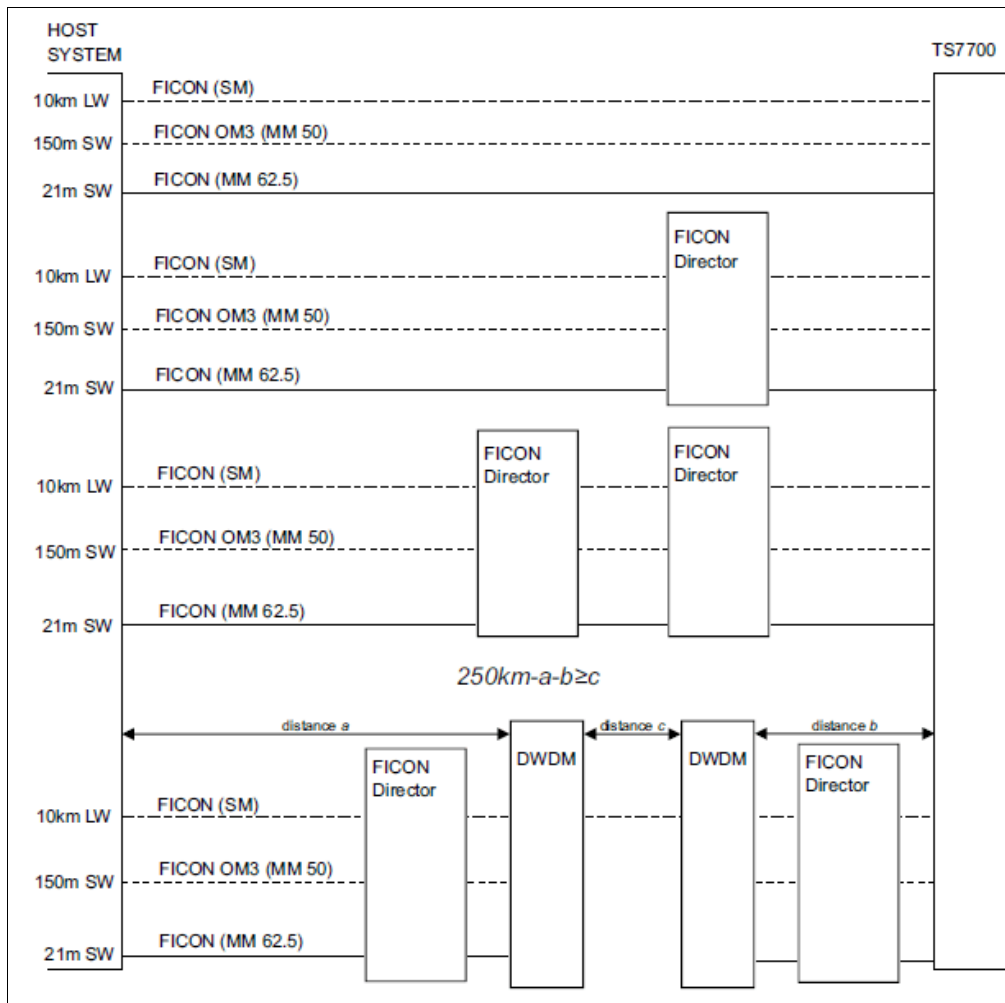


Figure 4-2 IBM Z host attachment to the TS7700 (at speed of 8 Gbps)

Table 4-14 Connection speed and distance by cable type

Cable type	Connection Speed	Maximum Distance
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.)
OM2	16 Gbps	35 m (115 ft.)
OM3	16 Gbps	100 m (328 ft.)
OM4	16 Gbps	130 m (426 ft.)

Note: Long wave cables attach only to long wave adapters and short wave cables attach only to short wave adapters. Intermixing cannot be done.

FICON Director support

All FICON Directors are supported for single and multi-cluster grid TS7700 configurations where code level 5.3 is installed with 2 Gbps, 4 Gbps, 8 Gbps, or 16 Gbps links. The components auto-negotiate to the highest speed allowed. The 16 Gbps ports cannot negotiate down to 2 Gbps links.

You cannot mix different vendors, such as Brocade (formerly McData, CNT, and InRange) and CISCO, but you can mix models of one vendor.

For more information about specific supported intermix combinations, see [System Storage Interoperation Center \(SSIC\)](#).

To view the FICON switch support matrix, see [Switches and Directors qualified for IBM zSystems FICON and FCP channels](#).

FICON channel extenders

FICON channel extenders can operate in one of the following modes:

- ▶ Frame shuttle or tunnel mode
- ▶ Emulation mode

By 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 end (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 IBM Z host and the TS7700 are in the same matrix as the FICON switch support. For more information, see the FICON Channel Extenders section in [Switches and Directors qualified for IBM zSystems FICON and FCP channels](#).

Cascaded switches

Consider the following general configuration rules for configurations with cascaded switches:

- ▶ Director Switch ID

This switch 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 were 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 address 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 4-3 on page 170 shows an example of host connection that uses DWDM and cascaded switches.

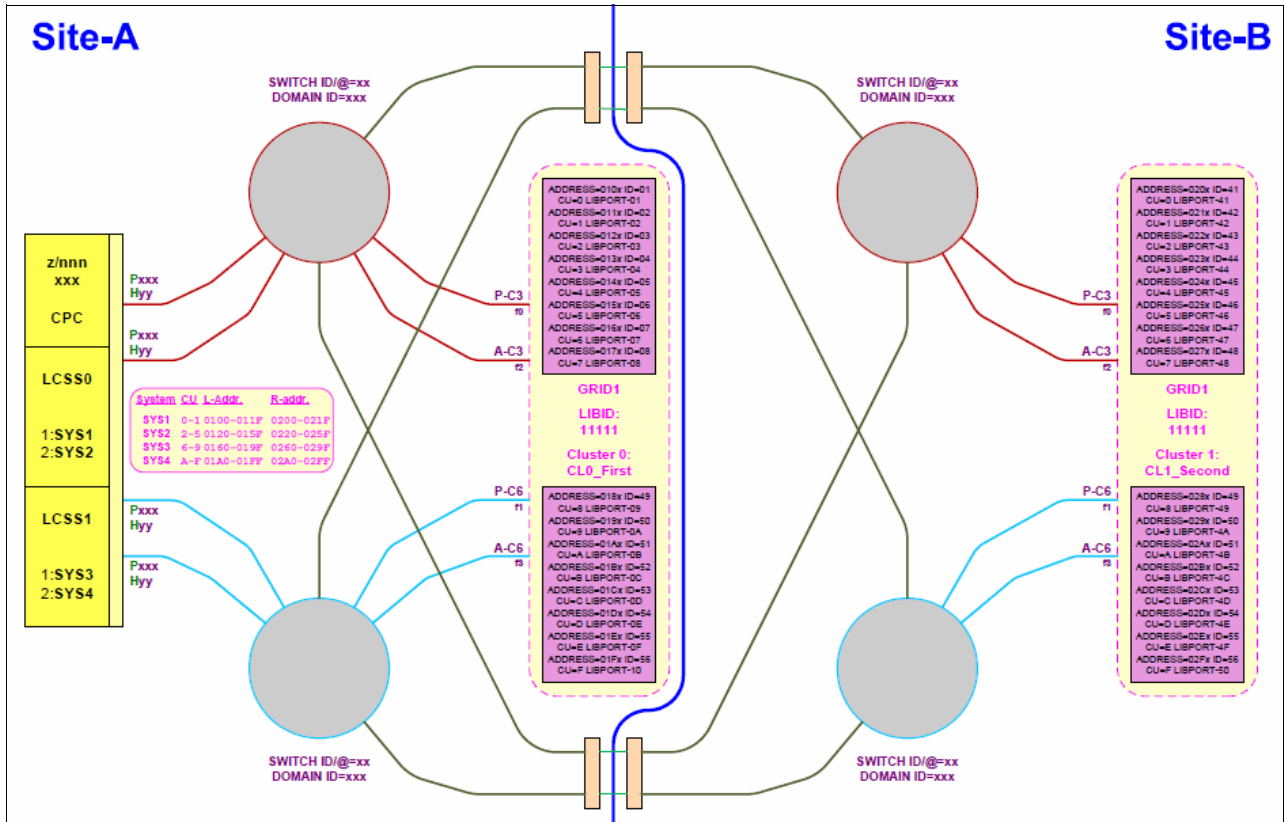


Figure 4-3 Host connectivity that uses DWDM and cascaded switches

4.1.6 Planning for LDAP

Depending on the security requirements in place, the user of the TS7700 can choose to have all 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. Within the LDAP-Policy definition, you can define that IBM Service personnel do not need to use LDAP-Credentials to access the system. This configuration is recommended.

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, the TS7700 clusters and the TSSC include native support for the LDAP server (Microsoft Active Directory [MSAD] is supported). Release 3.1 adds support for IBM Tivoli Directory Server for IBM Z along with IBM Resource Access Control Facility (RACF).

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, username, and password that you created for IBM SSRs and keep this information easily available in case you need it. 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 597.

4.1.7 Cluster time coordination

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

External NTP is required when any of the grid members are configured to use the Cloud Storage Tier because Time Synchronization is demanded for the cloud TS7700C interaction.

Figure 4-4 shows the NTP server configuration in grid.

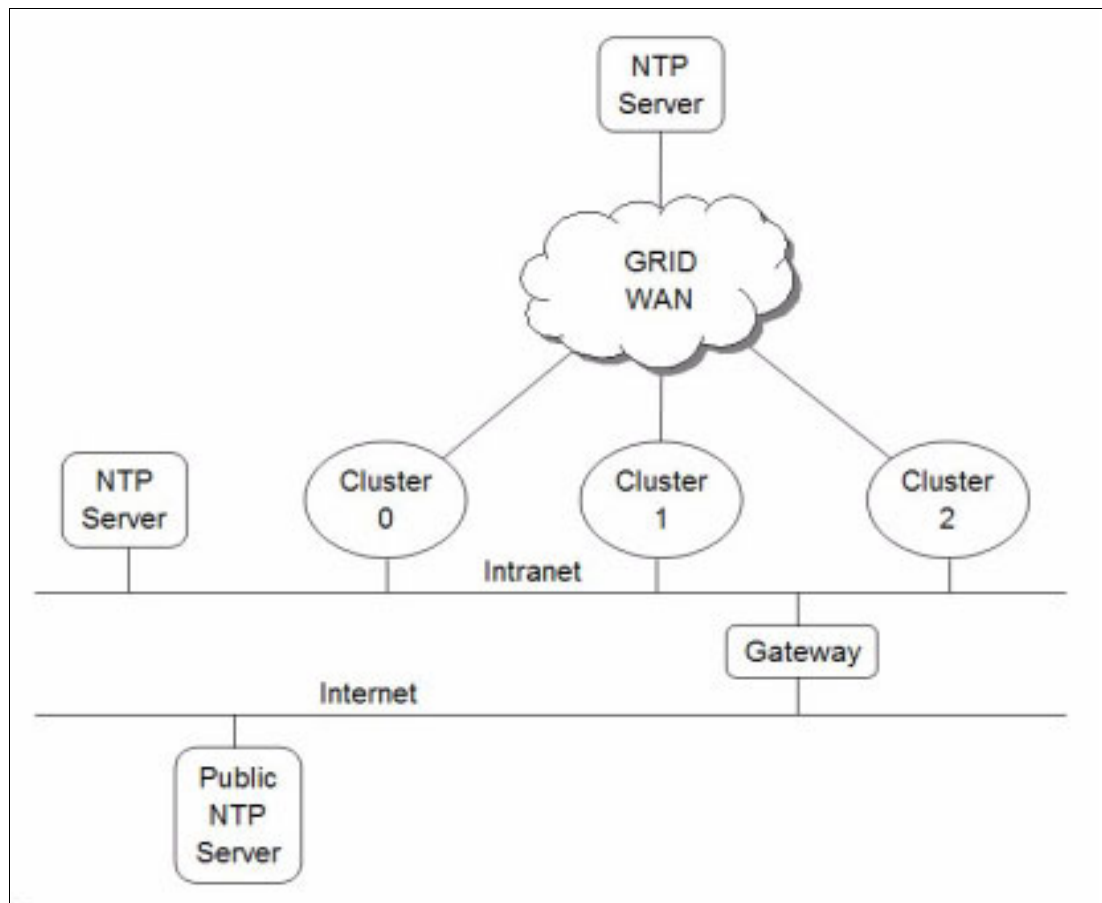


Figure 4-4 NTP server configuration

The NTP server address is configured into system VPD on a Grid-wide scope so that all clusters access the same NTP server. All the clusters in a grid must 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 (or the lowest-numbered available cluster in the grid).

4.2 Planning for a grid operation

The TS7700 grid provides configuration flexibility to meet various requirements. Those requirements depend on your business and 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 considerations

The Autonomic Ownership Takeover Manager (AOTM) is an optional function which, following a TS7700 cluster failure, automatically enables one of the methods for ownership takeover without operator intervention, which improves the availability of the TS7700. It uses the TS3000 System Console that is associated with each TS7700 to provide an alternative path to check the status of a peer TS7700.

Without AOTM, an operator must determine whether one of the TS7700 clusters failed, and then enable one of the ownership takeover modes. This operation is required to access the virtual volumes that are owned by the failed cluster.

Note: It is important that write ownership takeover is enabled only when a cluster fails, and not when a problem exists 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, which results in a corruption of the data. Although no data corruption issue exists with the read ownership takeover mode, it is possible that the remaining clusters might not have the latest version of the virtual volume and present previous data.

Even if AOTM is not enabled, it is advised that it be configured. Doing so 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 failed or is otherwise not operating. If a TS7700 suspects that the cluster that owns a volume it needs failed, it asks the TS3000 System Console to which it is attached to query the System Console that is attached to the suspected failed cluster.

If the remote system console can validate that its TS7700 failed, it replies 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 be enabled only by an operator.

To take advantage of AOTM, the customer must provide IP communication paths between the TS3000 System Consoles at the cluster sites. For AOTM to function correctly, it cannot share paths as the grid interconnection between the TS7700s.

AOTM can be enabled through the MI interface. It is also possible to set the default ownership takeover mode.

4.2.2 Defining grid copy mode control

When upgrading a stand-alone cluster to a grid, FC 4015, 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.

The following definitions of Copy Consistency Points in a multi-cluster grid are described:

- ▶ No Copy (NC): No copy is made to this cluster.
- ▶ Rewind Unload (RUN): A valid version of the virtual volume was copied to this cluster as part of the volume unload processing.
- ▶ Deferred (DEF): A replication of the modified virtual volume is made to this cluster after the volume is 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 88.
- ▶ Time Delayed: This policy enables better control of what data must be replicated to other clusters in the grid. For example, 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 Tape-attached cluster for later than the average Lifecycle of your data.

Then, most of the data expires before the time set for the delayed copies runs out, which avoided the processor burden that was introduced by the replication of archive or short retention data, and later the extra reclamation activity on the Tape-attached cluster. The time delay can be set to 1–65,535 hours. For more information, see the following web pages:

- [IBM TS7700 Series Best Practices - TS7700 Hybrid Grid Usage](#)
- [IBM TS7700 Series Best Practices - Copy Consistency Points](#)
- [IBM TS7700 Series Best Practices - Synchronous Mode Copy](#)

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 that are tailored to meet your requirements. The settings take effect for any mount requests that are received after you save the changes. Mounts that are 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 apply.

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 override 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 partitions (CP0) of a TS7700Ts for fast access, while other workloads can be directed to TS7700Ds or the cache partition (CPx) of TS7700Ts for archival purposes.

Clusters that are 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 SAA is enabled but not selected any cluster as SAA candidates in the Management Class, all clusters are treated as SAA candidates.

Understand that SAA influences only the mount behavior of the grid. Although other clusters can be selected as mount point if the original SAA clusters are not available or not configured to the host, they are not considered for the TVC selection. If all clusters that are specified in the Management Class as target are not available, the mount might be processed, but the job hangs afterward.

Before SAA is operational, the SAA function must be enabled in the grid by using the **LI REQ SETTING SCRATCH ENABLE** command.

4.2.4 Retain Copy mode

Retain Copy mode is an optional attribute, which is controlled by a Management Class construct configuration in which a volume's previously existing Copy Consistency Points are accepted rather than applying the Copy Consistency Points defined at the mounting cluster. This attribute applies to private volume mounts for read/write appends. It is used to prevent more copies of a volume from being created in the grid than wanted. This attribute is important in a grid with three or more clusters that includes 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 into families. By using cluster families, you can define a common purpose or role to a subset of clusters within a grid configuration. For example, the role that is assigned (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 clusters are used outside of its family.

4.2.6 TS7700 cache thresholds and removal policies

These thresholds determine the state of the cache as it relates to remaining free space.

Cache thresholds for a TS7700 cluster

There are three thresholds that define the capacity of CP0 in a TS7700T and the active cache capacity in a TS7700D. These thresholds determine the state of the cache as it relates to remaining free space.

The following thresholds are available in ascending order of occurrence:

- ▶ **Automatic Removal**

The policy removes the oldest logical volumes from the TS7700D and CP0 in TS7700T 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 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. You can also lower the removal threshold in the LI REQ. The default is 4 TB.

To perform removal operations in a TS7700T, the size of CP0 must be at least 10 TB. Consider the following points:

- You can disable automatic removal within any specific TS7700D cluster by using the following **LIBRARY REQUEST** command, as shown in the following example:

```
LIBRARY REQUEST,library-name,CACHE,REMOVE,{ENABLE|DISABLE}
```

- The default automatic removal threshold can be changed from the CLI by using the following library request command:

```
LIBRARY REQUEST,library-name,CACHE,REMVTHR,{VALUE}
```

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 less than 3 TB of free space is available 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 TS7700D 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 [TS7700 Cache thresholds and removal policies](#).

Clarification: Host writes to the TS7700 and inbound copies continue during this state.

► Out-of-cache resources

This state occurs when less than 1 TB of free space 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 TS7700D 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 be displayed:

- HYDME0997W
- HYDME1133W
- HYDME1201W

For more information about these messages, see [TS7700 Cache thresholds and removal policies](#).

Clarification: New host allocations do not choose a TS7700D Cluster in this state as a valid tape volume cache candidate. New host allocations that are issued to a TS7700D Cluster in this state choose a remote tape volume cache instead. If all valid clusters are in this state or cannot accept mounts, the host allocations fail.

Read mounts can choose the TS7700D Cluster in this state, but modify and write operations fail. Copies inbound to this cluster are queued as deferred until the cluster exits this state.

The start and stop thresholds for each of the active cache capacity states that are defined are listed in Table 4-15.

Table 4-15 Active cache capacity state thresholds

State	Enter state (free space)	Exit state (free space)	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	<ul style="list-style-type: none"> ▶ CBR3792E upon entering state ▶ CBR3793I upon exiting state
Out of cache resources (CP0 for a TS7700T)	<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	<ul style="list-style-type: none"> ▶ 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 TS7700 cluster in a grid configuration. These policies provide more control over the removal of content from a TS7700 cache as the active data reaches full capacity. To perform removal operations in a TS7700T cluster, the size of CP0 must be at least 10 TB.

To help ensure that data is always in a TS7700, 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 TS7700 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 TS7700D and for CP0 within a TS7700T. 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 TS7700 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 TS7700D cache, or for a TS7700T CP0.

Above this threshold, virtual volumes are removed from a TS7700 cache.

Note: Virtual volumes are removed only if another consistent copy exists within the grid.

Virtual volumes are removed from a TS7700 cache in the following 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 TS7700 continues to remove virtual volumes until the stop threshold is met. The stop threshold is a value that is the removal threshold minus 500 GB. A particular virtual volume cannot be removed from a 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 is not yet completed, the virtual 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 visible only if the selected cluster is a 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 in anticipation of a service mode event, or before a DR test where FlashCopy for DR testing is used.

Virtual 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, which prevents the removal of enough logical volumes. This scenario can quickly lead to the TS7700 cache reaching its maximum capacity.

The lower threshold creates more free cache space, which enables the TS7700 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 TS7700 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 TS7700 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 TS7700 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 TS7700 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 starting the temporary removal process does not lower its own removal threshold during this process.

4.2.7 Data management settings (TS7700T CPx in a multi-cluster grid)

The following settings for the TS7700 are optional, and they can be configured during the installation of the TS7700, or later by using the TS7700 **LIBRARY REQUEST (LI REQ)** command interface:

- ▶ Copies to Follow Storage Class Preference (COPYFSC)
- ▶ Recalls Preferred to be Removed from Cache (RECLPG0)

Note: For more information about the Host Console Request functions and their responses, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

Copies to follow Storage Class Preference (COPYFSC)

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 when a copy to physical media is completed (through the premigrate/migrate process). Consider the following points:

- ▶ *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. Preferred to stay in cache 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. Preferred to be removed from cache is also referred to as *Preference Group 0* (PG0).

Applied Cache Preference Group (PG0/PG1) to be applied depends on the Storage Class construct that is associated to individual virtual volumes. Therefore, it is possible to use different Cache Preference settings for different virtual volumes.

Note: Construct names are assigned to virtual volumes by the attached Host System. They are used to establish Data Management policies to be run by the TS7700 against specific volumes. Constructs (and associated policies) are defined in advance by using the TS7700 MI. (For more information, see “Defining TS7700 constructs” on page 593.) If the Host System assigns a construct name without first defining it, the TS7700 creates the construct with the default parameters.

In the default Storage Class Case, for a TS7700 running in a dual production multi-cluster grid configuration, virtual tape drives in both TS7700 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.

However, 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 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 business requirements, this behavior can be modified for a TS7700 by using the COPYFSC control, as shown in the following example:

```
LI REQ, <distributed-library>, SETTING, CACHE, COPYFSC, <ENABLE/DISABLE>
```

This control features the following characteristics:

- ▶ The default is set to disabled.
- ▶ When disabled, virtual volumes that are copied into the cache from a peer TS7700 are managed as PGO volumes (prefer largest files out of cache first), regardless of local definition of the associated Storage Class construct.
- ▶ When set to enabled, virtual volumes that are copied into the cache from a peer TS7700 are managed by using the actions that are defined for the Storage Class construct that is associated with the volume, as locally defined.

Recalls preferred for cache removal (RECLPG0)

Normally, a volume that is 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 the 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 the volumes to restore are resident in the cache. Therefore, some number of recalls are required. Unless you can explicitly control the sequence of volumes to be restored, it is likely that recalled volumes displace cached volumes that are not yet restored from, which results in more 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 are accessed so that they minimally displace other volumes in the cache.

Based on business requirements, this behavior can be modified by using the RECLPG0 setting, as shown in the following example:

```
LI REQ, <distributed-library>, SETTING, CACHE, RECLPG0, <ENABLE/DISABLE>
```

This setting features the following characteristics:

- ▶ When disabled, which is the default, virtual volumes that are recalled into cache are managed by using the actions that are defined for the Storage Class construct that is associated with the volume as defined in the local TS7700.
- ▶ When enabled, recalls are managed as PGO volumes (prefer out of cache first by largest size), regardless of the local definition of the associated Storage Class construct.

4.2.8 High availability considerations

High availability (HA) means providing continuous access to virtual 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) are 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 must be varied on in a planned or unplanned outage to enable production to continue.

- ▶ For the workload placed on the grid configuration, performance throughput must be sufficient to meet service level agreements (SLAs) when only one of the clusters is used. 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 the use of an MC whose Copy Consistency Point definition includes 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 a zero recovery point objective (RPO) at sync point granularity.

The copy is updated at the same time as the original volume, which keeps both instances of this logical volume synchronized at the record level. For more information, see Chapter 2, “Architecture, components, and functional characteristics” on page 15.

- ▶ 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 TS7700D grids. For more information, see 14.5, “Considerations for Virtual Device Allocation” on page 764.
- ▶ To minimize operator actions when a failure occurs 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 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 failed, an operator must determine whether a cluster failed and by using 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 approach 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 online) and an ownership takeover mode was established manually or through AOTM.

If it is necessary, access data to complete the job. For more information about AOTM, see 2.4.33, “Autonomic Ownership Takeover Manager” on page 97. 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) settings are defined at the system level. Therefore, if you modify them in one cluster, it applies to all clusters in that grid.

4.2.9 Planning for cloud operation

The following tasks must be performed by the cloud administrator:

- ▶ Choose the cloud provider and obtain its URL so it can be associated to a cluster in the grid to perform the premigrated tasks to cloud.
- ▶ Obtain the access credentials to configure TS7700 cloud accounts.
- ▶ Determine the cloud space to serve as cloud containers for the TS7700.

The following tasks must be performed by using the TS7760C management interface:

- ▶ Define the cloud pool
- ▶ Define the non-resident cache partition

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 4-16, 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 4-16 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 249. In the same manner, the distributed library name is linked to the distributed library ID, as shown in Figure 6-9 on page 250. Use names that are similar to those listed in Table 4-16.

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 simpler, 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, following 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 listed in Table 4-16 on page 183.

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 sequence number is the *distributed library ID*. It is arbitrary, can be selected by you and 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 of hexadecimal characters only. 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 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 that are participating in a grid have the same composite library ID. In the example that is described in “Distributed library ID”, the composite library ID starts with a “C” for this five hex-character sequence number.

The last four characters can be used to uniquely identify each composite library in a meaningful way. The sequence number must match the LIBRARY-ID 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 IBM Z 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.

The definitions of the LIBPORT-IDs in a multi-cluster grid are listed in Table 4-17 on page 185. For Cluster 0, 256 devices are 01–10 and 496 devices are 01–1F. LIBPORT-ID is always one more than CUADD.

Table 4-17 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'
6	0–1E	X'A1'–X'BF'
7	0–1E	X'E1'–X'FF'

Virtual tape drives

The TS7700 presents a tape drive image of a 3490 C2A, which is 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 V2R4 or later (earlier release level support must be done through the RPQ process).

There are no required host software APARs that are required for TS7700 support at Release 4.2 and higher, but there is some additional host support that is delivered through APARs. For more information, see [TS770 Virtualization Engine](#).

In general, install the host software support. For more information about software maintenance, see [IBM Support and Downloads web page](#). Also, refer to 8.2.1, “Field frame replacement migration for TS7700T” on page 309.

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 to which libraries the allocation can be directed.

In addition to defining new SMS classes in z/OS, the new SMS classes (constructs) must be defined in the TS7700 through the MI. This definition is necessary because each of the class types (DC, SC, MC, and SG) includes a corresponding meaning and use for the volume within the TS7700.

When a mount occurs for a volume and it is written to from the beginning of the tape (file sequence 1, DISP=NEW), the host passes the assigned SMS classes to the TS7700 and the TS7700 assigns those classes (constructs) to the volume. Figure 4-5 shows the default MC in the first line and another MC defined as described in the second line.

Select	Name	Secondary...	Description	Scratch moun...	Retain copy mo...	"Valkyrie[0]" (#BA...
<input type="checkbox"/>	-----		0 The default Management Class	C0, C1, C2, C3, C5, C6, C7	No	No Copy
<input type="checkbox"/>	CE32		32 Copy export MC	C0, C1, C2, C3, C5, C6, C7	No	Deferred

Figure 4-5 Default construct

4.3.4 Sharing and partitioning considerations

This section includes the following topics:

- ▶ Tape management system and OAM
- ▶ Partitioning the physical media in a TS7700T or TS7700C 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 and 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 can access a shared disk, an RMMplex is not necessary and is not recommended.

For more information about which RMM subsystem is best for your environment, see the following publications:

- ▶ *DFSMSrmm Primer, SG24-5983*
- ▶ *z/OS DFSMSrmm Implementation and Customization Guide, SC23-6874*

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 (VOLCATs), to manage volume and library entries.

If tape libraries are shared among hosts that are all intended to have access to the same volumes, they must all use a single TCDB on a shared disk. Each such TCDBplex must have a unique **DEVSUPxx** parmlib member that specifies library manager categories for each scratch media type, error, and private volumes.

It is always recommended to change the IBM-supplied DEVSUPxx parmlib member to have site-specific values. Even if this member is shared between all hosts, it is advisable not to use the default categories because the use of such categories complicates debugging of any cartridge entry problems.

Planning what categories are used by which hosts is an important consideration that must 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 between multiple hosts

The virtual drives and virtual volumes in a TS7700 can be partitioned just like physical drives and stacked volumes, although by way of different mechanisms. Any virtual volume can be directed to any physical stacked volume when you use a TS7700T. The TS7700 places no restrictions on the use and management of those resources. When you use a TS7700T, you can partition your stacked volumes in up to 32 separate pools by assigning a volume range to a pool and that pool to an SG before insertion time of the physical stacked volume.

4.3.5 Library Manager Category Usage Considerations

To partition a library among multiple TCDBplexes requires separation of the scratch pools; that is, each TCDBplex must have a separate library manager category for each scratch media type. For logical completeness, the error and private volume categories should also be unique to each TCDBplex.

To change the default category assignments, specify the categories in parmlib member DEVSUPxx. The category specification parameters enable the installation to change the default category assignments that are associated with a system or sysplex, or both.

It is the responsibility of the installation to ensure that all systems or sysplexes (or both) that are associated with the same TCDB (TCDBplex) use the same category assignments. For more information about the partitioning-related DEVSUPxx parameters, see *z/OS MVS Initialization and Tuning Reference*, [SA23-1380](#).

In a partitioned library, it is recommended that the installation uses DEVSUPxx to change the default categories that are associated with each TCDBplex. Therefore, because no TCDBplex uses the default categories, no volumes are in those categories.

If the DEVSUPxx parameters are inadvertently removed from one system, scratch mount requests are directed to the empty default categories and the mount requests fail. If a TCDBplex uses the default categories, volumes can be mounted by the system where the DEVSUPxx parameters were removed.

If a scratch volume from a default category is mounted on the system where the parameters were removed, it is not used because no tape volume record is in the TCDB. The volume is assigned to the error category with resultant disruption in library operations in the TCDBplex that owns the default categories.

4.3.6 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 resembles 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 the following 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 an IBM system-managed library means that all attached hosts have the same access to all volumes in the tape library. To achieve this sharing, you must share the host CDSs; that is, the TMS inventory and the TCDB, among the attached hosts.

Also, you must have the same categories that are defined in the DEVSUPxx member on all hosts. 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.

4.3.7 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 simplest 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. Historically, you needed to change the DEVSUPxx member and restart the system to modify a category value. Today, the use of the **DS QLIB,CATS** command 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 when a VOLSER range is defined. As a result, selective device protection applies retroactively to pre-existing data. For more information about a case study about sharing and partitioning the TS7700, see Appendix I, "Case study for logical partitioning of a two-cluster grid" on page 991.

4.3.8 Logical path considerations

The TS7700 attaches to the host system or systems through two or four FICON adapters. For 16 Gb and 8 Gb FICON adapters, each channel that is connected to the FICON adapter port supports 512 logical paths. The 4 Gb FICON adapter continues to support 256 paths per port. A four FICON (16/8 Gb adapter) configuration with dual ports enabled results in a total of 4,096 logical paths per TS7700, as shown in the following example:

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.

For more information about planning and implementing FICON channels, operating in FICON native (Fibre Channel [FC]) mode, and the FICON and FC architectures, terminology, and supported topologies, see *FICON Planning and Implementation Guide*, SG24-6497.

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.

4.3.9 Secure Data Transfer

Secure Data Transfer (SDT) provides a way to securely read and write logical volume data between clusters within a grid.

Logical volume copies are encrypted only when the encryption is enabled on both ends of a copy transaction. If one of the pairs does not have encryption enabled, the logical volume is not encrypted.

SDT uses OpenSSL software libraries with the TLS1.2 protocol following AES standards. AES-256 and AES-128 bit keys are supported. Logical volume data is encrypted within the TS7700 before transport, so special network requirements are not needed.

During logical volume transfers, each TS7700 can be used as a client or a server, depending on the direction of data travel. The client always starts the data transfer request.

As a server, AES-256 and AES-128 are always supported. As a client, the key size that is used depends on the selection that is made at the Secure Data Transfer on the TS7700 Management Interface page when SDT is enabled. During the key exchange process, the highest common key is used. SDT does not result in a performance degradation. SDT also supports the use of client-provided SSL certificates.

4.3.10 MI dual control

TS7700 offers MI dual security authentication to operatively reduce the risk of human error for accidental deletion and malicious damage. This function enhances the following robust security features that are engineered into each TS7700:

- ▶ Cybersecurity enhancements to help businesses avoid malicious activity from within the company, and pass audit and compliance requirements.
- ▶ Specific sensitive modifications can be double checked by a separate user.
- ▶ Expire hold enablement and duration-protected functions with a 'maker' and 'checker' approach.
- ▶ Solution for multi-tenancy clusters

For more information, see Chapter 9, "IBM TS7700 Management Interface operations: Part 1" on page 351, and Chapter 10, "IBM TS7700 Management Interface operations: Part 2" on page 449.

4.3.11 Planning for logical and physical volumes

As part of your planning process, you must 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 that is 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.

4.3.12 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 all storage hierarchies. It must also be unique across LPARs that are 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 is not 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 are 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.

4.3.13 Virtual volumes

Determine the number of virtual 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 (FC 5270), up to a total of 4,000,000 logical volumes.

Tip: For 3957-V06/VEA, the maximum is 2,000,000 virtual volumes, which is also the limit for the overall grid that contains clusters that correspond to one of those models.

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 virtual volumes

The TS7700 supports logical volumes with maximum sizes of 400 MiB, 800 MiB, 1000 MiB, 2000 MiB, 4000 MiB, 6000 MiB, 25000 MiB, and 65000 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 virtual 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 virtual volumes, you need more TS7700 virtual 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 that 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 MiB, 2000 MiB, 4000 MiB, 6000 MiB, 25000 MiB, or 65000 MiB sizes.

A virtual 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 the 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 virtual volume sizes does not affect the real space that is used in 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 are needed, first determine the number of private volumes that make up the current workload that you are migrating. One way to

determine this number is by looking at the amount of data on your current volumes and then matching that figure 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. Consider that fact when planning the number of logical volumes required. Consider the use of smaller volumes for applications, such as DFSMSHsm and larger volumes for backup and full-volume memory dumps.

If you plan to use 25000 MiB virtual volumes, a maximum size of 25000 MiB for virtual volumes is allowed without any restriction if all clusters in a grid operate at R3.2 or later level of Licensed Internal Code. Also, 65000 MiB virtual volumes are allowed in a grid operating at R5.4 or later. 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 a VTS installed, it is possible that you determined this number, and are therefore can 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:

$$V_v = C_v + T_r + (S_c)(S_i + 1)$$

The formula contains the following values:

- V_v** Total number of needed logical volumes.
- C_v** Number of logical volumes that is needed for current data rounded up to the nearest 10,000.
- T_r** 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.
- S_c** Number of of empty scratch volumes that are used per night, rounded up to the nearest 500.
- S_i** 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 must 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 S_i value in the formula.

If you define more volumes than you need, you can always eject the extra 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 number is the maximum in a stand-alone or grid configuration.

For more information about making 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 274.

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 virtual volumes in SCRATCH status to support the data that is written to the TS7700. This amount can be hundreds or thousands of volumes, depending on your workload. More than a single day’s worth of scratch volumes should be available at any time.

Return-to-scratch processing

Return-to-scratch processing involves running a set of tape management tools that identify the logical volumes that no longer contain active data, and then communicating with the TS7700 to change the status of those volumes from private to scratch.

The amount of time the process takes depends on the type of TMS that is used, how busy the TS7700 is when it is processing the volume status change requests, and whether a grid configuration is 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 are 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 the 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 is 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

TS7700T and TS7700C 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 the 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 are migrated and the PG1 threshold is still unrelieved does HSM operate on private PG1 volumes in LRU order.

Note: Define all scratch categories before the preferred migration enhancement is used.

4.3.14 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 DATACLAS.

TS7700 reporting volumes (BVIR) cannot be written in LWORM format. For more information, see 13.4, “Bulk Volume Information Retrieval” on page 687.

4.3.15 Physical volumes for TS7700T

This section describes the number of physical volumes that are required to accommodate the workload you are planning for the TS7700T. 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 specific host workload
- ▶ Average compression ratio that is achieved per workload
- ▶ Average usage rate of filling physical volumes
- ▶ Scratch physical volumes

Amount of data that is stored for a specific 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 percentage 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 questions that are described in this section 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 fewer than 15 or 50 volumes in your pool if these conditions are acceptable. You need at least two scratch physical volumes to avoid an out of scratch state.

The following suggested formula can be used 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 a 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 must plan on the following number of physical volumes in the TS7700T:

$$P_v = P_s + P_1 + P_2 = 15 + 8 + 16 = 39 \text{ physical volumes}$$

If you need dual copied virtual volumes in a single cluster, double the number of physical volumes for that workload. If a workload uses dedicated pools with the borrow/return sharing disabled, each workload must have its own dedicated extra 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 that are still in the physical library that are in the Copy Export state, and the empty, filling, and full physical volumes that eventually are set to the Copy Export state.

The default value can be adjusted through the MI (use the Copy Export Settings window) to a maximum value of 10,000. After your Copy Export operations reach a steady state, approximately the same number of physical volumes are returned to the library for reclamation because there are those being sent offsite as new members of the Copy Export set of volumes.

4.3.16 Data compression

Starting with R4.1.2, the following options are available for TS7700 data compression:

- ▶ FICON compression: The compression method that is embedded on the FICON adapter.
- ▶ LZ4 compression: Software compression that uses an LZ4 algorithm. This compression method is faster and uses less CPU than the ZSTD method, but it results in a lower compression ratio than ZSTD.
- ▶ ZSTD compression: Software compression that uses a Zstandard algorithm. This compression method results in a higher compression ratio than LZ4, but is slower and uses more CPU than LZ4.

No special feature code is needed for these options to be available, but all clusters in the Grid must run R4.1.2 or later (the presence of lower levels of code in the grid will prevent the use of this feature).

Different virtual volumes in the same grid can use different compression algorithms, depending on Data Class construct assignment (which contain the selected algorithm information). This feature implies this option is a grid-scope setting. The option to be selected depends on the wanted compression ratio/speed level

Only virtual volumes that are written from BOT can go through LZ4/ZSTD processing (if its associated Data Class is configured for that compression method). Previously written data in virtual volumes keeps the initially applied compression method, even if Data Class parameters are changed. A method is not available to convert old data to the new compression algorithms.

Note: The uncompressed data size that can be written to a single virtual volume includes a “logical” limit of 68 GB for channel byte counters tracking the amount of written data. This limit can be surpassed when compression rates are used that are equal or higher than 2.7:1 (with FICON traditional compression or the new enhanced algorithms) against volumes that are assigned to Data Classes configured for 25000 MiB volume size (after compression) and equal or higher than 1.05:1 for 65000 MiB volume size.

Considering that circumstance, Data Classes can now be configured to decide how to handle that event by using the new attribute 3490 Counters Handling with the following available options:

- ▶ Surface EOT: Set this option to surface EOT (End Of Tape) when channel bytes written reaches maximum channel byte counter (68 GB).
- ▶ Wrap Supported: Set this option to allow channel bytes written to exceed the maximum counter-value and present the counter overflow unit attention to the attached LPAR, which then can collect and reset the counters in the TS7700 by using the **RBL (X'24')** command.

If compression method is selected, the host compression definition is accepted (which is the same behavior that is shown with previous releases) when writing data to a virtual volume. Compression then 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.

4.3.17 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 because of security or business requirements.

TS7700T implements 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 entire tape before reuse. If a physical volume features encrypted data, the erasure is accomplished by deleting Encryption Keys on the physical volume, which renders 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, the TS7700T needs time to erase every physical tape. Therefore, the TS7700T needs 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 2.3.13, “Secure Data Erase function” on page 58, and “Defining physical volume pools in the TS7700T” on page 579.

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, which renders the encrypted data unrecoverable.

When compared to the normal or long erasure operation, EK shredding is 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.

4.3.18 Planning for tape encryption in a TS7700T

The importance of data protection became 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 TS7700T helps control the risks of unauthorized data access without excessive security management burden or subsystem performance issues.

Encryption on the TS7700T is controlled on a storage pool basis. SG and MC DFSMS constructs that are 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, were enhanced to include encryption characteristics.

The tape encryption solution in a TS7700T consists of the following components:

- ▶ The IBM Guardium Key Lifecycle Manager 4.1 and later as a central point from which all EK information is managed and served to the various subsystems.
- ▶ The TS1120, TS1130, TS1140, TS1150, or TS1160 encryption-enabled tape drives are the other fundamental piece of TS7700T tape encryption that provides hardware that runs the cryptography function without reducing the data-transfer rate.

- ▶ The 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 IBM Guardium Key Lifecycle Manager 4.1 and later by using Ethernet to communicate with the IBM Security Guardium Key Lifecycle Manager (or in-band through FICONs) to the tape drives. Encryption support is enabled with FC 9900.

Rather than user-provided key labels per pool, the TS7700T 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 TS7700T does not require any host software updates because the TS7700T 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 TS7700T.

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 TS7700T.

Note: The IBM Guardium Key Lifecycle Manager 4.1 and later is the only supported key manager. The Container edition for z/OS of the IBM Security Guardium Key Lifecycle Manager is also supported.

You must also 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 must have all 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 TS7700T when needed.

If the TS7700T 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 the cache.

IBM Guardium Key Lifecycle Manager 4.1 and later

IBM Security Guardium Key Lifecycle Manager, formerly known as *IBM Security Key Lifecycle Manager*, is the IBM strategic platform for the storage and delivery of encryption keys to encrypt storage endpoint devices.

IBM Security Guardium Key Lifecycle Manager can be used with the TS1160, TS1150, and TS1140 tape drives. Similar to the previous product, IBM Encryption Key Manager (EKM), IBM Security Key Lifecycle Manager, IBM Security Guardium Key Lifecycle Manager serves data keys to the tape drive.

It focuses on ease of use and provides a graphical user interface (GUI) to help with the installation and configuration of the key manager. It also allows for the creation and management of the key encrypting keys (certificates).

For more information, see [IBM Guardium Key Lifecycle Manager documentation](#).

Encryption capable tape drives

Because data is encrypted on the back-end tape drives, the TS7700T must be equipped with Encryption Capable tape drives, such as the following tape drives:

- ▶ TS1120 3592 E05 (Encryption Capable) tape drives; must be running in 3592 E05 native mode. TS1120 tape drives with FC 5592 or FC 9592 are Encryption Capable.
- ▶ TS1130 3592 E06 tape drives.
- ▶ TS1140 3592 E07 tape drives.
- ▶ TS1150 3592 E08 tape drives.
- ▶ TS1160 3592 60F tape drives

The TS7700T 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, ensure 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 TS7700T must be configured with the IP addresses and TCP/IP ports to find the encryption key managers in the network.

For more information about a comprehensive TS7700T encryption implementation plan, see “Implementing TS7700 Tape Encryption” in *IBM System Storage Tape Encryption Solutions*, SG24-7320.

4.3.19 Planning for cache disk encryption in the TS7700

This section describes cache disk encryption for the TS7760 and TS7770.

TS7760 disk cache encryption

The TS7700 cache models 3956-CSA/XSA support Full Disk Encryption (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 the following keys to protect HDDs:

- ▶ Data encryption key: This key is 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.
- ▶ Lock key or security key: This key is a 32-byte random number that authenticates the drive with the CSA Cache Controller by using asymmetric encryption for authentication. When the FDE drive is *secure-enabled*, it must authenticate with the 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 CSA cache controller and CSA cache expansion drawers.

Authentication occurs after the FDE disk has powered on, where it is in a locked state. If encryption was never enabled (the lock key is not initially established between the CSA cache controller and the disk), the disk is considered unlocked with unlimited access, as with a non-FDE drive.

The following feature codes are required to enable FDE:

- ▶ Feature Code 7404: Required on all 3956-CSA, and 3956-XSA cache drawers
- ▶ Feature Code 7333: Required on the 3952-F06/F07 base frame for a TS7760
- ▶ Feature Code 7334: Required on the 3952-F06/F07 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 TS7760-VEC (Encryption Capable Frames, as listed in the previous required feature code list).

Key management for FDE does not use the IBM Security Guardium Key Lifecycle Manager. Instead, the key management is handled by the disk controller 3956-CSA. No keys are available 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 that is 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.

TS7770 disk cache encryption

Disk encryption is available on a new order from manufacturing that ordered FC 5272, Disk Enabled Encryption, or FC 5276, Enable disk encryption - External Key Management. An order of FC 5272 or FC 5276 includes FC 7405, Encryption CSB (USB flash drives), which provides four USB sticks.

Lock keys are stored on USB sticks that are plugged into disk cache controllers. Locked DK is stored across all disks, which makes data accessibility possible only with a full string of DDMs and the USB sticks. USB sticks can be removed if the hardware is at risk, such as emergency evacuation.

An entire file system must be encrypted because a mixture of encrypted and non-encrypted arrays is not supported. All arrays in all strings must be encrypted. All strings in the cluster must also be encrypted.

All TS7770 configurations with 3956-CSB/XSB or CFC/XFC cache that include any encryption type that is enabled have always shipped with local key management enabled (FC 5272). This feature encrypts the data in the CSB/CFC processor and places that encrypted data onto regular disk drives.

The local encryption (FC 5272 Disk Enabled Encryption) is configured during the TS7770 initial installation by the service person. Consider the following points:

- ▶ FC 5272 Disk Enabled Encryption is not available for field Installation on the TS7770 and must be shipped from manufacturing for any Encryption.
- ▶ FC 7405 must be ordered on every 3956-CSB/CFC in the TS7770 configuration.
- ▶ FC 7405 provides four USB sticks per 3956-CSB/CFC that are used to store the local encryption keys.

Note: If external key management is later enabled, these USB sticks are no longer needed.

The External Key Encryption (FC 5276) must have FC 5272 installed on the TS7770 server before initial installation. Consider the following points:

- ▶ All TS7700 configurations with any encryption type enabled are always shipped with local key management enabled first.
- ▶ After a TS7770 with FC 5272 is configured in a customer environment and can communicate with an external key server, FC 5276 can be activated to migrate to external key management.

External key management

You can manage the external key for the disk drive modules (DDMs) externally. For external key management of encryption, the encryption must be enabled onsite by an IBM SSR.

The encryption key server is installed and configured on the network.

IBM Guardium Key Lifecycle Manager 4.1 and later is supported on TS7700 with 8.51.1.xx microcode.

The following functions are addressed:

- ▶ Enable the Feature Keys for Local Encryption (not for TS7770)
- ▶ Enable the Feature Keys for External Key management.
- ▶ Configure external key encryption (from SMIT only)
- ▶ Switch from internal key management to external key management (from SMIT only)
- ▶ Rekey (from SMIT and Management Interface)

Since Release 5.4, other EKM is supported

- ▶ Only VED disk only systems
- ▶ Only New systems, no conversion allowed
- ▶ Only Thales CipherTrust Manager 2.0 and later
- ▶ Only configurable by an IBM SSR by using SMIT

4.4 Tape analysis and sizing the TS7700

This section documents the process of the use of 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.

4.4.1 IBM tape tools

Although most IBM tape tools are available to you, some tools, such as *BatchMagic*, are available to IBM personnel and IBM Business Partners only. You can download the tools that are available from [IBM Software Tape Tools](#).

The web page features a list of .TXT, .PDF, and .XMI files. To start, open the OVERVIEW.PDF file to see a brief description of all the various tool jobs. The following files contain the program objects that should be installed to run a tool:

- ▶ IBMJCL.XMI: Job control language (JCL) for current tape analysis tools
- ▶ IBMCNTL.XMI: Parameters that are needed for job execution
- ▶ IBMLOAD.XMI: Load library for executable load modules
- ▶ IBMPAT.XMI: Data pattern library, which is needed only 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 negatively affect performance.

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 listed in Table 4-18 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 4-18 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).
15	EOV or CLOSE when open for writing (called “open for output” in reports).
21 ^a	Volume demount.
30 ^b	Address Space Record (contains subtypes 04, 05, 34, 35, and others).
34	Step End (Time Sharing Option [TSO]).
35	Job End (TSO).

a. Type 21 records exist for tape data only.

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 features 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 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 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. Based on input parameters, TAPEWISE can generate several reports that can help with the following items:

- ▶ Tape activity analysis
- ▶ Mounts and megabytes 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

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 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.Parm1ib 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 in the unload process are shown in Figure 4-6 on page 205:

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.

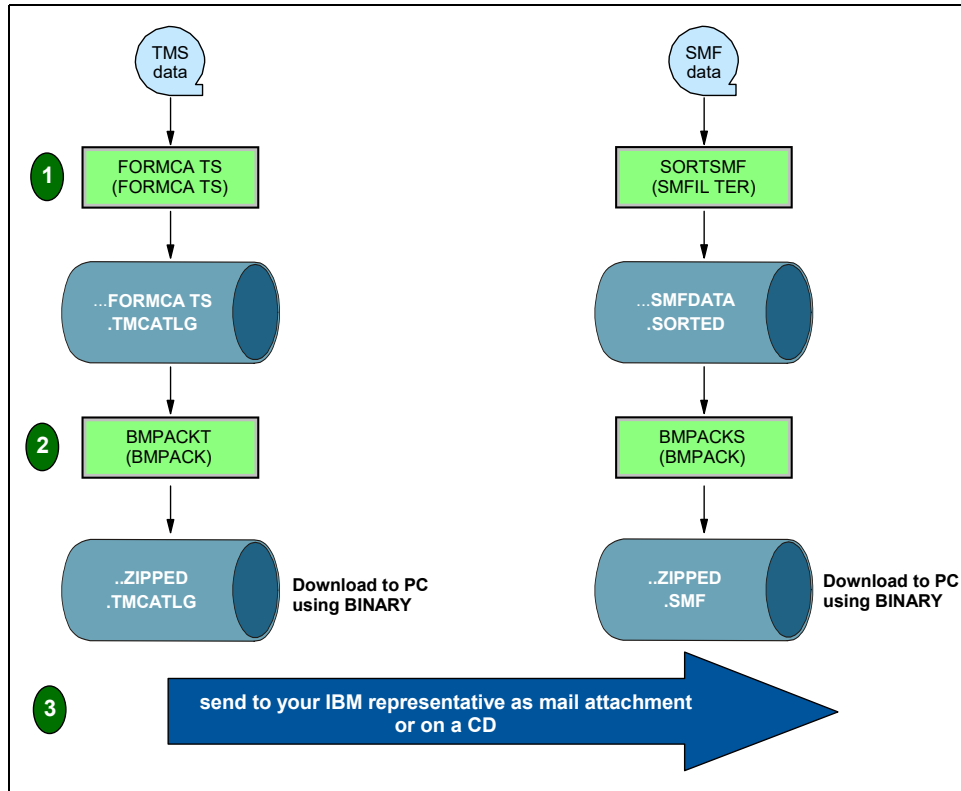


Figure 4-6 Unload process for TMS and SMF data

In addition to the extract file, the following information is useful for sizing the TS7700:

- ▶ Number of volumes in the 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, the number of volumes does need to be collected.
- ▶ 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 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

Similar to the number and type of tape CUs installed, this information 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.

4.4.2 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 that is 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 the cartridges that are required with new technology. It also models the operation of a TS7700 and 3592 drives (for TS7700T) 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.

Note: BatchMagic is specifically for IBM internal and IBM Business Partner use.

4.4.3 Workload considerations

The TS7700 appears as a group of 3490E subsystems, ranging 16–31 Virtual Control Units (depending on installed instances of FC 5275), with a maximum of 496 virtual devices attached per cluster. Any data that can be on a 3480, 3490, 3590, or 3592, previous 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 the following 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 and 16 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 periods occur during the day when your tape processing jobs are limited by drive availability, the TS7700 might help considerably in the area of processing.

The TS7700D enables access to multiple logical volumes directly from cache, at disk speed.

The design of the TS7700T enables access to multiple logical volumes on the same stacked physical volume because access to the logical volumes is solely through the TS7700T TVC. If access is needed to more than one logical volume on a physical volume, this access is provided without requiring any user involvement, unlike some alternatives, such as stacking by using JCL.

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 a drive earlier, so the next job that needs a tape drive can access it sooner because no rewind or unload and robotics time exist after closing the data set.

When a job attempts to read a volume that is not in the TS7700T 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 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 TS7700T.

Remember: The TS7700 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: Consider the following points:

- ▶ The term *local* refers to the TS7700 cluster that is running the logical mount to the host.
- ▶ The term *remote* refers to any other TS7700 that is participating in the same grid as the local cluster.

▶ **Scratch mount times**

When a program issues a scratch mount to write data, the TS7700 completes the mount request without having to recall the logical volume into the cache. With the TS7700D, all mounts are cache hit mounts. For workloads that create many tapes, a scratch mount significantly reduces volume processing times and improves batch window efficiencies.

The effect of the use of the scratch category on the TVC improves mount performance in the TS7700T because no physical mount is required. The performance for scratch mounts is the same as for TVC read hits.

Scratch mount times are further reduced when the optimal scratch allocation assistance function is enabled. This function designates one or more clusters as preferred candidates for scratch mounts by using a Management Class construct that is defined from the TS7700 Management Interface. The comparison between the time that is taken to process a mount request on a subsystem with cache to a subsystem without cache is shown in Figure 4-7 on page 207.

► 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 simple 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. For more information, see 2.3.14, “Copy Export function” on page 59.

► 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 TS7700T. Therefore, manual stacking of data sets onto volumes is no longer required. If you are planning for a new application that uses JCL to stack data sets onto a volume, the TS7700T makes this JCL step unnecessary.

Multi-file volumes that are moved to the TS7700T can also work without changing the stacking. However, the TS7700T recalls the complete logical volume to the TS7700T 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 TS7700T 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 TS7700T 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. For more information, see 2.3.14, “Copy Export function” on page 59.

► Grid network load balancing

For a TS7700 Grid link, the dynamic load balancing function calculates and stores the following information:

- Instantaneous throughput
- Number of bytes queued to transfer
- Total number of jobs queued on both links
- Whether deferred copy throttling is enabled on the remote node
- Whether a new job will be throttled (is deferred or immediate)

As a new task starts, a link selection algorithm uses the stored information to identify the link that most quickly completes the data transfer. The dynamic load balancing function also uses the instantaneous throughput information to identify degraded link performance.

The TS7700 provides a wide range of capabilities. Unless your data sets are large or require interchange or offsite storage, it is likely that the TS7700 is a suitable place to store your data.

4.4.4 Education and training

Plenty of information is available about the IBM TS7700 and IBM TS3500/TS4500 tape library in IBM Redbooks publications, operator manuals, IBM Documentation, and other resources. The amount of education and training your staff requires on the TS7700 depends on the following factors:

- ▶ If you are using a TS7700T, are you installing the TS7700T in an existing TS3500/TS4500 tape library environment?
- ▶ Are both the TS7700T 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 TS7700D solution?

A new TS7700T sharing an existing TS3500 or TS4500

When the TS7700T 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 TS7700T management interface (MI).

Also, you must cover how the TS7700T relates to the TS3500 or TS4500, which helps operational personnel understand the tape drives that belong to the TTS7700T, and which logical library and assigned cartridge ranges are dedicated to the TS7700T.

The operational staff must be able to identify an operator intervention, and perform the necessary actions to resolve it. They must also be able to perform basic operations, such as inserting new volumes in the TS7700T, or ejecting a stacked cartridge by using the MI.

Storage administrators and system programmers must be familiar with the operational aspects of the equipment and the following information:

- ▶ Understand the advanced functions and settings, and how they affect the overall performance of the subsystem (TS7700 or grid)
- ▶ How to use software choices, takeover decision, and library request commands and how they affect the subsystem
- ▶ Disaster recovery considerations

Storage administrators and system programmers must 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

For more information, see the following resources:

- ▶ Chapter 2, “Architecture, components, and functional characteristics” on page 15
- ▶ Chapter 6, “Implementing IBM TS7700” on page 239
- ▶ Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351
- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789

4.4.5 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 [IBM Services website](#).

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 4-19 on page 212 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 4-19 Main checklist

Task	Reference
Initial meeting	N/A
Physical planning	4.1, "Hardware installation and infrastructure planning" on page 146 and your IBM SSR
Host connectivity	4.1.5, "Host attachments" on page 166
Hardware installation	Specific hardware manuals and your IBM SSR
IP connectivity	4.1.3, "TCP/IP configuration considerations" on page 157
HCD	6.1, "Hardware configuration definition" on page 240
Maintenance PTF check	Fix category ptfs. 8.2.1, "Field frame replacement migration for TS7700T" on page 309
SMS	6.2, "Setting up the TS7700" on page 247
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, "IBM TS7700 Management Interface operations: Part 1" on page 351
Setting up the BVIR	13.4, "Bulk Volume Information Retrieval" on page 687
Specialist training	N/A
DR implications	Chapter 15, "Copy Export" on page 783
Functional/performance test	Chapter 14, "Performance considerations" on page 747
Cutover to production	N/A
Postinstallation tasks (if any)	14.3.1, "TS7700 components and task distribution" on page 753
Data migration (if required)	Chapter 8, "Migration" on page 303



Disaster recovery

This chapter describes the use of the TS7700 in disaster recovery (DR) and includes the following topics:

- ▶ 5.1, “TS7700 DR principles” on page 214
- ▶ 5.2, “Failover scenarios” on page 218
- ▶ 5.3, “Planning for DR” on page 219
- ▶ 5.4, “High availability and DR configurations” on page 222
- ▶ 5.5, “DR testing” on page 231
- ▶ 5.6, “A real disaster” on page 231
- ▶ 5.7, “Geographically Dispersed Parallel Sysplex for z/OS” on page 233

5.1 TS7700 DR 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
- ▶ How to use cloud-attached storage in your DR plan

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, or they might be accessible through a cloud store by any TS7700C cluster in the grid.

If a copy of the logical volume is not available at that TS7700 cluster (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.

Another possible scenario is if a copy of the logical volume is not available at the TS7700 cluster, but a copy is available at the cloud tier. Although the version in a remote TVC is the preferable source over a recall from the cloud, if for some reason no other cluster in the grid has a valid copy of the volume, the object that is stored in the cloud might be recalled and then accessed by any TS7700C in the grid.

Whether a copy is available at another TS7700 cluster in a multi-cluster grid depends on the following factors:

- ▶ 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 (Sync): Data that is written to the cluster is 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.
 - Time Delayed (TD): If the data expires before the Time Delayed setting is reached, no copy is made, otherwise a deferred copy is created.
 - No Copy (None): Data that is created on one cluster is not copied to the other cluster.
- ▶ Grid Cloud Awareness (TS7700C enabled Grids)

If you have at least one TS7700C cluster in the Grid and all clusters are at R5.1 or later, you can use Grid Cloud Awareness. Grid Cloud Awareness is a TS7700C cluster's ability to access data in the cloud independently of whether the cluster has a valid copy consistency point.

If a TS7700C is aware of a copy in the cloud, it can recall it if needed. If a cluster without cloud support is present, it can also access the data in the cloud by requesting any TS7700C cluster to recall the data into the TS7700C tape volume cache. After it is recalled, the original TS7700 cluster can access it remotely.

- ▶ Grid Cloud Awareness can also be used instead of grid-based consistency points for certain workloads. For example, a production TS7700C cluster can put a copy of a volume in an attached cloud. Although a remote TS7700C has “No Copy” configured, it still can access the data from the cloud, if needed. In essence, the cloud is used as the replication mechanism and grid copies do not need to occur.
- ▶ Any new TS7700C cluster that is introduced into a grid or a TS7700 that adds cloud support immediately can access all the data in the cloud that is put there by a TS7700C 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 Sync. 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 TVC, but might still be available if you use a cloud store. With Deferred and Time Delayed, a copy is available later; therefore, it might be available at the cluster that specified Deferred or it might not yet be completed.

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 to the device end, which can result in fewer copies of the data that is available than your copy policies specify.

The TS7700 Disk-Only solutions have a maximum storage capacity that is the size of their TVC, and TS7700T CP0 works similarly. Therefore, after the cache fills, the Volume Removal 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

In addition to a copy policy of No Copy, a Deferred Copy policy or Time Delayed Copy policy least affects the applications that are running on the host. In the case of Deferred Copy immediately after the volume is closed, the 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. For the Time Delayed Copy, after the timer setting is met, a copy of the volume is then put on the Deferred Copy Queue if the data is still valid.

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 because of 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, it might be acceptable that all the data did not get copied. However, for primary data, it is preferable to use a Sync 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 another cluster, the copy can be accessed through the original cluster or the other cluster.

However, a logical volume is owned by a single cluster at any time. This cluster is known as the *owning cluster*. The owning cluster controls access to the volume and changes to the attributes that are associated with the volume (such as category or storage constructs). The cluster that owns 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 own 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 cannot respond to requests from other clusters, the operation against that volume fails, unless more direction is given by way of Ownership Takeover in the TS7700 MI. Clusters do not automatically assume or assume ownership of a logical volume without being directed to do so, unless Autonomic Ownership Takeover Manager (AOTM) is confirmed and the TSSC components can confirm that the owning cluster failed.

This process is done to prevent the failure of the grid network communication paths between the clusters, which results in both clusters thinking that they own the volume. If more than one cluster owns a volume, that dual ownership might result in the volume's data or attributes being changed differently on each cluster, which results 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.

The following modes are set through the Management Interface (MI):

- ▶ Read Ownership Takeover (ROT): When ROT is enabled for a failed cluster, ownership of a volume can be taken from a cluster that failed. Only read access to the volume is allowed through the other clusters 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 can be taken from a cluster that is marked as failed. Full access is allowed through the other clusters in the grid. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster is restored. WOT should never be enabled for testing; only in the event of a true disaster to prevent multiple copies of the same volume from being created in the grid.
- ▶ Service prep/service mode: When a cluster is placed in service preparation mode or is in service mode, ownership of its volumes can 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.

Note: It is not possible to put a cluster into Service Prep if that cluster cannot communicate with the grid.

- ▶ In addition to the manual setting of one of the ownership takeover modes, an optional automatic method that is named AOTM is available when each of the TS7700 clusters is attached to a TS3000 System Console (TSSC) and a communication path 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. 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 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 receive 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 read or WOT, depending on what was set by the IBM SSR during AOTM configuration.

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 failed, which can take several minutes.

The following OAM messages can be displayed when AOTM enables the ownership takeover mode:

- ▶ 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.

When a cluster in the grid becomes inoperable, mounts that are directed to that cluster might fail with the following messages:

```
CBR4195I LACS retry possible for job OAM:
IEE763I NAME= CBRLACS CODE= 140394
CBR4000I LACS WAIT permanent error for drive xxxx.
CBR4171I Mount failed. LVOL=vvvvvv, LIB=libname, PVOL=??????, RSN=rc.
IEE764I END OF CBR4195I RELATED MESSAGES
```

CBR4196D Job OAM, drive xxx, volser vvvvvv, error code 140394. Reply 'R' to retry or 'C' to cancel.

If many jobs are failing the response to this message should be 'C' and the logical drives in the failing cluster should then be varied offline by using **VARY devicenumbr, OFFLINE** to prevent further attempts from the host to mount volumes on this cluster.

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 includes a valid copy of a logical volume, the jobs can be tried again.

If a logical volume is 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.

After the failed cluster comes back online and establishes communication with the other clusters in the grid, the following message is issued:

```
CBR3750I Message from library libname: G0011 Ownership takeover mode within  
library libname has been automatically disabled now that library libname has  
become available.
```

It is now possible to issue **VARY devicenumber, ONLINE** to bring the logical drives on the cluster for use.

After the cluster is operational again, the following message is issued if any volumes are in a conflicting state because they were accessed on another cluster:

```
CBR3750I Message from library libname: 0P0316 The TS7700 Engine has detected  
corrupted tokens for one or more virtual volumes.
```

If this issue occurs, see 10.5.2, “Repair Virtual Volumes window” on page 547 for more information about the process that is used for repairing the corrupted tokens.

The impact code and text that are issued with the CBR3750I messages now can be changed. For more information, see 12.2.1, “CBR3750I console message” on page 645.

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 during the initial system installation or at another time.

For more information about these scenarios, see [IBM TS7700 Series Grid Failover Scenarios Version 1.5](#), which is a white paper that was written to assist IBM specialists and clients in developing such testing plans.

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 DR

Although you can hope that a disaster does not occur, 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.

The following 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 DR site connectivity IODF considerations

If your production hosts feature 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 that are configured and offline to your production hosts helps to make it easier if a TS7700 failure occurs that requires FICON access to the DR clusters, which is distance-dependent and might not be appropriate for all configurations. This method also has the advantage that the DR host can share the IODF with the Production host, which makes failing over to a DR site more seamless.

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 process requires an IODF activation on the production hosts.

5.3.2 Grid configuration

With the TS7700, the following types of configurations can be installed:

- ▶ Stand-alone cluster

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, or a cloud object store was defined and CLOUD EXPORT is being used. In this case, no physical tapes are required. If the cluster goes out of service because of failures, whether the data is recoverable depends on the failure type.

If you do not have a TS7700C and a cloud store available, 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 also assumes that only a subset of the volumes is undamaged after the event. If the physical cartridges are destroyed or irreparably damaged, recovery is not possible, as with any other cartridge types.

In the case of TS7700C having access to the object store, Cloud Export and Recovery can be used as the DR solution.

For more information about the process, see the Chapter 17, “Cloud Storage Tier export, recovery, and testing” in *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.

Important: Consider the following points:

- ▶ Integrate the TS7700 recovery procedure into your current DR procedures.
- ▶ 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 or CLOUD EXPORT and Recovery from the object store in the cloud might take considerable time to complete. For those clients, the TS7700 grid provides a near-zero data loss and expedited recovery-time solution when implemented properly.

▶ Multi-cluster grid

With a multi-cluster grid configuration, up to eight 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. Also, if the two clusters are TS7700C, they can share access to the cloud object store.

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 it is likely that no connections exist 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. Also, two or more of the production clusters can share access to the cloud object store as the DR host.

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 and optionally to a cloud object store. The third and fourth (or more) remote clusters can hold deferred volume copies for DR and optionally have those copies in a cloud object store shared with the production clusters. 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.

With the addition of the TS7700C, the option to offload data to the cloud adds possibilities for recovery if a disaster occurs such that any of the clusters in the grid that are cloud-capable models can access the cloud object store.

The only connection between the production sites and the DR site is the grid interconnection. Normally, no host connectivity exists 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 or is made available in the cloud object store.

5.3.3 Planning guidelines

As part of planning a TS7700 grid configuration to address this solution, consider the following points:

- ▶ 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. Also, IPEX SAN42B-R switches and IP extension hardware are available 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 your network bandwidth and your cloud storage provider if you are running cloud-based recoveries. Ensure that the source location is the nearest possible to minimize latency on data transfer from cloud to your TS7700C.
- ▶ If you use one or more TS7700Cs in your production location and Cloud Export and Recovery as the recovery method, ensure that your DR cluster is configured to access the objects sent to the cloud by your production cluster. Your TS7700C cluster at the DR location needs access to the cloud object to complete the recovery.
- ▶ 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 no access exists 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 hold 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. The following scenarios can be set up:
 - Test that is based on all data from a TS7700?
 - Test that is based on the use of the Copy Export function and an empty TS7700?
 - What happens if the object store in the cloud becomes unavailable?
 - Test that is based on stopping production access to one TS7700 cluster and running production to another cluster?

For more information about DR Test, see Chapter 16, "Disaster recovery testing in a grid configuration" on page 819.

5.4 High availability and DR configurations

A few examples of grid configurations are described in this section.

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. They are typical examples that can be suited for cloud object usage. With up to eight cluster grids, many more ways to configure a grid might exist.

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 next. Other configurations are possible, and might be better suited for your environment.

DR 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 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. No host connectivity exists between the local hosts and the DR site cluster.

Figure 5-1 shows this configuration.

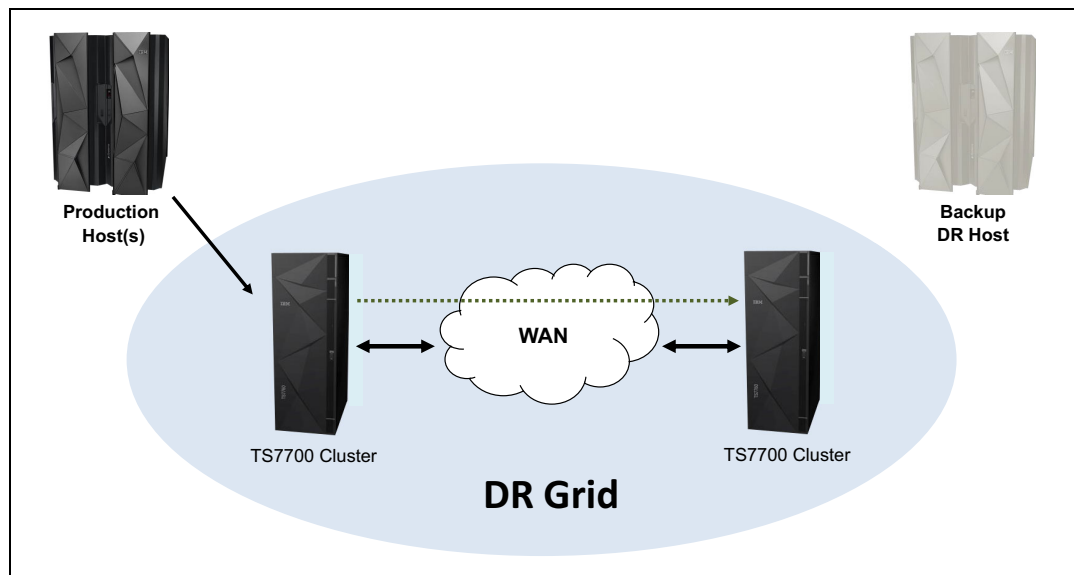


Figure 5-1 DR configuration

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 limited bandwidth is available between sites, copy critical data with a consistency point of RUN with the rest of the data that uses the Deferred or Time Delayed Copy Consistency Point. RUN or SYNC are only acceptable copy policies for distances less than 100 kilometers (62 miles). Distances that are greater than 100 km (62 miles) must rely on the Deferred or Time Delayed 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 the clusters becomes unavailable because it 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 points:

- ▶ 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.

For more information about AOTM, see 2.4.33, “Autonomic Ownership Takeover Manager” on page 97.

Figure 5-2 shows this configuration.




Figure 5-2 Availability configuration

Configuring for DR and high availability

You can configure a two-cluster grid configuration to provide 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 access can include 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 can 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, a difference in read or write performance can exist 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. In that use case, 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 shows this configuration.

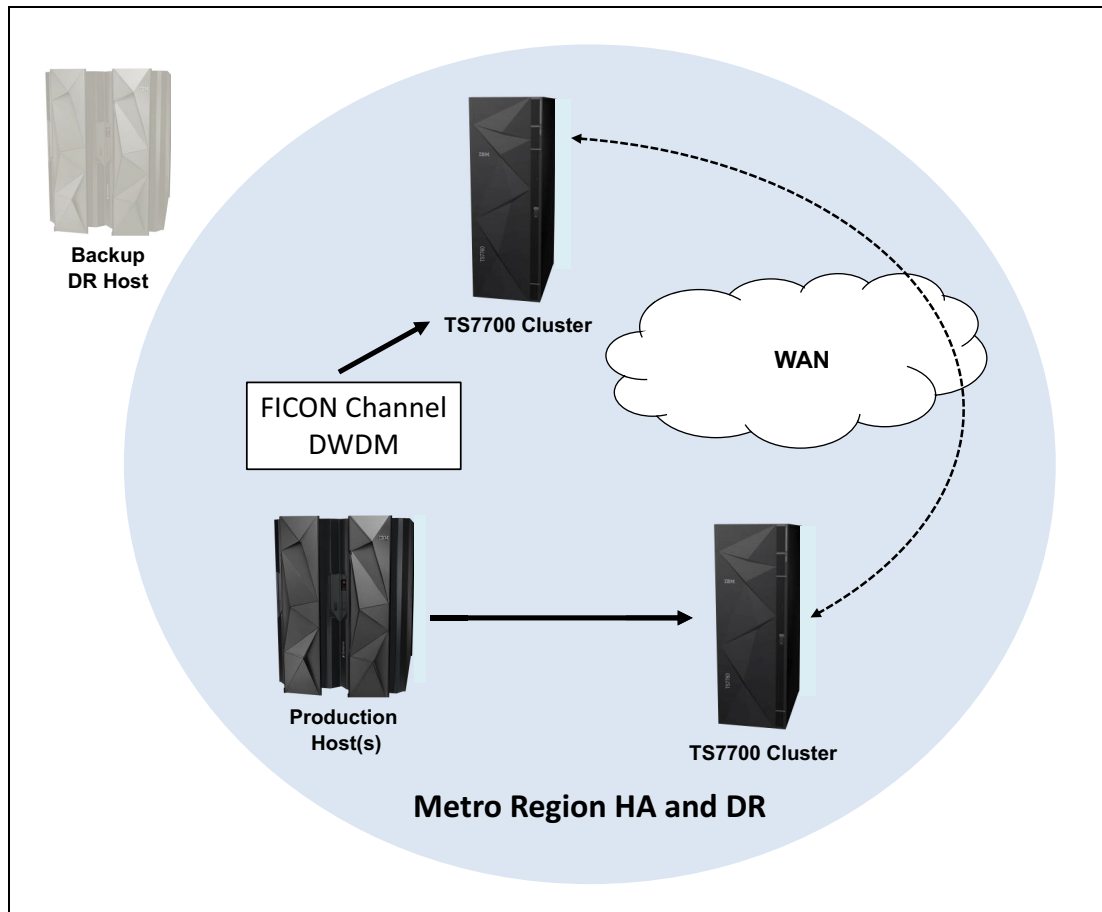


Figure 5-3 Availability and DR configuration

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 in this section. 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 DR

Figure 5-4 on page 226 shows 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 configuration 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.

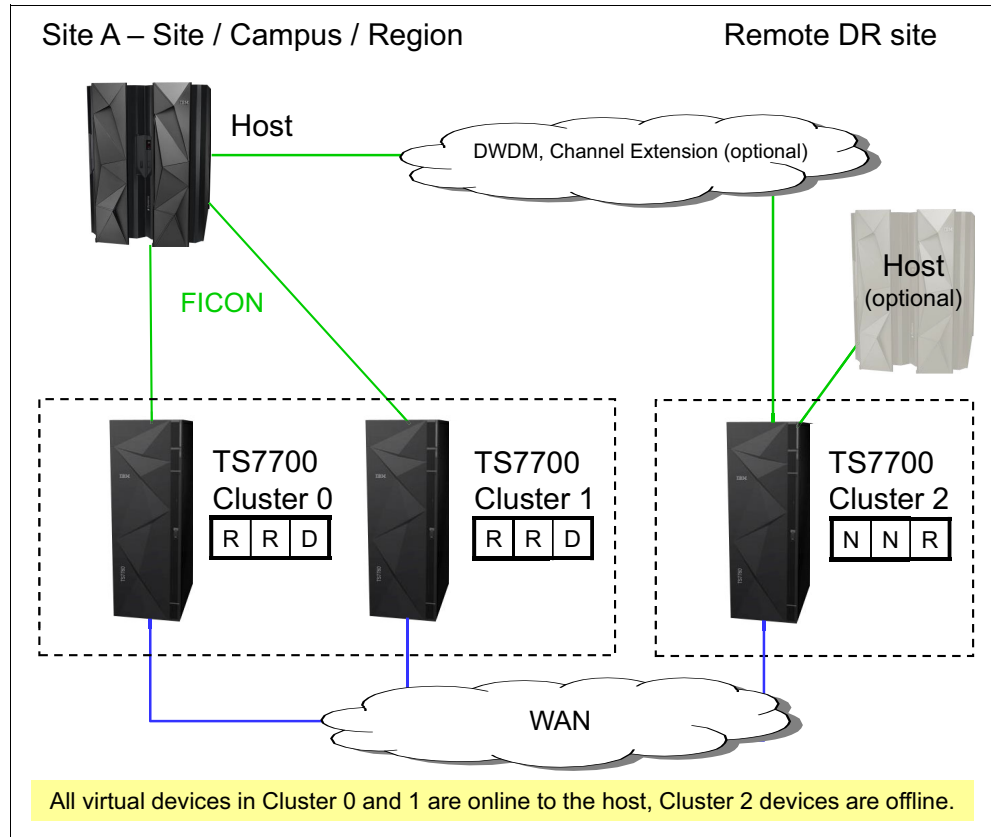


Figure 5-4 High availability and DR configuration

Dual production site and DR

Figure 5-5 on page 227 shows 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 configuration 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 DR

Three-cluster high availability production site and DR

This model was adopted by many clients. In this configuration, two clusters are in the production site (the same building or separate location within the 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 its local TS7700D peer and the remote TS7700. Optional copies in both TS7700D clusters provide high availability and cache access time for the host access. At the same time, the remote TS7700 provides DR capabilities and the remote copy can be remotely accessed, if needed.

This configuration, which provides a 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 shown in Figure 5-6.

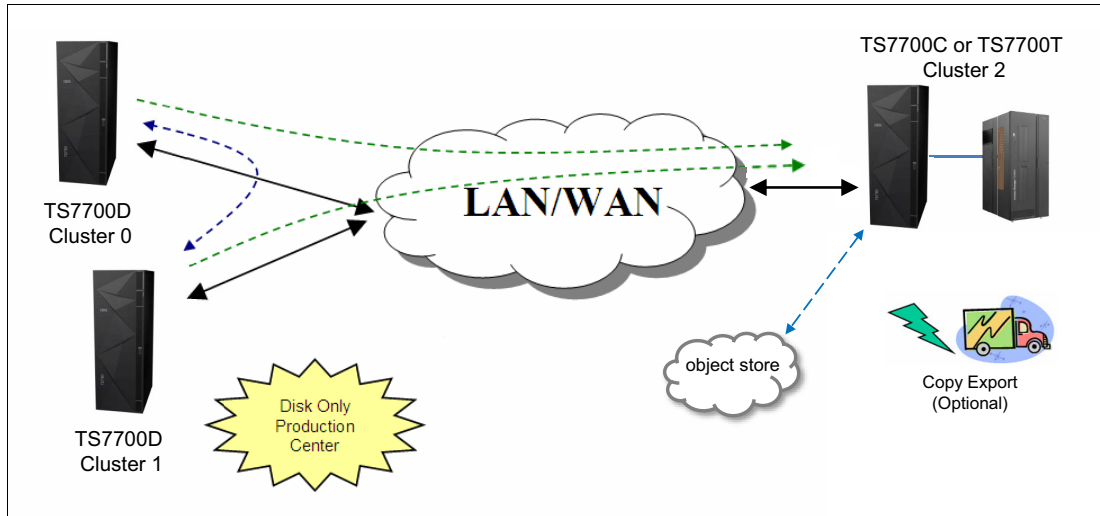


Figure 5-6 Three-cluster high availability and DR with two TS7700Ds and one TS7700T

Another variation of this model uses a TS7700D and a TS7700T or TS7700C for the production site (as shown in Figure 5-7), both replicating to a remote TS7700T or TS7700C with optional object store in the cloud shared by both sites.

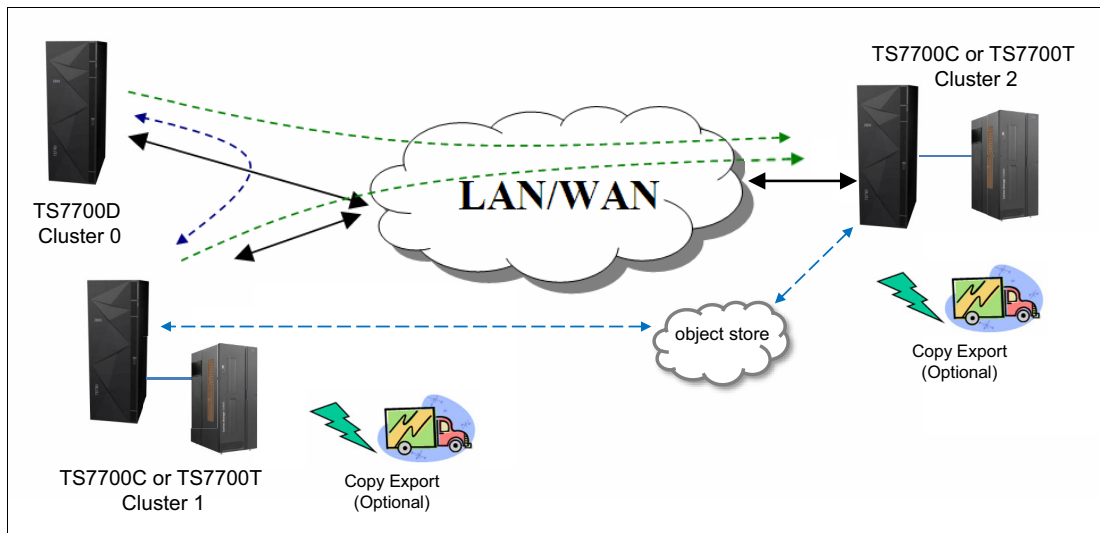


Figure 5-7 Three-cluster high availability and DR with two TS7700T/TS7700C and one TS7700D

In both models, if a TS7700D reaches the upper threshold of usage, the PREFER REMOVE data, which was replicated to the TS7700T or TS7700C, 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 TS7700T or TS7700C, and others that favor the TS7700D, suiting a specific workload to the cluster best equipped to perform it.

With the TS7700T, Copy Export (shown as optional in Figure 5-6 on page 228 and Figure 5-7 on page 228) can be used to have another copy of the migrated data from either or both sites, if required.

With the TS7700C, both sites can share a common Object Storage in the cloud that is accessible by either TS7700C.

Four-cluster grid

A four-cluster grid that can have both sites for dual purposes is described in this section. Both sites are equal players within the grid, and any site can play the role of production or DR, as required.

Dual production and DR at Metro Mirror distance

Dual production and DR sites are used in this model. Although a site can be labeled as a high availability pair or DR site, they are equivalent from a technology perspective and functional design.

In this example, two production sites are used within metro distances and two remote DR sites within metro distances between them. This configuration delivers the same capacity as a 2 two-cluster grid configuration, with the high availability of a four-cluster grid (see Figure 5-8).

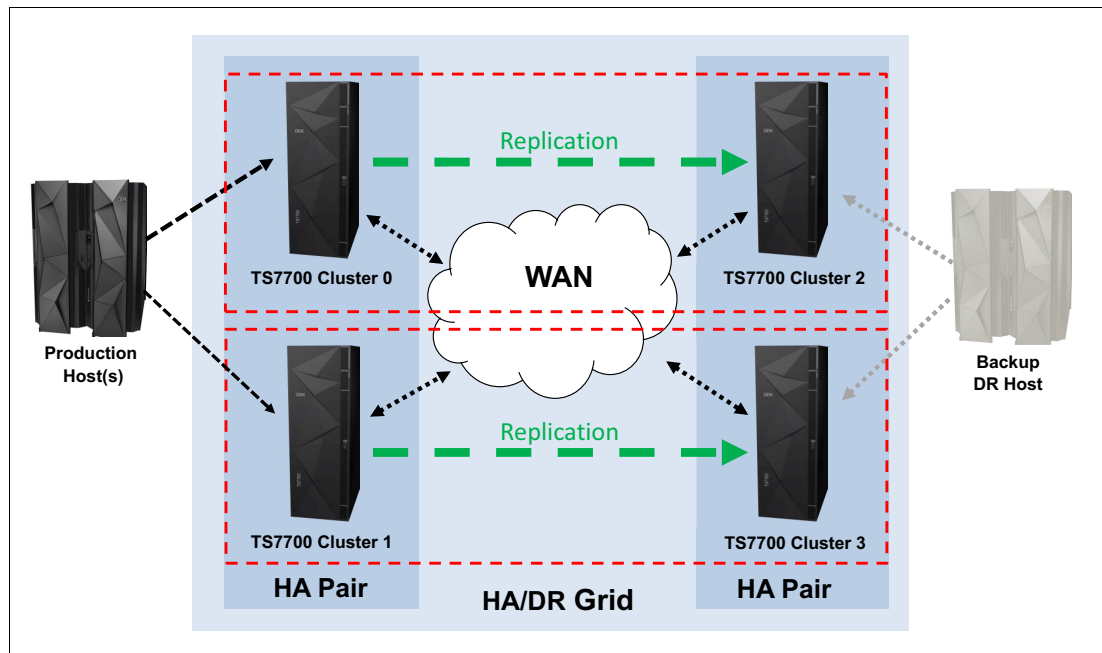


Figure 5-8 Four-cluster high availability and DR

You can have a host workload that is balanced across both clusters (Cluster 0 and Cluster 1 in Figure 5-8). The logical volumes that are written to a specific cluster are replicated to one remote cluster only. As shown 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, the copy mode for Cluster 0 is DNDN or SNSN and for Cluster 1 is NDND or NSNS.

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 fails because of 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.

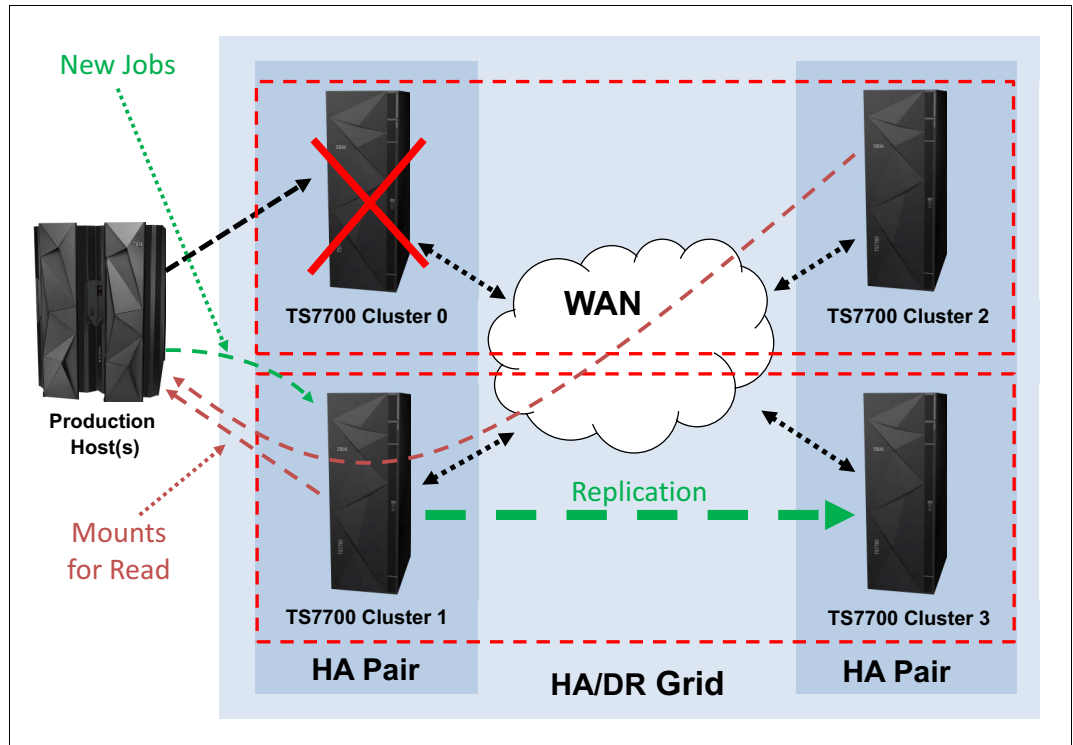


Figure 5-9 Four-cluster grid high availability and DR - Cluster 0 outage

During the outage of Cluster 0 in the example, new jobs for write use only one half of the configuration (the unaffected partition that is shown 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, which reestablishes 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 that are shown in Figure 5-9. In this case, copy policies can be DNDN for Cluster 2 and NDND for Cluster 3.

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 must be completed 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 IOCP statements necessary to include the recovery TS7700 tape drives in the IODF.

4. Define the library definitions in the source control data set (SCDS) with the Library IDs for the recovery TS7700 tape drives for both the composite and distributed libraries. Activate the IODF and the SMS SCDS.
5. If you use cloud as a data source on your recovery, ensure that the DR cluster can access the cloud object store and you have the cluster configured with the cloud information.

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 that are 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 DR testing

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. For more information about DR testing practices, see Chapter 16, “Disaster recovery testing in a grid configuration” on page 819.

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 disaster can be real. 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 use the DR TS7700, which makes your task easier. However, this easy-to-use capability does not mean that you have all the data copied to the DR TS7700.

If your copy mode is RUN, consider only in-flight tapes that are being created when the disaster occurs. Rerun all of these jobs to re-create tapes for the DR site.

Alternatively, if your copy mode is Deferred, tapes are available 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 in the copy queue. With this information, you can discover which data sets are missing, and rerun the jobs to re-create these data sets at the DR site by using your TMS.

Figure 5-10 shows an example of a real disaster situation.

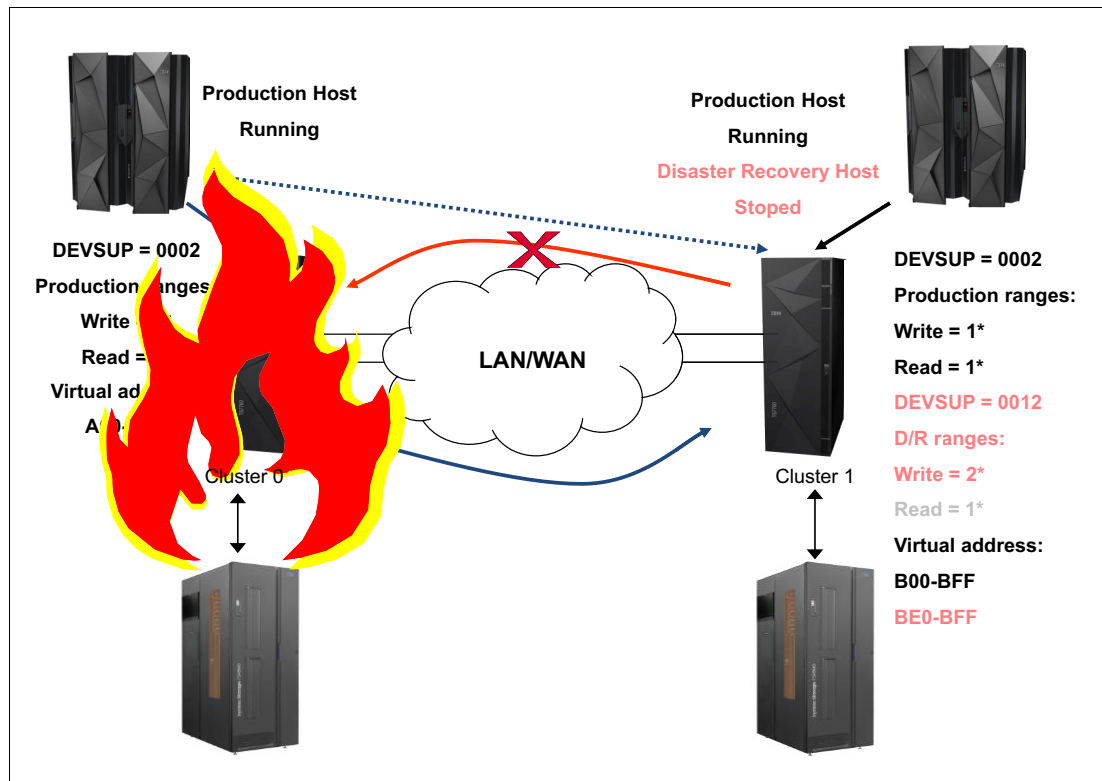


Figure 5-10 Real disaster situation

In a real disaster scenario, the entire primary site is lost. Therefore, you must start your production systems at the DR site. To complete this process, you must have a copy of all your information not only on tape, but all DASD data that is copied to the DR site.

After you can start the z/OS partitions, you must be sure that your hardware configuration definition (HCD) includes the DR TS7700 cluster. Otherwise, you cannot bring the TS7700 online.

You must also change ownership takeover. To perform this task, go to the MI interface and enable ownership takeover for read and write.

All the customizations that you made for DR testing are not needed during a real disaster. Production tape ranges, scratch categories, SMS definitions, RMM inventory, and so on, are in a real configuration that is in DASD and 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 the use of 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 must start the tasks that are required to recover your primary site or to create a site. The old DR site is now the production site, so you must create a DR site.

5.7 Geographically Dispersed Parallel Sysplex for z/OS

The IBM Z 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, which improves 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 TS7700

The hosts are directly attached to the TS7700 that is local to the host so that is your first consideration in directing I/O to a specific TS7700. Host channels from each site's GDPS hosts are also typically installed to connect to the TS7700 at the site that is remote to a host to cover recovery only when the TS7700 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 TS7700 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 TS7700's TVC (if required) to perform a mount operation unless override settings on a cluster are used.

To direct the TS7700 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 TS7700 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 TS7700 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 TS7700 Virtualization Engines.

Switching site production from one TS7700 to another one

The way that data is accessed by the TS7700 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 TS7700 grid environment with GDPS, the following scenarios can occur:

- ▶ GDPS switches the primary host to the remote location and the TS7700 grid is still fully functional. Consider the following points:
 - No manual intervention is required.
 - Logical volume ownership transfer is done automatically during each mount through the grid.
- ▶ A disaster occurs at the primary site, and the GDPS host and TS7700 cluster are down or inactive. Consider the following points:
 - Automatic ownership takeover of volumes, which are then accessed from the remote host, is not possible.
 - Manual intervention is required. Through the TS7700 MI, the administrator must start a manual ownership takeover. To do so, use the TS7700 MI and click **Service** → **Ownership Takeover Mode**.
- ▶ Only the TS7700 cluster at the GDPS primary site is down. In this case, the following manual interventions are required:
 - Vary online remote TS7700 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 TS7700 MI, start a manual ownership takeover. To do so, click **Service** → **Ownership Takeover Mode** in the TS7700 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 in 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 that is called *Bulk Volume Information Retrieval* (BVIR). If an unplanned interruption to tape replication occurs, 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 991.

5.7.3 Geographically Dispersed Parallel Sysplex implementation

Before implementing the GDPS support for TS7700, ensure that you review and understand the following topics:

- ▶ [IBM Virtualization Engine TS7700 Series Best Practices Copy Consistency Points](#)
- ▶ [IBM Virtualization Engine TS7700 Series Best Practices Synchronous Copy Mode](#)

The complete instructions for implementing GDPS with the TS7700 are available in the GDPS manuals.

Note: For DS8000 Object Store support, inbound objects are not affected by the state of any back-end tape library or object store. For more information about DS8000 Object Store, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).



Part 2

Implementation and migration

This part provides the required information to implement IBM TS7700 R5.4 in your environment, and to migrate from another tape solution to IBM TS7700.

This part includes the following chapters:

- ▶ Chapter 6, “Implementing IBM TS7700” on page 239
- ▶ Chapter 7, “Hardware configurations and upgrade considerations” on page 261
- ▶ Chapter 8, “Migration” on page 303



Implementing IBM TS7700

This chapter describes how to implement the IBM TS7700 on IBM Z platforms. From a software perspective, differences exist between the TS7700C, TS7700D, and TS7700T. If no specific differences are indicated, the implementation steps apply to all models. Otherwise, the differences are explained in each relevant step.

In this chapter, the various tasks that are needed to implement the TS7700 are described. These tasks include setting up for a tape library, cloud storage device, DS8000 Object Storage, and general TS7700 setup steps.

For more information about defining a tape subsystem in a DFSMS environment, see *IBM TS4500 R9 Tape Library Guide*, [SG24-8235](#)

For more information about the TS7770 Cloud Object Storage solution and how to implement and integrate this solution into your enterprise, see *IBM TS7700 R5.4 Cloud Storage Tier Guide*, [REDP-5573](#)

For more information about planning and implementing the function of using the TS7700 as an optional target for DS8000 Transparent Cloud Tier by using DFSMS, see:

- ▶ *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#)
- ▶ *IBM DS8000 Transparent Cloud Tiering: DS8000 Release 9.3*, [SG24-8381](#)

This chapter includes the following topics:

- ▶ 6.1, “Hardware configuration definition” on page 240
- ▶ 6.2, “Setting up the TS7700” on page 247
- ▶ 6.3, “TS7700 software definitions” on page 248
- ▶ 6.4, “Attaching the TS7700 to a physical tape library or cloud” on page 253

6.1 Hardware configuration definition

This section describes the process of defining the TS7700 by using the HCD interface. Usually, HCD definitions are made by IBM Z administrators. A helpful approach is to complete a table with all the definitions that the administrators 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

CHPID	CU	CUADD	Link	Devices	ADD	LIB-ID	Libport
		1A			00-0F		1B
		1B			00-0F		1C
		1C			00-0F		1D
		1D			00-0F		1E
		1E			00-0F		1F

6.1.1 Defining devices through HCD

The standard TS7700 configuration allows you to define 16 CUs with 16 devices each per cluster.

Purchasing multiple license feature code 5275 allows you to define a maximum of 31 CUs. For more information, see 7.2.2, “TS7700 nonconcurrent system component upgrades” on page 276.

Table 6-2, Table 6-3, Table 6-4 on page 242, and Table 6-5 on page 242 detail the mapping between LIBPORT-ID and CU for all TS7700 clusters.

Table 6-2 CU/LIBPORT-ID Definitions for the first 128(0-127) devices for cluster 0 through cluster 7

CU	1	2	3	4	5	6	7	8
CUADD	0x00	0x01	0x02	0x03	0x04	0x05	0x06	0x07
LIBPORT-ID Cluster 0	0x01	0x02	0x03	0x04	0x05	0x06	0x07	0x08
LIBPORT-ID Cluster 1	0x41	0x42	0x43	0x44	0x45	0x46	0x47	0x48
LIBPORT-ID Cluster 2	0x81	0x82	0x83	0x84	0x85	0x86	0x87	0x88
LIBPORT-ID Cluster 3	0xC1	0xC2	0xC3	0xC4	0xC5	0xC6	0xC7	0xC8
LIBPORT-ID Cluster 4	0x21	0x22	0x23	0x24	0x25	0x26	0x27	0x28
LIBPORT-ID Cluster 5	0x61	0x62	0x63	0x64	0x65	0x66	0x67	0x68
LIBPORT-ID Cluster 6	0xA1	0xA2	0xA3	0xA4	0xA5	0xA6	0xA7	0xA8
LIBPORT-ID Cluster 7	0xE1	0xE2	0xE3	0xE4	0xE5	0xE6	0xE7	0xE8

Table 6-3 CU/LIBPORT-ID Definitions for the second 128(128-255) devices for cluster 0 through cluster 7

CU	9	10	11	12	13	14	15	16
CUADD	0x08	0x09	0x0A	0x0B	0x0C	0x0D	0x0E	0x0F
LIBPORT-ID Cluster 0	0x09	0x0A	0x0B	0x0C	0x0D	0x0E	0x0F	0x10
LIBPORT-ID Cluster 1	0x49	0x4A	0x4B	0x4C	0x4D	0x4E	0x4F	0x50
LIBPORT-ID Cluster 2	0x89	0x8A	0x8B	0x8C	0x8D	0x8E	0x8F	0x90
LIBPORT-ID Cluster 3	0xC9	0xCA	0xCB	0xCC	0xCD	0xCE	0xCF	0xD0
LIBPORT-ID Cluster 4	0x29	0x2A	0x2B	0x2C	0x2D	0x2E	0x2F	0x30
LIBPORT-ID Cluster 5	0x69	0x6A	0x6B	0x6C	0x6D	0x6E	0x6F	0x70

CU	9	10	11	12	13	14	15	16
LIBPORT-ID Cluster 6	0xA9	0xAA	0xAB	0xAC	0xAD	0xAE	0xAF	0xB0
LIBPORT-ID Cluster 7	0xE9	0xEA	0xEB	0xEC	0xED	0xEE	0xEF	0xF0

Table 6-4 CU/LIBPxRT-ID Definitions for third 128(256-383) devices for cluster 0 through cluster7

CU	17	18	19	20	21	22	23	24
CUADD	0x10	0x11	0x12	0x13	0x14	0x15	0x16	0x17
LIBPORT-ID Cluster 0	0x11	0x12	0x13	0x14	0x15	0x16	0x17	0x18
LIBPORT-ID Cluster 1	0x51	0x52	0x53	0x54	0x55	0x56	0x57	0x58
LIBPORT-ID Cluster 2	0x91	0x92	0x93	0x94	0x95	0x96	0x97	0x98
LIBPORT-ID Cluster 3	0xD1	0xD2	0xD3	0xD4	0xD5	0xD6	0xD7	0xD8
LIBPORT-ID Cluster 4	0x31	0x32	0x33	0x34	0x35	0x36	0x37	0x38
LIBPORT-ID Cluster 5	0x71	0x72	0x73	0x74	0x75	0x76	0x77	0x78
LIBPORT-ID Cluster 6	0xB1	0xB2	0xB3	0xB4	0xB5	0xB6	0xB7	0xB8
LIBPORT-ID Cluster 7	0xF1	0xF2	0xF3	0xF4	0xF5	0xF6	0xF7	0xF8

Table 6-5 CU/LIBPORT-ID Definitions for the last 112(384-495) devices for cluster 0 through cluster7

CU	25	26	27	28	29	30	31
CUADD	0x18	0x19	0x1A	0x1B	0x1C	0x1D	0x1E
LIBPORT-ID Cluster 0	0x19	0x1A	0x1B	0x1C	0x1D	0x1E	0x1F
LIBPORT-ID Cluster 1	0x59	0x5A	0x5B	0x5C	0x5D	0x5E	0x5F
LIBPORT-ID Cluster 2	0x99	0x9A	0x9B	0x9C	0x9D	0x9E	0x9F
LIBPORT-ID Cluster 3	0xD9	0xDA	0xDB	0xDC	0xDD	0xDE	0xDF
LIBPORT-ID Cluster 4	0x39	0x3A	0x3B	0x3C	0x3D	0x3E	0x3F
LIBPORT-ID Cluster 5	0x79	0x7A	0x7B	0x7C	0x7D	0x7E	0x7F
LIBPORT-ID Cluster 6	0xB9	0xBA	0xBB	0xBC	0xBD	0xBE	0xBF
LIBPORT-ID Cluster 7	0xF9	0xFA	0xFB	0xFC	0xFD	0xFE	0xFF

Figure 6-1 and Figure 6-2 on page 244 show the two important windows that are used 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)

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 Connection (FICON) directors, the link address fields are blank.

3. Repeating the previous process, define the second through the 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.
4. 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. 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 the device number)
Timeout . . . . . 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. You can update the device parameters in this window.

```

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: Consider the following points:

- ▶ 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 IBM Z host is sharing the virtual drives in the TS7700, specify `SHAR-ABLE=YES`. This specification forces `OFFLINE` to `YES`. It is up to the installation to ensure the correct serialization from all attached hosts.
- ▶ Use the composite library ID of the TS7700 in your HCD definitions.
- ▶ The distributed library IDs are not defined in HCD definitions.

To define the remaining TS7700 3490E virtual drives, repeat this process for each CU in your implementation plan.

6.1.2 Activating the I/O configuration

Differences exist in the concurrent input/output definition file (IODF) activation process between a new tape library implementation and a configuration change that is made to a library. Changes to the virtual devices' address range of a library is an example of where concurrent IODF activation is useful.

Note: As an alternative to the procedures that are described next, you can 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, this installation is a new library from a host software definition perspective. 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 the devices in the library online. 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. Perform 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. For more information, see 12.1.2, "MVS system commands" on page 629.

Modifications to an existing tape library

When you are modifying a tape library so that device addresses can be changed, complete the following steps:

1. Activate an IODF that deletes all devices that correspond to the library.
2. Activate an IODF that defines again all of the devices of the modified library.
3. Run MVS console command **VARY ONLINE** to vary the devices in the library online. This command creates some of the control blocks. You see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```
4. Perform the final **ACTIVATE**.

Alternatively, you can use the **DS QL,nnnnn,DELETE** (where *nnnnn* is the LIBID) command to delete the library's dynamic control blocks. If IODF with LIBID and LIBPORT is already coded, 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, complete the following steps *after* you delete the library's dynamic control blocks:

1. Run MVS console command **VARY ONLINE** to vary on the devices in the library. This command creates some control blocks. 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 correctly defined.

6.2 Setting up the TS7700

The following sections describe the implementation and installation tasks to set up the TS7700. Specific names are used in this chapter if a task applies to only one model among supported models because slight differences exist between:

- ▶ TS7770C/TS7760C (TS7700C cloud storage tier configurations)
- ▶ TS7770D/TS7760D (TS7700D disk-only configurations)
- ▶ TS7770T/TS7760T (TS7700T tape-attached configurations)

For example, since the TS7700D and TS7700C do NOT support attachment to a tape library, then the implementation steps that are related to a physical IBM TS4500/TS3500 tape library are not applicable for those configurations. However, you can install the TS7760T or TS7770T with your existing or new TS4500/TS3500 tape library to serve as the physical back-end tape library.

6.2.1 TS7700 definitions

Use the TS7700 Management Interface (MI) for the following TS7700 Virtualization subsystem setup tasks:

- ▶ Defining scratch categories, including the data expiration settings for each category's virtual volumes.
- ▶ Defining TS7700 constructs

To use the Outboard Policy Management functions, you must define the following constructs:

- Storage Group (SG)
- Management Class (MC)
- Storage Class (SC)
- Data Class (DC)
- ▶ Defining Simple Network Management Protocol (SNMP) target.
- ▶ Defining a rocket-fast system for log processing ([RSYSLOG](#)) target.
- ▶ Defining Cluster Network Settings

- ▶ Enabling Secure Data Transfer, which requires FC 5281 installed in each TS7700 that uses this function.
- ▶ Defining Security Settings
- ▶ Inserting virtual volumes

For more information about MI definitions, see Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351, and Chapter 10, “IBM TS7700 Management Interface operations: Part 2” on page 449.

For more information about setting up the definitions, see 11.2.2, “TS7700 definitions” on page 578.

6.3 TS7700 software definitions

This section describes the software definitions that are required to implement the TS7700 in z/OS. The TS7700 must be defined as a new tape library with emulated 3490E Tape Drives from the host system.

To use the TS7700, at least one Storage Group (SG) must be created to enable the TS7700 tape library virtual drives to be allocated by the storage management subsystem (SMS) ACS routines. Only the composite library can be defined in the SG because all of the logical drives and volumes are only associated with the composite library.

For more information about host software implementation tasks for IBM tape libraries, see the following publications:

- ▶ *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, [SC23-6867](#)
- ▶ *IBM TS4500 R9 Tape Library Guide*, [SG24-8235](#)

If your Tape Management System (TMS) is DFSMS Removable Media Manager (DFSMSrmm), see the following publications for more information:

- ▶ *IBM z/OS DFSMShsm Primer*, [SG24-5272](#)
- ▶ *What is New in DFSMSrmm*, [SG24-8529](#)
- ▶ *z/OS DFSMSrmm Implementation and Customization Guide*, [SC23-5874](#)

If this installation is the first SMS tape library at this host, more steps are required. The full product documentation for your TMS must be consulted, in addition to the OAM PISA.

6.3.1 Defining volume catalogs

Consider the current size of the volume catalogs and the extra space that is required for the new volume and library entries. For a TS7700 with 100,000 volumes, at least 32 cylinders must be allocated.

To more precisely estimate the space allocation requirements in the tape volume catalog, complete the following steps:

1. 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.
2. 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](#).

6.3.2 Defining the TS7700 in a z/OS SMStape environment

Complete the following steps to define the TS7700 tape library in a z/OS SMStape environment:

1. Use the **ISMF Library Management** → **Tape Library** → **Define** window to create 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 250 show 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 window 1

Figure 6-7 shows the second window.

```
Command ==>                                TAPE LIBRARY DEFINE                                Page 2
SCDS Name . : MIKE.SCDS
Library Name : IBMC1
Media Type:                                Scratch Threshold
Media1: . . . . . 0      (0 to 9999999)
Media2: . . . . . 3000  (0 to 9999999)
Media3: . . . . . -      (0 to 9999999)
Media4: . . . . . 0      (0 to 9999999)
Media5: . . . . . 0      (0 to 9999999)
Media6: . . . . . 0      (0 to 9999999)
Media7: . . . . . 0      (0 to 9999999)
Media8: . . . . . 0      (0 to 9999999)
Media9: . . . . . 0      (0 to 9999999)
Media10: . . . . . 0     (0 to 9999999)
Media11: . . . . . 0     (0 to 9999999)
Media12: . . . . . 0     (0 to 9999999)
Media13: . . . . . 0     (0 to 9999999)
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 window 2

Figure 6-8 shows the third window.

```

Command ==>                                TAPE LIBRARY DEFINE                                Page 3
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 window 3

Figure 6-9 shows defining the distributed tape library.

```

Command ==>                                TAPE LIBRARY DEFINE                                Page 1
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 names that you defined at the TS7700 MI.
3. By using ISMF, create the SGs for the TS7700. Ensure that these defined construct names are the same as those names 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, the warning message CBR3017I is generated if a distributed library is defined to an SG. This message 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. Then, convert, test, and validate the ACS routines.
5. Activate the new Source Control Data Set (SCDS) by using the **SETSMS SCDS (scdsname)** command.
6. Restart the OAM address space by using the **F OAM,RESTART** command. SMS SCDS activation starts 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.

7. Define any required security profiles:

- RACF considerations for the VOLCAT

In general, tape users do not require any RACF access authority to the VOLCAT. During job processing, the updates to the VOLCAT are made by authorized system routines. However, the VOLCAT still needs a data set profile and must be defined with UACC(NONE).

Storage administrators that use ISMF must have READ access to STGADMIN.IGG.LIBRARY. IDCAMS users must have an access level to STGADMIN.IGG.LIBRARY that is suitable to the function that is performed.

For more information about the required RACF access level when IDCAMS is used, see “Required Security Authorization for VOLCAT Operations” in *z/OS DFSMS Access Method Services Commands*. SC23-6846.

8. Customize the tape management system:

- Customize your TMS to include the new volume ranges and library name. For DFSMSrmm, this process involves EDGRMMxx updates for OPENRULE, REJECT, or PRTITION statements as needed.
- Consider whether your current TMS database or control data set (CDS) includes sufficient space for the added library, data set, and volume entries. For more information about DFSMSrmm, see “Creating the DFSMSrmm CDS” in *DFSMSrmm Primer*, SG24-5983.
- Check and modify CBRUXxxx exits as required. For products other than DFSMSrmm, you might need to obtain modified exits from your supplier.

6.3.3 SYS1.Parmlib changes

Complete the following steps:

1. Modify the SYS1.Parmlib member DEVSUPxx for new categories, as described in 4.3.4, “Sharing and partitioning considerations” on page 186. The DEVSUPxx default categories must always be changed from the IBM-supplied defaults to prevent disruption of library operations. For more information, see “Changing the library manager category assignments in an ATLDS” in *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, [SC23-6867](#).
2. Add OAM to start automatically at IPL if it is not used.
3. If you want to vary virtual devices online after IPL, modify COMMNDxx.
4. In a Grid environment, you might want to update the ALLOCxx parmlib member SYSTEM TAPELIB_PREF to balance workloads. The default algorithm, EQUAL, works well but if the libraries under consideration include an equal number of online devices, the recommendation is to use BYDEVICES.
5. Change Missing Interrupt Handler (MIH) The TS7700 emulates 3490E devices and for clusters that are running microcode R4.1.2 and later, communicate a default Primary MIH Time OUT value of 45 minutes to the host operating system in the Read Configuration Data channel control word (CCW X'FA').

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.

The MIH timeout value applies only to the virtual 3490E drives and *not* to the real IBM TS1160/TS1150/TS1140/TS1130/TS1120/3592 drives that the TS7700T manages in the back end.

The commonly used MIH values are listed in Table 6-6. These values can be adjusted depending on specific operational factors.

Table 6-6 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 by using Re-wind Unload (RUN) copy policies	20 minutes

If adjustment is needed, the MIH value can be specified in the parmlib member IECIOSxx. Alternatively, you can set the MIH values by using the IBM Z operator command **SETIOS**. A user-specified MIH value overrides the default value and is effective 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:

- ▶ 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 the following command:


```
SETIOS MIH,DEV=(0A40-0A7F),TIME=45:00
```

 - To display the new settings, use the following command:


```
D IOS,MIH,DEV=0A40
```
 - To check the current MIH time, use the following 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** 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 Time OUT value is used for most I/O commands, but the *Secondary* MIH Time OUT value can be used for special operations, such as long-busy conditions or long-running I/O operations.

Whenever 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 Time OUT value. This design implies that if an MIH timeout value that is equal to the MIH default for the device class is specifically requested, IOS does not override the device-established Primary MIH Time OUT value. To override the device-established Primary MIH Time OUT value, you must specifically set a timeout value that is not equal to the MIH default for the device class.

6.3.4 Final steps

Complete the following steps:

1. Vary the composite library and distributed libraries online. For more information, see 12.1.1, “DFSMS operator commands” on page 626.
2. Vary the TS7700 virtual drives online, as described in 12.1.1, “DFSMS operator commands” on page 626.
3. Insert logical volumes: NB OAM must be started and the CBRUX* exits enabled.
4. Confirm that volumes are inserted correctly with correct Category on TS7700 (see 12.5, “Host cartridge entry processing” on page 650), volumes are in scratch status in the TCDB, and LI REQ display of Composite library shows the correct scratch count.
5. Run test jobs, and so on.

6.4 Attaching the TS7700 to a physical tape library or cloud

This section describes the optional attachment to an IBM Tape library or to an external cloud object store.

Note: Other cloud providers may be supported doing a SCORE request.

6.4.1 Defining cache partitions

The TS7700 supports using different data storage tiers. It also supports partitioning the free space in the cache storage subsystem for better management of workloads with the following use cases:

- ▶ TS7770D/T/C with installable optional Feature Code (FC) 5283, which enables the TS7700 to work as an Object Store for DS8000 series storage subsystem.
- ▶ TS7700T that writes copies of virtual volumes to an IBM TS4500/TS3500 physical tape library.
- ▶ TS7700C that writes copies of virtual volumes to a cloud repository.

For the TS7700T (tape attach) or TS7700C (cloud attach) cluster model, up to eight cache partitions (one CP0 cache resident and up to seven partitions) can be defined. See 9.8.1, “Cache Partitions” on page 404.

Note: The default size for nonresident cache partitions within a TS7700T (CP1-CP7) is 3 TB. The customer can modify that size by using the corresponding MI window according to specific business needs. This setting must be reviewed and adjusted if the TS7700 was replaced with a newer generation by way of the Field Frame Replacement method, which is also known as a Frame Roll.

6.4.2 TS4500/TS3500 tape library definitions

The section describes implementing a TS7700T attached to a TS4500/TS3500 tape library.

Your IBM Service Support Representative (IBM SSR) installs the TS7700T hardware, its associated tape library, and the frames. This installation does not require your involvement

other than the necessary planning. For more information, see Chapter 4, “Preinstallation planning and sizing” on page 145.

Clarification: The steps that are described in this section relate to the configuration of a new or existing IBM TS4500/TS3500 tape library with all the required features installed, such as Advanced Library Management System (ALMS).

If you are attaching to an IBM physical tape library that is also attached to Open Systems hosts or IBM Z hosts, for more information about extra steps that might be required, see *IBM TS4500 R9 Tape Library Guide*, [SG24-8235](#).

The following tasks apply to the TS4500/TS3500 library definition (For more information about the procedure, see 11.2.1, “Tape library with the TS7700T cluster” on page 561.):

- ▶ Defining a logical library:
 - Ensuring that ALMS is enabled
 - Creating a logical library partition with ALMS
 - Setting the maximum cartridges for the logical library
- ▶ Adding drives to the logical library

Each TS7700T is associated with only one logical library in the physical library, and each logical library can be associated with only one TS7700. A TS7700T requires a minimum of four installed tape drives to be operational, and a maximum of up to 16 is supported.
- ▶ Defining control path drives

Each TS7700T requires the definition of four tape drives (among installed drives) as so called “control path devices”.
- ▶ Defining the encryption method for the new logical library

The TS7700T supports only encryption when its attached logical library is configured to use the System-Managed encryption method.
- ▶ Defining Cartridge Assignment Policies (CAPs)

Defining CAP policies automates the assignment of physical volumes to the logical library that are associated with the target TS7700T.
- ▶ Inserting TS7770T or TS7760T physical volumes into the library

This task consists of the physical insertion of cartridges to be used by the TS7700. This process can be done by using the I/O station in the tape library, or by opening a tape library enclosure door and directly inserting cartridges into available slots. The inventory is needed only if the door is opened to insert cartridges.
- ▶ 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 declared in advance.
- ▶ Relabeling physical volumes

This procedure is necessary only if a mismatch exists between the internal and external labels of a physical tape cartridge. To correct such a condition, you can use the TS7700 library request command interface to submit a `LI REQ PVOL, <volser>, RELABEL, YES` request against the affected cartridge, which marks it for relabeling the next time that it is written from the beginning of tape.

Note: For more information about the Library Request commands and their responses, see *IBM TS7700 Series z/OS Host Command Line Request User's Guide Version 5.4*.

6.4.3 Definitions for TS7700T

The following tasks apply only to TS7700T:

- ▶ Defining VOLSER ranges for physical volumes
The physical VOLSER ranges in the TS7700T must match the cartridge assignments in the logical library of the attached TS4500/TS3500
- ▶ Defining physical volume pool properties:
 - Set the reclaim threshold.
 - Set up any Inhibit Reclaim schedules, if required.
- ▶ If the TS7700T is required to perform tape encryption (only for machines that have Feature Code 9900 installed), the following definitions are applicable:
 - Data at rest encryption configuration
This cluster setting requires the address of the IBM Security Guardium Key Lifecycle Manager server that is responsible for managing the encryption keys that are used for data that is copied to physical volumes.

Key exchanges for tape encryption are performed by using the IPP protocol. Starting with microcode R5.0, the TS7700T supports the use of the TLS1.2 protocol to encapsulate IPP traffic to target IBM Security Guardium Key Lifecycle Manager, which requires importing its corresponding server certificate into the TS7700T.

For more information about configuring related settings, see “TS7770 Data at Rest Encryption” on page 525.
 - Physical volume pool encryption settings.
Each physical volume pool can be configured to use encrypted data that is copied to physical volumes that are associated with the pool. If encryption is specified, encryption keys are required to be configured (which must be available in configured IBM Security Guardium Key Lifecycle Manager).

For more information about configuring related settings, see 10.1.1, “Physical Volume Pools” on page 450.

6.4.4 Cloud Tier Settings (TS7700C only)

Use the TS7700 MI for the setup tasks that are described in this section. Select **Cloud Tier Settings** in the Settings group, as shown in Figure 6-10.



Figure 6-10 Cloud Tier Settings option

To use TS7700C Cloud Storage Tier, you must define the following settings:

► Cloud pools

Similar to physical volume pools in TS7700T, a *cloud pool* is a data storage policy that is to be associated to virtual volumes that are intended to be copied to a cloud repository.

► Cloud accounts

This setting includes details about credentials for users who are authorized to access the target cloud repository and type of repository.

► Containers

A *Cloud Container* represents the specific space (within the target cloud repository) that is assigned to the TS7700C to store virtual volumes in the form of object files. The Cloud Container window in the MI is also used to declare or record the following information:

- Cloud URL of a supported S3 cloud repository service:
 - Amazon Web Services (AWS) (off premises)
 - IBM Cloud Object Storage (on premises)
 - IBM Cloud Public (off premises)
- Cloud URL to cluster association

Note: For more information about tasks that are required to implement a TS7770C or TS7760C, see *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-55733.

6.4.5 Object Store settings

In TS7700 code level 8.50.x.x, the DS8000 Object Store feature was first introduced (FC 5282). This feature allowed one or more DS8000 primary Storage Systems to connect to a TS7700 and move DFSMSHsm MIGRATE data and DFSMSdss Full Volume Dump data to the TS7700 in the form of object files and not a logical tape volume.

Those objects are stored on a TS7700 cluster in a logical cache partition. The cache partition provides a way to separate object data into a bucket where a specific amount of space can be allocated.

In TS7700 R5.0, the DS8000 targets one or two TS7700s in the same grid to provide synchronous copy support by way of the DS8000 function that is known as *forking*. However, those objects were not replicated further to any other cluster in the grid.

Starting with release R5.2.2 (code level 8.52.200.109 and higher), TS7770 machines with installed Feature Code 5283 (Advanced Object Store) now can catalog those objects and the containers in which they exist. Users can set policies so the objects can be managed for replication purposes, and to deploy copies to other object-enabled TS7770 machines that are also present in the same grid.

MI pages under the Object Store group (as shown in Figure 6-11 on page 257) allow the user to create, modify, and delete Object Policies to be linked to Object Stores.



Figure 6-11 Object Store Group

Use the TS7770 MI for the following Object Store setup tasks:

► Object Policy

Object policies are used to determine how to manage objects targeting TS7770 from an external host, such as a DS8000 running transparent cloud tiering (TCT) workloads. Object policies include replication settings to define object redundancy across a grid by choosing the Copy Mode which better suits business needs.

Object policies are grid scope and pertain to how an object is managed for each cluster in the grid. You can set up object policies for every object-enabled cluster in the grid by accessing this page from one of those clusters.

► Object Store

To use DS8000 TCT, you must create a z/OS DFSMS Cloud Network Connection Construct name (cloud name) in the ISMF cloud window. That same cloud name must also be defined on the TS7770 as an Object Store before DS8000 is configured and zSeries Host start transactions for that cloud name. If the cloud name does not exist on TS7770, the DS8000 configuration process or host TCT transactions fail.

The **Object Store** page is used to add cloud names to a TS7700 grid. Cloud names are used as object store targets for TCT, and each cloud name can be assigned to an existing object policy. Object policies enable the users to customize management of the objects within the TS7700 grid.

The object store page also lists the containers that are created by the z/OS host application. A container is a data storage bucket that contains 0–*n* data objects.

The **Object Store** page is grid-scope. Object stores that are created on any object-enabled cluster in the grid are replicated on the other object-enabled clusters.

Notes:

1. For more information about tasks that are required to implement a TS7770 Object Store, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).
2. Also, users of the DS8000 Object Store capability are encouraged to join the **DS8000 Transparent Cloud Tiering** public forum in the IBM Storage Community (the online community where IBM Storage Users meet, share, discuss, and learn).

6.4.6 Defining fence actions

Usually, the TS7700 can detect problems that can impact its availability. However, there are specific conditions that are difficult to detect, caused by the intermittent nature of the failure, the frequency at which it occurs, its duration, or other circumstances that can escape the normal mechanism for local detection. Under these circumstances, the failing cluster can still report itself as being “healthy”, even if it is not.

This type of failure is known as “Sick But Not Dead” (SBND). If only a single cluster is in this condition, It can affect the performance or function of other clusters that are present in the same grid.

The consequences for this type of problem can be prevented by a “fence” mechanism that is available when all machines in the grid are running microcode R4.1.2 or later. This mechanism isolates the failing box while repair actions are performed while the other clusters in the grid continue to run unaffected.

In an SBND scenario, the failing box cannot recognize that a problem exists and keeps attempting to take part in grid transactions. The other clusters in the grid recognize this condition and run a “remote fence” action against the sick cluster.

The following new LIBRARY REQUEST FENCE options are available to customize the fencing process, depending on business requirements:

► Remote Cluster Fence Function enablement:

```
LI REQ, <composite-library>, FENCE, <ENABLE|DISABLE>
```

The default is DISABLE. This control allows or prevents the execution of remote cluster fence actions from the perspective of the entire grid.

► Primary fence action:

```
LI REQ, <distributed-library>, FENCE, ACTION, PRI,  
<NONE|ALERT|OFFLINE|REBOOT|REBOFF>
```

This command features the following options:

- NONE: This setting is the default. Other clusters in the grid cannot perform remote fencing against the specified cluster (“distributed-library”) if it fails. Therefore, the system depends on the ability of the local cluster to detect that something is wrong.
- ALERT: Other clusters surface an alert only when they believe that the specified grid member is having problems.
- OFFLINE: Other clusters trigger a force-offline sequence against the specified cluster to exclude it from grid transactions if the corresponding thresholds are met. If this primary action fails, a secondary fence action (if enabled) is started.
- REBOOT: Other clusters trigger a restart sequence against the specified cluster if the corresponding thresholds are met. Sometimes, restarting the SBND cluster is the best solution and faster recovery. The cluster automatically starts online processing after the restart. If the machine fails to be restarted, a secondary fence action (if enabled) is started.
- REBOFF (restart and stay offline): Other clusters trigger a restart sequence against the specified cluster when the corresponding thresholds are met. This sequence prevents it from attempting to come online again.

Use this option if health checks and troubleshooting are required the system comes back online again. If the machine fails to be restarted, a secondary fence action (if enabled) is started.

► Secondary fence action:

```
LI REQ, <distributed-library>, FENCE, ACTION, SEC, <ENABLE|DISABLE>
```

Referred to as “Isolate from grid network”, this secondary action is triggered against the specified cluster only if it was previously enabled (default is DISABLE), and the primary fence action failed. In this case, other clusters automatically run a reconfiguration of the intercluster network to logically exclude the failing cluster from any grid transactions. This action is available for offline or restart primary actions only.

- ▶ Enable/Disable Capturing IBM AIX dump at restart action:

LI REQ, <distributed-library>, FENCE, ACTION, AIXDUMP, <ENABLE|DISABLE>

An AIX dump is a data capture mechanism that fetches information that is related to the IBM AIX kernel state at the time it is taken. This information can then be used by authorized IBM Service or Support personnel during the diagnostic process if a failure occurs. This AIX dump is taken only when AIXDUMP is enabled (default is DISABLE) and the selected primary fence action is REBOOT or REBOFF.

- ▶ Remote Cluster Fence Thresholds:

LI REQ, <composite-library>, FENCE, THRESHLD,
<TMO|ERR|SCRVOAVG|PRIVOAVG|VCAVG|TOKAVG|EVALWIN> <value>

The THRESHLD keyword provides a way to select how and when a specified fence action occurs if an SBND condition is detected. These settings result in a grid-wide effect, and the following options can be adjusted, depending on your business needs:

- Threshold concepts (the third keyword and the allowed values in the fourth keyword) in the LI REQ command):

- TMO (timeout event counts): The number of internal system timeouts that must occur before the primary fence action is started. This amount relates to the combined sum of timeout incidents for scratch mounts, private mounts, volume close, and token handshakes.

This fourth keyword of the LI REQ allows values is 0–1000 seconds (the default is 20).

- ERR (error event counts): The number of specific failures that must occur before the primary fence action is started. This amount relates to the combined sum of specific error incidents for scratch mounts, private mounts, volume close, and token handshakes.

This fourth keyword of the LI REQ allows values is 0–1000 seconds (the default is 20).

- SCRVOAVG (scratch volume open average): The average peer handshake time to open or mount a scratch virtual tape that must be exceeded before the primary fence action is started.

This fourth keyword of the LI REQ allows values is 0–1200 seconds (the default is 180 seconds).

- PRIVOAVG (private volume open average): The average peer handshake time to open or mount a private virtual tape that must be exceeded before the primary fence action is started.

This fourth keyword of the LI REQ allows values is 0–1200 seconds (the default is 180 seconds).

- VCAVG (volume close average): The average peer handshake time to close or unmount a tape that must be exceeded before the primary fence action is started. This value excludes the RUN (Rewind and UNload) copy duration.

This fourth keyword of the LI REQ allows values is 0–180 seconds (the default is 120 seconds).

- TOKAVG (token handshake average): The average peer handshake time to process a miscellaneous token handshake that must be exceeded before the primary fence action is started.

This fourth keyword of the LI REQ values is 0–180 seconds (the default is 120 seconds).

- EVALWIN (evaluation window): The referred lapse in minutes that is used to calculate counts and averages for target concepts.

This fourth keyword of the LI REQ allows values is 1–30 minutes (the default is 7 minutes).

► Showing current Fence Status and Settings:

LI REQ, <composite-library>, FENCE, SHOW

Note: For more information about applicable fence actions, its corresponding configuration, associated Library Request commands, and their responses, see [IBM TS7700 Series Grid Resiliency Improvements User's Guide V1.1](#).



Hardware configurations and upgrade considerations

This chapter includes descriptions of the hardware configurations and upgrade considerations for the TS7700.

This chapter includes the following topics:

- ▶ 7.1, “TS7700 hardware components” on page 262
- ▶ 7.2, “TS7700 component upgrades” on page 274
- ▶ 7.3, “TS7700 upgrade to Release 5.4” on page 288
- ▶ 7.4, “Adding clusters to a grid” on page 289
- ▶ 7.5, “Removing clusters from a grid” on page 299

7.1 TS7700 hardware components

IBM TS7700 Release 5.4 microcode runs on only a 3957/3948 model VED with the necessary hardware installed. The model VED is an IBM Power9 processor-based server that includes attached redundant I/O expansion drawers that contain Peripheral Component Interface Express (PCIe) adapters.

Systems that are shipped from manufacturing with R5.3 and later are shipped as updated machine type 3948. Systems do not change machine type labels when upgraded to R5.4 in the field. Systems cannot mix MTs, so the hardware upgrades in the field must be of the same machine type. Manufacturing continues to ship field upgrades for 3957 and 3956 systems. All Feature Codes that are supported for previous MTs are supported for new MT 3948. The new MT 3948 is included in the Expert Care Services.

MT 3948 for all components:

- ▶ 3952-F07 (frame), becomes 3948-F07
- ▶ 3957-VED (server), becomes 3948-VED
- ▶ 3956-CSB,XSB,CFC,XFC (cache controller and expansions), becomes 3948-CSB,XSB,CFC,XFC

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 TS7770, which can be attached to an IBM TS4500 Tape Library containing IBM 3592 tape drives. Attachment is through 16 Gbps fiber switches. As of this writing, the attachment of the TS7770 is mutually exclusive to an IBM tape library or to a cloud storage tier.

DS8000 Object Store (FC 5282) was introduced in 8.50.x code level and is supported on the TS7770 VED model only.

TS7700 Advanced Object Store (FC 5283) was introduced in the 8.52.2xx code level and is supported on TS7770 VED models. FC 5283 replaces the previously implemented FC 5282. However, FC 5282 continues to be supported.

7.1.1 Common components for the TS7700 models

This section lists the components for the TS7700 models.

3952/3948 Tape Frame

The 3952/3948 Tape Frame houses a TS7700 server, controllers, and their components. The 3952/3948 Tape Frame that is used with the TS7700 contains the following components:

- ▶ Ethernet switches
- ▶ TSSC Server

- ▶ Optional components:
 - 3-Phase power distribution unit by way of RPQ 8B3722
 - 3-Phase power cord by way of RPQ 8B3723
 - Front/Rear and Side Panel Locking Procedure by way of RPQ 8B3669
 - Door lock kit rear by way of feature code 2749
 - Top cable exit by way of feature code 2750
- ▶ TS7700 Server:
 - TS7770 Server (3957/3948-VED)
- ▶ I/O drawers
- ▶ 3956/3948 Cache controller:
 - TS7770 Cache Controller (3956/3948-CFC)
 - TS7770 Cache Controller (3956/3948-CSB)
 - TS7760 Cache Controller (3956/3948-CSA)
- ▶ 3956/3948 Optional cache expansion drawers:
 - TS7770 Cache Drawer (3956/3948-XFC)
 - TS7770 Cache Drawer (3956/3948-XSB)
 - TS7760 Cache Drawer (3956/3948-XSA)

The 3952/3948 Tape Frame can be designated as a TS7700 Storage Expansion Frame when ordered with FC 7336. A total of two Storage Expansion Encryption-capable Frames is supported.

Any lock on the 3952/3948 Tape Frame prevents access to the front 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. Also, the emergency contact label that is included with the Installation Instruction RPQ 8B3669 (Front/Rear and Side Panel Locking Procedure), P/N 00VJ478, 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 3952/3948 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, contact your IBM service representative to order a new lock and key set (FRU 12R9307).

Available by feature code 2749 is a door lock and key that prevents access to the rear components of the TS7770.

For more information about 3952/3948 Tape Frame specifications, see Table 4-2 on page 147.

The IBM 3952/3948 Model F07 Tape Base Frame provides up to 36EIA units of usable space. The rack units contain the components of the defined tape solution. In 2021, IBM announced a new Tape Frame (3592-F07), which is slightly shorter compared to the old F06 frame (1.28 m compared to 1.42 m). The F07 frame is similar to the DS8000 frame.

The 3952/3948 Tape Base Frame is not a general-purpose frame. The 3952/3948 Model F07 is designed to contain the components of the TS7770. Only components of one solution family can be installed in a 3952/3948 Tape Frame. The Tape Frame is configured with Dual AC Power Distribution units for redundancy.

Note: Consider the following points:

- ▶ Feature code 2750 allows cabling to exit from the top of the 3952/3948 F07 frame.
- ▶ RPQ 8B3749 enables the installation of the TS7770 with solid-state cache and components into a customer-supplied rack that requires 18U of space. When this option is used, a maximum of 640 TB of cache can be configured.
- ▶ RPQ 8B3721 enables the installation of the TS7770 with hard disk drive cache and components into a customer-supplied rack that requires 18U of space. When this option is used, a maximum of 780 TB of cache can be configured.

Ethernet switches

Primary and alternative switches are used in the TS7700 internal network communications for redundancy.

The communications to the external network use a set of dedicated Ethernet ports on adapters in the 3957/3948 server.

Figure 7-1 shows the Ethernet switch that is used in the TS7770.



Figure 7-1 Ethernet switch

TS7700 grid adapters

The connection paths between multiple TS7700 clusters in a grid configuration are grid adapters. The dual-port 1 gigabit per second (Gbps) Ethernet adapters can be copper RJ45 or 10 Gb optical fiber (longwave [LW]) that include an LC duplex connector. In the TS7770, the grid adapters are now in the 3957/3948-VED server.

Depending on your bandwidth and availability needs, the TS7700 can be configured with two or four 1 Gb links. Feature Code (FC) 1034 is needed to enable the second pair of ports in the grid adapters. These ports are copper.

Optionally, an LW dual-port Optical Ethernet adapter (FC 1038) is available for two or four 10 Gbps 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).

TS7770 Disk Encryption

The TS7770 must be requested with encryption when it is ordered, which results in the CSB/CFC Distributed RAID (DRAID) being encrypted at creation time because encryption cannot be enabled later. For countries in which encryption cannot be requested, a TS7770 can be ordered without encryption.

For local encryption, the following points apply:

- ▶ Local encryption includes four USB drives to store encryption keys. Two of the drives must be attached to the CSB/CFC during start, and two are used as backup keys. If the USB drives are removed and any CSB controller is restarted, it fails to come up correctly.

- ▶ Rekey for Local encryption needs support to assist because USB drives must be moved during the process.

For external encryption, the following points apply:

- ▶ IBM Security Guardium Key Lifecycle Manager KMIP with TLS 1.2 support (distributed IBM Security Guardium Key Lifecycle Manager only) TLS V1.2 must be configured on the Key Manager Server.
- ▶ The customer must exchange certificates and trust them on the Key Manager server and TS7770 by using the Management Interface (MI).
- ▶ After external encryption is enabled, the customer can perform a rekey from MI.
For more information, see Appendix J, “Configuring externally managed encryption” on page 1013.
- ▶ Thales CipherTrust Manager 2.0 and later is supported on new systems with no conversion allowed. For more information about the Thales Cipher Trust manager, see [IBM TS7700 Disk Encryption for CSB and CFC models: Local, GKLM and Thales](#) and “Thales CipherTrust Manager Configuration” on page 1038.

The TS7770 cache model 3956/3948-CSB and 3956/3949-CFC support data-at-rest encryption, which is defined by the following characteristics:

- ▶ Data-at-rest means that the data is encrypted on the end device (drives) by way of software and hardware.
- ▶ The algorithm that is used is the Advanced Encryption Standard (AES) a US government standard from 2001.
- ▶ Encryption of data-at-rest complies with the Federal Information Processing Standard 140 (FIPS-140) standard, but is not certified.
- ▶ Ciphertext stealing XTS-AES-256 is used for data encryption.
- ▶ AES 256 is used for master access keys.
- ▶ The algorithm is public; the only secrets are the keys.
- ▶ A symmetric key algorithm is used. The same key is used to encrypt and decrypt data.

The following types of keys are used to encrypt the drives and are defined in the system:

- ▶ Master access key:
 - The master access key is created when encryption is enabled.
 - The master access key can be stored on USB flash drives or key servers. One master access key is created for each enabled encryption key provider.
 - It can be copied or backed up as necessary.
 - It is not permanently stored anywhere in the system.
 - It is required at start to unlock access to encrypted data.
- ▶ Data encryption keys (one for each encrypted object):
 - Data encryption keys are used to encrypt data. When an encrypted object (such as an array, pool, or child pool) is created, a new data encryption key is generated for this object.
 - Managed disks (MDisk) that are not self-encrypting are automatically encrypted by using the data encryption key of the pool or child pool to which they belong.
 - MDisks that are self-encrypting are not reencrypted by using the data encryption key of the pool or child pool they belong to by default. You can override this default by manually configuring the MDisk as not self-encrypting.

- Data encryption keys are stored in secure memory.
- During cluster internal communication data encryption keys are encrypted with the master access key.
- Data encryption keys cannot be viewed or changed.
- When an encrypted object is deleted, its data encryption key is discarded (secure erase).

Note: If all master access key copies are lost and the system must cold restart, all encrypted data is gone. No method exists, even for IBM, to decrypt the data without the keys. If encryption is enabled and the system cannot access the master access key, all SAS hardware is offline, including unencrypted arrays.

The TS7700 can use external key management with the 3956/3948-CSB/CFC cache types. The 3956/3948-CSB/CFC cache requires FC 5272,5276 and FC 7405 to enable external management.

Systems that are set up with local encryption key management can be converted to external key management. For more information, see Appendix J, “Configuring externally managed encryption” on page 1013.

The internal disks within the Power9 Server are storing system data and metadata only; therefore, they are not encrypted.

Note: Only the distributed IBM Security Guardium Key Lifecycle Manager supports external disk and tape (TS1140, TS1150, and TS1160 tape drives) encryption. The settings for Encryption Server are shared for tape and external disk encryption.

Although the IBM Security Guardium Key Lifecycle Manager for z/OS supports TS7700 physical tape, it does not support TS7700 disk encryption. The only valid option is the Container Edition for z/OS of IBM Security Guardium Key Lifecycle Manager, which is used with TS7770 disk encryption.

Back-end drive 16 Gb fiber switches

Release 5.4 or later includes support for TS7770T to attach to a TS4500 Tape Library. R5.4 supports the 16 Gbps fiber Switches that are used to communicate with the back-end tape drives. The TS7700 16 Gbps fiber Switches can be housed in a 3584-LXX or 3584-DXX frame. The switches can be in a frame that contains some of the associated back-end drives, or can be 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.

A frame that contains the back-end switches can still house up to 12 or 16 drives (based on a TS4500). FC 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 TS7770T 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 was removed.

Note: The TS7770T does not support 4 Gbps and 8 Gbps fiber switches for connection to the back-end drives. Currently, a TS7770T must use the 16 Gbps fiber switches to connect to the back-end drives.

7.1.2 TS7770 components

This section describes the TS7770 components.

Figure 7-2 shows the frame layout for a TS7770 Storage Expansion Frame.

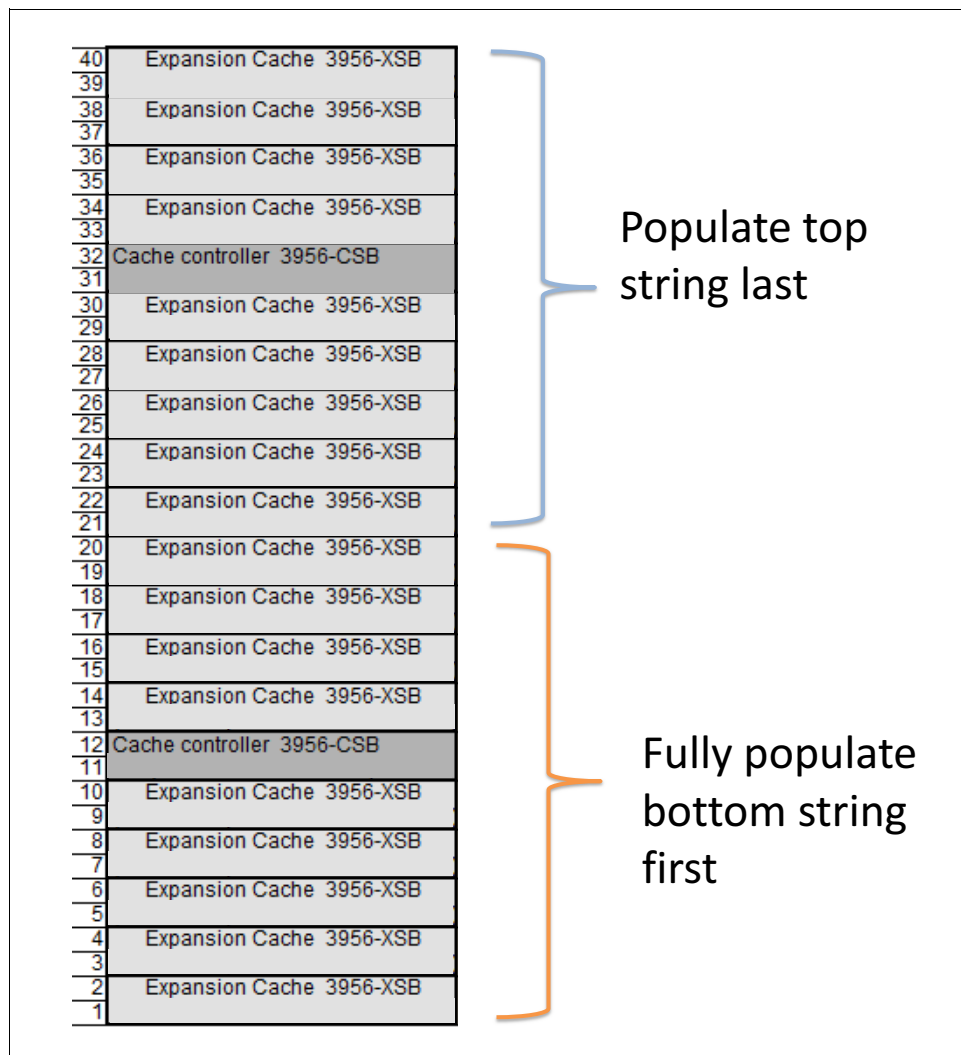


Figure 7-2 Layout of a TS7770 Storage Expansion Frame with 3956-CSB and 3956-XSB

The TS7770 Storage Expansion Frame is a 3952 Tape Frame that is designated as a cache expansion frame for use with a fully configured TS7770 Base Frame. A TS7770 Storage Expansion Frame enables expansion of the TS7770 Cache by attaching up to two more storage frames. Each frame contains two TS7770 Cache Controllers and optional TS7770 Cache Drawers. The cache controllers are model CSB, and cache drawers are model XSB.

The distance between a TS7770 Storage Expansion Frame and the TS7770 Base Frame cannot exceed 10 meters (32.8 feet). This distance enables connection of the frames by using a 30-meter (98-foot) fiber cable.

A total of two Storage Expansion Frames are supported, which adds up to a total capacity of 3.94 PB of usable capacity using CSB cache strings.

TS7700 Server model (3957-VED)

The TS7770 Server comprises a server and two I/O expansion drawers for PCIe adapters. The TS7770 Server controls virtualization processes, such as host connectivity, 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-VED) offers the following features:

- ▶ Two 10 core processors each at 3.8 GHz
- ▶ 128 GB of DDR4 2400/2133 MHz memory default since R5.3
- ▶ FC 3479, 64 GB memory upgrade (that is, 128 GB or 256 GB supported)
- ▶ 6 SFF 600 GB SAS 10k RPM internal drives that use RAID 0 or 6 SFF 775 GB SSD internal drives that use RAID 0
 - Two more SFF SSDs when FC 8083 is installed, which is a prerequisite for FC 5283 TS7700 Advanced Object Store for DS8000.
- ▶ 1 Gbps Quad-port Copper or 10 Gbps Dual-port LW Optical Ethernet adapters for grid network (up to four ports total)
- ▶ Four USB ports:
 - Two front USB 3.0 ports
 - Two rear USB 3.0 ports
- ▶ 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
- ▶ Integrated service processor

Each Expansion I/O adapter drawer offers the following features:

- ▶ Six extra hot-pluggable PCIe cartridge style slots (used to house FICON adapters for host attachment and Fibre Channel adapters for cache drawer attachment)
- ▶ Redundant AC power
- ▶ Redundant cooling
- ▶ Concurrent maintenance of:
 - PCIe or PCI-X adapters
 - I/O drawer
 - Two power supplies
 - Two fans

The TS7770 can also be installed in a customer-supplied, 19-inch rack.

Figure 7-3 shows the TS7770 server unit.



Figure 7-3 TS7700 Server (3957-VED) System Unit (front view)

Figure 7-4 shows the TS7770 rear locations.




Figure 7-4 TS7700 Server (3957-VED) System Unit (rear view)

Figure 7-5 shows the redundant I/O drawer adapter locations.




Figure 7-5 TS7700 Server Expansion Unit I/O drawer (rear view)

Note: The top two ports on the adapter in Slot 4 provide connectivity for the first Storage Expansion Frame. The lower two ports provide connectivity for the second Storage Expansion Frame.

TS7770 Cache Controller

The TS7770 Cache Controller is a self-contained 2U enclosure that mounts in the 3952 Tape Frame. The TS7770 Cache Controller allows for the creation of Pools, MDisk (DRAIDS), and Volumes. Each pool is made up of 1 MDisk with each MDisk being made of 1 Enclosure and 12 drives. These Distributed RAID arrays provide an equal of RAID6 level of protection.

The TS7770 Cache Controller offers the following features:

- ▶ Two Canisters (node) each containing 32 GB of RAM with 6-core Broadwell-DE Xeon running at 1900 MHz and system flash
- ▶ Two active 16 Gbps FC connectivity to server per canister (x4 = 64 Gbps)
- ▶ Two Drive Expansion connectors per canister for a total of 4 cables x 12 Gbps transfer rate = 48 Gbps
- ▶ 12 hot swap 10 TB NL-SAS 7.2k RPM drives for Model CSB or
- ▶ 24 hot swap 3.84 TB SAS SSD drives for Model CFC
- ▶ Two battery backup units (one per canister)
- ▶ Two AC power supplies with embedded enclosure cooling units
- ▶ One 10/100/1000 Mbps Ethernet technician port per canister
- ▶ Two 1/10 Gbps Ethernet ports per canister
- ▶ One USB 2.0 port per canister

Figure 7-6 shows the TS7770 Cache Controller front view.



Figure 7-6 TS7770 Cache Controller 3956-CSB (front view)

Figure 7-7 on page 271 shows the TS7700 Cache Controller rear view.




Figure 7-7 TS7770 Cache Controller 3956-CSB (rear view)

TS7770 Cache Drawer

The TS7770 Cache Drawer is a self-contained 2U enclosure that mounts in the 3952 Tape Frame. The TS7770 Cache Drawer expands the capacity of the TS7770 Cache Controller by providing more dynamic disk pools-protected disk storage. Figure 7-8 shows the front view of the TS7770 Cache Drawer.



Figure 7-8 TS7770 Cache Drawer 3956-XSB (front view)

Figure 7-9 on page 272 shows the rear view of the TS7770 Cache Drawer.




Figure 7-9 TS7770 Cache Drawer 3956-XSB (rear view)

TS7700 High-Performance Tape/Cloud Controller (C07 replacement)

This FC 5999 hard bundles multiple Feature Codes into one code as a replacement for C07 tape controllers. Only a subset of all Feature Codes are allowed. This FC can be Cloud or Tape attached.

The following features are included in FC 5999:

- ▶ No capacity expansion, 1 CFC cache drawer only
- ▶ Enable FICON 2nd port
- ▶ Grid Enablement
- ▶ 60 TB of cache enablement
- ▶ 4 GBps of incremental throughput
- ▶ 4 TB of premigration
- ▶ Can be installed in customer-supplied rack RPQ 8B3749 (CSB cache)
- ▶ Can be installed in customer-supplied rack RPQ 8B3721 (CFC cache)

7.1.3 TS7700 tape library attachments, drives, and media

In TS7700 configurations, the TS7770T is used with an attached tape library. The TS7770T must have their own logical library within the TS4500 tape library, with dedicated tape drives and tape cartridges.

Tape libraries

A TS7770T attached to a TS4500 Tape Library interfaces directly with tape drives in the library. 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. A pair of Fibre Channel switches routes the data to and from the correct tape drive.

Note: TS1140 (EH7) tape drives, TS1150 (EH8) tape drives, and TS1160 (60F) are used with the TS4500 Tape Library.

Tape drives

The 3592 Tape Drives are supported for use with the TS7770 include:

- ▶ TS1160 Tape Drive (As of this writing only TS7700 running code level 8.52.101.12 and later supports the TS1160 tape drive)
- ▶ TS1150 Tape Drive
- ▶ TS1140 Tape Drive
- ▶ TS1130 Tape Drive
- ▶ TS1120 Tape Drive (in native mode and emulating 3592 J1A Tape Drives)
- ▶ 3592 J1A Tape Drive

Tape media

The TS7770 supports 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 Advanced Type E Data (JE)
- ▶ 3592 Economy Tape Cartridge (JJ)
- ▶ 3592 Advanced Type K Economy (JK)
- ▶ 3592 Advanced Type L Economy (JL)
- ▶ 3592 Advanced Type E Economy (JM)

For more information, see “Tape drives and media support (TS7700T)” on page 153.

7.1.4 TS3000 Total System Storage Console

The TS3000 Total System Storage Console (TSSC) connects to multiple Enterprise Tape Subsystems, including TS4500 tape libraries, 3592 Controllers, and the TS7700. The TS3000 TSSC is a required component for the TS7770. The new P9 server based console is installed in the TS7700 3952 F07 Tape Base Frame. Starting with R5.3 pga1 (late summer 2023) the TSSC/TS3000 Service Console is part of the 3948-VED server ship-group. This means the TSSC/TS3000 specific feature codes within the F07 are discontinued and will not be listed anymore in a configuration.

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 is enabled.

7.1.5 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 TS7770.

A TS7770 Server with the FICON Attachment features (FC 3402 or FC 3403) can attach to FICON channels of IBM Z components by using FICON cable features that are ordered on the TS7700 Server.

One cable must be ordered for each host system attachment by using the following cable features:

- ▶ FC 3402 (16 Gb FICON Short 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 16 Gbps speed, the total cable length cannot exceed the following lengths:
 - 130 meters (426.5 feet) by 50-micron OM4 (47000 MHz*km) Aqua blue-colored fiber
 - 100 meters (328 feet) by 50-micron OM3 (500 MHz*km) Orange-colored fiber
 - 35 meters (114.8 feet) by 50-micron OM2 (500 MHz*km) Orange-colored fiber
- ▶ FC 3403 (16 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 (6.2 miles). Each FICON attachment can support up to 512 logical channels.

Requirement: FC 3401 (Enable 8 Gb or 16 Gb FICON dual-port) enables the second port on each installed 8 Gb or 16 Gb FICON adapter.

Extra cables, fabric components, and cabling solutions

Conversion cables from SC Duplex to LC Duplex are available as features on IBM Z 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.

For more information about Fibre Channel cable planning, see *IBM Virtualization Engine TS7700 Introduction and Planning Guide*, GA32-0568.

If Grid Enablement (FC 4015) is ordered, Ethernet cables are required for the copper or optical adapters to attach to the communication grid.

7.2 TS7700 component upgrades

Several field-installable upgrades give a 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 onsite TS7700:

- ▶ Secure data transfer FC 5281

Secure data transfer offers next-level cybersecurity with SP 800-131A-compliant encryption for Secure File Transfer between TS7700 clusters within a grid and logical volume data encryption, including remote reads, writes, and replication. Secure data transfer supports TLS 1.2 by using AES128 or AES256 bit encryption and uses the Power9 (VED) in-core cryptographic acceleration. Encryption settings can be enabled and disabled through the management interface concurrently.

No performance impact is expected when SDT is used.

► Capacity on Demand

This component allows clients to enable disk capacity in 20 TB and 100 TB increments concurrently. Consider the following points:

- CSB/CFC only: Every cache size from 20 TB to 3.9 PB can be ordered. The usable capacity can then be enlarged in steps of 20 TB.
- At least 1 TB of licensed capacity must be enabled on the next set of drawer pairs (through FC 5262 and 5263) before the next drawer pair can be installed. Adding physical drawer pairs without at least 1 TB of licensed capacity that is enabled is not supported.
- Feature numbers 5262 and 5263 should be ordered on 3957-VED
- At least one feature number 5262 for 20 TB cache enablement increment minimum must be plant-installed.

► Enable 1 TB Active Premigration Queue capacity (FC 5274)

Each instance of FC 5274 (Enable 1 TB Active Premigration Queue Capacity) enables up to 1 TB of data to be pending migration to physical tape or cloud storage tier. Up to 10 instances can be installed.

► Enable 5 TB Active Premigration Queue Capacity (FC 5279)

Each instance of FC 5279 (Enable 5 TB Active Premigration Queue Capacity) enables up to 5 TB of data to be pending migration to physical tape or cloud storage tier. Up to 10 instances can be installed after 10 instances of FC 5274 are installed for an active premigration queue size to 60 TB and higher.

► Enable 16 Gb FICON second port (FC 3401)

► Incremental data throughput (FC 5268)

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. The maximum number of features 5268 is 39. This installation is performed through the TS7700 MI by entering the license key that is obtained with the purchase of FC 5268 (100 MBps increment).

► 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.

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 a grid, the joining cluster must meet or exceed the currently supported number of logical volumes of the grid.

When merging one or more clusters into a 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 the grid connection adapter in the following TS7700 Server configurations:

- On a 3957-VED when FC 1034 1 Gbps quad-port copper connection is present
- On a 3957-VED when FC 1038 or FC 1041 10 Gbps dual-port optical LW connection is present

Note: A nonconcurrent installation is required to change grid adapters to a different configuration.

► Disk encryption

On a TS7770, you must encrypt the cache at installation time except for those countries that do not allow encryption and it is ordered from manufacturing with no encryption. Use FC 5272 for local encryption or FC 5276 for external encryption in which the user must secure IBM Security Key Lifecycle Manager V2 software entitlements for all cache hard disk drives that are installed in the TS7770 system. Also, Thales CipherTrust Manager 2.0 and later is supported on new systems only, no conversion allowed.

► TS7700 Storage Expansion frame

You can add up to two cache expansion frames to a fully configured TS7770 by using FC 9336 (Expansion frame attachment) and applying FC 7336 (TS7700 Encryption-capable expansion frame) to a 3952 F07 Tape Frame.

7.2.2 TS7700 nonconcurrent 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, nonconcurrent 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 following component upgrades must be made nonconcurrently to a TS7700:

- TS7770 optional Feature Code 5283 enables TS7700 Object Store for DS8000. For more information, see *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583-02.
- TS7770 optional features for tape attach

For TS7770 server VED hardware, use the following features to upgrade and enable the server to attach to a tape library:

- FC 5273 (TS7770 Tape Attach enablement)
- FC 5274 (Enable 1 TB Pending Tape Capacity)
- FC 5279 (5 TB Active Pre-Migration Queue)

Before installing the features, ensure that all hardware and microcode prerequisites are met. In addition, plan the future usage of the TS7700T and prepare the necessary information ahead. After the upgrade is installed, update the following information on the MI:

- Number of tape partitions: Depending on the usage decision, you might want to introduce multiple tape partitions for different customers, LPARs (such as prod, test, and development), or workload types.
- Size of the tape partitions: CP1 is the only tape partition that is assigned during the upgrade. The origin tape partition size is 3 TB. Consider that currently all data is still in CP0. If you do not redirect the workload to the CP1, the origin value is acceptable. However, 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 the workload to the Tape partitions, consider whether you want to introduce 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. You might also want to introduce a Copy export pool. If so, update the appropriate Management Class.
- Set an 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). It is also strongly advised to update the ALERT settings to ensure that alerts are 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 that is 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 more information about the **LI REQ PARTRFSH** command, see [IBM TS7700 Series/OS Host Command Line Request User's Guide](#).

► 16 Gbps FICON adapters

TS7770 model VED supports 16 Gbps FICON adapters. You can install up to four 16 Gbps FICON adapters and enable the second port with FC 3401 on each adapter for a total of eight host ports. An intermix of adapters cannot be used. Use FC 3402 (16 Gbps FICON Short Wavelength Attachment) or FC 3403 (16 Gbps FICON Long Wavelength Attachment) for this installation.

► Ethernet adapters for grid communication:

- LW fiber Ethernet

You can add a longwave fiber Ethernet adapter for grid communication between TS7700s.

For a 3957-VED use FC 1038 1041, a 10 Gbps dual-port optical LW connection is needed to achieve this upgrade.

Consideration: These 10 Gbps adapters cannot negotiate down to run at 1 Gbps. They must be connected to a 10 Gbps capable network connection.

– Copper Ethernet

You can add a 1 Gbps copper Ethernet adapter for grid communication between TS7700s. On a 3957-VED, use FC 1039 1 Gbps grid quad-port copper connection to achieve this upgrade.

Clarification: Intermixing different types of Ethernet adapters within one cluster is not supported. On a TS7700, you can have two 1 Gbps copper Ethernet adapters or two 10 Gbps LW fiber Ethernet adapters. Up to four ethernet ports can be enabled.

– TS7700 Server dual copper/optical Ethernet Adapter Card Conversion

You can convert a quad-port copper grid Ethernet adapter in a TS7700 Server for a dual-port adapter of the opposite type by ordering FC 1038 (dual-port LW optical), or vice versa. When you upgrade four 1 Gbps grid links to two 10 Gbps grid links, consult with your IBM SSR.

Each grid link can communicate to only 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 are no longer used and do report link degradation hardware messages. Consider disabling these unused grid links until all the clusters upgrade to matched grid links.

▶ TS7700 Server Fibre Channel host bus adapter installation

You can install two Fibre Channel interface cards in the TS7700 Server (3957-VED) to connect the TS7700 Server to the disk arrays in the TS7700 Storage Expansion Frame. Use FC 5243, Quad-port 16 Gb Fibre Channel HBA to achieve this installation.

Note: The adapter installation (FC 5243 Fibre Channel HBA) is nonconcurrent.

▶ Additional Virtual Devices (FC 5275)

Each instance of FC 5275 enables up to 16 extra virtual devices (virtual tape drives). A maximum of 15 features are supported for a total of 496 virtual devices. If FC 5275 is installed on a cluster in a mixed GRID microcode level, the extra devices are unavailable 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 must be taken offline and back online. If this process was done previously while the cluster member had FC 5275 installed, the cluster does *not* need to be taken back offline again.

To gain access to the new devices, the IODF must be modified. More control units with new libport-IDs must be defined. This process can be done before or after the installation of the feature codes. For more information, see Appendix G, “IBM TS7700 parameter examples” on page 959.

After the installation and activation of the FCs and the activation of the new IODF, you must restart the OAM address space on all attached LPARs where the extra drives are

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 include the new Library Port IDs into the Access port groups by using the TS7700 MI. If you do not include these Library Port IDs 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 TS7770 cache upgrade options

This section describes the TVC upgrade options that are available for the TS7770. Encryption is configured during the manufacturing process and is not field implemented.

The TS7770 Base frame minimum cache configuration is one 3956-CSB Cache Controller that includes 12 10 TB 7.2k rpm hard disk drives and one 3956-XSB Cache Drawer that includes 12 10 TB hard disk drives up to a maximum of nine 3956-XSB Cache Drawers. The usable cache capacity is 157 TB per pair of drawers for a total usable capacity of 789 TB in a fully configured base Frame.

The TS7770 Storage Expansion Frame consists at maximum configuration two 3956-CSB Cache Controllers each with nine 3956-XSB Cache Drawers that are attached to each controller for a total of 2.37 PB inclusive of the base frame.

A second fully populated expansion frame expands the usable capacity to 3.2 PB. The cache expansion is concurrent, except for the addition of the two 16 Gb Fibre Channel adapters that must be installed to support the first storage expansion frame.

For SSD-based cache models, the minimum cache configuration is one 3956-CFC cache controller that includes twenty-four 3.84 TB SSDs, which is 60 TB of usable capacity.

In addition to the base CFC cache controller drawer, a maximum of 9 XFC cache expansion drawers are supported, which is 640 TB of usable capacity.

7.2.4 Upgrading drive models in a TS7700T

This section describes upgrading back-end tape drives in a TS7700T cluster with data. You might want to upgrade the back-end tape drives to a higher model to have more capacity from the media because the drives are not encryption capable, or for any other reason.

TS7700T supports the following tape drives:

- ▶ 3592-J1A
- ▶ TS1120 (3592-E05)
- ▶ TS1130 (3592-E06/EU6)
- ▶ TS1140 (3592-E07)
- ▶ TS1150 (3592-E08)
- ▶ TS1160 /(3592-60G)

Note: Consider the following points:

- ▶ 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.
- ▶ As of this writing, TS1160 tape drives are supported only for TS7700 code level 8.52.101.12 (release R5.2.1 pga1). TS7700 release R5.2.2 (code level 8.52.200.109 and later) do not currently support TS1160 tape drives.

Hardware configuration and limitations

For more information about tape drives and supported media, see 4.1.2, “TS7700 specific limitations” on page 155.

Note: Throughout this section, the term *TS7700* refers to the TS7760/TS7770 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 media restrictions apply when a library attached to a 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 are 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 TS1160 Tape Drives are installed, media reclamation reduces (instead of increases) the number of available scratch volumes. Reclamation of this 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 JK physical volumes are 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 JK physical volumes to occur. If the system contains TS1150 or TS1160 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 MI after their active data is reclaimed onto newer media.

Note: JK and JC media must not be inserted as scratch into the library 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, the IBM Security Guardium Key Lifecycle Manager must be used.
- ▶ 3592 EU6 Tape Drives cannot be converted to TS1140 Tape Drives.

Restrictions for use with TS1150 and TS1160 Tape Drives (Homogeneous Configuration)

TS1160 Tape Drives are supported in new TS7700 orders from manufacturing and with existing TS7700s attached to a library. The following media restrictions also apply when a library that is attached to a TS7700T contains TS1160 Tape Drives:

- ▶ JA, JJ, and JB media are not supported.
- ▶ IBM Security Guardium Key Lifecycle Manager must be used if using tape encryption for this drive type.
- ▶ TS1140 Tape Drives cannot be converted to TS1150 and TS1160 Tape Drives.

Restrictions for use with TS1150 and TS1160 Tape Drives (Heterogeneous Configuration)

TS1160 Tape Drives are supported in new TS7700 orders from manufacturing, and with TS7700s attached to a library. The following media restrictions also apply when a library attached to a TS7700 Tape Attach contains TS1150/TS1160 Tape Drives:

- ▶ TS1150 and TS1160 Tape Drives can be intermixed with one other TS11xx drive type in a library that is attached to a 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 are 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 and TS1160 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 JA physical volumes are 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, TS1150, or TS1160 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 MI 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 TS11xx Tape Drive. If encryption is used, the IBM Security Guardium Key Lifecycle Manager must be used.
- ▶ TS1140 Tape Drives cannot be converted to TS1150 or TS1160 Tape Drives.

Considerations for upgrading tape drives

This section describes considerations when you upgrade your back-end tape drives.

Upgrading to a homogeneous TS1150 or TS1160 tape drive configuration

Homogeneous tape drive configurations include the following limitations:

- ▶ A maximum of 16 tape drives are supported.
- ▶ A minimum of 4 tape drives are required.
- ▶ JJ, JA, and JB (if written in J1A/E05/06 format) cannot be read or written by TS1150 or TS1160 drives. These media are called *sunset media* and to read them, the target TS7700 still requires availability of previous-generation drives installed in the configuration.
- ▶ Existing JC and JK cartridges (originally written in E07 format) can still be accessed but become READ-ONLY. They are also marked as *sunset media* when upgrading to TS1160.
- ▶ TS7700 does not go online if a logical volume is in sunset media and no read-compatible drives are available.

It is possible to replace all tape drives with TS1150/TS1160 tape drives if all active data on tape is on JK or JC cartridges, which implies that the tape drives that were installed are TS1140s. If the TS7700T had 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 TS7700T to complete the following steps:

1. Change the Storage Class definitions to point to the resident-only partition.
2. Run a **PARTRFSH** command to move the logical volumes from the tape-attached partition to the resident-only partition.
3. Recall all of the migrated data from the older tapes into the cache-resident partition and allow the MES to have TS1150 or TS1160 tape drives installed.
4. After the MES, you can again change the Storage Class by running a **PARTRFSH** command, and push the logical volumes back to backend tapes.

After the first TS1150/TS1160 is installed and configured, the TS7700 detects the new drive generation during the online process and acknowledges the cartridge types. Ensure that all of the logical volumes in sunset media, such as JJ, JA, or JB, are migrated to a supported media before TS1150/TS1160 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 the 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 that correspond to older media that were used by older drive types are marked as sunset read-only after TS7700 comes online. This media cannot be used as scratch tapes.

When a TS7700 has all TS1160 tape drives, the following 3592 media types can be used as scratch media for read/write.

- ▶ Cartridges with E08 format for pools that are marked as “Copy Export” with the following media types:
 - JK - IBM 3592 Advanced Type C Economy Cartridge
 - JC - IBM 3592 Advanced Data Type C Cartridge
 - JL - IBM 3592 Advanced Type D Economy Cartridge
 - JD - IBM 3592 Advanced Data Type D Cartridge
- ▶ Cartridges with 60F format (native) with the following media types:
 - JM - IBM 3592 Advanced Economy Type E Cartridge
 - JE - IBM 3592 Advanced Data Type E Cartridge

Empty sunset media, such as JC or JK, 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 3592-60F 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

A heterogeneous tape drive configuration includes the following limitations:

- ▶ A maximum of 16 tape drives are supported.
- ▶ Up to 14 TS1160 tape drives are allowed.
- ▶ A minimum of four TS1160 tape drives are 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 that are stored in JA, JB, and JJ media.
- ▶ JA, JB, and JJ cannot be read or written by TS1160. They cannot be used as scratch media.
- ▶ At least 15 empty scratch media are necessary to support data migration from sunset media to newer media.

If running an intermix of TS1160 and TS1150 tape drives, consider the following points:

- ▶ The JD media are written with the new TS1160 tape drives with 15 TB capacity and are also only read by the TS1160 tape drives.
- ▶ The existing JD media with active data, which were written with the TS1150 tape drives, remain in 10 TB format and are only read by the TS1150 tape drives.
- ▶ When one of the existing JD media, which was written with the TS1150 tape drives, goes to “scratch”, it is written by the new TS1160 tape drives with the 15 TB capacity.

Support for limited heterogeneous tape drives seamlessly helps customers move from older media and drives to supported media and TS1160 tape drives. This option enables you to add TS1160 tape drives to the TS7700 so that all new workloads can be directed to the TS1160 tape drives. It also leaves at least one of the legacy tape drive generations (3592-J1A, TS1120, TS1130, or TS1140) to manage legacy media.

Note: Only one generation of older tape drives can be included with the newly installed TS1160.

During the TS7700 online processing, media types JA, JB, and JJ 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 that at least 15 scratch media are inserted to run reclamation of sunset media. After the reclaim process moves the data to TS1160 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, read-only drives are used for writing data in the following exceptions:

- ▶ 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 TS7700T is part of a grid, vary online the logical drives in the other clusters to provide tape resources to the host. If this cluster is a stand-alone cluster, plan for a 3–4 hour outage.
2. Vary offline 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 MI, the drives must be unassigned from the TS7700T Logical Library.
3. Drives can now 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. If a TS7760T is used, 16 Gbps switches are necessary.
4. New drives are assigned back to the 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 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 step completes the drive upgrade change. Now, the 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 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 11-2 on page 562) by using a specific command in the TS7700T by the IBM SSR. The media format is handled as described in the previous scenario.

Migrating TS7700T data from sunset media type after upgrading heterogeneous drive configuration

Consideration: Restrictions apply to some configurations. For more information about the valid combinations of media types and drive models within the TS7700T, see “Tape drives and media support (TS7700T)” on page 153.

This procedure can be helpful when upgrading your tape drives to the TS1160 3592-60F tape drives, or when replacing the media cartridges with a newer type to increase the storage capacity of your library.

The 60F drives cannot read JA, JB, or JJ tapes; therefore, you must have JC, JK, JL, JD, JL, or JM media for the TS7700T to support new logical volumes that are written to TS1160 drives. You must also have at least two drives of a sunset generation that can read logical volumes from tapes that were supported.

In this scenario, coming from a previous 3592 tape drive model to the 60F, all JA, JB, JJ, JC, or JK (written on E07 or older format) media are *sunset*, which means that after reclaiming the active logical volumes still contained in it, the 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, JL, JD, JE, and JM. Your data, formerly in a JA, JB, or JJ media, is forcibly migrated into the supported media.

Two alternatives are available to introduce the new media: 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 TS7700T for the new media, as shown in Figure 11-4 on page 564.
2. Create a Cartridge Assignment Policy (CAP) for the new range and assign it to the correct TS7700T logical partition (LPAR), as described in “Defining Cartridge Assignment policies” on page 571.
3. Insert the new cartridges in the TS3500/TS4500 tape library, as described in “Inserting TS7700T physical volumes” on page 572.
4. Assign an existing pool or pools of physical volumes in the TS7700T to the new media type, as described in “Defining physical volume pools in the TS7700T” on page 579.
5. Modify the Storage Group (SG) in the TS7700T constructs to point to the new cartridge pools, as described in “Defining TS7700 constructs” on page 593.
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 579.

These settings cause the TS7700T to start using the new media for stacking newly created logical volumes. Existing physical volumes are reclaimed into the new media and become empty.

You might prefer to keep your pool definitions unchanged throughout the media change process. In this case, you complete steps 1–3. No other change is necessary if you are migrating from JJ/JA/JB cartridges to JD or JM cartridges.

If you are migrating from JC/JK/JD to JD/JM cartridges, 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 579. 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 (Re-CaLI SUNset) to expedite the replacement of the sunset media with newer media. In this case, ensure that the common scratch pool includes the new media type that is 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 uses the following format:

```
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 that are in use. No changes 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 is less than the threshold, the sunset media is reclaimed and active data is migrated to the newer media.

Figure 7-10 shows how the Sunset Media Reclaim Threshold Percentage is set in the Physical Volume Pool properties.

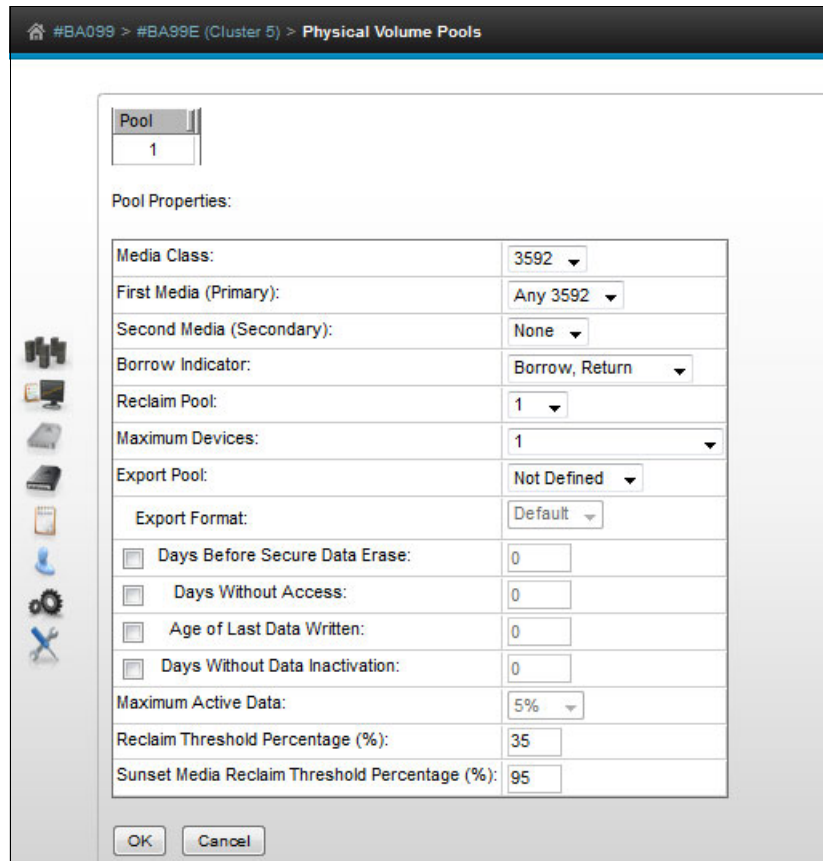


Figure 7-10 Sunset Media Reclaim Threshold Percentage

If you want to limit the service resource of data migration, a new Host Console Request command 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" (the default value). The library request command is shown in the following example:

```
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 fell 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 nonsunset drives and the new keywords are for sunset drives.

Clarification: The low sunset drive alert includes the following parameters:

► **SETTING2,ALERT,PDRVSLOW**

Shows 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**

Shows 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*.

The following library request command is available:

```
LI REQ, <lib_name>, SETTING2, ALERT, [PDRVSLOW|PDRVSCRT], <number_of_drives>
```

7.2.5 Frame replacement of old hardware with new hardware

This scenario describes the frame replacement of a TS7760T (VEC) with a TS7770T (VED).

You might want to replace old hardware VEC with a VED to have significant performance improvement and increased disk cache capacity over the old ones. The source TS7700 must be at a minimum 8.33.x.x. The target TS77000 must be on the recommended code level.

To replace the old hardware, all of the logical volumes must be pre-migrated to physical volumes, transferred to the new frame, and recalled from the physical volumes.

If the frame you want to replace is a TS7760T (VEC):

- All the *private* logical volumes in CP0 must be moved to CPx before frame replacement.
- All the *resident* logical volumes with delayed premigration setting must be premigrated before frame replacement.

All of the physical volumes that you want to migrate to the new frame must be readable by the new tape drives. For more information, see 4.1.2, “TS7700 specific limitations” on page 155.

7.3 TS7700 upgrade to Release 5.4

Release 5.4 and later can be installed on a previous VED-based TS7700 and requires 128 GB of memory. New VED's come with 128 GB by default.

7.3.1 Planning for the upgrade

A multi-cluster GRID configuration can enable almost all changes or upgrades to be concurrent from a client's standpoint, putting one individual cluster into service at a time. The Release 5.4 Licensed Internal Code upgrade is a disruptive activity in a stand-alone cluster. A Licensed Internal Code (LIC) update is done by an IBM service support representative (IBM SSR) or by a remote code load team member. 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. 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 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 Host software support before you perform the LIC upgrade and have the IBM service representative apply any required maintenance that is associated with the LIC level you are installing.

Important: A fix category (fixcat) identifies groups of software fixes (PTFs). For example, a fix category might identify software fixes that are required to support a particular hardware device, support compatibility with a new software release, or provide a software function.

Contact your z system software admin and tell them which fixcat categories to be looking out for such as IBM.Device.Tape.TS7700-3948 and IBM.Device.Tape.TS7700-3957 before performing the LIC upgrade. For more information about fixcat and how your z system admin can find the PTFs that you need, see [IBM Fix Category Values and Descriptions](#).

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 up to eight 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.

For more information about distances that are supported by different grid adapters and cabling options, see "TS7700 grid and cloud LAN/WAN requirements" on page 157.

For separated sites or three or more cluster grids, be sure that the network is prepared at the time that the migration starts. The TS7770 provides two or four independent 1 Gbps copper (RJ-45) or 10 Gbps optical LW fiber Ethernet links (single or dual-port) for grid network connectivity. For more information, see 4.1.3, "TCP/IP configuration considerations" on page 157.

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. Many combinations of code levels and configurations exist 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.

7.4.2 Considerations when adding a cluster to the configuration

Figure 7-11 shows an example of joining or merging a new cluster to a stand-alone configuration.

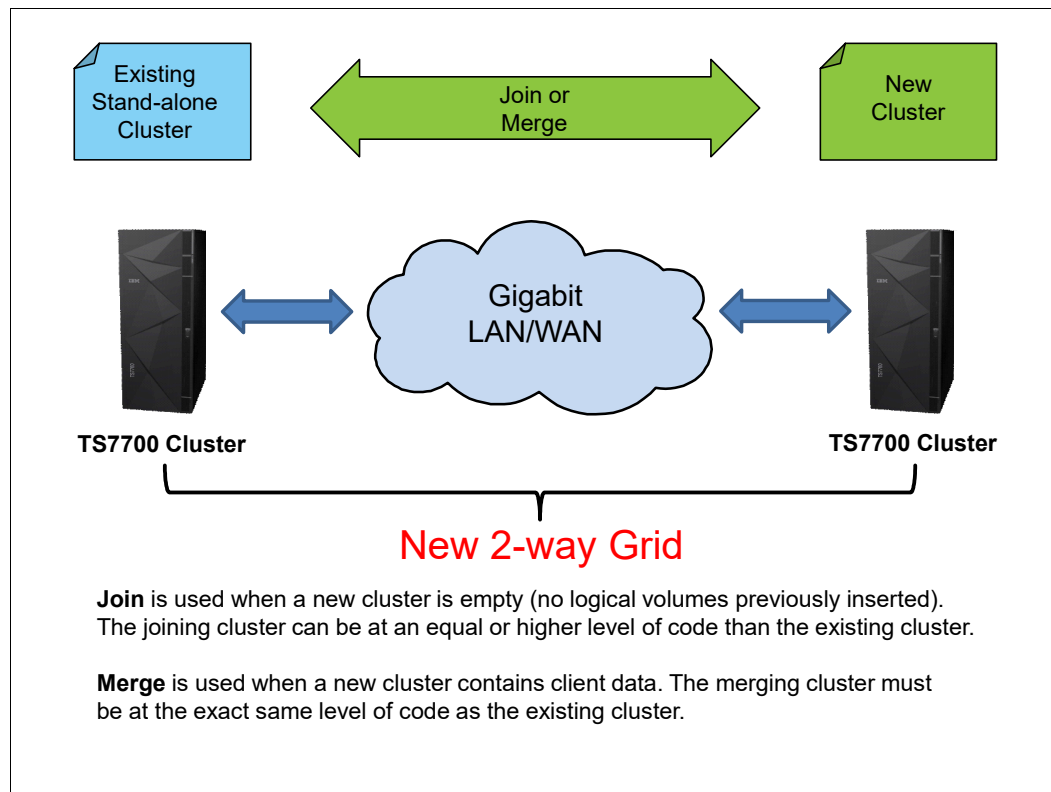


Figure 7-11 Example of a join or merge of a new cluster

Figure 7-12 shows an example of merging or joining a new cluster to a grid configuration (in this example, a 5-cluster grid).

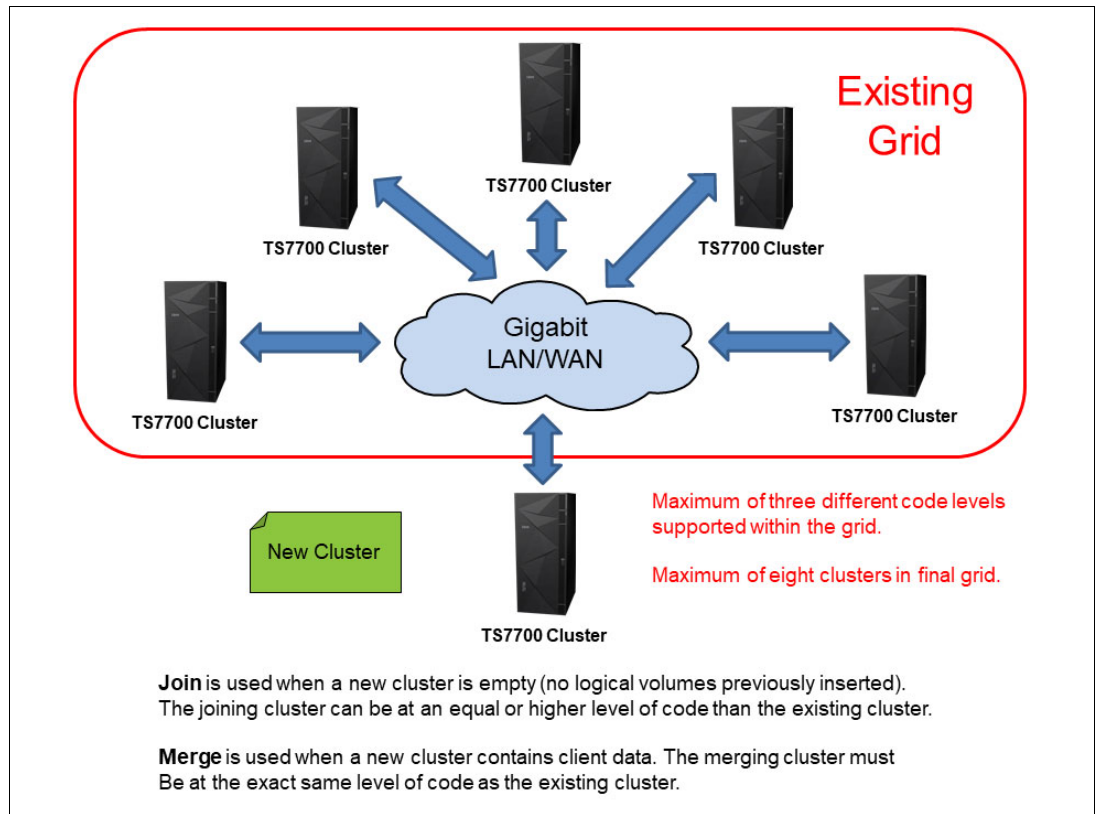


Figure 7-12 Join or merge a new cluster to a multi-cluster grid

Preparation

When performing a join, the 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 a grid, the process is started 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 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 182 and Chapter 6, “Implementing IBM TS7700” on page 239. Consider the following points:

- ▶ 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 TS7700T Cluster join is complete:

- ▶ Define stacked volume ranges
- ▶ Define a reclaim threshold percentage

Logical volume considerations

Ensure that the joining cluster has at least the same number of FC 5270 components that are installed as in the existing cluster or grid.

Licensed Internal Code supported levels and feature code for join

Release 5.4 supports the ability to have VED clusters (new from manufacturing or empty through a manufacturing clean-up process) join a grid with a restricted mixture of Release 5.1 minimum. Three total code level differences can exist across both targets and the joining system during the MES where R5.1 can be the lowest of the three levels.

When you join one cluster to a cluster in a 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 addressed the following restrictions for the join process:

- ▶ The joining cluster must be empty (contain no data, logical volumes, and constructs).
- ▶ If the existing cluster to be joined to is a member of a grid, it must have the highest code level present among grid members. 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.

Other considerations

The existing clusters must be *write-protect disabled* and *FlashCopy disabled*.

Join steps

Complete the following steps to join the cluster:

1. Arrange for the following 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. 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.

3. Change the 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.

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.

Note: For DS8000 Object Store support, inbound objects are not affected by the state of any back-end tape library or object store. For more information about DS8000 Object Store, see *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-55833.

7. Review your family definitions and decide whether the cluster must 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 **PMTHLVL** setting and set it to the amount of installed FC 5274; for example, 6 x FC 5274 = 6000.
For more information about these settings, see Chapter 14, "Performance considerations" on page 747.
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 14, "Performance considerations" on page 747.
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 an MC for BVIR purposes 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 process can be done in different ways. In general, the new rules of the management policies must be retrieved for every logical volume that is 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:

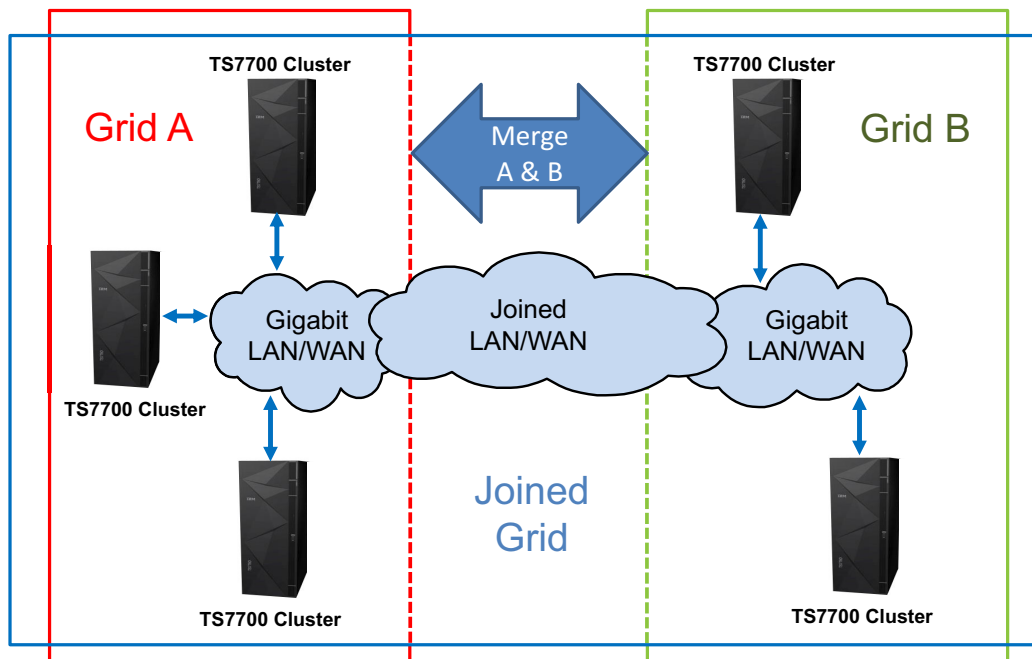
- ▶ Mount or demount to an LVOL
- ▶ **COPYRFSH**, a **LI REQ** command, based on a single logical volume

To produce a new copy, the data must be in the cache. If your source cluster is a TS7700T, consider sorting the logical volumes in a copy order that maps to the physical volume layout. This sorting improves 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 13, “Monitoring” on page 667. The tools are available at [this website](#).

7.4.3 Considerations for merging an existing cluster or grid into a grid

Figure 7-13 shows a grid merge scenario that involves a two-cluster grid and a three-cluster grid being merged into a five-cluster grid.



- All clusters must be at a common code level
- Grid merge requires all clusters within one grid to be inaccessible
- Other grid must have all clusters online
- Maximum of eight clusters in final grid
- Maximum of 4 million logical volumes. Logical volumes must be unique within grid.

Figure 7-13 Grid merge example

Preparation

You can add a TS7700 Cluster to another 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 a grid, or merge two grids.

You can merge two 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, the data is not 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 started 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 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, you must plan which LPAR can access which clusters and device ranges in the grid in advance. These changes must be prepared in each LPAR (HCD, SMS, and TCDB). Consider the following points:

- ▶ 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.

Conditions exist during a **MERGE** of two grids where the newly merged clusters are not correctly recognized in the new grid and do not come online to the host. This issue is typically only seen whether the newly merged grid reuses one of the 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, the situation can be resolved by issuing the console command **DS QL,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 in the existing cluster, are updated with the content of the existing cluster.
- ▶ Constructs in a cluster, which are in the existing cluster but not in the merging cluster, are copied.
- ▶ Constructs in a cluster, which exist in the merging cluster but not in the existing cluster, are kept, but not copied to the existing cluster or grid.

Figure 7-14 shows the MC definition of the merging cluster and the two-cluster grid before the merge.

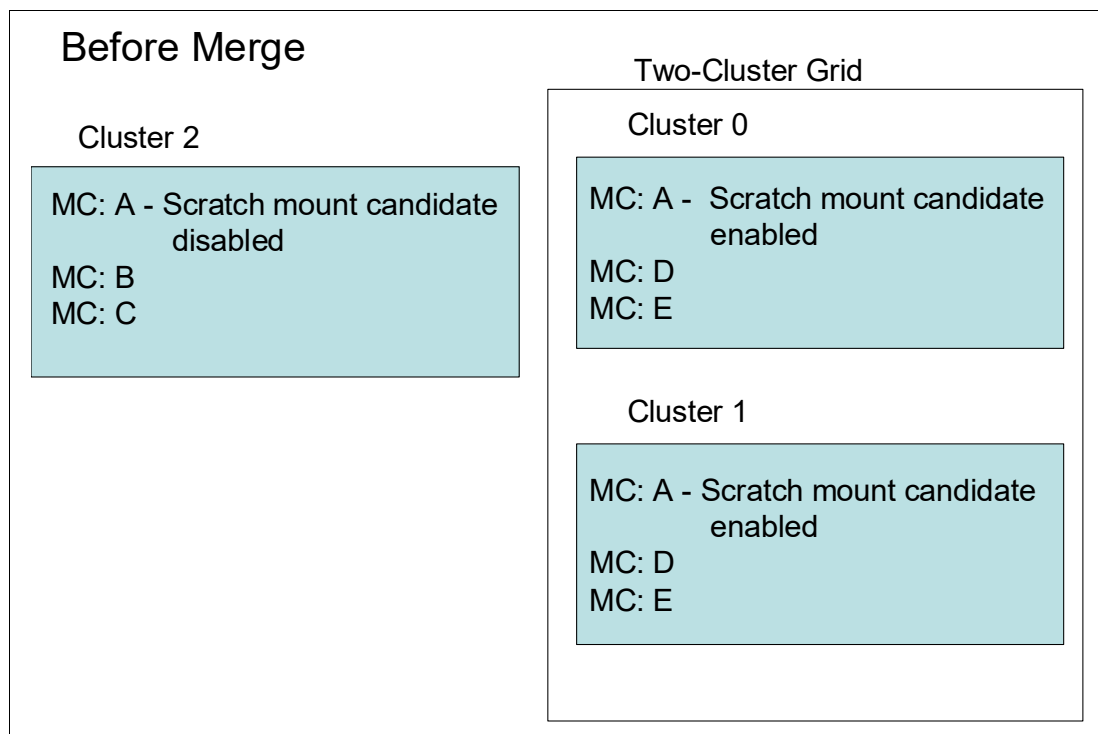


Figure 7-14 Management Class definition before merge

Figure 7-15 shows the MC definition of the merging cluster and the two-cluster grid after the merge.

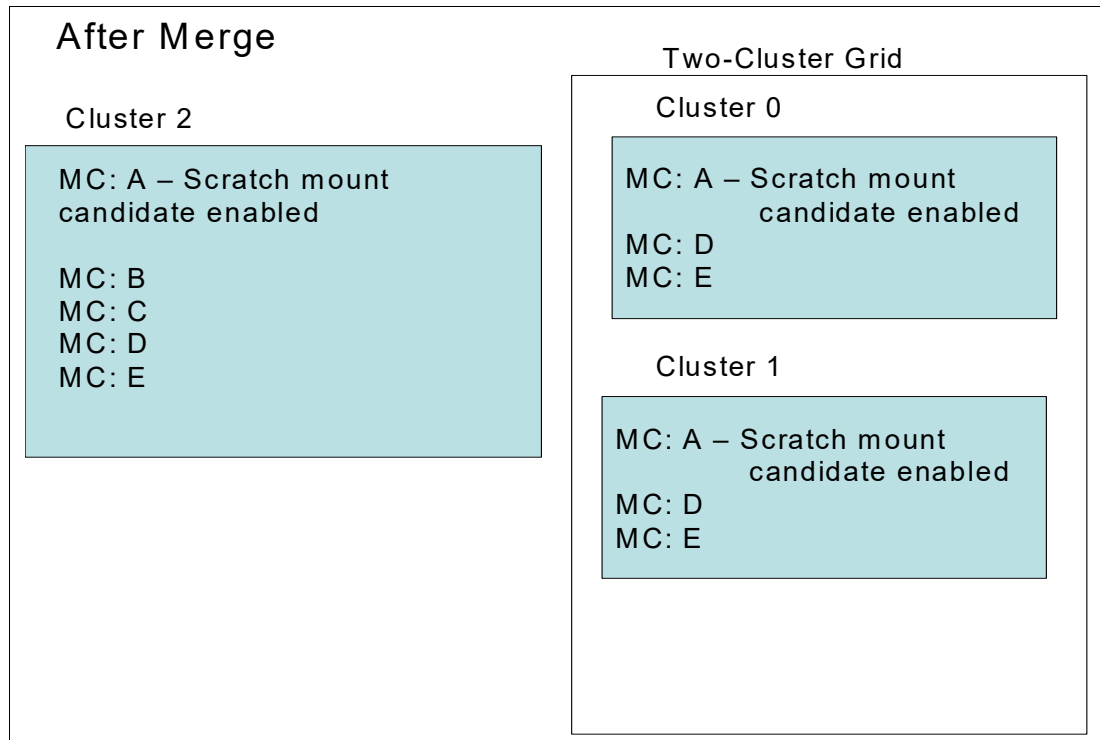


Figure 7-15 MC definition after merge

If categories and constructs are defined on the merging cluster, verify that the total number of each category and construct that exist in the grid does not exceed 256. If necessary, delete 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. 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. 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 are 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, more inserts are not allowed until the feature counts on those clusters are increased.

Minimum Licensed Internal Code level and feature code for merge

Before a cluster or a grid is merged 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.
- ▶ Only a grid-to-grid merge is supported and both grids must operate at the same Licensed Internal Code level.
- ▶ Merges are supported only when all clusters in the resulting grid are at the same code level.
- ▶ FC 4015, Grid enablement, must be installed on all TS7700 clusters that operate in a grid configuration.
- ▶ Existing clusters and merging clusters must contain enough features to accommodate the total resulting volume count post merge, or the merge fails.
- ▶ The merging cluster must contain FC 5271 if the cluster to be merged has it installed.
- ▶ If the merging cluster includes FC 1035 installed, the client's infrastructure must support 10 Gb.
- ▶ If categories and constructs are defined on the merging cluster, verify that the total number of each category and construct that exist in the grid following the merge does not exceed 256. If necessary, delete existing categories or constructs from the merging clusters before the grid join 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

Merge steps

Complete the following 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 182 and Chapter 6, "Implementing IBM TS7700" on page 239.
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. Now, define the newly 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 part of an existing grid, you might need to delete the services control blocks of that cluster's devices by 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 2.4.5, "Copy consistency points" on page 70.
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. For more information, see "Population of a new cluster (COPYRFSH)" on page 294.

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 from a grid configuration. It can be used for removing one cluster from a two-cluster to eight-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), which keeps 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). Additionally the client must determine how to handle object data on the cluster to be unjoined. This process must 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 that is 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.

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 by 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.

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 FC 4017.

7.5.2 High-level description of the process

The following high-level preparation activities occur before a cluster is removed from a domain:

- ▶ Determine whether any volumes exist that are available only on the cluster to be removed (for example, MCs that are defined to have only a copy on one cluster, or auto removal from a TS7700). Before the removal, you must create consistent copies on other clusters in the domain.

For more information about the BVIR Copy Audit function, see [IBM TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#).

- ▶ 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 12.1.7, “Ejecting virtual volumes” on page 643).
 - 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 volumes do not exist in the damaged category. You can use the Repair Logical Volumes menu under the MI window to repair them.
- ▶ Modify MCs so that the removed cluster is no longer the target for copies.
- ▶ If 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 no Management Class exists on any other cluster where NO COPY is the only remaining copy consistency policy.
- ▶ Adjust VEHSTATS runs to reflect the new configuration.

Note: Consult with your IBM SSR for the code version prerequisites for FC 4016.

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 perform this check 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 that are associated with the removed cluster.



Migration

This chapter explains aspects of migrating to a TS7700 environment from an IBM Virtual Tape Server (VTS) or other tape drive technologies. It also 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 simpler are highlighted.

This chapter includes the following topics:

- ▶ 8.1, “Migration to a TS7700” on page 304
- ▶ 8.2, “Migration between TS7700s” on page 305
- ▶ 8.3, “Moving data for host-based migration” on page 318
- ▶ 8.4, “Moving data out of the TS7700” on page 324
- ▶ 8.5, “Migrating DFSMSHsm-managed data” on page 327
- ▶ 8.6, “DFSMSRmm and other tape management systems” on page 336
- ▶ 8.7, “IBM Spectrum Protect” on page 338
- ▶ 8.8, “DFSMSDss” on page 341
- ▶ 8.9, “Object access method” on page 343
- ▶ 8.10, “Database backups” on page 344

8.1 Migration to a TS7700

This section covers various aspects of migrating from an existing tape technology to the TS7700. Host-based Migration is used to move customer data from older technology or other vendor tape solutions to the TS7700 by using host-based programs.

Hardware migration from tape technology to the TS7700 includes the following aspects:

- ▶ Software changes in storage management subsystem (SMS), hardware configuration definition (HCD), tape configuration database (TCDB), and tape management system (TMS)
- ▶ Connectivity for the new configuration
- ▶ Migration of the database from source TS7700 to the newer target TS7700

Information is provided about the TS7700 family replacement procedures that are available with the new hardware platform and the TS7700 R5.0 Licensed Internal Code (LIC) level. With the availability of the new generation hardware, an upgrade path is provided for TS7700 users to migrate to this new hardware.

This section covers upgrading tape drive models in a TS7700T to gain capacity from your existing media, or to provide encryption support. It also describes the hardware upgrade procedure and the cartridge migration aspects.

8.1.1 Host-based migration

Host-based migration is a migration method that uses host programs. It is used in the following cases:

- ▶ No compatibility exists between the physical tape media that are used by source and target solutions:
 - No physical tape media compatibility between IBM 3590 tape drives and IBM 3592 tape drives (TS7700T supports only IBM 3592 series of tape drives)
 - No internal data format compatibility between physical tape media that is written by native tape drives and ones that are written by TS7700
- ▶ Target library does not include any attached back-end tape drives and it cannot join a grid with the source tape library
- ▶ New tape library cannot join a grid with the source library

Although both migration source and target feature physical tapes, migration from VTS with 3590 tape drives or native tape drives to a TS7700T always requires host-based migration to copy the data into the TS7700. This requirement is because no data compatibility exists between the physical tape media that is used by migration source and target solutions.

Migration from VTS with 3590 tape drives or native tape drives to a TS7700D can be performed by using only host-based migration because the TS7700D does not have any attached back-end tape drives. Therefore, data must be copied into the TS7700D by using host programs.

A migration between two TS7700s might require a host-based solution if your environment cannot accept another cluster in the grid; for example, when no network infrastructure is available to create a grid links between source and target tape libraries.

For more information about the methods you can use for host-based migration, see 8.3, “Moving data for host-based migration” on page 318.

Host-based migration is also available for migration from a VTS with 3592 tape drives to a TS7700T; however, tape-based migration might be a better choice because TS7700T includes attached back-end tape drives and you do not have to use host cycles in the process. For more information, see 8.1.2, “Tape-based migration” on page 305.

8.1.2 Tape-based migration

Migration from a TS7700T to another by using physical volumes, also called *outboard VTS migration*, is possible depending on the target configuration. It provides an upgrade path for old VTS models to a newer 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 considerations that are 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 migration of data from an older TS7700T to a newer model

Work with your IBM service support representative (IBM SSR) for more information 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 TS7700T that is 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 TS7700T is attached to another tape library, the process is called *data migration with tape move*. When a source VTS is migrated to a TS7700T, or two VTSs are migrated to the same target TS7700T, the process is called *merge*.

8.2 Migration between TS7700s

The method that is described in 8.1.1, “Host-based migration” on page 304 also can be used for data migration between TS7700s. However, other preferred methods are available for this purpose that do not need to involve a host for data movement. These methods are explained in the following sections.

The user should select the most efficient approach, depending on the case:

- ▶ Host-based migration: Uses host cycles and operation involvement for data movement.
- ▶ Frame replacement migration for TS7700T:
 - Option to replace an older TS7700T model with an entirely new machine.
 - New (“target”) TS7700T takes the identity of the old one (“source”), so the replacement operation is transparent to other cluster members of the same grid.
 - It does not involve a host for data movement.
 - It depends on all data that is in the source machine to include completed copies to physical tape in advance to actual frame replacement.

- ▶ Join and Copy Refresh processing:
 - Used for data migration to empty TS7700 inside the same grid
 - Does not involve host for data movement
 - Does not use physical tape volumes
- ▶ Migration service offering:
 - Used for data migration to another TS7700 inside the same grid
 - Does not involve host for data movement
 - Does not use physical tape volumes
 - Multiple concurrent migrations are allowed inside the same grid
- ▶ Copy Export and Copy Export Recovery/Merge:
 - Used for data migration from one TS7700T to another TS7700T cluster
 - Target TS7700T can be a stand-alone or part of a different grid
 - Does not involve host for data movement
 - Uses physical tape volumes to export data to target TS7700T
- ▶ Grid to Grid Migration:
 - Used for data migration from a grid to another grid
 - Does not involve host for data movement
 - Does not use physical tape volumes
- ▶ Cloud-Based Migration:

The following specific use cases are pertinent to the use of cloud storage tier:

- Partition refresh for cloud migration

Virtual volumes that were originally stored in the resident cache partition (CP0) of the TS7700C can be moved to a cache partition that is intended to use the Cloud Storage Tier in the same cluster. Consider the following points:

 - Requires modifying attributes of target virtual volumes in corresponding constructs (storage group and storage class).
 - Mount/demount operations are not needed for new attributes to be applied if the library request command (**LI REQ PARTRFSH, <volser>, MMOUNT**) is used.
- Cloud ghost copy processing:

Data that was premigrated to cloud storage tier by one cluster can be made accessible to a newly joined TS7700C. Consider the following points:

 - Requires modifying copy mode attributes of target virtual volumes (management class), and ensuring corresponding storage groups and storage classes in the target machine are correctly configured to use the cloud storage tier.
 - Data that existed in the cloud is recognized by the new machine without copying data through grid links if library request command (**LI REQ COPYRFSH, <volser>, CLDGHOST**) is used. This command updates the library database so that the grid copies that are delayed or set as PGO skips the specific copy processing through the grid links after verifying that the content is accessible in the cloud.

- Cloud Migration and Grid Cloud Awareness:
 - Volumes that are pre-migrated or migrated to the cloud storage tier are automatically made available to a newly joined cluster if it is configured to access the cloud, as shown in Figure 8-1.




Figure 8-1 Cloud Attached Clusters with Grid Cloud Awareness

- TS7700C clusters configured as No-Copy (by way of management class) are now allowed to do recalls of premigrated/migrated content from the cloud when other copy sources in the grid are not available. For example: consider a scenario where copy modes for a virtual volume were set to (C0,C1)=(Deferred,No Copy). In this case, if C1 is a TS7700C configured to access the same cloud pool, it can access the volume immediately after the volume is pre-migrated or migrated to the cloud by C0.
- After adding a TS7700C cluster to the grid, you do not need to wait for all of the data to be replicated. Instead, set up the cloud and container information and you can read any tape from the cloud.

Note: For Grid Cloud Awareness, all clusters within the grid must be Cloud Attached TS7700C and be at R5.1 or later.

- Cloud Export and Recovery:
 - Similar to the basic Copy Export function, Cloud Export and Recovery provides the capability to export a backup into a cloud tier versus physical tapes; therefore, the backup can be used later to restore an empty cloud-attached cluster (TS7700C).

- In a tape migration project, Cloud Export and Recovery can be used as a preferable method as opposed to Copy Export. It has the benefit of not needing to manage physical ejects, shipment, and re-insertions. Also, no physical drives limitation exists when recalls or the prestage process for the recovery are done.
- The cloud object store is an extension of the TS7700C storage capabilities. Unlike a physical tape library, cloud provides the ability to access object data from any TS7700C cluster in a grid directly from a cloud storage container if that cluster has a physical network path and is configured with credentials to access the object.
- A mixed environment can be used that features Tape Attached (TS7700T), Disk Only (TS7700D), and Cloud Attached (TS7700C), but only the Cloud Attached clusters can access the cloud object store and perform read/write transactions against it. Corresponding volume tracking information is still shared with all clusters in the grid (an example is shown in Figure 8-2).




Figure 8-2 TS7700 Grid with Cloud and Tape Attached

- Cloud Export requires an Export List File Volume with information about supported options to be used for the operation. That information must be written in a specific format, and then the Cloud Export is triggered by using the **LIBRARY EXPORT <volser>** command.

This command starts the backup task and sends it to the cloud store. For more information about formatting in the Export List File Volume, and supported options for the operation, see [TS7700 Cloud Storage Tier Export, Recovery, and Testing Guide](#).

- After issuing the **LIBRARY EXPORT** command, the cluster monitors pending pre-migration tasks and ensures they all complete. After it has a consistent point with all volumes in the cloud, it starts the database backup process.

- For the recovery, after installing the new empty TS7700C in your target location, configure TS7700C cloud pool and cloud vault credentials. You need the container ID, backup ID, and the serial number of the source tape library that started the export to the cloud. This information must be entered on the TS7700C MI.
- TS7700C locates the backup and imports the database backup from cloud store and perform the restore.

Note: For Cloud Export and Recovery all clusters must be at R5.1 or later, older releases can coexist in the same grid. For more information, see [TS7700 Cloud Storage Tier Export Recovery and Testing Guide](#).

If data is moved inside the same grid (after a join of a cluster or a merge), and cloud-based migration is not an option, COPYRFSH is the preferred method. For more information, see “Population of a new cluster (COPYRFSH)” on page 294.

8.2.1 Field frame replacement migration for TS7700T

This procedure must be done by an IBM SSR. It is an option to upgrade from an older TS7700T generation to a newer one, which provides benefits to the user in terms of performance, disk cache capacity, and new features.

This operation implies the substitution of a “source” TS7700T by using a newer target generation TS7700T, which takes the role of the machine being replaced (from a grid perspective). If host-related feature codes that are installed in the target machine are equivalent to the codes that are in the original source machine, no configuration changes from the connected host side are required.

A physical tape library that is attached to the original source machine, and its entire inventory of physical volumes, is kept for the new target TS7700T.

The following high-level steps are performed during this process by the IBM SSR:

1. Ensure pre-MES compatibility.

This step is a precheck, which can be run concurrently to customer operations, and should be done before initiating the outage window, which will be required for the subsequent non-concurrent steps of the frame replacement. The objective is to look for potential mismatches or incompatibilities between the source and target machines so that such issues can be resolved in advance. The following aspects are evaluated:

- Physical media compatibility.
- Installed feature codes.
- Data that might exist in the resident cache partition 0. The TS7700 field frame replacement process can only be performed if it is verified that all data in the source machine was copied to physical tape in advance and is located in a tape partition.
- If the premigration delay function is used, it should therefore possibly be deactivated some time before the field frame replacement is performed by the IBM SSR. Also, manually trigger the delayed premigration data in the system for premigration to tape. Consult with an appropriate VTS consultant from IBM or with a Business Partner who is familiar with this particular function.

2. Perform a backup of the database, internal file systems, and software configuration information that corresponds to the source TS7700T.

3. Remove the source machine.

4. Activate the target machine and restore information from backup.

Note: The cache on the new frame is empty and if host jobs must access the data, it requires a recall from back-end tape. Customers might need to run the *prestige* tool to stage frequently accessed volumes back in cache to help with system performance after the replacement.

The *prestige* tool is part of the IBM Tape Tools package for z/OS. This package can be downloaded free of charge and installed on your z/OS host. You can download this tools package and other stand-alone tape tools from the [IBM tool website](#).

8.2.2 Join and Copy Refresh processing

If you want to move to a new data center or conduct a technical refresh, use Join and Copy Refresh processing to migrate the data to a new cluster without the use of host-based migration. This method can be used only when the data migration is done inside a grid. Grid migration is the fastest proven migration method, and includes the following steps:

1. Join a new cluster to a grid.

For more information about how to join a new cluster to a grid, see 7.4, “Adding clusters to a grid” on page 289.

2. Change the Copy Mode in Management Class (MC) to allow copies to the new cluster.

MC can be configured from MI. For more information, see 10.2, “Constructs icon” on page 476.

For more information about Copy Mode, see 4.2.2, “Defining grid copy mode control” on page 173.

3. Generate a list of logical volumes that need to be copied to the new cluster.

A user can list the status of logical volumes by following the Bulk Volume Information Retrieval (BVIR) command. A user can process its output and create a list of logical volumes that must be copied to the new cluster by using **VOLUME STATUS**.

The following IBM Tape Tools are available. Combining the tools can list logical volumes that must be copied to a new cluster and generate Copy Refresh commands list for them easily:

- BVIRMES
- VESYNC
- BVIRVTS
- COPYRFSH

4. Run Copy Refresh to each logical volume in the list that was created in Step 3 to produce a new copy of the data in the new cluster. Copy Refresh can be run by using the **LI REQ COPYRFSH** command.

5. If you are removing a cluster, make sure that no volumes use that specific cluster as the only valid source for a copy. At least one other consistent copy must exist in the grid before the removal.

The **COPY AUDIT INCLUDE <Old Cluster Lib ID>** command can be used to determine copy consistency gaps (do not use COPYMODE if you are certain copies must exist that are independent of the configured copy modes).

6. Remove the old cluster from the grid. For more information about how to remove the old cluster from the grid, see 7.5, “Removing clusters from a grid” on page 299.

While the Copy Refresh command is submitted from a host, the data is copied internally through the gridlinks. No Host I/O exists through the FICON adapters, and all data in the TCDB and tape management remains unchanged. This method can be used only if the data migration is inside a grid. The BVIR AUDIT parameter provides a simple method to ensure that all data is copied.

For more information about the available tools, see 13.5, “IBM Tape Tools” on page 703.

For more information about the commands that are described in this section, see [IBM TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#).

For more information about the **LI REQ COPYRFSH** command, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

Generating Copy Refresh commands list

IBM Tools BVIRMES, VESYNC, and BVIRVTS can help the user to build a list of logical volumes that should be copied over to the new cluster. This process generates a **COPYRFSH** command list that is used to automate the process.

Complete the following steps:

1. Make sure that Management Class values are in place.
2. Run BVIRMES to the new cluster that was joined to a grid to distribute output of BVIR VOLUME STATUS to all logical volumes on the cluster. The output is stored in MESFILE.
3. Run VESYNC to get the list of logical volumes that must be copied to the new cluster from MESFILE gotten at step 1.
4. If the copy source cluster is a tape-attached model (TS7700T), run BVIRVTS (TYPE=VOL) to get VOLFILE. It is needed only for tape-attached source.
5. Run BVIRVTS (TYPE=CACHE) to the copy source cluster to get CACHEFILE. If CACHEFILE is also used for tape-attached source, TVC resident logical volumes are grouped so a physical volume does not need to be mounted to recall them to cache for Copy Refresh.
6. Run **COPYRFSH** to generate a Copy Refresh commands list from information of logical volumes that are gotten at steps 3, 4 and 5.
7. Run Copy Refresh commands in the list and wait for them to complete.
8. Monitor the status queue of the Copy Refresh commands by using the TS7700 MI under the performance page. For more information, see 13.3.1, “TS7700 Management Interface: Performance page” on page 672.

Figure 8-3 shows logic flow of the process.

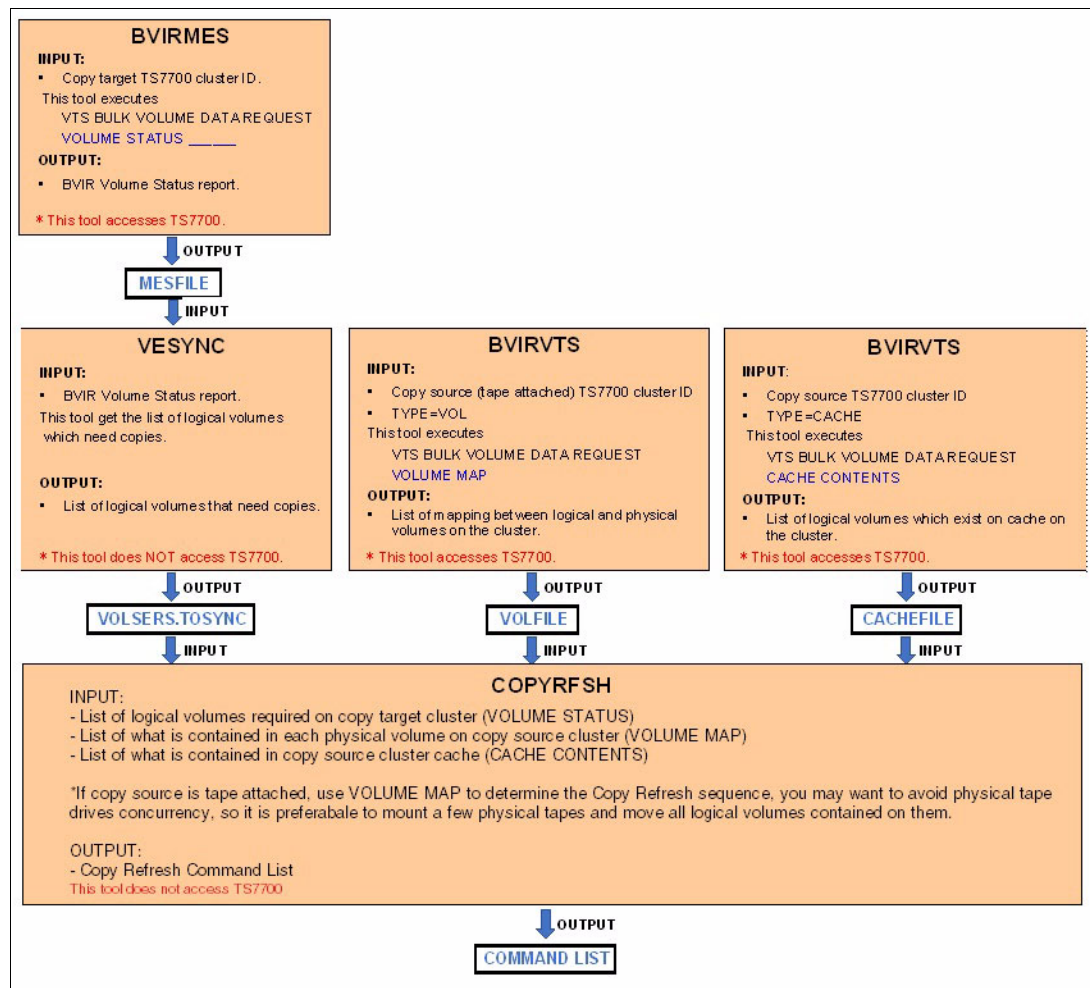


Figure 8-3 Generating Copy Refresh command list by using the IBM Tape Tools

8.2.3 Migration service offering

Migration service is a fee-based offering from the IBM Technology Expert Labs team that substantially automates the Copy Refresh process that is used to copy or move data between TS7700 clusters.

The offering includes the following benefits:

- ▶ A method to simplify hardware refreshes and changes, which is an alternative to the host-based method.
- ▶ Runs on the TS7700 grid, and no host access is required.
- ▶ Most migrations are done remotely by LBS personnel by using the remote support (“call-in”) function that is provided by the TSSC (TS3000 System Console).
- ▶ Uses TS7700 grid functions, which are the same as functions that are used by the host-based method.
- ▶ Separate from TS7700 code, can be installed on a cluster independently.

- ▶ Customer can see results of migration with existing interfaces or by asking for updates from IBM Technology Expert Labs.
- ▶ The customer provides the following requirements for LBS to run the migration:
 - Which and how many cluster migrations are running concurrently
 - Priorities can be specified by management class, category, and volume serial range
 - Rate of migration:
 - Limit the number of copies in the queue
 - “Black out” windows for copy processing
 - Ability to pause and restart migrations

Use the migration service offering to automate virtual volume migrations that are related to grid changes.

8.2.4 Copy Export and Copy Export Recovery/Merge

Copy Export function can export a copy of the selected logical volumes that are stored in physical tapes that are managed by TS7700T. Exported physical tapes can be removed from TS7700T and taken offsite. Typically, exported data that is stored in physical tapes is used for disaster recovery. However, the user can move data into another TS7700T by using exported data.

To restore exported data, the TS7700T must have physical tape drives that can read the exported physical tape volumes. As an example, in this case when the target TS7700T cluster has TS1150 tape drives only, the source TS7700T must use JK or JC cartridge with TS1140 or TS1150 tape drives for export. The user can restore exported data into another TS7700T by using one of the following methods:

- ▶ Copy Export Recovery
 - Can move exported data into a stand-alone empty TS7700T.
- ▶ Copy Export Merge:
 - Can merge exported data into another active TS7700T that has data.
 - Target clusters can be stand-alone or part of a grid.

Copy Export Merge is available through a service offering only, which provides an IBM storage specialist to help plan and implement the merge of logical tape volumes that are copy-exported from another IBM TS7700T cluster.

For more information about Copy Export and Copy Export Merge, see Chapter 15, “Copy Export” on page 783.

8.2.5 Grid-to-Grid Migration tool

The Grid-to-Grid Migration (GGM) tool is a service offering from IBM where you can copy logical volumes from one grid to another grid while both grids feature 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 IP addresses for the gridlinks (see Figure 8-4). No host I/O exists with the FICON adapters.

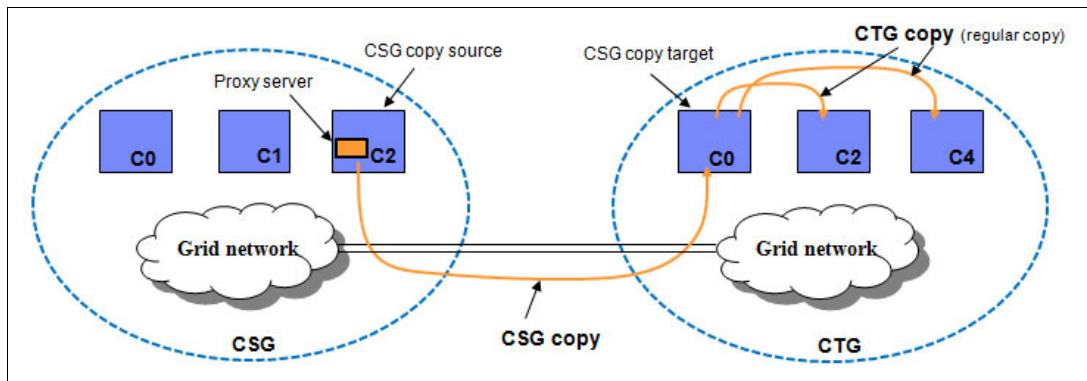


Figure 8-4 TS7700: Grid to Grid Migration overview

The following main components are shown in the GGM overview that is shown in Figure 8-4:

CSG	Copy Source Grid
CTG	Copy Target Grid
CSG Copy Target Cluster	Cluster that receives the CSG copy in the CTG
Proxy server	A binary program that is installed in the CSG through <code>vt_d_exec</code> , which enables the CTG to communicate with the CSG

The GGM tool should be considered whether the following situations are true:

- ▶ Eight clusters exist in the installed grid.
- ▶ The Join and Copy Refresh processing cannot be used (floor space requirements, microcode restrictions, and other considerations must be met).
- ▶ CSG and CTG are maintained by different providers.

Considerations for GGM

Consider the following points regarding the GGM tool:

- ▶ GGM is based on **LI REQ** commands. All commands are run on the CSG Copy Target Cluster. The chosen target cluster in the CTG must have at least R3.3 installed.
- ▶ It is single logical volume based. **LI REQ** commands (which belong to the M2 command class) must be run for **each** logical volume that must be copied. The quantity of requests (50 is the maximum) that can be concurrently processed in the M2 class level is limited. For more information about related constraints and how to correctly manage the corresponding queue, see 12.1.3, “Host Console Request function” on page 632.
- ▶ A volume to be copied by GGM must be in the *Cluster Source Grid* (CSG). If no copy is in this cluster, the logical volume cannot be copied, even if other clusters in the source grid include 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), more copies might be produced in the CTG.

- ▶ GGM traffic between the two participating grids is unencrypted. This condition persists even if Secure Data Transfer (SDT) is locally enabled for each CSG and CTG grids.
- ▶ 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 must run actions manually.
- ▶ After a logical volume is accessed for update (by an append or a housekeeping process), the logical volume cannot be copied again.
- ▶ GGM does not include any automatic interaction to the TMS. Depending on the scenario, the workload profile, amount of data, and used TS7700 models, IBM recommends the use of different approaches to select the sequence of the copies.
- ▶ GGM is based on a *deferred copies* mechanism. Therefore, you must review your CPYCNT parameter for the CTG because the Copy Source Grid (CSG) might throttle the GGM copies by using the values that are defined for DCOPYT and DCTAVGTD. The throttling might affect the performance of the GGM copies.
- ▶ The data placement in the CSG affects the performance of GGM. Recalls from physical tape can slow down the process of GGM. Whenever possible, use a TS7700D as the CSG.
- ▶ The GGM tool provides different options of how new data (new device categories) and old data (keep or delete the data in the source grid) are treated.
- ▶ GGM from any grid to the grid that supports the new compression method is allowed. In such a case, CTG can read logical volumes that are written by R4.1.1 or earlier code level in addition to volumes that are written by newer code levels.
- ▶ GGM from the grid that supports the new compression method to the grid that does not support it is not allowed. In such a case, CTG cannot read logical volumes that are migrated from CSG.
- ▶ GGM can also benefit from *Grid Cloud Awareness*, which was introduced at R5.1. It allows CTG to access volumes that are transferred from CSG to a cloud tier without passing through the grid network. After CSG completes data transfer to the cloud tier, CTG can immediately access it. When the CSG is at R5.1 and supports Grid awareness of volumes function, the CTG also needs to be at R5.1 and is configured with cloud-attached capabilities to get the benefit of Grid awareness.

In general, there are two major logical scenarios for GGM.

Data movement with GGM

The 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 must 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, consider that a multivolume file can be accessed only if all logical volumes belonging to the multivolume file are in the same grid.

Switch the production workload to the new grid before you start the GGM process. This switch includes several benefits:

- ▶ The CSG is not subject to Deferred copy throttling because no host I/O is processed.
- ▶ If the CSG is a TS7700T CPx, pre-staging the data speeds up the GGM copy processing.
- ▶ Data with a short lifecycle can expire in the source grid without needing to be copied at all to CSG, reducing the amount of data that must be copied.

Data duplication with GGM

This scenario is applicable if dedicated LPARs or the workload of an LPAR must be separated from an environment, which is 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 grid until a time 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, consider the use of the lifecycle information (expiration, retention) as input to avoid data with a short lifetime being copied.

To ensure that during tests of the new environment the original copied data to the new grid is only read but not modified, 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, GGM copies *cannot* be performed.

Copy Source Grid considerations

Whenever possible, use a TS7700D or the CP0 in TS7700T/TS7700C models in the CSG. This configuration simplifies the migration process because all source data is in cache, and no concerns exist about your back-end resources. In this case, the user can concentrate on data lifecycle considerations to prioritize data to be copied, or source data being needed by application purposes, especially for multivolume data sets.

If only a TS7740 can be chosen or a cache partition that is associated to a tape or cloud storage tier, also consider the back-end resources and the available cache for the GGM copy processing. In this case, prestage the data, and ensure that the data is copied based on the 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 is recalled by the GGM process, the recall process might take 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 be with the original Storage Class settings in the cluster. Otherwise, recalled data might be migrated before GGM can begin copy processing.

For more information, see [IBM TS7700 Series Grid To Grid Migration User's Guide](#) or contact your IBM SSR. It contains a detailed description of the GGM process, requirements, and user setup and preparation, and examples and **LI REQ** commands that are used in the GGM process.

In addition, several supporting tools to create the necessary input control statements, and the necessary TCDB entry changes and TMC entry changes, are available at the [IBM Tape Tool page](#).

8.2.6 Cloud-based migration

With the implementation of the Cloud Storage Tier capability, the TS7700 product now offers the following options for data migration:

► Partition refresh for cloud migration

This option is available only when all clusters in the grid are running microcode R5.0 or later. This requirement applies to cases in which a previously installed disk only TS7700D is upgraded to TS7700C so that it is now enabled to use the Cloud Storage Tier. The newly upgraded machine can include a significant number of private virtual volumes that are in the default cache partition and depending on business requirements, it can be desirable to migrate a portion, or all of them, to the assigned cloud storage system.

In this case, the following steps must be completed:

- a. Change the Cloud Premigration Rank in Storage Group (SG) and Partition in the Storage Class (SC) constructs to allow copies to be sent to the selected cloud repository.
- b. Determine the list of existing private virtual volumes that are intended to be reallocated to the cloud storage layer.

A user can list the status of virtual volumes by using the **BVIR VOLUME STATUS** command. The output of the resulting report can be filtered and then used to create a list of virtual volumes to send to the cloud.

- c. Use the library request interface to process each virtual volume in the volume list so that it can be moved to the cloud storage system: **LI REQ PARTRFSH,<volser>,MMOUNT**

This option updates the cloud attributes of target virtual volumes by mimicking a mount/demount, so the user does not need to specifically exercise such operation for new attributes to occur.

► Cloud ghost copy processing

If the grid includes a TS7700C and virtual volumes were premigrated to the cloud storage repository, cloud copies can be made directly accessible to a new TS7700C cluster joining the grid, a need does not exist for explicit copy processing through the grid links. The following steps must be completed:

- a. Join a new cluster to a grid.

For more information about how to join a new cluster to a grid, see 7.4, “Adding clusters to a grid” on page 289.

- b. Change the Copy Mode in Management Class (MC) to allow copies to the new cluster.

MC can be configured from MI. For more information, see 10.2.2, “Management Classes window” on page 478.

For more information about Copy Mode, see 4.2.2, “Defining grid copy mode control” on page 173.

- c. Generate a list of private virtual volumes that were premigrated to the cloud, which must be copied to the new cluster.

A user can list the status of virtual volumes by using the **BVIR VOLUME STATUS** command, and verify which volumes were premigrated to the cloud by using the **BVIR CLOUD VOLUME MAP** command. The output of the resulting reports can be filtered and then used to create a list of virtual volumes that are copied to the new TS7700C cluster.

The following IBM Tape Tools are available to support this process:

- BVIRMES
- VESYNC
- BVIRVTS
- COPYRFSH

9. Run **Copy Refresh** on each logical volume in the list that was created in Step 3 to produce a new ghost copy of the data in the new cluster. Copy Refresh with “ghost” option can be run by using the **LI REQ COPYRFSH, <volser>, CLDGHOST** command.

The **Copy Refresh** command can be submitted from a host or the MI, and the data is copied only in a logical level through the gridlinks. No Host I/O is done through the FICON adapters, nor bandwidth consumption through grid links for copies that are required, and all data in the TCDB and tape management remains unchanged.

This method can be used only if the data migration is inside a grid with target volumes previously premigrated to the cloud storage repository. The BVIR AUDIT parameter provides a method to help ensure that all data is copied.

For more information about the available tools, see 13.5, “IBM Tape Tools” on page 703.

For more information about the commands that are described in this section, see [IBM TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#).

For more information about the **LI REQ COPYRFSH** command, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

8.3 Moving data for host-based migration

When moving data into the TS7700, it is not possible to move the cartridges out of a VTS and insert them into a stand-alone TS7700T cluster, 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.

For more information about these publications, see the [IBM Redbooks website](#).

In all other scenarios, migrating data into the TS7700 requires that the TS7700 and the existing environment remain installed in parallel until the data is migrated through the host that is 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 into a TS7700. You can start using the TS7700 by moving data into it. The best method depends on the application that you want to manage with the TS7700.

The following methods are available:

- ▶ **Phased method**

This method consists of the use of the TS7700 with new allocations. Migrating data takes longer, but it can be more controlled and flexible.

- ▶ **Quick method**

Use this method when you want to move data into the TS7700. This method is considered quick because it swiftly puts all of the data that you want to move under TS7700 control.

For more information about how to move data out of the TS7700, see 8.4, “Moving data out of the TS7700” on page 324. 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.

8.3.1 Phased method of moving data

The techniques that are described in this section to move data depend more on changing parameters or adjusting routines and procedures to meet the objectives rather than a straightforward data movement.

Selecting the data

If you select DFSMSHsm-owned data, you can group your data as listed according to any or all the following items:

- ▶ 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 that are provided in 8.5, “Migrating DFSMSHsm-managed data” on page 327 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 copies, 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 use 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. For more information, see 8.6, “DFSMSrmm and other tape management systems” on page 336.

Watching the data move to the TS7700

Data movement that uses this option does not involve deliberate 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 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.

8.3.2 Quick method of moving data

The steps that are outlined in this section involve specific actions on your part to move data into the TS7700. As with the techniques that are described in 8.3.1, “Phased method of moving data” on page 319, 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 know the VOLSER where the data is stored. These applications feature special considerations. If you change the VOLSER on which the data is stored, the application has no way of knowing the location of the data. For more information, see Appendix B, “IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments” on page 883.

A simple method is to obtain information from the TMS database. Reports can give you details about the data you have in the tape shop. These details help you select the input volumes.

If DFSMSrmm is used, you can easily acquire data from a Removable Media Management (RMM) EXTRACT file, which is normally created as part of the regular maintenance. Then, by 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 can 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 SMS tape is used, the first step is to update the ACS routines to direct all new data to the TS7700. With this change, new data on tapes is created in the TS7700 so that moving it again later is not necessary.

If you move OAM-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. If you move DFSMSHsm-owned data, you must use the RECYCLE command to move incremental backups and migration volumes. Use a COPYDUMP job to move DFSMSdss data to the TS7700. The utility to use depends on the data selected. Usually, it is sequential data that can be copied by using the IEBGENER utility, DITTO/ESA. If you have DFSORT, ICEGENER and ICETOOL perform better.

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 multfile 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. Then, begin moving these tapes and leave the short-term retention tapes to the last phase of data movement. They likely are moved by the everyday process.

Updating the tape management system with the correct retention information

When the manual copy operation is 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/write date
- ▶ Job name

Optionally, you can also update the following items:

- ▶ Stepname
- ▶ DDname
- ▶ Account number
- ▶ Device number

In RMM, this step can be done by using a CHANGEDATASET command that has special authority to update O/C/EOV recorded fields. For more 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 process can be by using **IDCAMS DELETE NOSCRATCH** or within TSO with the **U** command followed by **DEFINE NONVSAM**. 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 was correctly copied. Also verify that the retention and catalog information is correct.

By using this quick-method sequence, you can copy every type of tape data, including generation data groups (GDGs), without modifying the generation number.

In an RMM environment, you can use RMMCLIST variables, RMM REXX variables (or both), 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 a TMS is used other than RMM, see the appropriate product functions to obtain the same results.

Migrating data inside the TS7700 can be made simpler 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 for more information:

- ▶ 8.5, “Migrating DFSMSHsm-managed data” on page 327
- ▶ 8.7, “IBM Spectrum Protect” on page 338

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. This process changes the device class definitions for the selected data to put in the TS7700 and then starts 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, which uses TS7700 virtual volumes as output.

In general, Db2 image copies and Archlog are not retained for long. After all new write activity goes to the TS7700, you can expect that this data is moved by the everyday process.

8.3.3 Products to simplify the task

You might want to consider the use of a product that is designed to copy data from one medium to another. The first choice is the IBM offering that interacts with DFSMSrmm and is called *Tape Copy Tool* (see Table 8-1). 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 that are 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. Choose another Tape Copy Tool from Table 8-1.

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 the use of the product for other functions, such as duplicate tape creation
- ▶ Ease of use
- ▶ Ability to:
 - Create a pull list for any manual tape mounts.
 - Handle multivolume data sets.
 - Handle volume size changes, whether from small to large, or large to small
 - Review the list of data sets before submission.
 - Handle failures during the copy operation, such as input volume media failures
- ▶ Audit trail of data sets that were 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.

Several common tape copy products are listed in Table 8-1. You can choose one of these products or use your own utility for tape copy. The use of any product item in Table 8-1 is optional, and the user might want to evaluate the convenience of working with any of them based on particular needs or business requirements.

Table 8-1 Selection of tape copy tools

Product name	Vendor name	For more information
Tape Copy Tool/ DFSMSrmm	IBM	Contact your IBM SSR. Do not confuse this offering with the Tape Analysis Tools that are described in 13.5.2, "IBM Tape Tools installation" on page 708, which can be downloaded from IBM for no extra fee.
Tape Optimizer	IBM	https://www.ibm.com/products/tivoli-tape-optimizer-for-zos?mhsrc=ibmsearch_a&mq=Tape%20optimizer

Product name	Vendor name	For more information
CA-1/TLMS Copycat	Broadcom, Inc.	https://techdocs.broadcom.com/us/en/ca-mainframe-software/performance-and-storage/ca-1-tape-management-system/14-0/utilities-and-reports/ca-copycat-utility.html
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	https://seasoft.com/products/solutions-for-system-z/tape-management/zela
FATScopy	BMC	https://www.bmc.com/it-solutions/bmc-ami-storage-migration.html

In addition to using one of these products, consider the use of IBM Technology Expert Labs to assist you in planning and moving the data into the TS7700.

8.3.4 Combining methods to move data into the TS7700

You 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. This data can be moved into the TS7700 by using the quick method only. Decide how much of this data is to be moved into the TS7700. One way to decide is to examine expiration dates. You can then set a time when all volumes (or a subset) are copied into the TS7700. These volumes might not need to be copied if they are going to expire in two months. You can save yourself some work by enabling these volumes to go to SCRATCH status.

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. These volumes do not need to be copied if they are going to expire soon.

8.4 Moving data out of the TS7700

You might want to move data out of the TS7700 for many reasons. The most common reason is for disaster recovery or data interchange. You can move data out of the TS7700 by using one of the following methods:

- ▶ Host-based copy tools
- ▶ Copy Export and Copy Export Recovery/Merge
- ▶ Cloud Export and Recovery
- ▶ DFSMSHsm aggregate backup and recovery support

8.4.1 Host-based copy tools

You can use a host-based tool to copy the data from the TS7700 to the target.

By using this method, the data is reprocessed by the host and copied to another medium. This method is described in 8.3.1, “Phased method of moving data” on page 319. The only difference is that you must address the TS7700 as input and the non-TS7700 drives as output.

8.4.2 Copy Export and Copy Export Recovery/Merge

You can use the Copy Export function to export a copy of the selected logical volumes that are stored in physical tapes on TS7700. Exported physical tapes can be removed from TS7700 and taken offsite. Typically, exported data is used for disaster recovery, but it can also be used for data migration between TS7700s. For more information, see 8.2, “Migration between TS7700s” on page 305.

For more information about Copy Export and Copy Export Merge, see Chapter 15, “Copy Export” on page 783.

8.4.3 Cloud Export and Recovery

Similar to the Copy Export function but without a secondary copy of the data, Cloud Export performs the following steps:

1. Tracks queued premigration (copies of virtual volumes to the cloud object store) for cloud pools that are specified to be exported per user-supplied instructions.
2. After the pending premigration tasks to target cloud pools are completed, TS7700 performs a database backup covering virtual volumes that are stored in those pools.
3. Database backup that is generated during a cloud export can then be used as a restore point by performing a process denominated “Cloud Export Recovery”. This process makes copies that are on the cloud usable for a new, empty TS7700C library that was granted access to cloud containers where the objects are stored.

The TS7700Cs must be configured with valid cloud object store account information, authentication credentials, and container information.

Cloud Export is started by regularly running the **LIBRARY EXPORT** command against a virtual volume that contains an Export List File, which includes cloud-related options to run the operation.

Backups are retained based on Cloud Pools retention settings. Therefore, if a cloud retention is modified on MI, the new retention is considered when daily cleanup process runs.

All TS7760-VECs and TS7770-VEDs in the Grid must be at R5.1 to use this function. For more information about Cloud Export, see Chapter 17, “Cloud Storage Tier export, recovery, and testing” in *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.

8.4.4 DFSMSHsm aggregate backup and recovery support

The third method 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 SMS-managed and non-SMS-managed data that is cataloged on DASD, migration level 1, migration level 2, and tape.

By using the DFSMSHsm ABARS function, group the data you want to move outside the TS7700. Then, start addressing other tape drives that are 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.

This process requires performing the following tasks:

1. Identify the data sets to be backed up by any of these options:
 - Create a selection data set.
 - Define an aggregate group.
2. Run the ABACKUP VERIFY command.
3. Run the ABACKUP EXECUTE command.

These tasks are described next.

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.

By 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 that is 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 windows.

The *aggregate group* lists the selection data set names, instruction data set names, 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 process is the same as performing a test run of aggregate backup. The following example shows the **ABACKUP** command:

```
HSEND ABACKUP aggname 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 EXECUTE command

When you are ready, start the backup by using the following command:

```
HSEND ABACKUP aggname EXECUTE UNIT(non_TS7700_unit) PROCESSIONLY(USERTAPE)
```

When you enter the **ABACKUP** command with the **EXECUTE** option, the following tape files are created for later use as input for aggregate recovery:

- ▶ Data file: Contains copies of the data sets that were 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 by using other physical tapes. Consider the use of the Copy Export function, which enables you to move a copy of the original logical volume 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 15.2, "Implementing and running Copy Export" on page 800.
- ▶ For Copy Export Recovery, see 15.3, "Using Copy Export Recovery" on page 810.
- ▶ For the use of the DFSMSShm ABARS function, see the *z/OS DFSMSShm Storage Administration*, [SC23-6871](#).

8.5 Migrating DFSMSShm-managed data

DFSMSShm is an application that can use the full cartridge capacity. However, you might want to consider the use of the TS7700 rather than native physical drives for DFSMSShm data for various reasons. 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 that are 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 frequent recall takeaway activity and *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 result in a positive effect by reducing the number of mounts that are required to run several recalls. 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 be performed in parallel, given enough recall tasks. In addition, the persistence of the virtual tape in the Tape Volume Cache (TVC) after it is 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 DFSMShsm auxiliary tasks (MASH) rather than increasing the number of functional tasks within the started tasks. Each DFSMShsm started task can support up to 15 AUTOBACKUP tasks.

Another reason for the use of the TS7700 with DFSMShsm is 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 suited for the TS7700 because of the appropriate tailoring of those parameters that can affect DFSMShsm performance. This tailoring is described next.

For more information, see the *z/OS DFSMShsm Storage Administration*, [SC23-6871](#).

8.5.1 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 a 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. Consider the following points:

- ▶ 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.

The maximum data set sizes that are supported by DFSMSHsm in z/OS environments are listed in Table 8-2.

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 recycling.

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, 25000, and 65000 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 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 later, this data set can fit on 800 MiB volumes. Any single user data set larger than 3.81 TiB in z/OS 1.13, or 15.875 TiB in z/OS 2.1 and later 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 the use of 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 the 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 that share 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 commands can apply to all HOSTs in an HSMplex.

Each instance of DFSMSHsm can have a unique MIGUNIT specified. For example, 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 stored. The other DFSMSHsm hosts can then migrate and back up SGs that contain 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 by 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 a migration/backup volume. As the volume nears full, it handles spanning of data sets, as described in “Tape spanning” on page 333. 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 did not reach its physical capacity. By marking the last volume full, the next time processing starts, DFSMSHsm uses a new volume, starts a new multivolume set, and enables the use of a new DC and other SMS constructs.

If the volume is not marked full, the multivolume set continues to grow and use the old constructs.

Use the SETSYS PARTIALTAPE MARKFULL option because it reduces the number of occasions in which DFSMSHsm appends to a partial tape. This use 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 use case is relevant to Outboard policy management and the implementation of different logical volume sizes. If all volumes are marked full, you can 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.

8.5.2 TS7700 implementation considerations

This section summarizes DFSMSHsm implementation considerations regarding the TS7700.

Mount wait time

You can direct DFSMSHsm data into a TS7700. For a TS7700T, modify your DFSMSHsm mount wait timer to be 12 minutes. This modification enables extra time that might be needed on specific mounts for the TS7700T to stage the data back into the 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 the use of large logical volumes, such as 6000 MiB, 25000 MiB, or 65000 MiB, for backup and smaller logical volumes for migration, especially if the TS7700T is used. If a high recall rate from ML2 exists, you might not even want to use the entire capacity of an MEDIA1 or MEDIA2 virtual volume.

Installations in which recalls from ML2 are rare, and installations in which 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 are based on the optimum volume size.

When you customize the ARCCMDxx **SETSYS** parameters, see Table 8-3 on page 334 for more information. HSM is aware of the large virtual volume capacity; it is not necessary to use high PERCENTFULL values to tune the capacity of tapes from a DFSMSHsm perspective. 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 **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 *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 that is based on the default volume size for the media type. However, the volume usage (or a similar parameter) shows how much data was 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. By using a global scratch pool, you benefit from a fast mount time by establishing your scratch categories, as explained in "Inserting virtual volumes" on page 589. Consider the use of 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 a 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 DFSMSShsm might leave unused while it tries to eliminate the spanning of data sets. That is, 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 tradeoff: It reduces the occurrences of a user data set spanning tapes in exchange for writing less data to a specific tape volume than its capacity otherwise enables. The amount of unused media can vary from 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 experience an excessive number of spanning data sets must consider specifying a larger value in the **SETSYS TAPESPANSIZE** command. The use of a high value reduces tape spanning. In a TS7700, this value reduces the number of virtual volumes that must be recalled to satisfy DFSMSShsm 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.

8.5.3 DFSMSShsm task-related considerations

To better understand the use of DFSMSShsm with TS7700, this section summarizes the DFSMSShsm functions that use tapes and analyzes the benefits of tape virtualization for these functions.

Backups of DFSMSShsm control data sets

DFSMSShsm 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 TS7700 is used as output for the DFSMSShsm AUTODUMP function, *do not* specify the following parameters:

```
DEFINE DUMPCLASS(dclass STACK(nn))  
BACKVOL SG(sgname) | VOLUMES(volser) DUMP(dclass STACK(10))
```

These parameters were introduced to force DFSMSShsm 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 (DFSMSShsm Migration Level 2)

When a TS7700T is used as DFSMSShsm 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 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 it is possible that all migrated data sets were removed from the TVC and must be recalled from tape.

Backup and recovery

Unlike the DFSMSShsm RECALL operation, RECOVERY often has a lower frequency in a DFSMSShsm environment. Therefore, the use of TS7700 for DFSMSShsm backup and recovery functions benefits you without affecting DFSMSShsm performance. However, review your DFSMSShsm performance requirements before moving DFSMSShsm BACKUP to the TS7700.

TAPECOPY

The DFSMSShsm TAPECOPY function requires that original and target tape volumes are of the same media type and use the same recording technology. The use of a TS7700 as the target for the TAPECOPY operation from an original volume that is not a TS7700 volume might cause problems in DFSMSShsm because TS7700 virtual volumes feature different volume sizes.

Use the information that is listed 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%
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 DFSMSShsm 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 is copied completely.

However, the use of 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 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 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, two valid instances from which the data set can be retrieved are always available: one backup and one ML2 version. After the second copy is written at rewind-unload time, two copies of the ML2 data 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 TS7700T recall process. If you have a TS7700T with four physical drives, use a maximum of two concurrent DFSMSHsm RECYCLE tasks. If you have a TS7700T with six physical drives, use no more than five concurrent DFSMSHsm RECYCLE tasks.

Select the RECYCLEPERCENT and consider the following points:

- ▶ You 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 example, you might want to limit RECYCLE to only your old technology, and exclude the newer tape technology from RECYCLE until the conversion is complete:

- ▶ 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.

For more information about implementing DFSMSHsm, see *IBM z/OS DFSMSHsm Primer*, [SG24-5272](#).

8.6 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. 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 the following messages:

- ▶ CBR3660A
- ▶ CBR3792E
- ▶ CBR3794A

Note: Before APAR OA49373, the CBR3660A message was deleted when the scratch count was 2X+1 above the threshold. With this APAR (z/OS V2R1 and later), when the CBR3660A is deleted it can be customized by 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 is never to 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 used for only stacked volumes that were 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 was 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 CBRSP LCS. 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; instead, they are managed exclusively by the TS7700T. Do not enable any host to 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 parmli b 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. No similar REJECT parameters are used 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 no RMM entry exists 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 that are defined in its CDS. An extra 1,000,000 volumes need 1,000 MB of space. However, do not add all the volumes initially.

For more information, see “Inserting virtual volumes” on page 589.

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 more information about calculating the correct size of the RMM CDS, see *z/OS DFSMSrmm Implementation and Customization Guide*, [SC23-6874](#).

Consider the use of 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.

Also, some TMSs do not enable the specification of tape volumes with alphanumeric characters or require user modifications to do so. For more information, see the corresponding product documentation.

In 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 was done by EDGINERS or IEHINITT.

8.7 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 is referenced. IBM Tivoli Storage Manager, such as DFSMSHsm, can automatically fill a native 3592 cartridge. It can use the tape up to EOV, independent of the media type.

Released in 2009, Tivoli Storage Manager 6.1 did not include 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.

Introduced with Version 6.3, IBM Tivoli Storage Manager for z/OS Media and IBM Tivoli Storage Manager for z/OS Media Extended Edition enable IBM Tivoli Storage Manager V6.3 servers that are running on IBM AIX and Linux on IBM Z 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. For more information, see [Tivoli Storage Manager branding transition to IBM Spectrum Protect](#).

For more information about the current Tivoli Storage Manager supported levels for Linux on IBM Z, see [Overview - IBM Storage Protect Supported Operating Systems](#).

Figure 8-5 shows a high-level data flow in a Tivoli Storage Manager for z/OS Media environment.

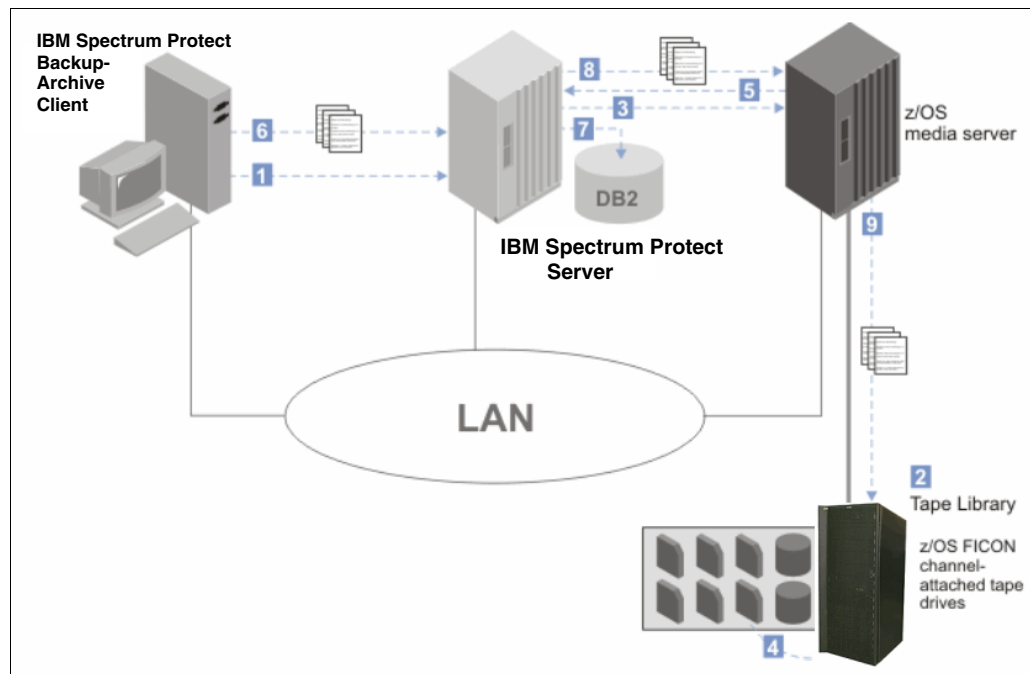


Figure 8-5 Data flow from the backup-archive client to z/OS media server storage

The following numbers correspond to the numbers that are shown in Figure 8-5 on page 338:

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 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. 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 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, similar to 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 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 does not expire for a longer period. IBM Tivoli Storage Manager reclamation must be scheduled for off-peak hours.

- ▶ Use colocation to reduce the number of TS7700 volumes required for a full restore. IBM Tivoli Storage Manager includes a colocation function to group IBM Tivoli Storage Manager client data onto a minimum set of tapes to provide a faster restore and a separation of client data onto separate physical tapes. Colocation with TS7700T does not minimize the physical tapes that are used, but minimizes the number of logical volumes that is used.

Colocation with TS7700T can improve the restore time for large amounts of data. TS7700T does not ensure physical tape separation when colocation is used because separate logical volumes can be on the same physical tape.

- ▶ Use TS7700T for IBM Tivoli Storage Manager database backups that are to be used for recovery from local media, and use 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 the TS7700 is used, do not be concerned about the unused capacity of logical volumes.

- ▶ Use 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 TS7700T because it is not guaranteed 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 TS7700T. The copy storage pools for offsite backups must use native drives or TS7700T at the recovery site.

- ▶ Use TS7700 in server-to-server configurations for multiple IBM Tivoli Storage Manager server implementations. If you use 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 files are ideal candidates for a TS7700.

8.7.1 Native or virtual drives

When only one stand-alone TS7700T is available, you might choose native drives for data that is 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 include compatible devices for the EXPORT media. Use native drives for IBM Tivoli Storage Manager EXPORT, unless you have multiple TS7700T tape drives and can IMPORT/EXPORT between the TS7700T tape drives.

8.7.2 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 that is available for IBM Tivoli Storage Manager use likely can 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 parameter avoids the need to keep the virtual volume mounted and frees a virtual drive for other users.

- ▶ MAXCAPacity (device class definition)

By using 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. The use of 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 [Installing and upgrading the server](#).

8.8 DFSMSdss

This section describes the uses of DFSMSdss with the TS7700.

8.8.1 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 are not 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 TS7700T, you fill the stacked cartridge without changing JCL by using multiple virtual volumes. The 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 or 65000 MiB logical volume through your SMS DATACLAS. This method prevents unneeded multivolume files.

By 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.

8.8.2 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 an IBM Z 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 DFSS 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 when you can and start using the enhanced Stand-Alone Services.

To use Stand-Alone Services, create a stand-alone core image that is suitable for IPL by using the new **BUILDSA** command of DFSMSdss. Create a 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 432.

Complete the following 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 are 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 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 mounts in “Virtual tape drives” on page 412.
4. Load the Stand-Alone Services program from the device that 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 the use of 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. 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 that is shown in Figure 9-6 on page 361, 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.

8.9 Object access method

With OAM's object support, which is used to store unstructured (byte-stream) data, objects can be stored on disk (in Db2 or in a file system [zFS or NFS]), on optical, tape, and with OA55700 directly to the cloud. Storing objects to tape (such as the IBM TS7700 Virtualization Engine) provides a low-cost storage medium for storing primary and backup copies of OAM objects. When storing objects to the TS7700 transparently to OAM, those same objects can be later migrated to physical tape or to the cloud by using policies that were established within the TS7700 MI.

Critical to OAM is the ability to retrieve an object on the storage medium that is most suitable to the access requirements of the data. Typically with OAM objects, the data is first stored in OAM's disk tier (in Db2 or in a file system) and then, as the data ages, OAM can be instructed to move that data to another tier in its storage hierarchy through SMS policies.

Through its policies, the TS7700 can also be instructed to move OAM object data from its disk cache to tape or to the cloud as the data ages. Although virtual tape (while in the disk cache of the TS7700) is faster than physical tape, virtual tape is still serially accessible and multiple read requests for objects on the same virtual volume still need to wait until the previous request is satisfied.

Regardless of where the data is stored in the OAM storage hierarchy, you want to ensure that it is on storage that can meet the current access requirements of the data. For objects that are stored on virtual tape, the ability is also available in OAM to recall (make a copy of an object) to the disk layer of the OAM storage hierarchy for days.

Because the retrieval of an OAM object is often what is most important, consider the use of smaller logical volume sizes when the primary copy of an OAM object is stored on virtual tape. Then, because OAM also supports up to two OAM managed backup copies, larger logical volume sizes can be used when an OAM managed backup copy is stored on virtual tape.

With the TS7700, OAM obtains the size of the logical volume from the library. Therefore, the use of the TAPECAPACITY parameter (in the CBROAMxx parmlib member) is not needed. Although the TS7700 can replicate OAM's object data across clusters, the recommendation is to maintain at least one OAM-managed backup copy to ensure that a separately written backup copy is available.

For object data that can be frequently accessed, policies must be set up within the TS7700 so that the primary copy of an OAM object is available in the disk cache of the TS7700 for fast access. An OAM managed backup copy can then be set up to take advantage of the other cache partitions of the TS7700 and migration capabilities to physical tape or to the cloud, depending on recovery needs. If an OAM primary copy is not available, OAM can automatically access a backup copy of an object.

As with DFSMSHsm, the Synchronous mode copy option that is provided by the TS7700 can be used to replicate your object data to another cluster in the grid. This replicated copy (and other copies in the grid) is then managed by the TS7700. The replication capabilities in the grid can be used in addition to any separately managed OAM backup copies.

When the Synchronous mode copy option is used, it is recommended to set the CBROAM parmlib option, **DISABLEOPENTYPEJ (INOUT)**, on the **SETOAM** command. This setting is important if during OAM object processing the Synchronous mode copy operations go deferred (CBR3730E).

To optimize the use of the TS7700 when storing OAM object data, consider the following points:

- ▶ Review the **MAXTAPERETRIEVETASKS** and **MAXTAPESTORETASKS** parameters at the global and storage group level to ensure that they are optimized based on read requests for the data. For example, if object storage groups exist that get more read requests, ensure that they have more devices available to them. Other parameters, such as **DEMOUNTWAITTIME**, **TAPEPERCENTFULL**, and **TAPEFULLTHRESHOLD**, also might need to be reviewed when the TS7700 is used to store OAM object data.
- ▶ Review the **MOUNTWAITTIME** parameter in the CBROAMxx parmlib member to see whether the default (5 minutes) is sufficient. If the primary copy of an OAM object is no longer available in the disk cache of a TS7700 (it might be on physical tape or in the cloud), the **MOUNTWAITTIME** accounts for the time that is needed to recall the volume back into the disk cache of the TS7700 and also any queue time that is spent within the TS7700.
- ▶ Review and consider other TS7700 policy settings: size of the logical volume for primary and for backup copies, whether logical WORM (LWORM) is needed, delete expire hold settings, copy policies for the data, disk cache and migration settings for frequently accessed data, changes in policy settings as the data ages, and so on.

8.10 Database backups

The use of a TS7700 as output confers several advantages to database backups. This section provides more information about these benefits for database products, such as Db2.

8.10.1 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 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 DASD or tape. When archiving to tape, a scratch tape volume is requested each time.

Archive logs contain unique information that is 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 and DFSMSHsm CPU cycles, and simplifies the process. The disaster recovery copy can be created by using Copy Export capabilities in the 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–1 GB. The size of a virtual volume on a TS7700 can be up to 65000 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 issue is not a problem with the TS7700 because it requires no manual operation. Even with the 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*, which is 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 features 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 full and incremental image copies. A full image copy contains an image of the whole table space at the time that 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 must be allocated at the same time. This requirement can tie up many tape drives and in smaller installations, prevent other work from being run. With one TS7700 and 496 virtual drives, this potential problem is not an issue.

The large number of tape drives is also important 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.

8.10.2 CICS and IMS

As with Db2, IBM CICS® and IMS use tapes to store logs and image copies of databases.

CICS is only a data communication product. IMS includes the data communication and the database function (IMS-DL/1). IBM 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, 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 helps to provide a 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 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 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 case implies the use of many input logs that ARE 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. The use of one virtual volume per database does not waste space because the TS7700 then groups these copies into a stacked volume.

unlike IBM Db2, IMS includes a batch function that works with databases through an IMS batch region. If you do not use logs when an IMS batch region is open, 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. By 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.

8.10.3 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 large if your environment does not use TMM or if you do not have DFSMSHsm installed. All of this data benefits from the use of the TS7700 for output.

With TS7700, the application can write one file per volume by using only part of the volume capacity. The 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 must move the data offsite, you must address a device that is outside of the local TS7700, or use other techniques to copy TS7700 data onto other movable tapes, as described in 8.4, “Moving data out of the TS7700” on page 324.



Part 3

Operations

This part describes daily operations and the monitoring tasks that are related to the IBM TS7700. It also provides you with planning considerations and scenarios for disaster recovery, and disaster recovery testing.

This part includes the following chapters:

- ▶ Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351
- ▶ Chapter 10, “IBM TS7700 Management Interface operations: Part 2” on page 449
- ▶ Chapter 11, “IBM TS7700 common operations and procedures” on page 557
- ▶ Chapter 12, “IBM z/OS host console operations” on page 625
- ▶ Chapter 13, “Monitoring” on page 667
- ▶ Chapter 14, “Performance considerations” on page 747
- ▶ Chapter 15, “Copy Export” on page 783
- ▶ Chapter 16, “Disaster recovery testing in a grid configuration” on page 819
- ▶ Chapter 17, “RESTful API” on page 855
- ▶ Chapter 18, “IBM TS7700 support for zTape Air-GAP” on page 871



IBM TS7700 Management Interface operations: Part 1

This chapter and Chapter 10, “IBM TS7700 Management Interface operations: Part 2” on page 449, provide information about how to configure and operate the IBM TS7700 by using the Management Interface (MI). The MI is the web-based Graphical User-Interface (GUI) of the IBM TS7700.

This chapter includes the following topics:

- ▶ 9.1, “Overview” on page 352
- ▶ 9.2, “User interfaces” on page 354
- ▶ 9.3, “Tape library management GUI” on page 355
- ▶ 9.4, “TS7700 Management Interface” on page 357
- ▶ 9.7, “Monitor icon” on page 387
- ▶ 9.8, “Virtual icon” on page 403
- ▶ 9.9, “Object Store icon” on page 443

9.1 Overview

Note: R5.4 can be installed on either the TS7770 VED models, 3957-VED, through a code upgrade in the field, or on the M/T 3948-VED as a new installation from manufacturing.

Existing 3957-VED models must have 128 GB of RAM before upgrading to R5.3 or later.

The 3957-VEC models (or previous models) are not supported by R5.3 code or later.

The TS7770 Advanced Object Store for DS8000 (FC 5283) was introduced with R5.2 phase2, and brought many enhancements to the original implementation of the DS8000 Object Store function (FC 5282), as for instance:

- ▶ All object enabled clusters in a grid become aware of all objects in their grid.
- ▶ All objects can be accessed from any object-enabled cluster in the grid, regardless whether a local copy of the object exists.
- ▶ Policy management for the DS8K objects is now available. The user can define what clusters receive a copy and if copies are synchronous or deferred between clusters.
- ▶ DS8K Multi-Cloud support, meaning that each DS8900 can target up to eight unique object stores (R9.2x multi-cloud support required for DS8K). The TS7770 supports up to 256 “Cloud Name Virtual Object Stores” in a grid. The same Cloud Name Virtual Object Store can be targeted by one or more DS8Ks configured with the same Cloud Name.
- ▶ Statistical information about the use of the Object cache partition is supported. These statistics are further broken down by Cloud Name, which can provide a more granular view of workloads.

TS7770 Advanced Object Store for DS8000 introduces a new navigational icon in the TS7770 Management Interface - the “Object Store” icon. Two new MI pages are now available under this icon, the “Object Policy” and “Object Store” pages. The Object Policy page is used to manage object policies and the Object Store page is used to assign object policies to Cloud Name Virtual Object Stores. Both are grid scope pages that allow a user to manage these items for the entire grid.

The TS7770 Advanced Object Store for DS8000 (FC 5283) is supported only on TS7770-VED and requires Feature Code 8083 (3.84 TB solid-state drives, SSDs). This FC allows the field or plant installation of 2 x 3.84 TB SSDs into empty slots of a new or existing VED server. This feature is not available for VEC-based TS7760 systems. This extra SSD storage space is used for internal object management purposes to support Advanced Object Store requirements.

Note: Clusters with FC 5283 (Advanced Object Store) cannot coexist within the same grid with clusters that have FC 5282 (DS8000 Object Store) installed.

For more information, see Chapter 2, “Architecture, components, and functional characteristics” on page 15, and 9.9, “Object Store icon” on page 443.

In the TS7770 SSD-based cache system for user data, introduced with R5.2, the code originally was limited to a maximum of two SSD drawers with a total of 120 TB of usable capacity.

With R5.3 code, the SSD-based cache system can be increased to a maximum of 4 drawers, delivering 260 TB (or 240 TiB) of usable data capacity without considering the compression. For more information, see Figure 9-4 on page 359, and Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Note: SSD-based cache is available for 3957-VED as 3956-CFC model (and XFC expansion) or in new installations of 3948-VED as 3948-CFC and XFC.

The Control Unit Initiated Recovery (CUIR) and Grid Resilience functions have been enhanced in the R5.2 (8.52.100.x code [phase1]):

- ▶ Grid Resilience (a function introduced in R4.1.2) monitors thresholds and other indicators to determine whether a TS7700 cluster in a grid is not healthy, and fences the cluster to preserve grid wellness.
- ▶ CUIR provides a way for a TS7700 cluster to notify attached LPARS with ZOS installed as the main operating system to vary offline devices before a planned outage. CUIR for healthy clusters was also introduced in R4.1.2.

With R5.2, when a cluster is fenced as part of the Grid resilience functions, any other cluster that is connected to the same host as the unhealthy cluster notifies the attached hosts about the sick cluster. This notification causes the hosts to automatically vary offline the devices that belong to the unhealthy cluster.

A single cluster can also notify its own attached hosts that it is sick, and auto-fence is being applied. Post outage, when the sick cluster is recovered, it notifies attached hosts to auto vary the devices back online.

For more information about how to fence and unfence a cluster, and how to enable CUIR, see 9.6.2, “Actions menu on the Grid Summary page” on page 368, and [TS7700 Series CUIR User Guide V1.2](#) or later.

Beginning with the R5.1 level of code, the TS7700 Management Interface introduced another layer of security by implementing the Dual Control. When enabled, Dual Control requires two people to complete a sensitive change in the MI settings by introducing a new attribute to the user, the checker authority.

Any user with privileges to modify category or cloud pool settings can propose a setting change, and another user with checker authority must approve it before the change is affected. A user with checker authority can propose a change, but it takes a different user who also is a checker to approve it.

Note: Enable or disable dual control along with user settings of checker authority is a grid-wide setting, not a cluster individual setting.

The implementation of Dual Control protects the following settings to prevent accidental or intentionally harmful changes that can pose a risk for client’s data:

- ▶ Modify category
- ▶ Delete category
- ▶ Modify logical volume version retention for a Cloud Pool
- ▶ Object Store and Object Policy pages

After Dual Control is enabled, it also covers the following actions, which avoids the bypassing of the protection:

- ▶ Disable Dual Control
- ▶ Modify User Password
- ▶ Modify Local User
- ▶ Enable Local/Remote Security Policy
- ▶ Modify Remote Security Policy User and Group Mapping

MI Dual Control options are available under Access Icon, as described in 10.3, “Access icon” on page 492.

For more information about TS3500 or TS4500 tape libraries, see the following publications:

- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *IBM TS4500 R8 Tape Library Guide*, SG24-8235

9.2 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 the following groups:

- ▶ Logical view
- ▶ Physical view

Each component and function belong to only one view.

The logical view is named the *host view*. From the host allocation perspective, only one library exists, called the *composite library*. The logical view includes virtual volumes and virtual tape drives.

The composite library can have up to 3,968 virtual addresses for tape mounts, considering an eight-cluster grid with support for 496 virtual devices in each cluster (available with FC 5275 and z/OS APAR OA44351). For more information, see 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 following configured physical tape library and tape drives:

- ▶ The TS4500 Tape Library with supported tape drives TS1150 (3592 EH8), and TS1160 (3592 60F) models
- ▶ The TS3500 tape library, which supports the TS1140, TS1150, and TS1160 (3592 60G) tape drive models
- ▶ TS1160 (Jag6) support included and JE/JM media support with R5.3.

Note: TS7770-VED tape attach configurations must use IBM supplied 16 Gbps switches.

The following operator interfaces are available for providing information about the TS7700:

- ▶ 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 12, “IBM z/OS host console operations” on page 625.
- ▶ Web-based management functions are available through web-based user interfaces (UIs). The following browsers can be used to access the web interfaces:
 - Firefox ESR: 60.xChrome: 77.x
 - Microsoft Internet Explorer: 10.x, and 11 Microsoft Edge: 44

Enable 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, the following web-based functions are available:

- The tape library GUI, which enables management, configuration, and monitoring of the configured tape library in tape attach configurations. The TS4500 and TS3500 are the supported tape libraries for the TS7700 implementation.
 - The TS7700 MI, which 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.
 - ▶ R5.3 of code introduces support to the Representational State Transfer (RESTful, or also REST) API, which provides a standardized and platform-independent user interface to the TS7700. Refer to Chapter 17, “RESTful API” on page 855 for more information on REST API. All user operations that are supported by any of the previously available interfaces should also be available in the new REST API.

Note: TS7700 RESTful API currently supports only the GET method for all the resources except for authentication where the POST method is used.

Find more about RESTful API on R5.4 Customer IBM Docs, available at TS7700 Management Interface Help documentation or online at:

<https://www.ibm.com/docs/en/ts7700-virtual-tape/5.4.1?topic=reference-restful-api>

9.3 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. The TS7700 tape attachment 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.

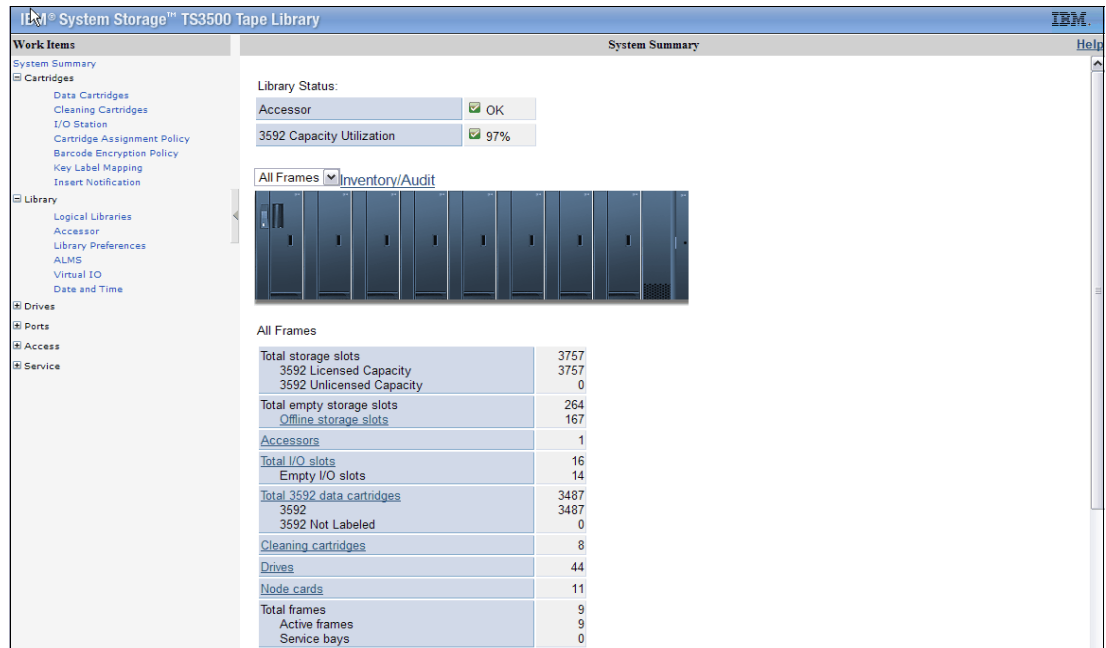


Figure 9-1 TS3500 tape library GUI initial window

Figure 9-2 shows the aspect of the Summary window of the TS4500 Tape Library GUI.



Figure 9-2 TS4500 Tape Library management GUI

The tape library management GUI windows are used during the hardware installation phase of the TS7700 tape attach models. For more information about installation tasks, see 11.2.1, “Tape library with the TS7700T cluster” on page 561.

9.4 TS7700 Management Interface

The TS7700 MI is the primary interface to monitor and manage the TS7700. The TS7700 GUI is accessed through TCP/IP, by entering the TS7700 IP address in your web browser. The current TS7700 graphical user interface (GUI) implementation features an appearance and feel that is similar to other MI that was adopted in other IBM Storage products.

9.4.1 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
```

If a local name server is used, where names are associated with the virtual IP address, the cluster name rather than the hardcoded address can be used for reaching the MI.

Note: If HTTPS exclusive access is enabled, you cannot access the Management Interface of a remote cluster. You can access only the Management Interface of the accessing local cluster.

4. The login page for the Management Interface loads, as shown in Figure 9-3 on page 358. The default login name is `admin` and the default password is `admin`. Starting with R5.0, the first time that an administrator logs in to the MI, the application reports the password as expired and requests the user to change it.

Note: New passwords must be 6–16 characters in length and composed of alphanumeric characters. Passwords include the following restrictions:

- ▶ At least five letters and one number must be used
- ▶ A number cannot be in the first or last position
- ▶ No spaces can be included
- ▶ The username cannot be included

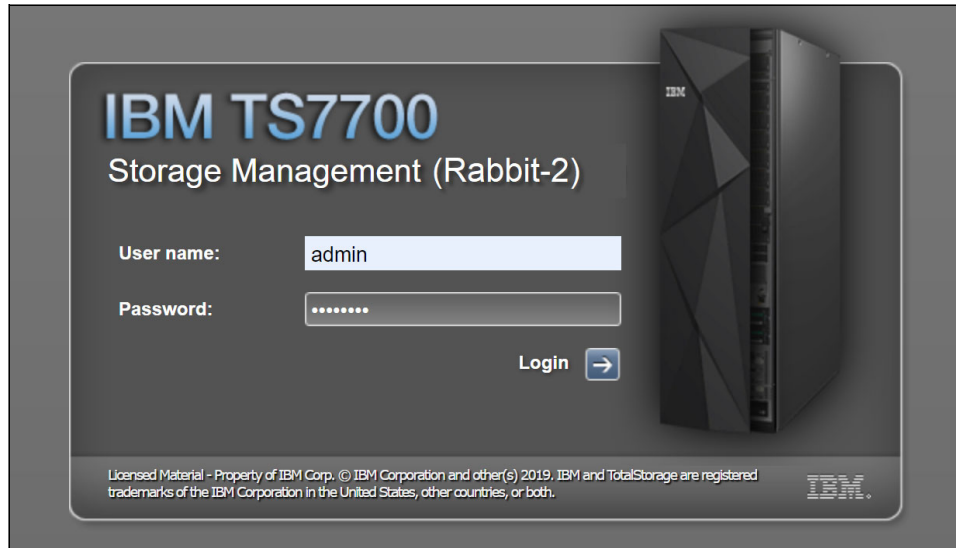


Figure 9-3 TS7700 MI login

When logging in, the Grid Summary page opens, as shown in Figure 9-4 on page 359.

If security policies are implemented locally at the TS7700 cluster or by using centralized role-based access control (RBAC), a unique user identifier and password can be assigned by the administrator. The user profile can be modified to provide only functions that are applicable to the role of the user. All users may not have access to the same functions or views through the MI.

For more information, see 10.3, “Access icon” on page 492.

Every user (regardless of the role) can reset their own expired password from the MI login page (as shown on Figure 9-3) by entering the expired password and updating it.

Figure 9-4 on page 359 shows an example of Grid Summary window of a TS7700 Grid. It shows a three-cluster grid, the members of the grid, and the health status of the components. The composite library is shown as a data center, with all members of the grid on the raised floor. Also shown is the navigation icons and options and the new Object navigational icon for Advanced Object Store for DS8000 (grid scope) is highlighted.

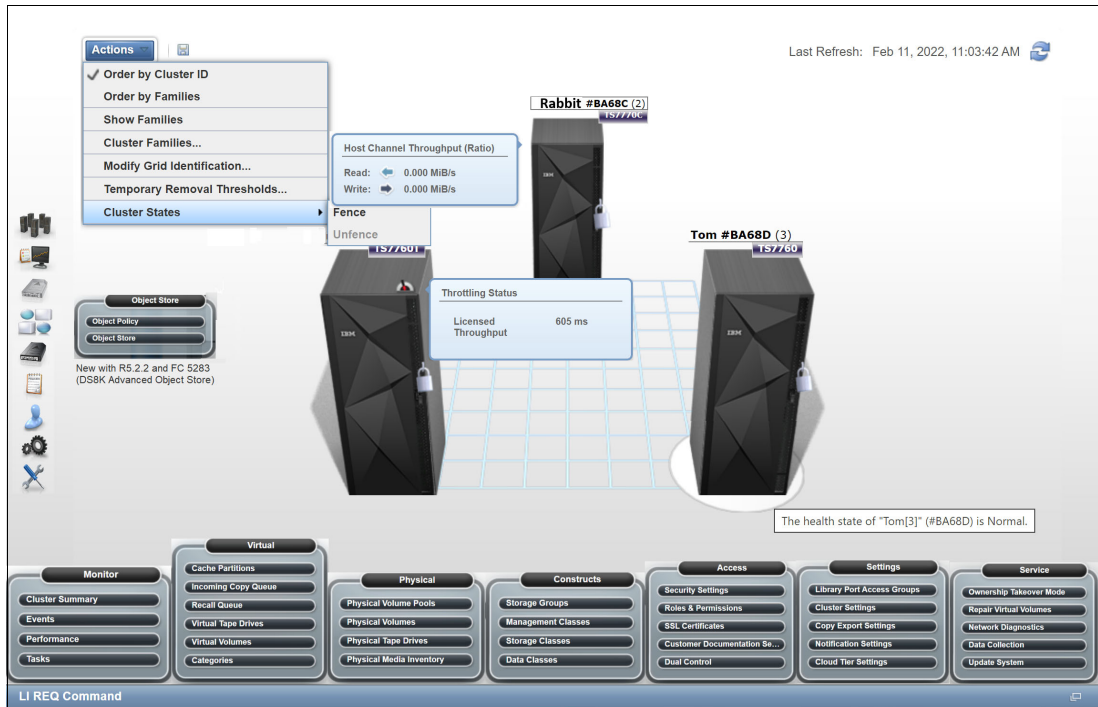


Figure 9-4 MI Grid summary

Each cluster is represented by an image of the TS7700 frame, which displays 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 interval 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, which is in the upper-right, 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 (see Figure 9-4).

Hovering the mouse pointer over the components on the page displays more information about the operational status of the component.

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. For more information about a specific cluster, click that component's icon. In the example that is shown in Figure 9-4, the clusters that are on the right side of the figure have a Warning (yellow) indication on them.

9.4.2 Using the TS7700 Management Interface

This section describes how to use the TS7700 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 more information about the MI. The banner is across the top of the MI web page, and allows a secondary navigation scheme for the user.

Figure 9-5 shows an example of the TS7700 MI banner element.

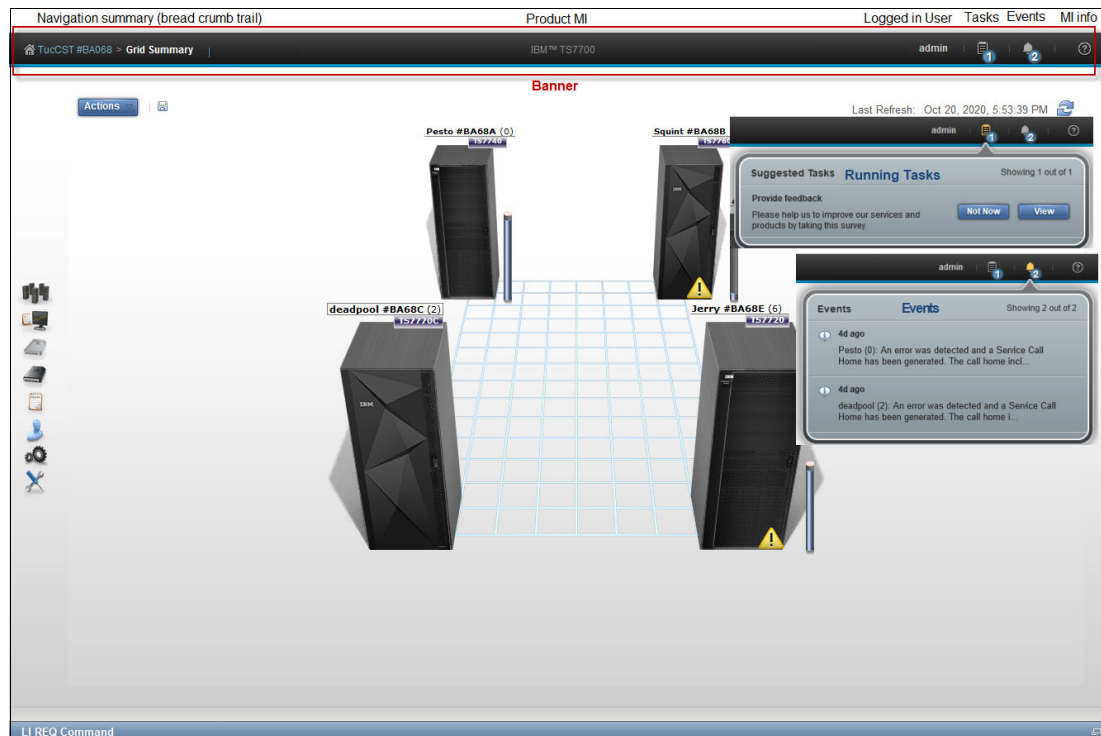


Figure 9-5 Management Interface Banner

The left field in the banner shows the sequence of selections that are made by user in the TS7700 MI website hierarchy (that is, the “bread crumbs” trail). The user can navigate directly to a different page by hovering the mouse over that field and clicking to select a different page.

At the right side of the banner (showing admin in Figure 9-5), the user that is logged in MI is shown. Hovering the mouse over it gives you the choices of logging out, changing the user password, and turning low graph mode on or off.

The suggested Task in the example that is shown in Figure 9-5 is an invitation to take an IBM Net Promoted Score (NPS®) Survey as a way to enhance the communications with the user by collecting feedback and suggestions regarding the implemented functions and the TS7700 Management Interface usability.

The last field to the right of the banner (question mark symbol) provides information about current MI window. In addition, you can start learning tutorials, IBM Documentation, and check the level of the installed IBM Documentation by hovering over it and clicking the wanted option.

Status and event indicators

Status and alert indicators are shown at the top 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 of the window footer and include tasks and events.

Figure 9-6 shows some examples of status and events that can be displayed from the Grid Summary window.

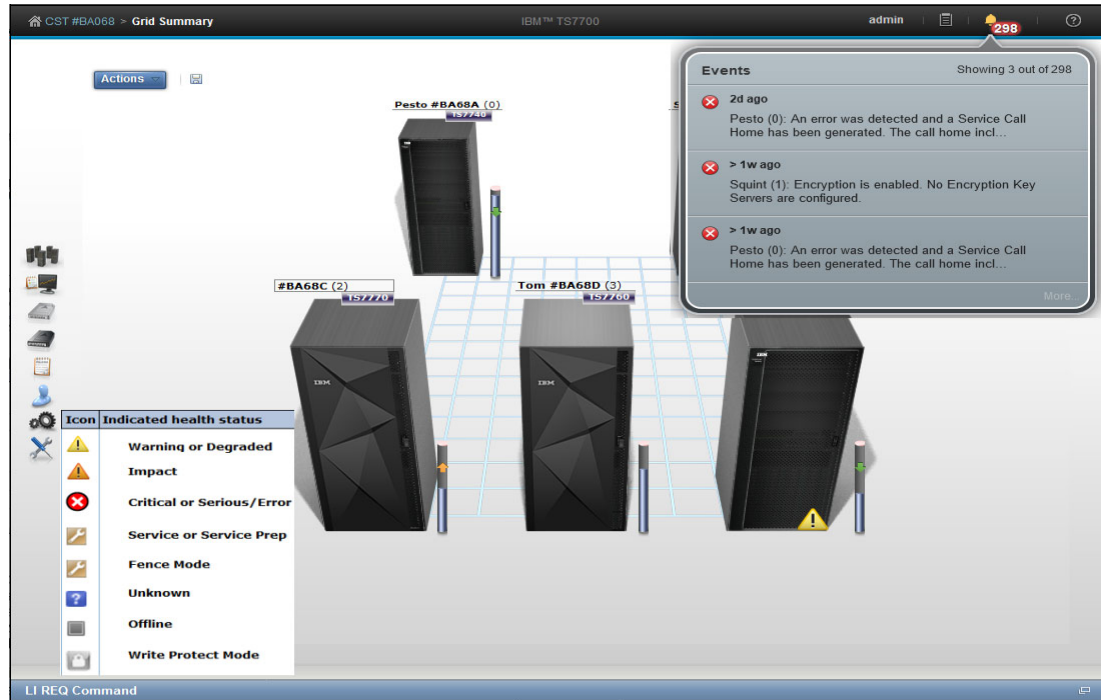


Figure 9-6 Status and Events indicators in the Grid Summary panel

All cluster indicators provide information for the accessing cluster only, and are displayed only on MI windows that have a cluster scope. 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 more information about tasks, the behavior of health and status icons, and a description of how to optimize the table presentations, see this IBM Documentation [web page](#).

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 window, 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 window is minimized and docked at the bottom of the MI window. The user must click it (at the lower right) only to open the LI REQ command pane. Figure 9-7 shows the new LI REQ command window and operation.

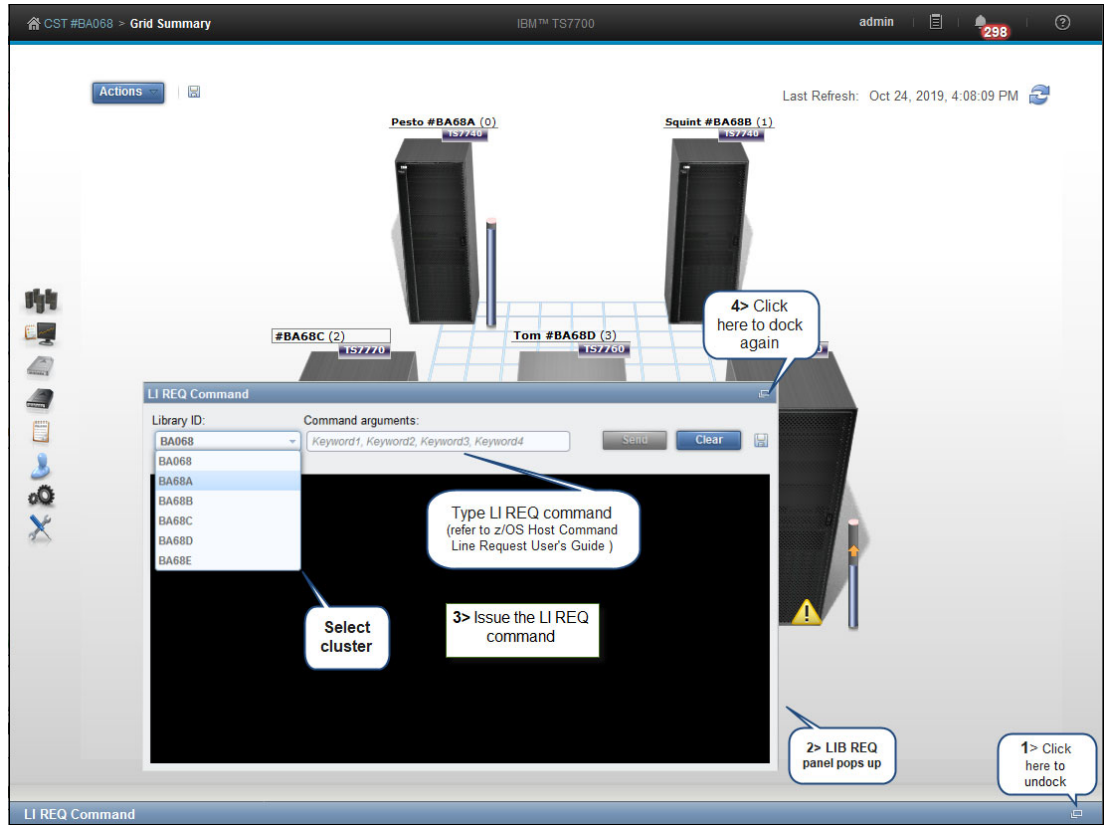


Figure 9-7 LI REQ Command window and usage

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 is shown only in the bottom of the MI windows for users with the Administrator role.

Figure 9-8 shows an example of a library request command that is reported in the Tasks list. It also shows how to get more information about the command by selecting **Properties** and then clicking **See details** in the MI window.

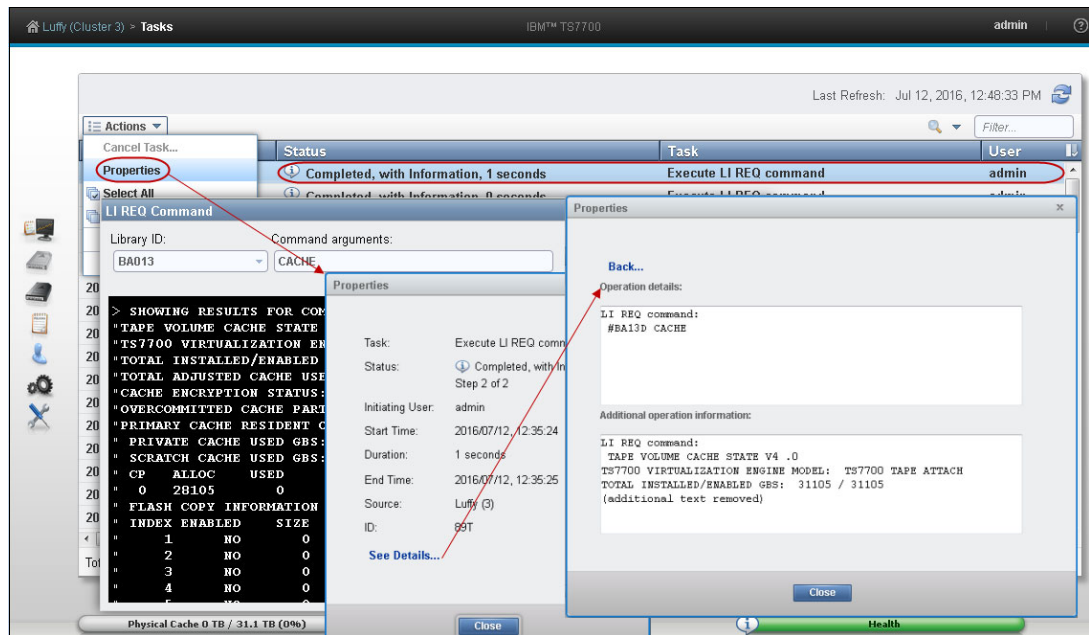


Figure 9-8 LI REQ command log and information

Important: LI REQ commands that are issued from this window are not presented in the host console logs.

For more information about available LI REQ commands, their usage, and respective responses, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

Standard navigation elements

This section of the TS7700 MI provides functions to manage and monitor the health of 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.

The following 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.
- Sphere/box icon** Information that is related to the Advanced Object Store for DS8000.
- Dark cartridge icon** Information that is related to physical cartridges and the associated tape library are in this window.
- Notepad icon** This window contains the construct settings.
- Blue man icon** All security-related settings are grouped under the Access icon.

- 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.



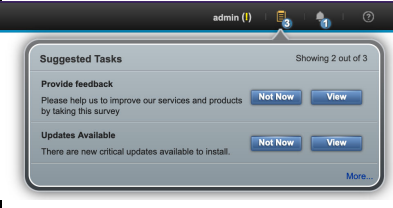
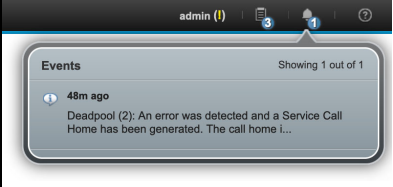
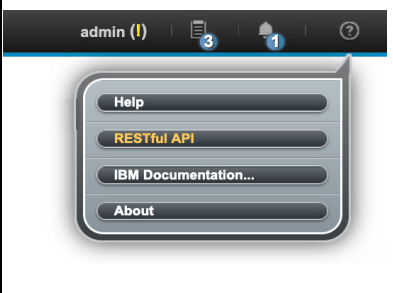

9.5 MI Navigation

A summary of the TS7700 MI Navigation is shown in Table 9-1.

Table 9-1 TS7700 MI Navigation

MI Function icon	Icon name	Usage description
	<p>System</p> <p>9.6, "Systems icon" on page 367</p>	<p>GRID Summary:</p> <ul style="list-style-type: none"> ▶ Investigate the details and health of a specific cluster ▶ Change the order of clusters displayed ▶ Show or hide cluster family relationship ▶ Change cluster family configurations ▶ Change grid nickname or description ▶ Change temporary Removal Threshold ▶ Fence/unfence a specific cluster ▶ Check host throughput values ▶ Check copy queue status for a cluster ▶ Identify throttling behavior for a cluster ▶ Download a file containing the grid information <p>Cluster Summary:</p> <ul style="list-style-type: none"> ▶ Change Cluster State ▶ Modify Cluster Identification ▶ Vary Devices Online for the cluster ▶ Fence/unfence cluster ▶ Cluster health and details ▶ Cluster components and alerts ▶ Cache capacity tube ▶ Port details and health ▶ Library details and health ▶ Expansion frame details and health
	<p>Monitor</p> <p>9.7, "Monitor icon" on page 387</p>	<ul style="list-style-type: none"> ▶ Cluster Summary ▶ Events ▶ Performance <ul style="list-style-type: none"> – Historical summary – Virtual mounts – Physical mounts – Host throughput – Cache throttling – Cache utilization (single or multiple partitions) – Grid Network throughput – Pending updates ▶ Tasks

MI Function icon	Icon name	Usage description
	<p>Virtual</p> <p>9.8, "Virtual icon" on page 403</p>	<ul style="list-style-type: none"> ▶ Cache Partitions ▶ Incoming copy queue ▶ Recall queue ▶ Virtual tape drives ▶ Virtual volumes ▶ Categories
	<p>Object Store</p> <p>9.9, "Object Store icon" on page 443</p>	<ul style="list-style-type: none"> ▶ Object policy ▶ Object store
	<p>Physical</p> <p>10.1, "Physical icon" on page 449</p>	<ul style="list-style-type: none"> ▶ Physical volume pools ▶ Physical volumes ▶ Physical tape drives ▶ Physical media inventory
	<p>Constructs</p> <p>10.2, "Constructs icon" on page 476</p>	<ul style="list-style-type: none"> ▶ Storage Groups ▶ Management Classes ▶ Storage Classes ▶ Data Classes
	<p>Access</p> <p>10.3, "Access icon" on page 492</p>	<ul style="list-style-type: none"> ▶ Enabling external disk encryption ▶ Security settings ▶ Roles and permissions ▶ SSL certificates ▶ Customer Documentation Settings ▶ Dual Control
	<p>Settings</p> <p>10.4, "Settings icon" on page 512</p>	<ul style="list-style-type: none"> ▶ Cluster network settings ▶ Notification settings ▶ Cloud Tier settings ▶ Feature licenses ▶ SNMP ▶ Library port access groups ▶ Cluster settings ▶ Copy export settings ▶ Secure Data Transfer

MI Function icon	Icon name	Usage description
	<p>Service</p> <p>10.4, “Settings icon” on page 512</p>	<ul style="list-style-type: none"> ▶ Ownership takeover mode ▶ Repair virtual volumes ▶ Cloud export ▶ Cloud export recovery ▶ Network diagnostics ▶ Data collection ▶ Copy export recovery ▶ Copy export recovery status ▶ Update system
	<p>Banner</p> <p>“Banner” on page 360</p>	<ul style="list-style-type: none"> ▶ Help items: <ul style="list-style-type: none"> – Learning and Tutorial – Information center ▶ User items: <ul style="list-style-type: none"> – Logoff – Change Password ▶ Bread crumbs
	<p>Tasks</p>	<ul style="list-style-type: none"> ▶ Active tasks (hyperlink to Tasks page)
	<p>Events</p>	<ul style="list-style-type: none"> ▶ Current Events (hyperlink to Events page)
	<p>Information</p>	<ul style="list-style-type: none"> ▶ Information about: ▶ Current TS7700MI page help ▶ Restful API ▶ IBM TS7700 Customer documentation R5.4 ▶ TS7700 MI, Cluster S/N, model, and microcode level.
	<p>Status Pods</p> <p>Figure 9-6 on page 361</p>	<ul style="list-style-type: none"> ▶ Grid scope Tasks and Events ▶ Cluster scope Cache capacity, Copy queue, and Health

9.6 Systems icon

The TS7700 MI windows that are under the Systems icon can help to quickly identify cluster or grid properties, and assess the cluster or grid “health” at a glance.

Tip: The Systems icon appears only when the accessed TS7700 cluster is part of a grid.

9.6.1 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 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 9-9 shows an example of a Grid Summary window, including the windows.

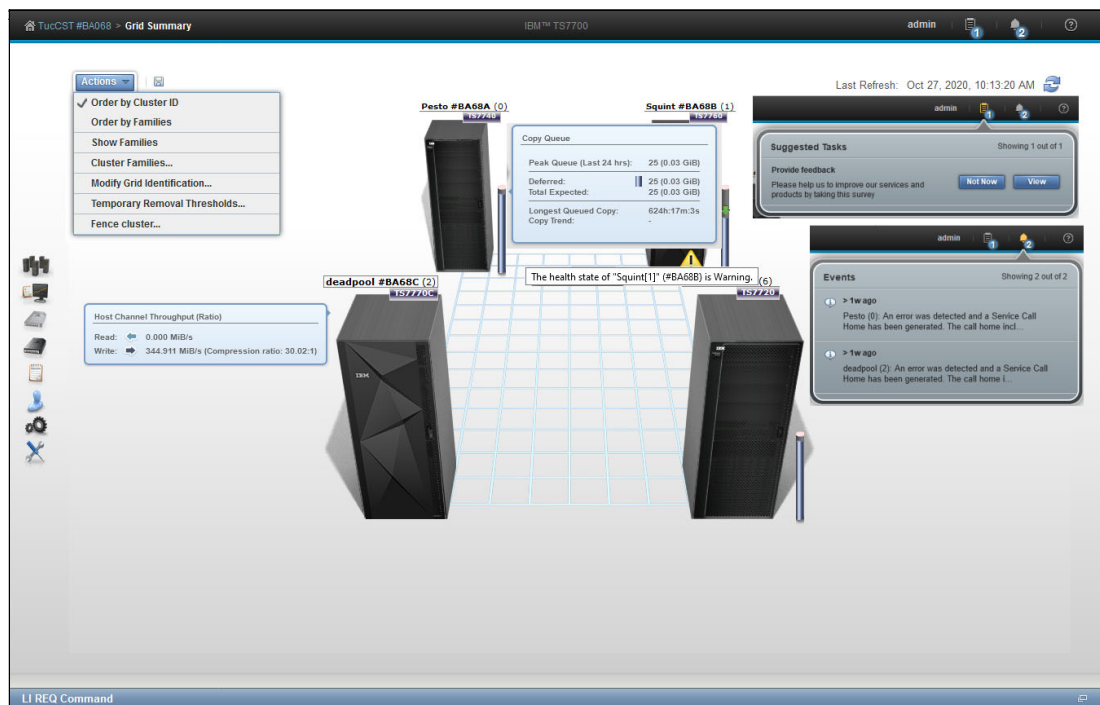


Figure 9-9 Grid Summary and windows

The Grid Summary window includes the following information:

- ▶ Cluster health and details
- ▶ Cluster families
- ▶ Throttling status
- ▶ Action menu
- ▶ Events
- ▶ Host throughput rate (sampled before compression by host adapters within cluster)

- ▶ Copy queue size and type
- ▶ Running tasks and events

A diskette icon is on the right side of the Actions button. Clicking the icon saves a CSV-formatted file with a summary of the grid components information, such as:

- ▶ Composite Library Sequence number
- ▶ Distributed Library Sequence number
- ▶ Grid nickname
- ▶ Cluster name
- ▶ Families
- ▶ Model of all clusters
- ▶ Microcode level for all clusters
- ▶ Disk Encryption and Key creation data
- ▶ Installed VTD_EXEC per cluster
- ▶ SDT status (on/off, protocol, version, key size, certificates) per cluster
- ▶ Serial number of the DDMS per Cache enclosure per cluster

9.6.2 Actions menu on the Grid Summary page

Use the Actions menu (see Figure 9-9 on page 367) 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 (CP0) of a TS7700T (tape attach) or a TS7700C (Cloud attach) clusters.

The following options are available:

- ▶ 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 and 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 that are used in the grid.

- ▶ Modify Grid Identification

Use this option to change grid nickname or description.

- ▶ Temporary Removal Thresholds

This option is used to temporarily change the removal thresholds of the disk-only clusters in the grid.

► Vary Devices Online

Select this option to vary devices online for the selected cluster. A blue informational icon is shown on the lower left corner of the cluster image if logical devices of that cluster must be varied online.

Figure 9-10 shows the vary cluster devices' online status. Notice the information icon on the cluster that is reporting devices offline.

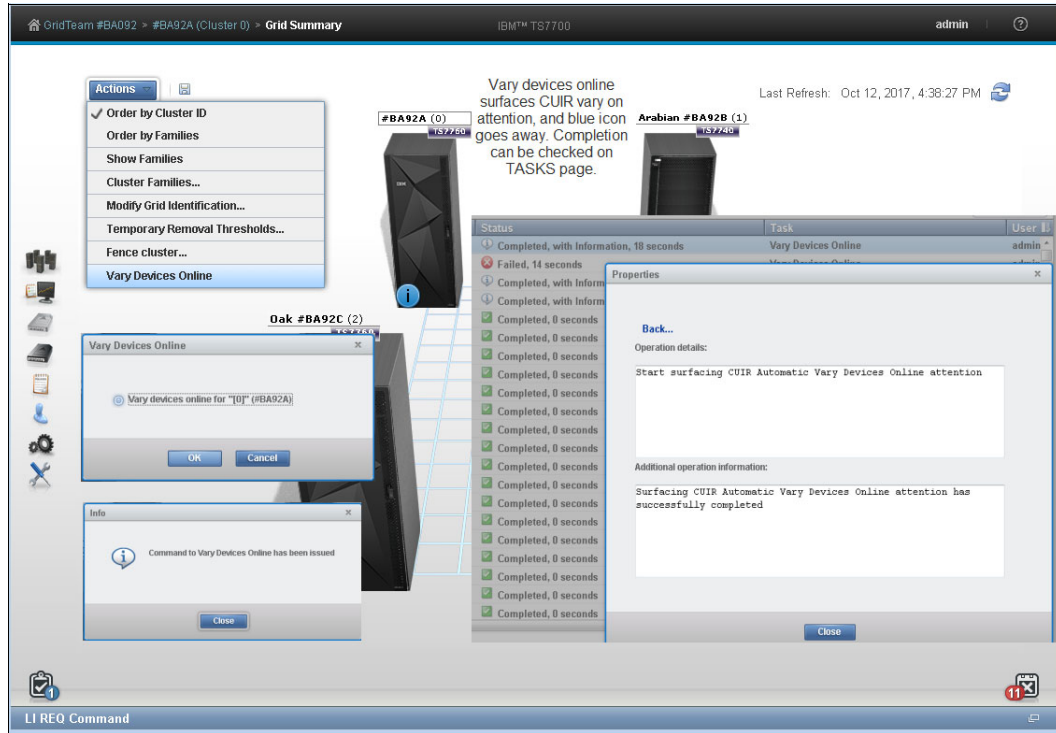


Figure 9-10 Vary cluster devices online

This menu option is available only if control unit-started reconfiguration (CUIR) is enabled by a **LI REQ** command and the Automatic Vary Online (AONLINE) notification is disabled. For more information about the use of the LI REQ commands, see Chapter 12, “IBM z/OS host console operations” on page 625.

► Fence Cluster, Unfence Cluster

Select **Fence Cluster** to place a selected cluster in a fence state. If a cluster is in a fence state, the option Unfence Cluster is shown instead.

Select **Unfence Cluster** to unfence a selected cluster that is in a fence state. These functions are part of the grid resilience improvements package. The Fence Cluster option in the Actions menu allows the user Administrator (default) to manually remove (fence) a cluster that was determined to be not functioning properly from the grid. Fencing a cluster isolates it from the rest of the grid. The administrator can fence the local cluster (the one being accessed by MI) or a remote cluster in the grid from this window.

Note: *Remote cluster fence* is enabled only when all clusters in a grid are at R4.1.2 (or later) code level.

The user can decide which of the following actions is taken by the sick cluster after the fence cluster action is selected:

- Options for the local cluster:
 - Forced offline
 - Restart
 - Restart and stay offline
- Options for a remote cluster (from any other cluster in the grid except for the cluster that is under suspicion):
 - Send an alert
 - Force cluster offline
 - Restart
 - Restart and stay offline or isolate from the grid

Figure 9-11 shows the TS7700 MI sequence to manual fence a cluster.

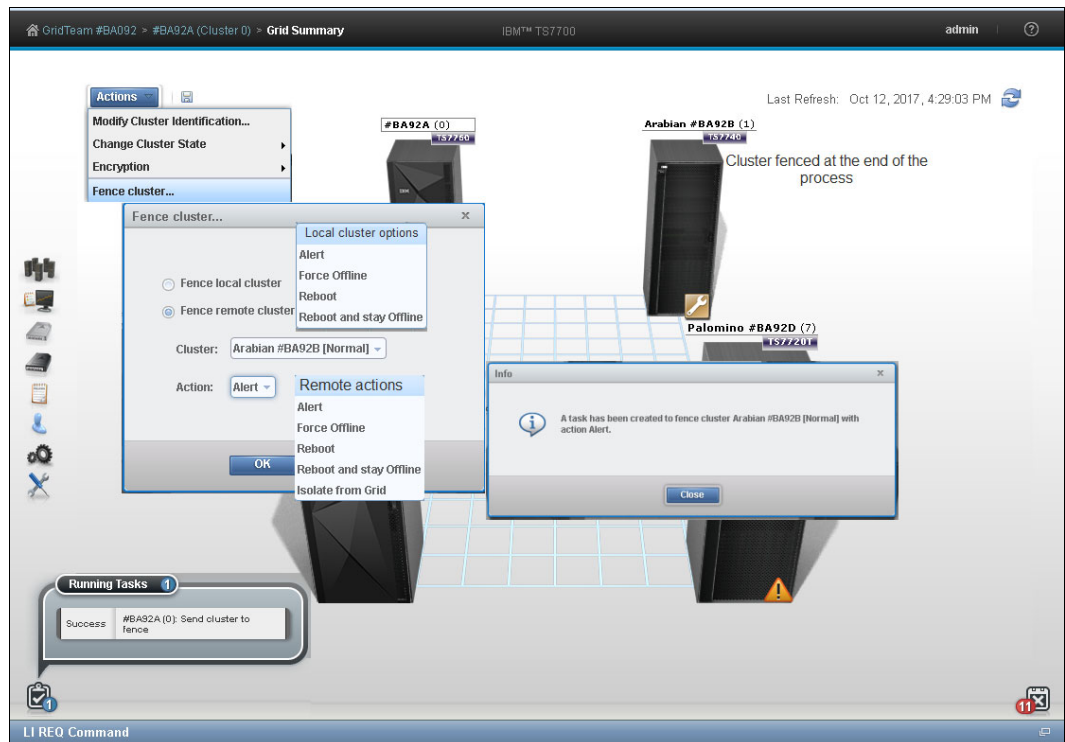


Figure 9-11 Fence a cluster operation.

For more information about cluster fence function and proper usage, see 2.4.36, “Grid resiliency functions” on page 101.

Figure 9-12 shows how to manually unfence a cluster by using the TS7700 MI.




Figure 9-12 Unfence cluster sequence

9.6.3 Cluster Families window

To view information and run actions that are related to TS7700 cluster families, use the window that is shown in Figure 9-13.

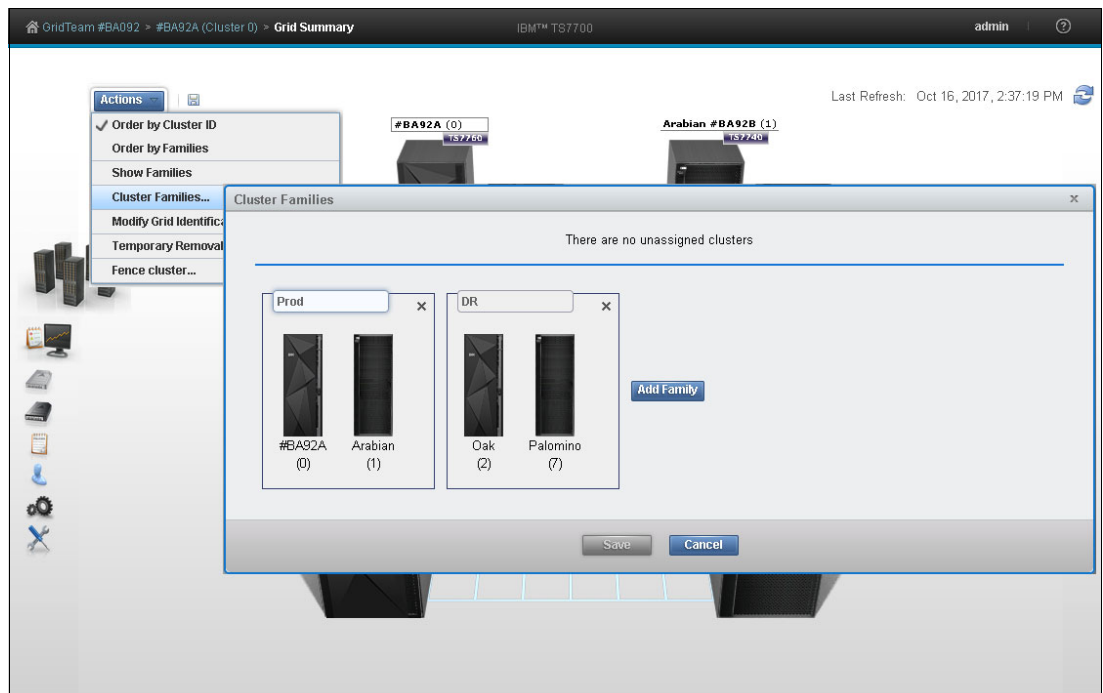


Figure 9-13 MI Add Cluster Families: Assigning a cluster to a family

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 cluster family. A new cluster family placeholder is created to the right of any 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 include 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 next.

► Move a cluster

One or more clusters can be moved by dragging to a new cluster family from the Unassigned Clusters area between existing cluster families, or to the Unassigned Clusters area from an existing cluster family:

- Select a cluster: A selected cluster is identified by its highlighted border. Select a cluster from its resident cluster family or the Unassigned Clusters area by using one of the following 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: A cluster family cannot be moved within the Cluster Families window.

► Delete a family

To delete a cluster family, click the **X** in the upper-right corner of the cluster family. If the cluster family to be deleted includes 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.

9.6.4 Grid Identification properties window

Use this option to view and alter identification properties for the TS7700 grid. In a multigrid environment, use this window to identify clearly a particular composite library, which makes it simpler to distinguish, operate, and manage this TS7700 grid (which avoids operational mistakes because of ambiguous identification).

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 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.

9.6.5 Lower removal threshold

Select **Actions** → **Temporary Removal Threshold** in the Grid summary view to lower the removal threshold for any disk-only cluster or cache-resident partition of a cloud or tape attach cluster in a grid that possess a physical tape library.

For more information about removal policies, see 4.2.6, “TS7700 cache thresholds and removal policies” on page 175.

Figure 9-14 shows the Temporary Removal Threshold window.

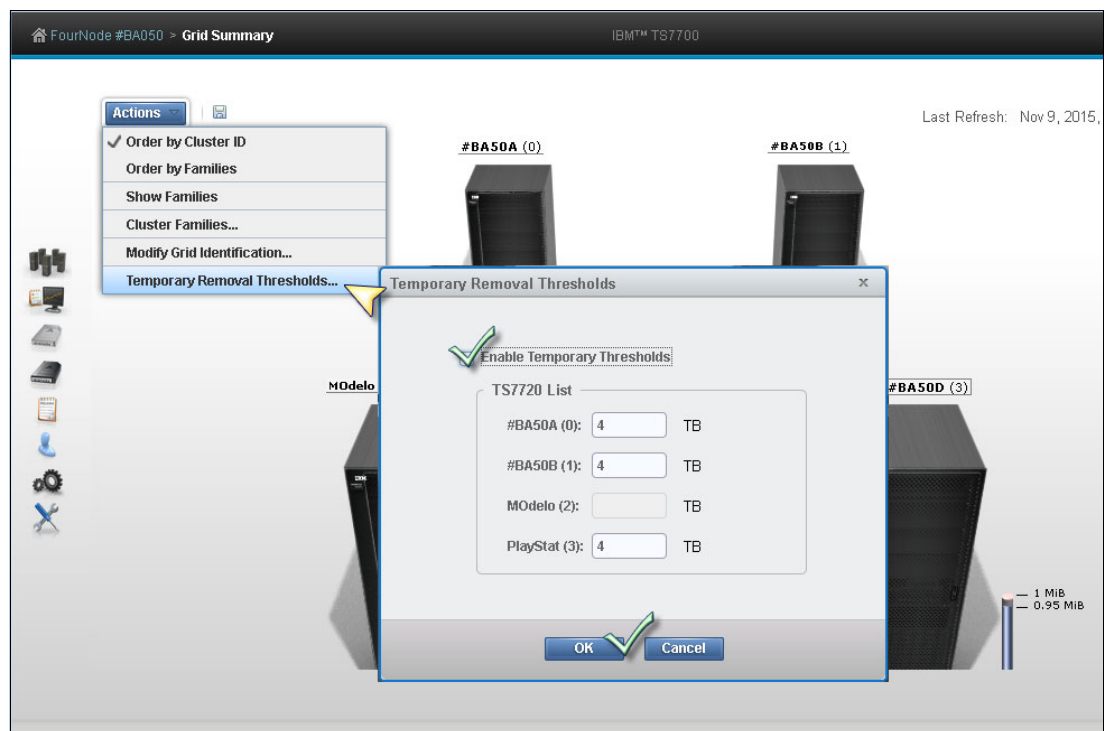


Figure 9-14 Setting the Temporary Removal Threshold









9.6.6 Grid health and details


In the Grid Summary view, the cluster is in a normal state (healthy) when no warning or degradation icon is displayed in 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, which confirms 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. More information about the status can be viewed by hovering your cursor over the icon.

The icons and how they are displayed in the Grid Summary page are listed in Table 9-2.

Table 9-2 Health status icons

Icon	Meaning	Reason
	Warning or Degraded	<ul style="list-style-type: none"> ▶ Out of Empty Stacked Volumes ▶ Copies Disabled by the System (out of physical scratch or out of cache space) ▶ Copies Disabled by Host by way of LI REQ ▶ Immediate Deferred copies ▶ All Storage Cells Full in Physical Library ▶ Limited Free Space in Cache ▶ Out of Cache Resources
	Impact	<p>No immediate impact to operations: Early warning for serious type interventions including:</p> <ul style="list-style-type: none"> ▶ Performance degradation ▶ Copies are degraded or not being performed ▶ Ejects cannot proceed ▶ Input station processing problems
	Critical or Serious/Error	<p>Operations in the reporting TS7700 or Library are stopped. All allocated jobs might fail. One or more jobs might fail, some virtual volumes might be inaccessible, mounts are suspended, or performance might be degraded.</p>
	Service or Service Prep	A cluster is in the Service or Service Prep mode
	Fence	A cluster is in the Fence Mode
	Unknown	The status of the cluster is unknown; its status has not been updated for at least 15 minutes.
	Offline	<p>A cluster is considered offline because of the following conditions:</p> <ul style="list-style-type: none"> ▶ Being in Service or Service Prep ▶ Being in Fence mode ▶ An active hnode or all vnodes are offline ▶ The cluster is in a pending online state
	Write Protect Mode	Write Protect Mode is enabled on the cluster. Clicking this icon opens the Write Protect Mode panel for the cluster.

Icon	Meaning	Reason
	Cluster Fenced or Devices Require Manual Vary Online	Hover over the icon to determine which informational state it represents: <ul style="list-style-type: none"> ▶ The cluster is fenced because of an Alert or Isolate from Grid condition and must be manually unfenced. ▶ The cluster surfaced CUIR attentions to vary offline devices, and LI REQ AONLINE is disabled, and the cluster is waiting to surface CUIR online attentions. Manually request the devices to vary online from the actions menu from the cluster summary panel.

For more information about icons and meanings, see this IBM Documentation [web page](#).

In the Grid Summary window, an alert icon indicates throttling activity on a cluster within the grid. Throttling can severely affect the overall performance of a cluster and the operation schedule. It might also result in job execution delays (see Figure 9-15).

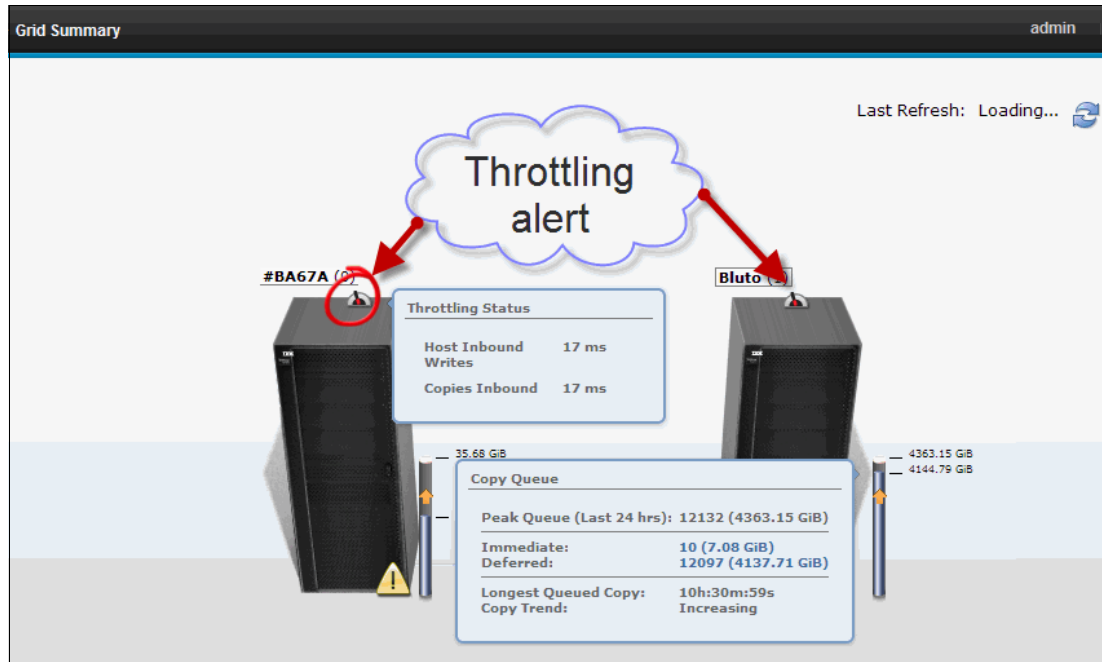


Figure 9-15 Clusters throttling in a two-cluster grid

For more information and how to avoid this issue, see Chapter 13, “Monitoring” on page 667.

For more information about throttling in a TS7700 grid, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

9.6.7 Cluster Summary window

The Cluster Summary window can be accessed 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. In a stand-alone configuration, this icon is the first icon that is available in the MI.

Figure 9-16 shows an example of the Cluster Summary window.

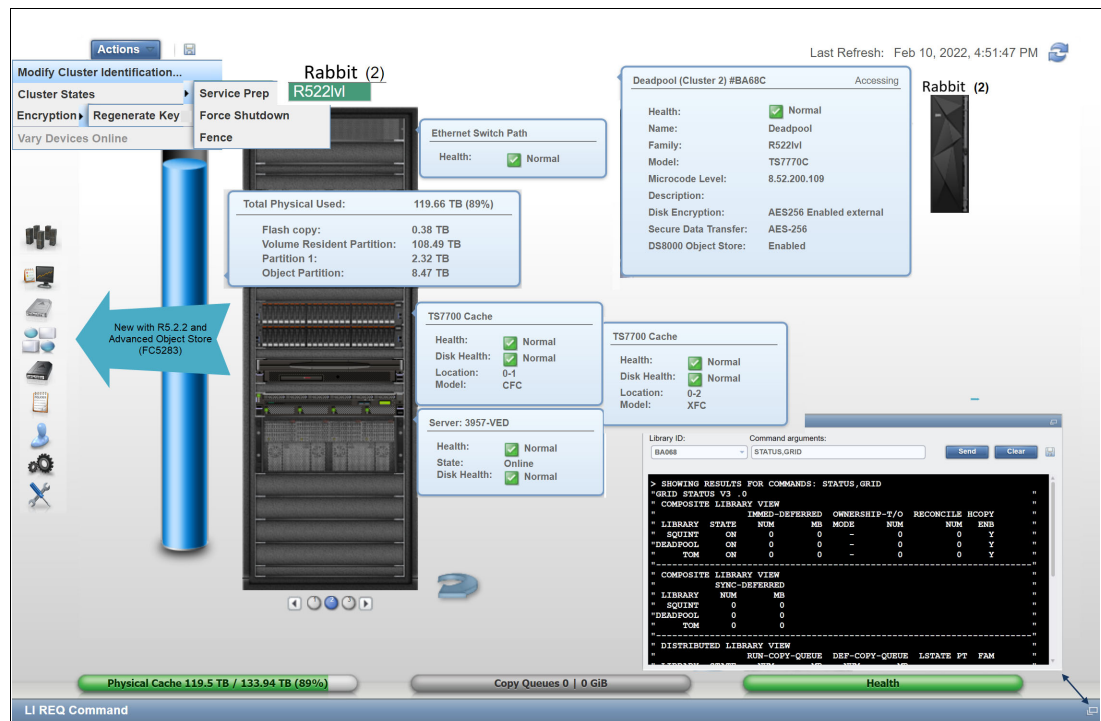


Figure 9-16 Cluster Summary with an SSD-based TS7700 and TS7700 Advanced Object Store

The Cluster Information can be displayed by hovering the cursor over the components, as shown in Figure 9-16. The following information is available:

- ▶ Cluster components health status
- ▶ Cluster Name
- ▶ The 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)

A diskette icon is shown to the right of the Actions button. Clicking this icon downloads a CSV-formatted file with information about that cluster, such as:

- ▶ Distributed Library Sequence number
- ▶ Cluster name
- ▶ Family
- ▶ Model
- ▶ Microcode level
- ▶ Disk Encryption and Key creation data
- ▶ Installed VTD_EXEC
- ▶ SDT status (on/off, protocol, version, key size, certificates)
- ▶ Serial number of the DDMs per Cache enclosure

9.6.8 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.

The available options to change the state of a cluster are listed in Table 9-3.

Table 9-3 Options to change the cluster state

If the current state is	Selection	Restrictions and notes
Online	Service Prep	<p>All of the following conditions must be met first:</p> <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. <p>Caution: If only one other cluster remains online, a single point of failure exists when this cluster state becomes service prep mode.</p> <p>Select Service Prep to confirm this change.</p>
	Force Shutdown	<p>Select Force Shutdown to confirm this change.</p> <p>Important: After a shutdown operation is initiated, it cannot be canceled.</p>
	Fence	<p>Select Fence to manually fence the accessing cluster</p>
Service Pending	Force Service	<p>Use this option when an operation stalls and is preventing the cluster from entering Service Prep.</p> <p>Select Force Service to confirm this change.</p> <p>All but one cluster in a grid can be placed into service mode, but it is advised that only one cluster is 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.</p>
	Return to Normal	<p>Select this option to cancel a previous service prep change and return the cluster to the normal online state.</p> <p>Select Return to Normal to confirm this change.</p>
	Force Shutdown	<p>Select Force Shutdown to confirm this change.</p> <p>Important: After a shutdown operation is started, it cannot be canceled.</p>
Shutdown (offline)	User interface not available	<p>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.</p>
Online-Pending or Shutdown Pending	Menu disabled	<p>No options to change state are available when a cluster is in a pending state.</p>

Going offline and coming online considerations

Whenever a member cluster of a grid goes offline or comes back online, it must 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 *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, which represents all of the logical volumes that exist in the grid, and those copies are kept updated at the same level by the grid mechanism. When coming back online, a cluster must reconcile its own collection of tokens with the peer members of the grid, ensuring that it represents the status of the grid inventory. This reconcile operation is also referred to as *token merge*.

Consider the following items 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 that is 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, an option is available to skip the merge step.

Click **Skip Step** to skip the merge operation. This option is available only 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 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 also is 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 helps ensure 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, AOTM 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 of which it need not take ownership. 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. Also, manual ownership takeover can be enabled, if necessary.

► Shutdown restrictions

To shut down a cluster, the user must be logged in directly to that cluster. To shut down a different cluster in the grid, log out of the current cluster and log in directly to the cluster to be shut down. For more information, see 9.6.10, “Cluster Shutdown window” on page 382.

9.6.9 Service mode window

Use the window shown in Figure 9-17 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. For more information, see Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Remember: Service mode is possible for only clusters that are members of a grid.

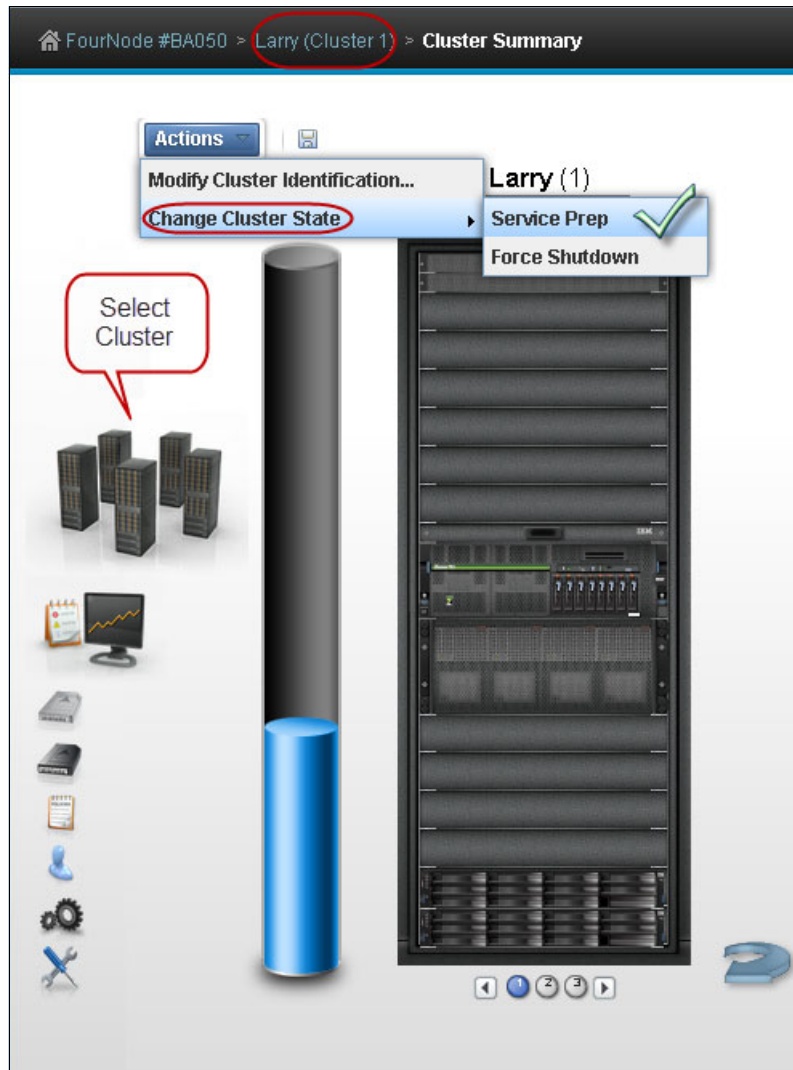


Figure 9-17 Cluster Summary: Preparing for service

Service mode allows 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 the client. Host operations continue automatically by using the other members of the grid if consistent copies of the logical volumes exist in those clusters and the host has access to the clusters.

Before CUIR, the user needed to vary offline all logical drives that were associated to the cluster going into service before changing cluster state to service. This process had to be done across all LPARs and system plexes that were attached to the cluster

Note: On the host side, vary logical drives online on the remaining clusters of the grid to ensure mount points enough for the system to continue operation before setting a cluster in Service mode.

The Control Initiated Reconfiguration (CUIR) helps to alleviate client involvement and simplify the necessary process to start service preparation in a grid member. The CUIR function is enabled from the CLI by using the following command:

```
LIBRARY REQUEST,library-name,CUIR,SETTING,SERVICE,ENABLE
```

R5.2 introduced improvements to the original CUIR and Grid Resiliency functions, building upon them to improve synergy, and obtain better autonomic responses.

For more information, see [IBM TS7700 CUIR User's Guide](#).

Any long-running jobs that used these pending offline devices continue to run to the end of the task. Therefore, the user issues **SWAP** commands to these jobs, which cause them to move to a different logical drive in a different cluster of the grid.

After the cluster maintenance is completed and IBM CSR cancels service for the cluster, the user needed to vary online all the devices again across all LPARs and system plexes. When service-prep is started on the cluster, a Distributed Library Notification is surfaced from the cluster to prompt the attached host to vary off the devices automatically after the following conditions are met:

- ▶ All clusters in the grid must be at 8.41.200.xx or a later microcode level for CUIR Service Vary functions to be available.
- ▶ The z/OS host must include APAR OA52376 with code level V2R2 and above for Service Vary-based CUIR.
- ▶ All clusters in the grid or stand-alone must be at 8.52.100.32 or later for the CUIR Unhealthy Vary function to be available.
- ▶ The z/OS host must include APAR OA60929 with code level V2R3 and later for Unhealthy Vary-based CUIR support.

For more information about CUIR, see Chapter 12, “IBM z/OS host console operations” on page 625.

Important: Forcing Service Mode causes jobs that are mounted or use resources that are provided by the 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 automatically enters 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 in execution, and ensures that no remote mounts are targeting local TVC. When all running tasks end and no other 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 enters in Service mode, the attempt fails.

After service prep completes for one cluster and that cluster enters 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 set in Service mode at the same time within a grid, best practice is having only one cluster in service mode at a time.

When multiple clusters are in service mode simultaneously, they need to be brought back to Normal mode at the same time. Otherwise, the TS7700 does not get to ONLINE state, awaiting until the remaining clusters also leave service mode. Only then, those clusters merge their tokens and rejoin the grid back to ONLINE state.

For a CP0 partition or a disk-only TS7700 cluster in the grid, the Lower Threshold option can be used to reduce the required threshold at which logical volumes are removed from cache before set a peer cluster in service. This option can be useful to prevent reaching maximum cache capacity in CP0 or disk-only cluster, if an extended outage is expected to the cluster being set in service. For more information about the Temporary Removal Threshold, see “Temporary removal threshold” on page 175.

The following operational modes of a cluster (Cluster State) are available:

- ▶ Normal: The cluster is in a normal operation state. Service prep can be started 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. The following options are available:

- ▶ 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 available only 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 the new Cluster State. Click **Cancel** to stop the change process.

9.6.10 Cluster Shutdown window

Use the window that is shown in Figure 9-18 to shut down remotely a TS7700 cluster for a planned power outage or in an emergency.

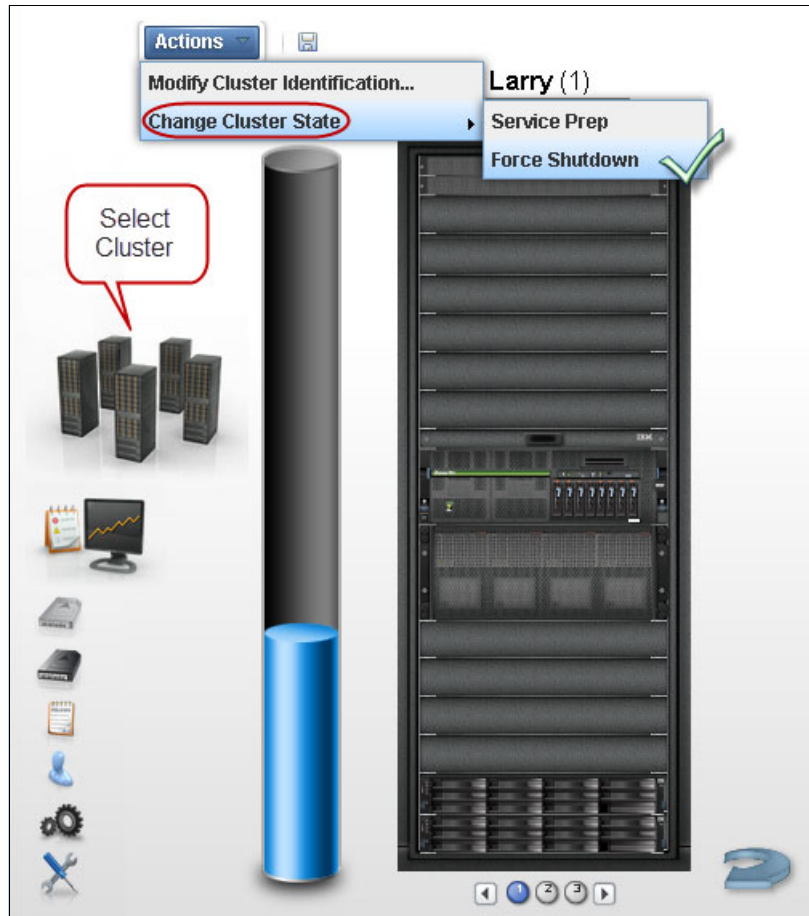


Figure 9-18 MI Cluster: Forcing a cluster to shutdown

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 shown.

Note: After a **shutdown** or **force shutdown** action, the targeted cluster (and associated cache) are powered off. Manual intervention is required onsite (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. Ensure that no jobs or copies are targeting or being sourced from this cluster during the shutdown.

Jobs that use this specific cluster are affected, but copies also are canceled. Eligible data that is not yet 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 initiated from the TS7700 MI also shuts down the cache, powering it off. The cache must be powered on and complete its initialization before the TS7700 server is powered up to resume cluster operation.

If the Shutdown option is selected from the action menu for a cluster that is still online (as shown at the top of Figure 9-18 on page 382), a message alerts the user to put the cluster in service mode first before shutting down, as shown in Figure 9-19.

Note: For normal situations, set the cluster into service mode before shutdown is always recommendable.

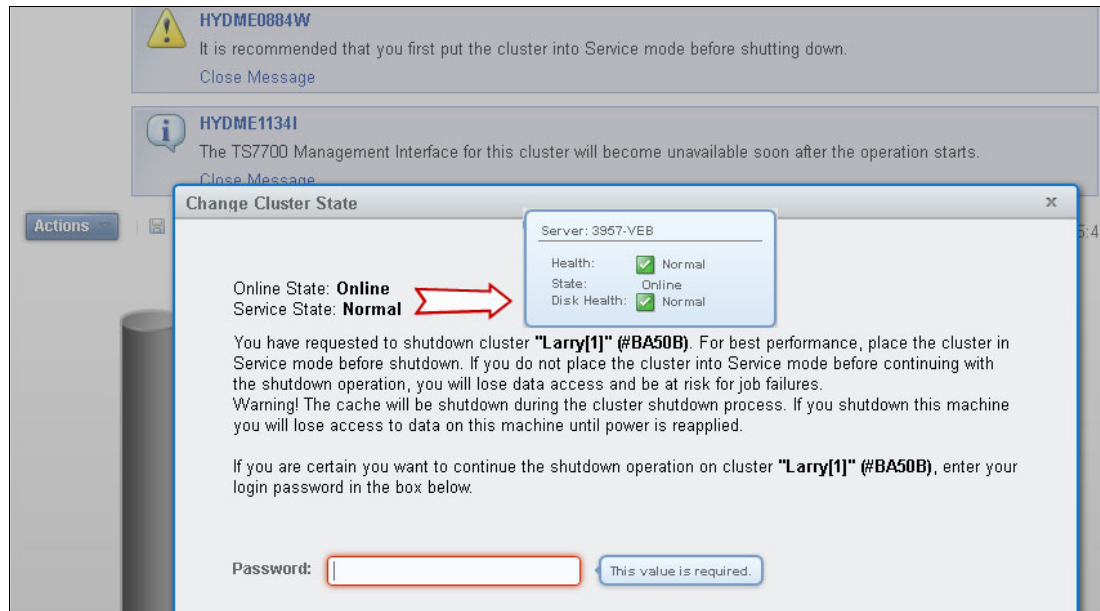


Figure 9-19 Warning message and Cluster Status during forced shutdown

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.

9.6.11 Cluster Identification Properties window

Select this option to view and alter cluster identification properties for the TS7700.

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 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.

9.6.12 Cluster health and detail

The health of the system is checked and updated periodically 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.

As shown in Figure 9-16 on page 376, the Cluster Summary page shows a TS7760T with a TS4500 Tape Library attached. Within the cluster front view page, the cluster badge (shown at top of Figure 9-16 on page 376) brings a general description about the cluster, such as model, name, family, Licensed Internal Code level, cluster description, and cache encryption status.

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 cluster server along with the internal disks (the middle section). The summary of cluster health is shown in the lower-right status bar and in the badge health status (over the frame).

Figure 9-20 shows the back view of the cluster summary window and health details.

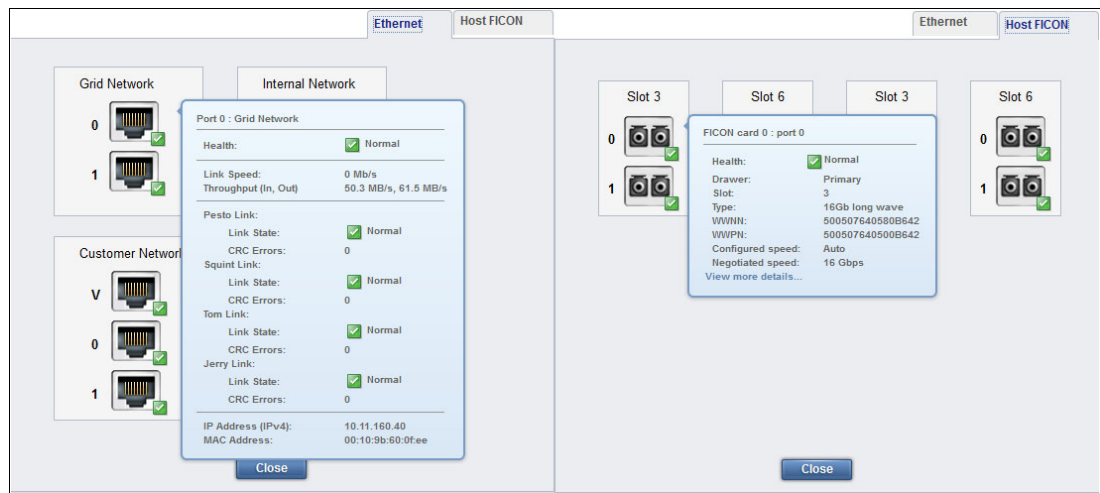


Figure 9-20 Back view of the TS7770-VED cluster summary with health details

The components that are shown 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 the ports that are dedicated to the external (client) network. The assigned IP addresses are also displayed.

Information about the ports is 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 9-20 on page 384. The available information includes the adapter position and general health for each port.

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 TS7700 cluster. To open the expansion frame view, click the small image that corresponds to a specific expansion frame that is below the Actions button.

Tip: The expansion frame icon is displayed only if the accessed cluster includes an expansion frame.

Figure 9-21 shows the Cache Expansion frame details and health view through the MI.

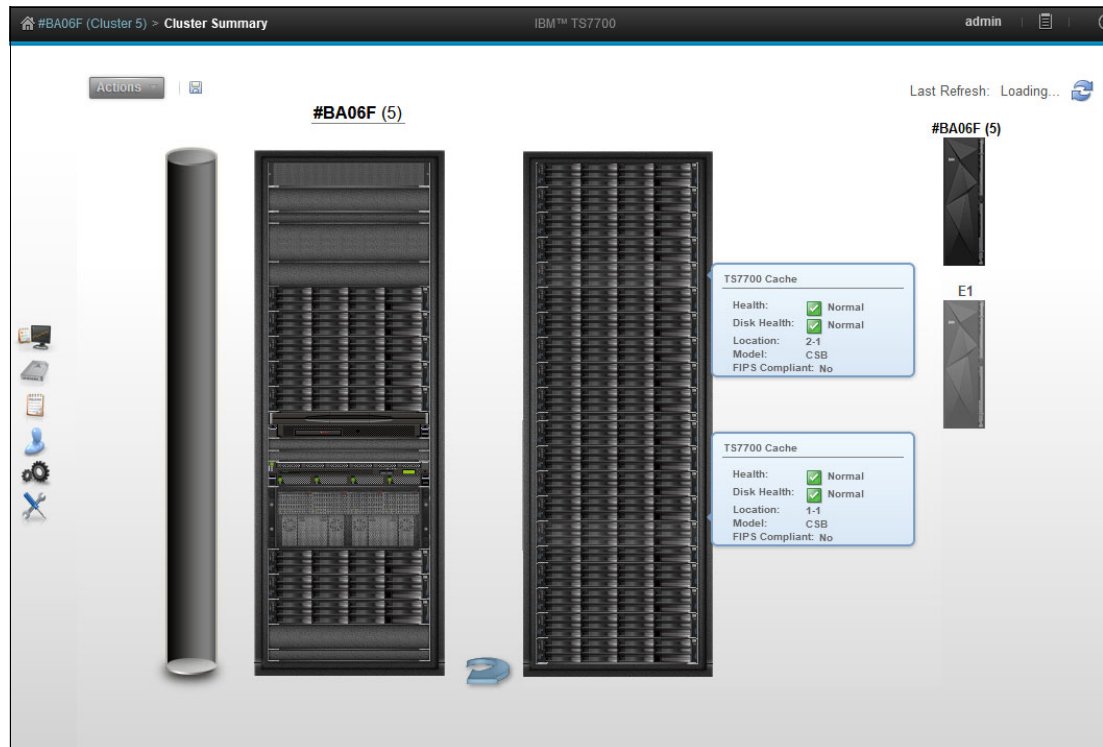


Figure 9-21 TS7700-VED Cache expansion frame details and health

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 9-22 shows a TS4500 Tape Library that is attached to a TS7760 cluster.



Figure 9-22 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 listed in Table 9-4.

Table 9-4 Library health details

Detail	Definition
Physical library type - virtual library name	The type of physical library (type is always TS3500), which is accompanied by the name of the virtual library that is 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

Detail	Definition
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 shows all the specific details about a physical tape drive, such as its serial number, drive type, whether the drive includes a cartridge mount on it, and for what it is mounted. To see the same information (such as drive encryption and tape library location) about the other tape drives, select a specific drive and click **Select Action** → **Details**.

9.7 Monitor icon

By using the items under the Monitor icon in the MI, users can monitor tasks, events, and performance statistics within the TS7700. Figure 9-23 shows the Monitor icon in the TS7700 MI.

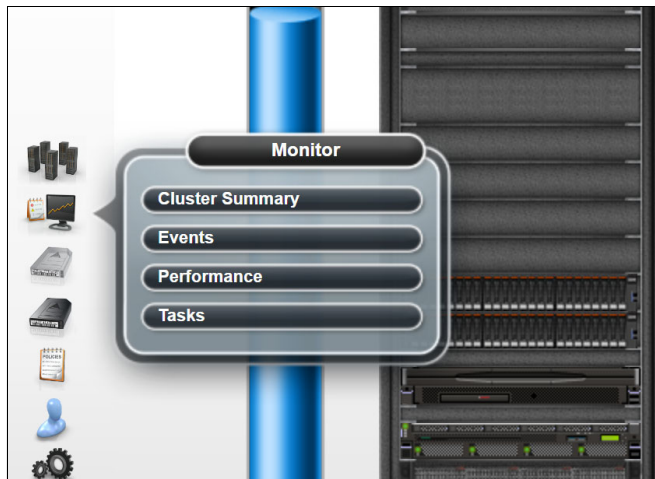


Figure 9-23 Monitor Icon

9.7.1 Events

Use this window that is shown in Figure 9-24 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 malfunction alert. Operator information is displayed on the Events table for 30 days after the operation stops or the event becomes inactive.

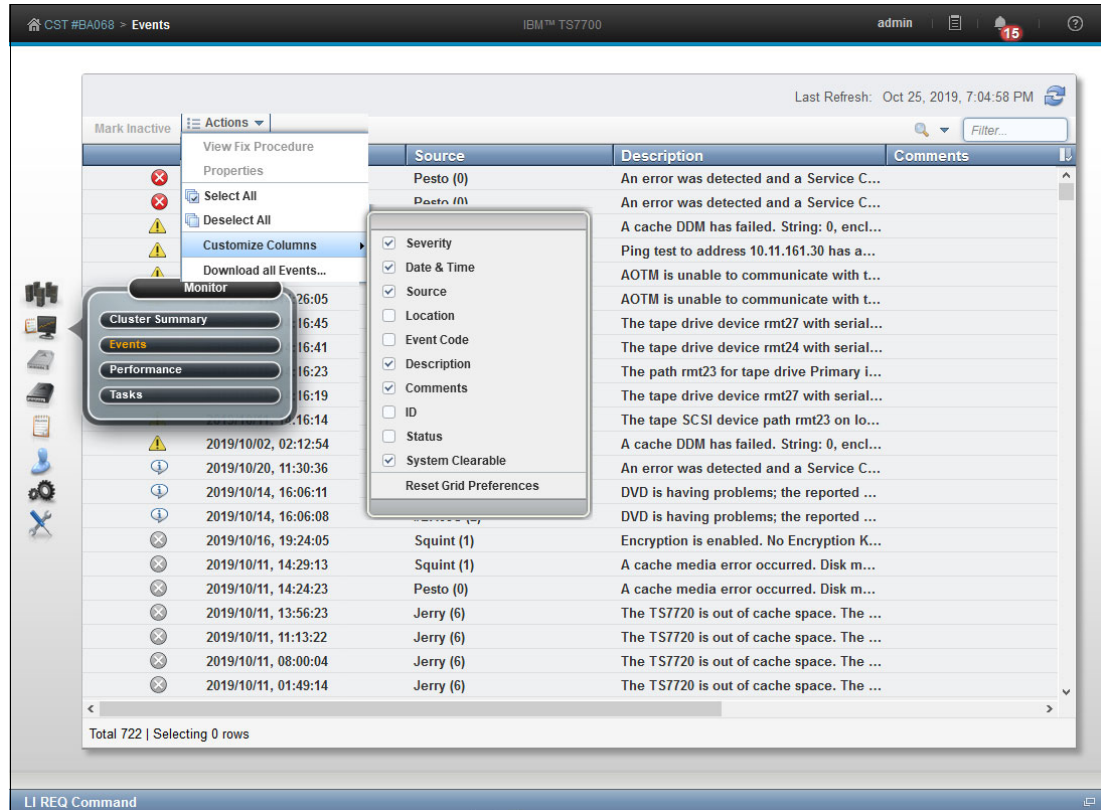


Figure 9-24 TS7700 MI Events window

Note: The timestamps of the events shown in the “Start Time” column of the **Events** page are generated by the TS7700 by using its internal clock, which might or might not be set to an NTP time source. The time zone for these timestamps is always Coordinated Universal Time. The values are sent to the user’s web browser labeled as Coordinated Universal Time, whereupon the web browser converts them to the user’s local time zone based on the user’s operating system settings.

In contrast, the “Last Refresh” time value in the upper right of the screen is generated locally at the user’s workstation, based on last time that the page was reloaded. This will always be in the user’s local time zone and reflects the current time as known by the user’s workstation.

This may cause discrepancies or confusion if the TS7700 internal clock and the user’s workstation clock are not synchronized to the same time, which can be corrected by having both the TS7700 and the user’s workstation configured to use the same or equivalent NTP time sources.

The Event window can be customized to meet your needs. Select which columns to show in the Events window in the selection box that is shown in Figure 9-24 on page 388.

For more information about the Events page, see IBM Documentation for TS7700, which is available locally by clicking the question mark symbol at the right of the banner on the TS7700MI, or at this [web page](#).

9.7.2 Performance

This section presents information for viewing IBM TS7700 Grid and Cluster performance and statistics.

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 move back up to 90 days, which covers three months of operations.

For more information about achieving peak performance, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

Historical Summary

Figure 9-25 shows the Throughput View for the Historical Summary in **Monitor** → **Performance** MI operation in a tape-attached cluster. The performance data for a specific TS7700 cluster can be viewed in this page.

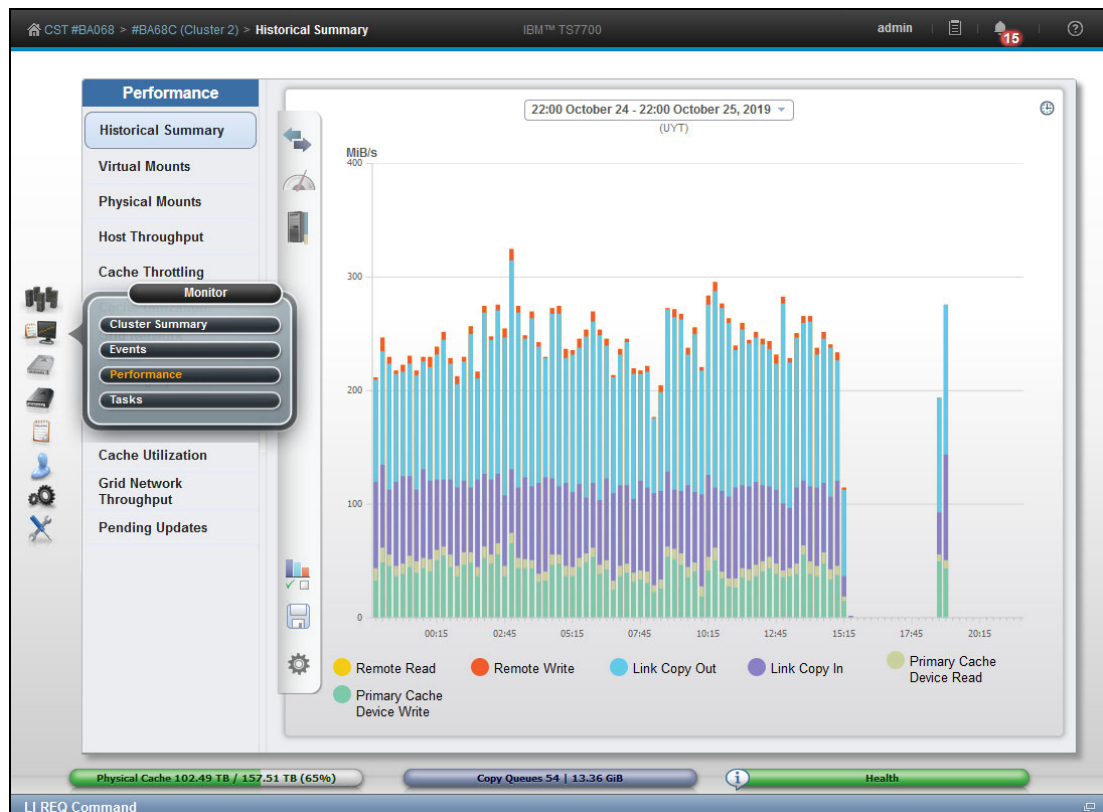


Figure 9-25 Performance window operation, throughput view

Figure 9-26 shows the Performance Historical Summary and related chart selections that are available for this item.

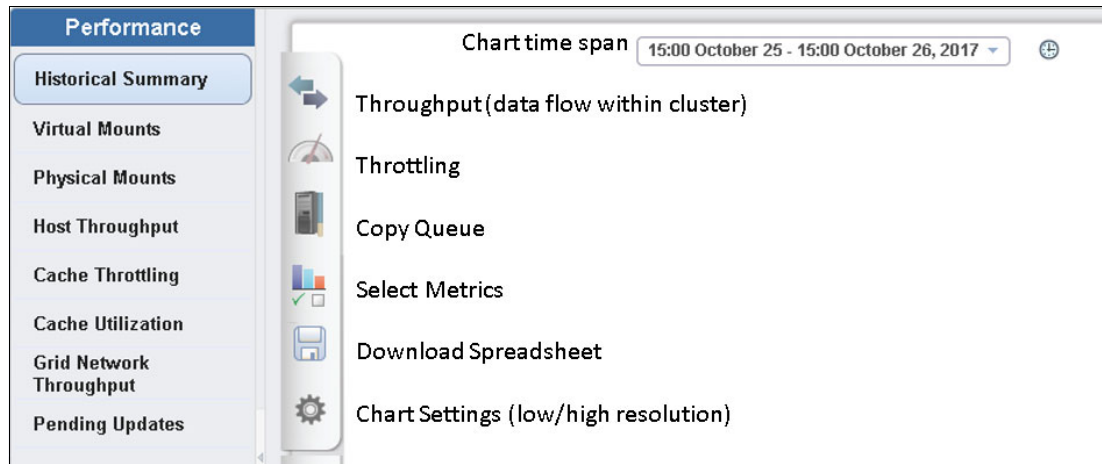


Figure 9-26 Performance options and chart selections

Figure 9-27 show samples of graphs that use the options that are shown in Figure 9-26.

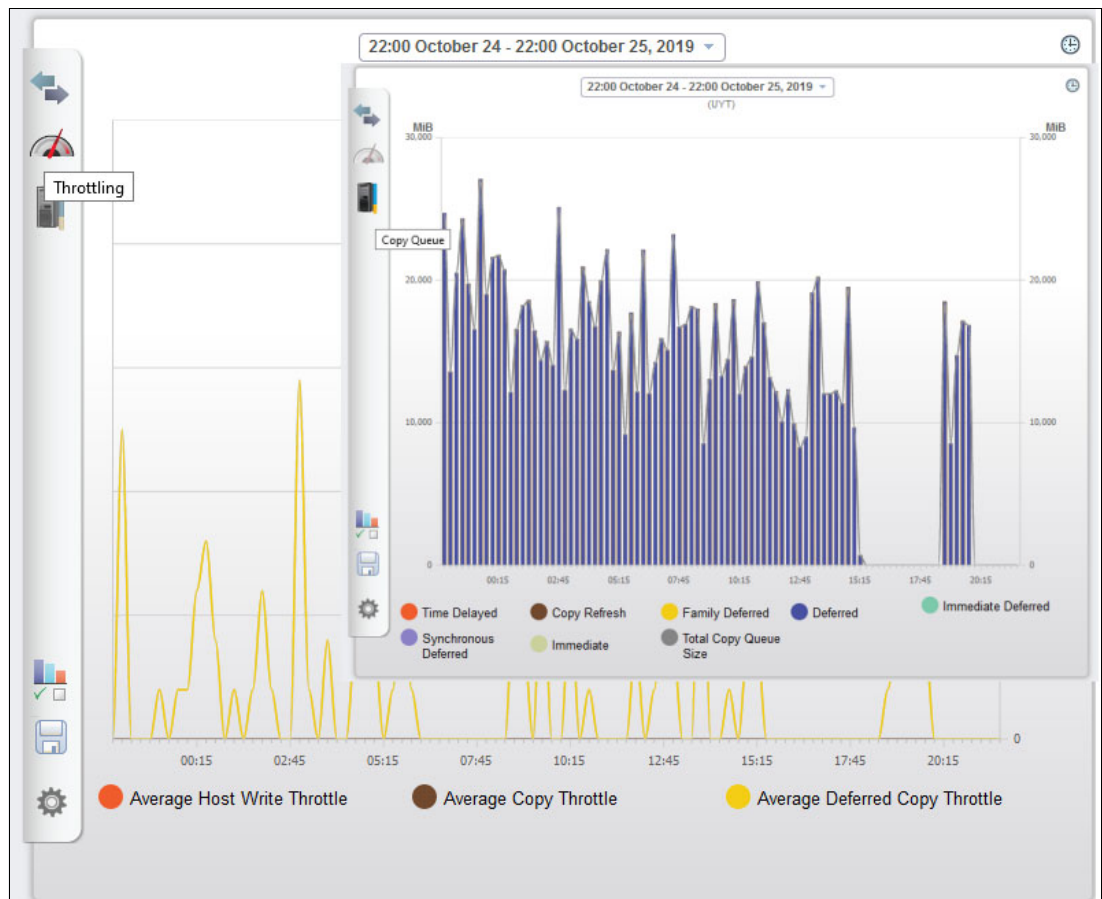


Figure 9-27 Example of Throttling and Copy Queue graphs from the Historical Summary page

Table 9-5 lists the chart elements that can be viewed or changed on this page.

Table 9-5 Elements of the Historical Summary chart

Element	Description
Y axes	The left vertical axis measures throughput (MiBps) or copies (MiB), depending on the selected data sets. The right vertical axis measures the number of mounted virtual drives, reclaimed physical volumes, and data set size to copy. Possible measurements include MiB, GiBs, milliseconds, seconds, hours, and percentage.
X axis	The horizontal axis measures in hours the period for the data sets that are shown. Measurements are shown in 15-minute increments by default. Click the time span (at the top-center of the chart) to change the display increments. The following values are available: <ul style="list-style-type: none"> ▶ 1 day ▶ 12 hours ▶ 1 hour ▶ 30 minutes ▶ 15 minutes ▶ Custom
Last 24 hours	Click the icon in the upper right corner of the page to reset the time span that is shown on the chart to the past 24-hour period. A change to the time span does not alter the configuration of data sets that are displayed in the main area of the chart.
Data sets	Data sets that are displayed in the main area of the chart are shown as lines or stacked columns. Data sets that are related to throughput and copy queues can be grouped to better show relationships between these sets. For more information about all of the data sets, see Table 9-6 on page 393.
Legend	The area below the X axis lists all data sets that are selected to display on the chart with their identifying colors or patterns. The legend displays a maximum of 10 data sets. Click any data set shown in this area to toggle its appearance on the chart. Note: To remove a data set from the legend, clear it by using the Select metrics option.
Time span	The period from which the displayed data is drawn. This range is shown at the top of the page. Note: Dates and times that are displayed reflect the time zone in which your browser is located. If your local time is not available, these values are shown in Coordinated Universal Time. Click the displayed time span to modify its start or end values. Time can be selected in 15-minute increments: <ul style="list-style-type: none"> ▶ Start date and time: The default start value is 24 hours before the present date and time. You can select any start date and time within the 90 days that precede the present date and time. ▶ End date and time: The default end value is 24 hours after the start date or the last valid date and time within a 24-hour period. The end date and time cannot be later than the current date and time. You can select any end date and time that is between 15 minutes and 24 hours later than the start value.

Element	Description
Presets	<p>Click one of the Preset buttons at the top of the vertical toolbar to populate the chart by using one of three common configurations:</p> <p>Throughput: Data sets in this configuration include the following information:</p> <ul style="list-style-type: none"> ▶ Remote Read ▶ Remote Write ▶ Recall from Tape ▶ Write to Tape ▶ Link Copy Out ▶ Link Copy In ▶ Primary Cache Device Read <p>Throttling: Data sets in this configuration include the following information:</p> <ul style="list-style-type: none"> ▶ Average Host Write Throttle ▶ Average Copy Throttle ▶ Average Deferred Copy Throttle <p>Copy queue: Data sets in this configuration include Copy Queue Size.</p> <p>The established time span is not changed when a preset configuration is applied.</p> <p>Note: The preset options that are available depend on the configuration of the accessing cluster. For more information about restrictions, see Table 9-6 on page 393.</p>
Select metrics	<p>Click the Select metrics button on the vertical toolbar to add or remove data sets that are displayed on the Historical Summary chart. For more information about all data sets, see Table 9-6 on page 393.</p>
Download spreadsheet	<p>Click the Download spreadsheet button on the vertical toolbar to download a comma-separated (.csv) file to your web browser for the period shown on the graph. Time is shown in 15-minute intervals in the .csv file.</p> <p>Note: The time that is reported in the .csv file is shown in Coordinated Universal Time. You might find time differences if the system that you use to access the Management Interface is configured for a different time zone.</p>
Chart settings	<p>Click the Chart settings button on the vertical toolbar to enable the low graphics mode for improved performance when many data points are displayed. Low graphics mode disables hover-over tool tips and improves chart performance in older browsers. If cookies are enabled on your browser, this setting is retained when you exit the browser.</p> <p>Note: Low graphics mode is enabled by default when the browser is Internet Explorer, version 8 or earlier.</p>

Click the **Select metrics** button to open the Select metrics window to add or remove data sets displayed on the Historical Summary chart.

The Select metrics window organizes data sets by sections and categories.

The user can select up to 10 data sets, as listed in Table 9-6, to display on the Historical Summary chart.

Table 9-6 Data set descriptions

Metrics section	Metrics category	Data set	Description
Throughput	I/O	Channel R/W MiBps	Transfer rate (MiBps) of host data on the FICON channel, which includes the following information: <ul style="list-style-type: none"> ▶ Host raw read: Rate that is read between the HBA and host. ▶ Host raw write: Rate that is written to the virtual drive from the host.
Throughput	I/O	Primary Cache Read	Data transfer rate (MiBps) read between the virtual drive and HBA for the primary cache repository.
Throughput	I/O	Primary Cache Write	Data transfer rate (MiBps) written to the primary cache repository from the host through the HBA.
Throughput	I/O	Remote Read	Data transfer rate (MiBps) to the cache of the accessing cluster from the cache of a remote cluster as part of a remote write operation. This data set is visible only when the accessing cluster is part of a grid.
Throughput	I/O	Remote Write	Data transfer rate (MiBps) to the cache of a remote cluster from the cache of the accessing cluster as part of a remote read operation. This data set is visible only if the access cluster is part of a grid.
Throughput	Copies	Link Copy Out	Data transfer rate (MiBps) for operations that copy data from the accessing cluster to one or more remote clusters. This is data that is transferred between legacy TS7700 Grid links. This data set is visible only if the access cluster is part of a grid.
Throughput	Copies	Link Copy In	Data transfer rate (MiBps) for operations that copy data from one or more remote clusters to the accessing cluster. This is data that is transferred between legacy TS7700 Grid links. This data set is visible only if the access cluster is part of a grid.
Throughput	Copies	Copy Queue Size	The maximum size of the incoming copy queue for the accessing cluster, which is shown in MiBs, GiBs, or TiBs. The following incoming copy queue options are available: <ul style="list-style-type: none"> ▶ Immediate ▶ Synchronous-deferred ▶ Immediate-deferred ▶ Deferred ▶ Family deferred ▶ Copy refresh ▶ Time delayed ▶ Total This data set is visible only if the accessing cluster is part of a grid.
Throughput	Copies	Average Copy Life Span	The average age of virtual volumes to be copied to the distributed library for the accessing cluster. The following options are available: <ul style="list-style-type: none"> ▶ Immediate Mode Copy ▶ Time Delayed Copy ▶ All other deferred type copies This data set is visible only if the accessing cluster is part of a grid.

Metrics section	Metrics category	Data set	Description
Storage	Cache	Cache to Copy	<p>The number of GiBs that are in the incoming copy queue of a remote cluster, but are destined for the accessing cluster. This value is the amount of data that is being held in the cache until a copy can be made.</p> <p>This data set is visible only if the accessing cluster is part of a grid.</p>
Storage	Cache	Cache Hit	<p>The number of completed mount requests where data is resident in the TVC.</p> <p>If two distributed library access points are used to satisfy a mount with synchronous mode copy enabled, this count is advanced only when the data is resident in the TVC for both access points. For this reason, this data set is visible only if the accessing cluster is a TS7700 Tape Attach, or is part of a grid that contains a TS7700 Tape Attach.</p>
Storage	Cache	Cache Miss	<p>The number of completed mount requests where data is recalled from a physical stacked volume.</p> <p>If two distributed library access points are used to satisfy a mount with synchronous mode copy enabled, this count is advanced when the data is not resident in the TVC for at least one of the two access points. For this reason, this data set is visible only if the accessing cluster is a TS7700 Tape Attach, or is part of a grid that contains a TS7700 Tape Attach.</p>
Storage	Cache	Cache Hit Mount Time	<p>The average time (ms) to complete Cache Hit mounts. This data set is visible only if the accessing cluster is attached to a tape library. If the cache is partitioned, this value is displayed according to partition.</p>
Storage	Cache	Cache Miss Mount Time	<p>The average time (ms) to complete Cache Miss mounts.</p> <p>This data set is visible only if the accessing cluster is attached to a tape library. If the cache is partitioned, this value is displayed according to partition.</p>
Storage	Cache	Primary Cache Total Physical Used	<p>The amount of used cache in a primary cache. If the cache is partitioned, this value is the total physical cache that is used in all partitions.</p>
Storage	Cache	Partitions	<p>If the accessing cluster is a TS7700 attached to a tape library, a numbered tab exists for each active partition. Each tab displays check boxes for the following categories:</p> <ul style="list-style-type: none"> ▶ Cache Hit ▶ Cache Miss ▶ Mount Time Hit ▶ Mount Time Miss ▶ Data in Cache
		Primary Used	<p>The amount of used cache in a partitioned, primary cache, according to partition.</p> <p>This data set is only visible if the selected cluster is a TS7700T.</p>

Metrics section	Metrics category	Data set	Description
Storage	Cache	Data Waiting for Premigration	The amount of data in the cache that is assigned to volumes that are waiting for premigration. This data set is visible only if the selected cluster is a TS7700T.
Storage	Cache	Data Migrated	The amount of data in the cache that was migrated. This data set is visible only if the selected cluster is a TS7700T.
Storage	Cache	Data Waiting for Delayed Premigration	The amount of data in the cache that is assigned to volumes waiting for delayed premigration. This data set is visible only if the selected cluster is a TS7700T.
Storage	Virtual Tape	Maximum Virtual Drives Mounted	The greatest number of mounted virtual drives. This value is a mount count.
Storage	Physical Tape	Write to Tape	Data transfer rate (MiBps) that is written to physical media from cache. This value typically represents premigration to tape. This data set is not visible when the selected cluster is not attached to a library.
Storage	Physical Tape	Recall from Tape	Data transfer rate (MiBps) that is read from physical media to cache. This value is recalled data. This data set is not visible when the selected cluster is not attached to a library.
Storage	Physical Tape	Reclaim Mounts	Number of physical mounts that are completed by the library for the physical volume reclaim cache operation. This value is a mount count. This data set is not visible when the selected cluster is not attached to a library.
Storage	Physical Tape	Recall Mounts	Number of physical mounts that are completed by the library for the physical volume reclaim operation. This data set is not visible when the selected cluster is not attached to a library.
Storage	Physical Tape	Premigration Mounts	Number of physical mount requests that are completed by the library that are required to satisfy pre-migrate mounts. This data set is not visible when the selected cluster is not attached to a library.
Storage	Physical Tape	Physical Drives Mounted	The maximum, minimum, or average number of physical devices of all device types that are concurrently mounted. The average number displays only when you hover over a data point. This data set is visible only when the selected cluster attaches to a library.

Metrics section	Metrics category	Data set	Description
Storage	Physical Tape	Physical Mount Times	<p>The maximum, minimum, or average number of seconds required to complete the execution of a mount request for a physical device. The average number displays only when you hover over a data point.</p> <p>This data set is visible only when the selected cluster attaches to a library.</p>
System	Throttling	Average Copy Throttle	<p>The average time delay as a result of copy throttling, which is measured in milliseconds. This data set contains the averages of nonzero throttling values where copying is the predominant reason for throttling.</p> <p>This data set is visible only if the selected cluster is part of a grid.</p>
System	Throttling	Average Deferred Copy Throttle	<p>The average time delay as a result of deferred copy throttling, which is measured in milliseconds. This data set contains the averages of 30-second intervals of the deferred copy throttle value.</p> <p>This data set is visible only if the selected cluster is part of a grid.</p>
System	Throttling	Average Host Write Throttle for Tape Attached Partitions	<p>The average write overrun throttle delay for the tape-attached partitions. This data set is the average of the nonzero throttling values where write overrun was the predominant reason for throttling.</p> <p>This data set is visible only if the selected cluster is a TS7700 attached to a tape library.</p>
System	Throttling	Average Copy Throttle for Tape Attached Partitions	<p>The average copy throttle delay for the tape-attached partitions. The value that is presented is the average of the nonzero throttling values where copy was the predominant reason for throttling.</p> <p>This data set is visible only if the selected cluster is a TS7720 or TS7760 attached to a tape library.</p>
System	Throttling	Average Deferred Copy Throttle for Tape Attached Partitions	<p>The average deferred copy throttle delay for the tape-attached partitions. This value is the average of 30-second intervals of the deferred copy throttle value during the historical record.</p> <p>This data set is visible only if the selected cluster is part of a grid and is a TS7700 attached to a tape library.</p>
System	Utilization	Maximum CPU Primary Server	The maximum percentage of processor use for the primary TS7700 server.
System	Utilization	Maximum Disk I/O Usage Primary Server	The maximum percentage of disk cache I/O uses as reported by the primary server in a TS7700.

Refer to Chapter 13, “Monitoring” on page 667 for more information on how to use the graphs described in this section.

For more information about the window and available settings, see TS7700 R5.4 IBM Documentation:

- To view TS7700 R5.4 IBM Documentation locally on the TS7700 MI, click the question mark icon in the upper right corner of the window.

- ▶ To view TS7700 R5.4 IBM documentation on the web, see [Performance](#).

Also, the paper *IBM Virtualization Engine 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. In addition, it explains throttling mechanisms and available tuning options for the subsystem to achieve peak performance. Available online at [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

Virtual mounts

This page is used to view the virtual mount statistics for the TS7700 Grid. Virtual mount statistics are displayed for activity on each cluster during the previous 15 minutes. These statistics are presented in bar graphs and tables and are organized according to the following sections:

- ▶ Number of virtual mounts

This section provides statistics for the number of virtual mounts on a specific cluster during the most recent 15-minute snapshot. Snapshots are taken at 15-minute intervals. Each numeric value represents the sum of values for all active partitions in the cluster. The following information is displayed:

- Cluster: The cluster name.
- Fast-Ready: The number of virtual mounts that were completed by using the Fast-Ready method.
- Cache Hits: The number of virtual mounts that were completed from cache.
- Cache Misses: The number of mount requests that are unable to be fulfilled from cache.

Note: This field is visible only if the selected cluster possesses a physical library.

- Total: Total number of virtual mounts.

- ▶ Average mount times

This section provides statistics for the average mount times in milliseconds (ms) on a specific cluster during the most recent 15-minute snapshot. Snapshots are taken at 15-minute intervals. Each numeric value represents the average of values for all active partitions in the cluster. The following information is displayed:

- Cluster: The cluster name.
- Fast-Ready: The average mount time for virtual mounts that were completed by using the Fast-Ready method.
- Cache Hits: The average mount time for virtual mounts that were completed from cache.
- Cache Misses: The average mount time for mount requests that cannot be fulfilled from cache.

Note: This field is visible only if the selected cluster possesses a physical library.

Physical mounts

The physical mounts statistics for the last 15 minutes of activity are displayed in bar graphs and table format per cluster: One for number of mounts by category and one for average mount time 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 page is not visible on the TS7700 MI if the grid does not possess a physical library (no tape-attached member).

The following information is available on this page:

- ▶ Cluster: The cluster name
- ▶ Pre-Migrate: The number of pre-migrate mounts
- ▶ Reclaim: The number of reclaimed mounts
- ▶ Recall: The number of recalled mounts
- ▶ Secure Data Erase: The number of secure data erase mounts
- ▶ Total: The total number of physical mounts
- ▶ Mount Time: The average mount time for physical mounts

Host throughput

Use the page shown in Figure 9-28 on page 399 to view statistics for each cluster, vnode, host adapter, and host adapter port in the grid. At the top of the page is a collapsible tree that allows you to view statistics for a specific level of the grid and cluster. The following links are available:

- ▶ Grid hyperlink: Displays information for each cluster.
- ▶ Cluster hyperlink: Displays information for each vnode.
- ▶ vnode hyperlink: Displays information for each host adapter.
- ▶ Host adapter link: Displays information for each of its ports.

The host throughput data is displayed in two bar graphs and one table. The bar graphs represent 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 following types of data are available. The letters (A, B, C, D) correspond to the data flow shown in Figure 9-28 on page 399:

- ▶ Cluster/vnode/Host Adapter/Host Adapter Port: Cluster or cluster component for which data is being displayed
- ▶ Compressed Read (A): Amount of data that is read between the virtual drive and HBA
- ▶ Raw Read (B): Amount of data that is read between the HBA and host
- ▶ Read Compression Ratio: Ratio of raw data read to compressed data read
- ▶ Compressed Write (D): The 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 raw data written to compressed data written

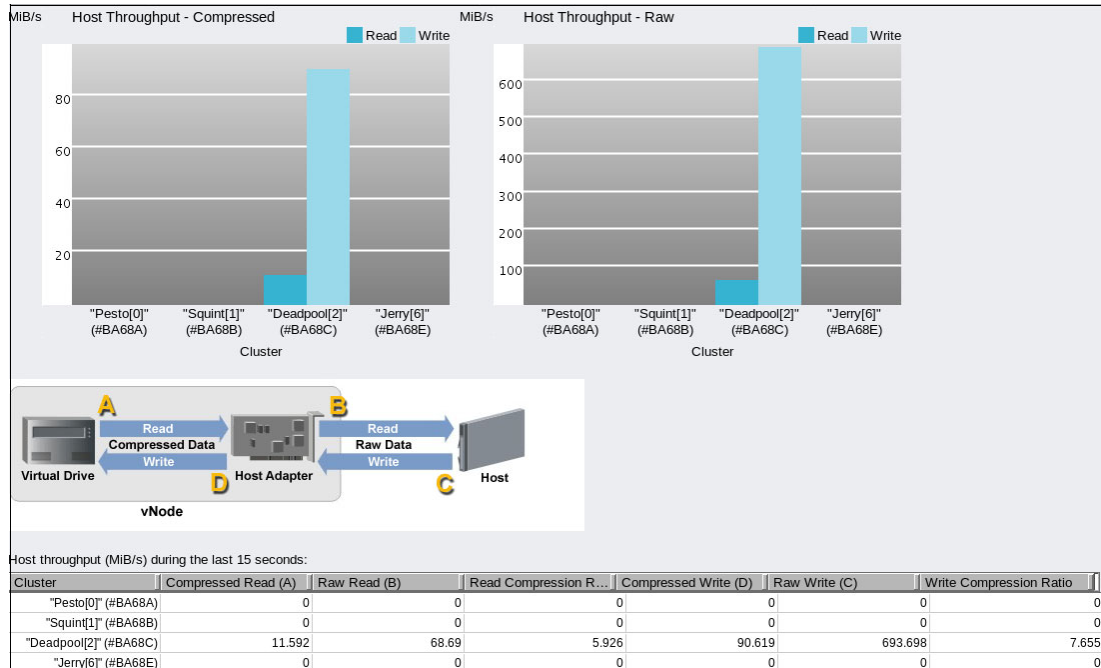


Figure 9-28 Host Throughput page

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.

Throttling refers to the intentional slowing of data movement to balance or [prioritize](#) system resources in a busy TS7700. Throttling can be applied to host write and inbound copy operations. Throttling of host write and inbound copy operations limits the amount of data movement into a cluster. Throttling is typically done for one of the following reasons:

- ▶ The amount of unused cache space is low.
- ▶ The amount of data in the cache that is queued for premigration exceeded a threshold.

Host write operations can also be throttled when RUN copies are being used and it is determined that a throttle is needed to prevent pending RUN copies from changing to the immediate-deferred state. A throttle can be applied to a host write operation for the following reasons:

- ▶ The amount of unused cache space is low
- ▶ The amount of data in the cache that must be pre-migrated is high
- ▶ For RUN copies, an excessive amount of time is needed to complete an immediate copy

The Cache Throttling graph displays the throttling that is applied to host write operations and to inbound RUN copy operations. The delay represents the time delay, in milliseconds, per 32 KiB of transferred post-compressed data. Each numeric value represents the average of values for all active partitions in the cluster. The following information is shown:

- ▶ **Cluster:** The name of the affected cluster.
- ▶ **Copy:** The average delay, in milliseconds that is applied to inbound copy activities.
- ▶ **Write:** The average delay, in milliseconds that is applied to host write operations locally and from remote clusters.

For more information about achieving peak network performance, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

Cache utilization statistics

Cache utilization statistics are presented for clusters that feature a cache single partition and for clusters with caches with multiple partitions. Models TS7760 or TS7770 disk-only clusters have only one resident partition, which accounts for the entire cache.

Note: *Cache Utilization* page is available for TS7700 Disk-only or Object store cluster. Page not visible for TS7700 attached to a cloud or tape library).

Figure 9-29 shows an example of a multiple partitioned Cache Utilization window; as in, a TS7770 VED with Advanced Object Store (FC 5283) that shows usage by logical volumes and objects.

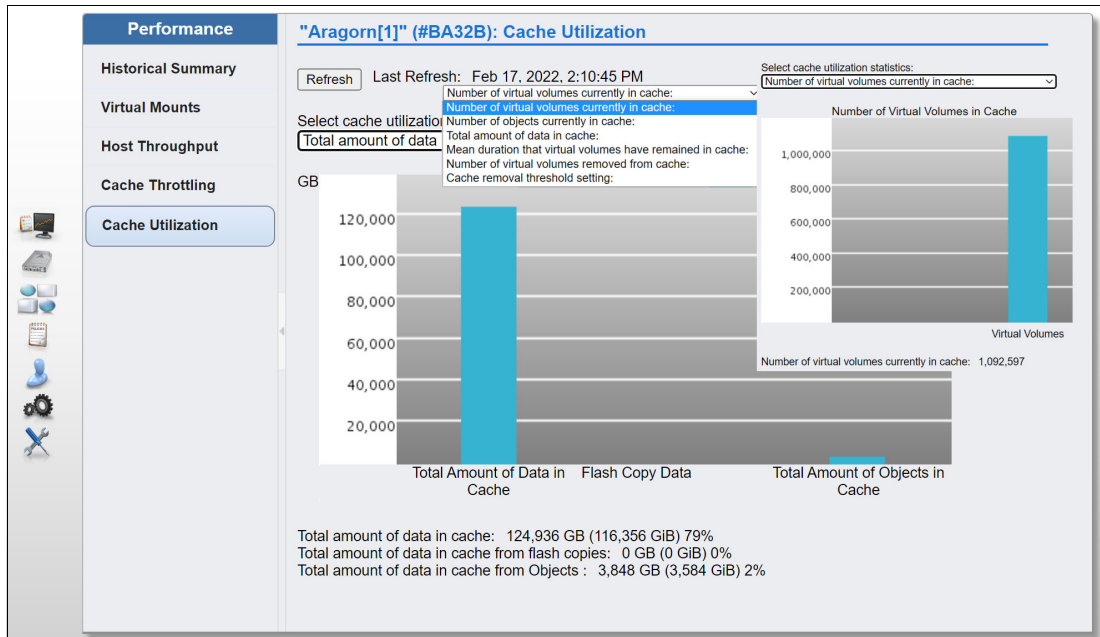


Figure 9-29 TS7700 Cache Utilization window

Cache Partition

The Cache Partition window presents the cache use statistics for the TS7700T or TS7700C models, in which the cache is made up of multiple partitions. Figure 9-30 on page 401 shows a sample of the Cache Partitions (multiple partitions) window, which can be accessed by clicking the **Monitor** or **Virtual** icon (both direct to the same window). In this window, the user can display the existent cache partitions, create or reconfigure a partition, or delete a partition as needed.

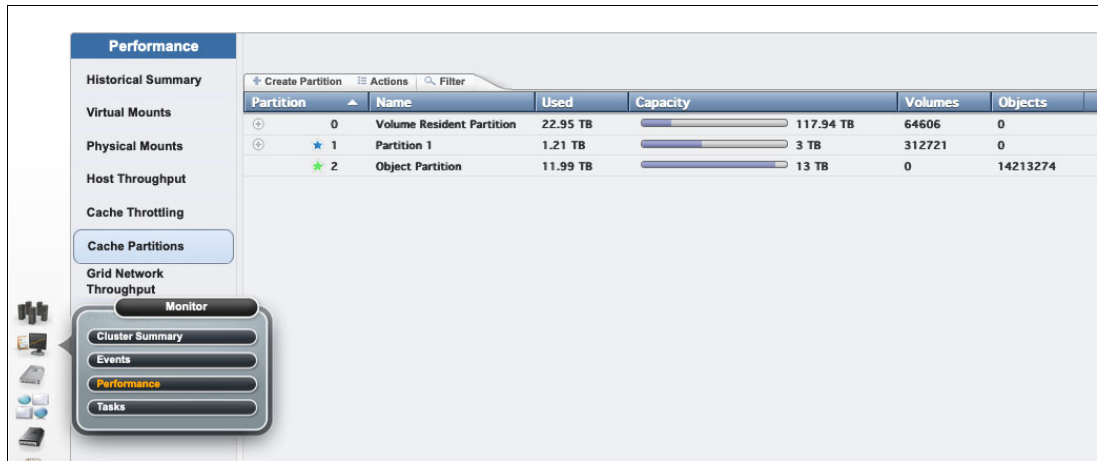


Figure 9-30 Cache Partitions window

Tip: Consider limiting the MI user roles that are allowed to change the partition configurations through this window.

For more information about the window and available settings, see TS7700 R5.4 IBM Documentation:

- ▶ To view TS7700 R5.4 IBM Documentation locally on the TS7700 MI, click the question mark icon in the upper right corner of the window.
- ▶ To view TS7700 R5.4 IBM Documentation on the web, see [Management interface help](#).

Grid Network Throughput

This window is available only 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 in the bar graph chart and tables. The following information is shown:

- ▶ Cluster: The name of the cluster
- ▶ 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
- ▶ Total: Total data transfer rate for the cluster

For more information about data flow within the grid and how those numbers vary during the operation, see in Chapter 13, “Monitoring” on page 667.

Pending Updates

The Pending Updates window is available only if the TS7700 cluster is a member of a grid. The Pending updates window can be used to monitor the 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 might also 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.

The Download button at 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 that cannot merge became "hot").

Stolen volume describes a volume whose ownership was taken over during a period in which the owner cluster was in service mode or offline, if an unexpected cluster outage occurs when the volume ownership is taken over under an operator's direction, or by using AOTM.

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 more information about this MI function, see [Pending updates](#).

9.7.3 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 at the top of the table; the completed tasks are listed at the bottom. Figure 9-31 on page 403 shows an example of the Tasks window. Notice that the information on this page and the task status pods are of grid scope.

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.

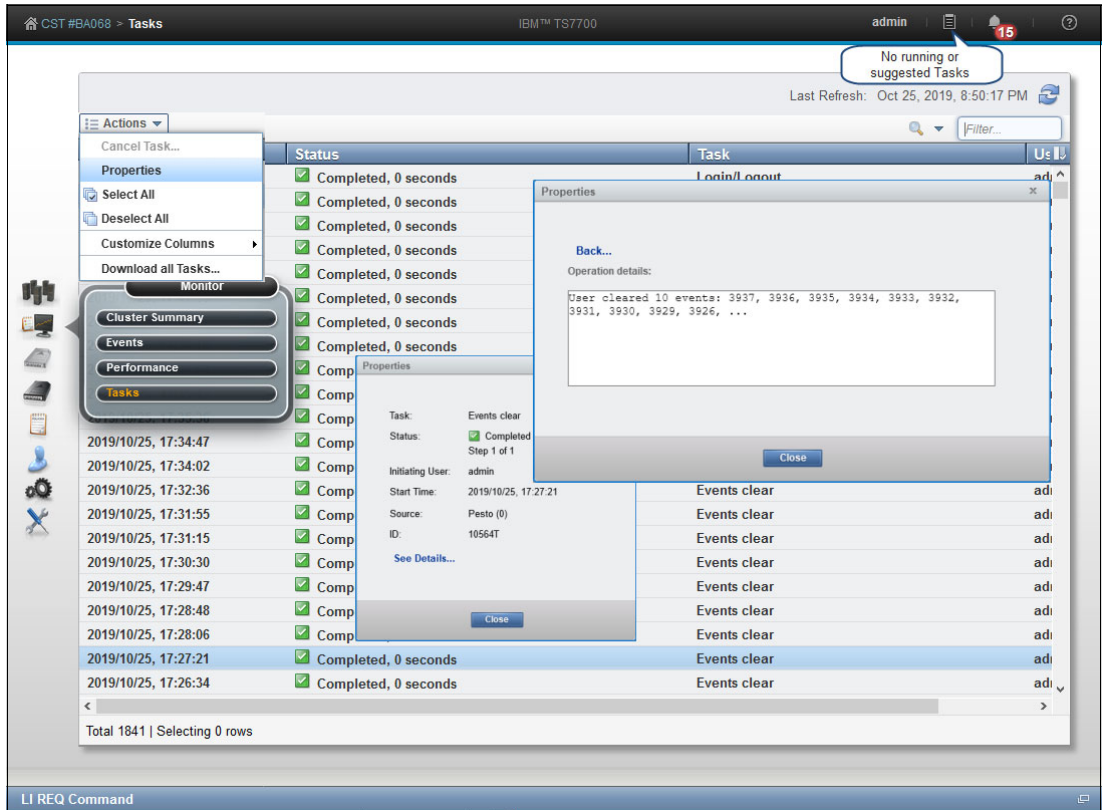


Figure 9-31 Tasks window: displaying the properties of one specific task

9.8 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 categories. Cache Partitions also is under the Virtual icon, which you can use to create, modify, or delete cache partitions.

Figure 9-32 shows the Virtual icon and the available options. The Cache Partitions option is available for the TS7700T, TS7700C, or TS7700 that is enabled with the DS8000 Object Store (FC 5282). The Incoming Copy Queue option appears in grid configurations only.

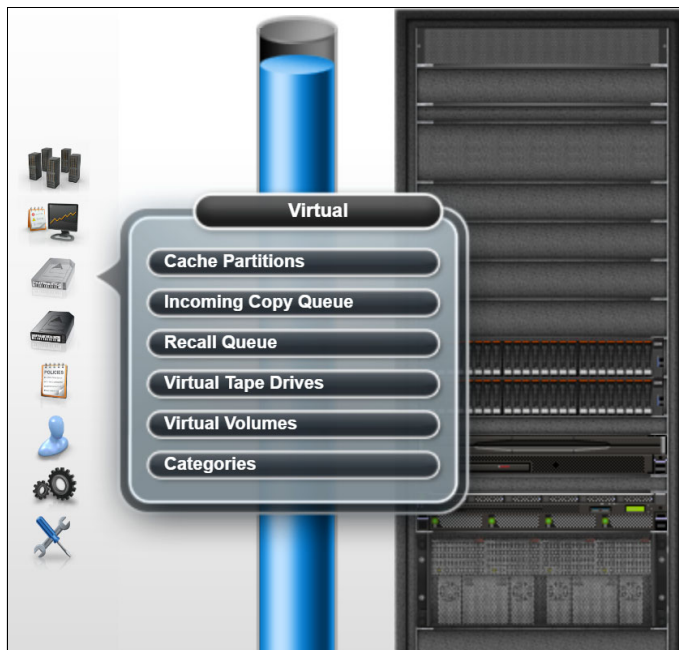


Figure 9-32 Virtual icon and options

The available items under the Virtual icon are described next.

9.8.1 Cache Partitions

In the Cache Partitions window in the MI, you can create a cache partition, or reconfigure or delete a cache partition for the TS7700T (tape-attached, FC 5273 Tape Attach Enablement) or TS7700C (cloud-attached, FC 5278 Cloud Enablement) models. Also, you can use this window to monitor the cache and partitions occupancy and usage.

Figure 9-33 on page 406 shows a sequence for creating a partition. As many as eight partitions can exist, from Resident partition (partition 0) to Partition 7, if needed. The partition-allocated size is subtracted from the resident partition (CPO) capacity if at least more than 2 TB of free space is available in it. For more information about rules and allowed values in effect for this window, see the TS7700 R5.2.2 IBM Documentation.

The TS7700 Transparent Cloud Tiering allows the TS7700 to offload data to public or private clouds. When data is stored in the cloud by a cluster, all clusters in the grid that are cloud-attached enabled can access that object store instance. IBM Cloud Object Storage (on-premises cloud) and AWS S3 (public cloud) are supported. The TS7700 Transparent cloud tiering is enabled through the FC 5278 (Enable Cloud Storage Tier) and it is mutually exclusive with Tape Attach function (FC 5273).

Note: Cloud Enablement requires a minimum of 64 GB of RAM, which is the default in VED models. Use FC 3466 to upgrade VEC memory size.

VED models must have 128 GB of RAM before upgraded to R5.3

Advanced Object Store for DS8000

TS7700 Advanced Object Store for DS8000 (FC 5283) implemented several enhancements to the original DS8K Object Store function, FC 5282.

Note: FC 5282 is no longer available for distribution and is replaced by Feature Code 5283 TS7700 Advanced Object Store. FC 5282 is still supported but can be upgraded to FC 5283 on existing TS7700 systems. For more information on upgrading your FC 5282, contact your IBM Representative.

Advanced Object Store requires DS8900 and TS7770 VED to work, and cannot coexist in the same grid concurrently with the previous function DS8K Object Store (FC 5282).

Consider the following points about Advanced Object Store:

- ▶ All TS7700 object-enabled clusters (FC 5283) are aware of all objects in the grid.
- ▶ All objects can be accessed from any object that is enabled clusters in the grid, whether the cluster has a local copy.
- ▶ Object policy management is available. A user can determine whether object copies exist, where they are within the grid, and if copies are performed synchronously or asynchronously.
- ▶ The grid provides automatic healing of changes during cluster outages.
- ▶ DS8K multi-cloud is supported.

FICON virtual volumes and DS8K objects can coexist in the same TS7770 cluster. The function uses existing cache storage and grid Ethernet adapters.

The Advanced Object Store (FC 5283) is supported on TS7770 VED with FC 8083 only, which installs two SSDs in the VED server for the use of the new function.

Note: FC 5083 is required on the TS7700-VED clusters that are a DS8900 TCT target.

When Advanced Object Store (FC 5283) is enabled, one of the CP1 through CP7 partitions is reserved exclusively for DS8000 objects and can be sized dynamically, as with all other partitions.

Note: New, empty clusters with FC 5283 cannot be joined to a grid that contains FC 5282 on any cluster. Also, FC 5283 cannot be activated on a VED if any other TS7700 in the grid has FC 5282 on it.

For more information about Advanced Object Store, see *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583.

For an example of a Cache Partition page, see Figure 9-38 on page 408. Notice that Objects and Volume counts are reported.

Considerations: No partition can be created if Resident-Only (CP0) has 2 TB or less of free space. Creating partitions is blocked by a FlashCopy for DR in progress, or by one of the existing partitions being in an overcommitted state.

Figure 9-33 shows creating a cache partition.

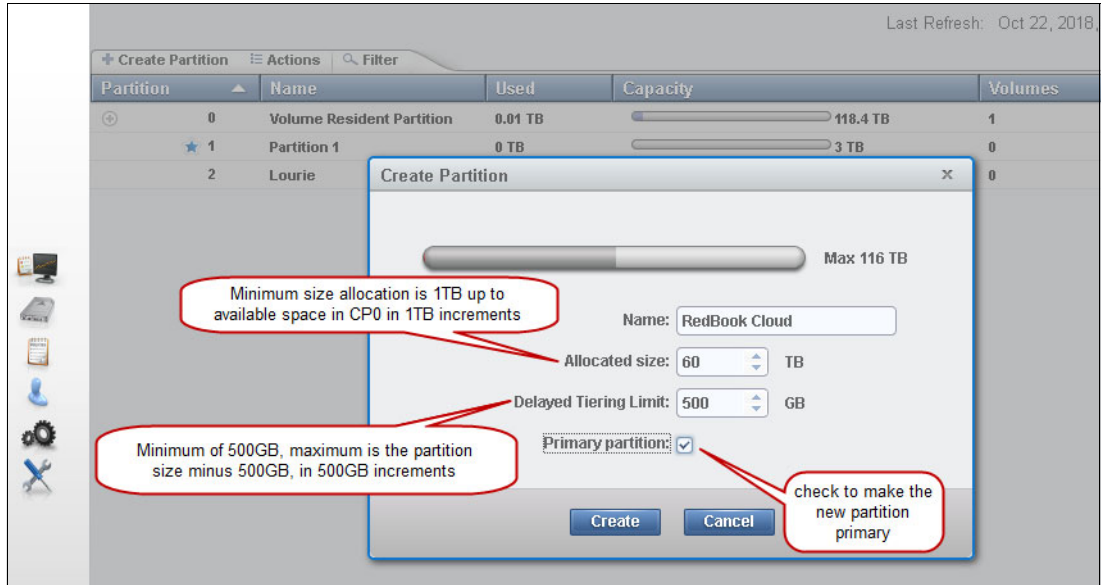


Figure 9-33 Creating a cache partition

Figure 9-34 shows an example of the successful creation in the upper half of the window. The lower half of the window shows an example in which the user failed to observe the amount of free space that was available in CP0.

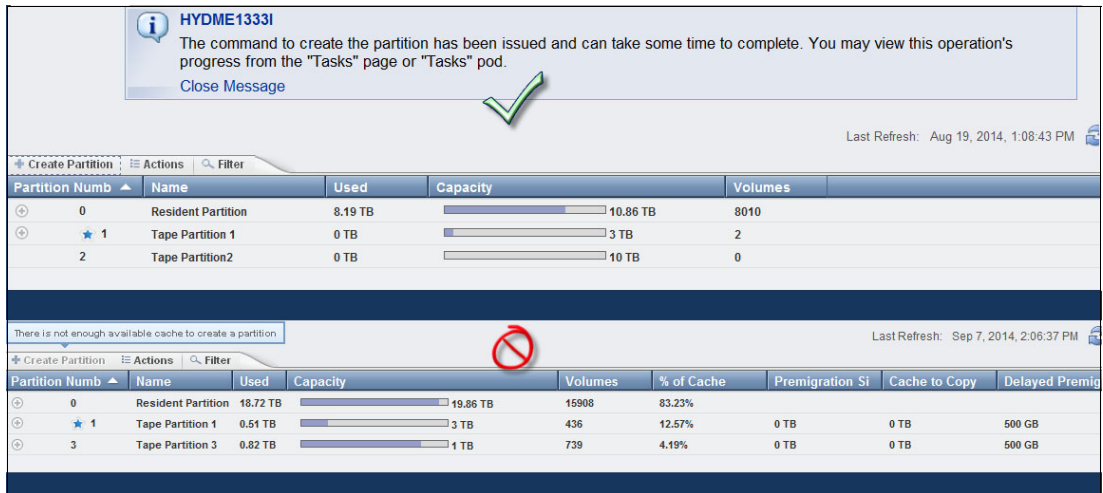


Figure 9-34 Example of a success and a failure to create a new partition

Notice that redefining the size of existing partitions in an operational TS7700T might create unexpected load peak in the overall premigration queue, which causes host write throttling to be applied to the partitions.

For example, consider an example in which a tape-attached partition is downsized, and becomes instantly overcommitted. In this example, the TS7700T 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 9-35 shows the previous scenario, TS7700T operating and 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 9-35 Partition 1 operating with 12 TB cache

Figure 9-36 shows Partition 1 being downsized to 8 TB. Note the initial warning and subsequent overcommit statement that shows up when resizing the 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 9-36 Downsizing Partition 1, and the overcommit warning

Accepting the Overcommit statement starts the resizing action. If this time is not best-suited for the partition resizing (as during the peak load period), the user can click **No** and decline to take the action and resume it at a better time. Figure 9-37 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

Last Refresh: Oct 6, 2014, 8:01:31 AM

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 9-37 Resizing message and resulting cache partitions window

Figure 9-36 on page 407 and Figure 9-37 on page 407 show an example of downsizing a partition in a TS7700T (tape attach), but the same behavior is expected if it was a TS7700C (cloud attach) partition instead. Partitions function the same for TS7700C and TS7700T configurations.

Note: The Tape Attach (FC 5273) and Cloud Storage Tier (FC 5278) functions are mutually exclusive within a cluster. Nonetheless, a tape attach and a cloud attach cluster can be part of the same grid. DS8000 Object Store (FC 5282) can exist in a tape-attach, cloud-attach, or disk-only VEC/VED cluster.

In all cases, CP0 remains dedicated for resident logical volumes only and does not support direct movement of the resident volumes to or from an object store or tape storage.

Note: DS8000 objects are in an object cache partition (only one is allowed) and cannot be migrated to physical tapes or cloud. With R5.2.2 and Advanced Object Store, objects can be replicated to other clusters in the grid under policies for object management.

Figure 9-38 shows an example of the Advanced Object Store object partition (CP1 in the example) in a cloud-enabled cluster (FC 5283).

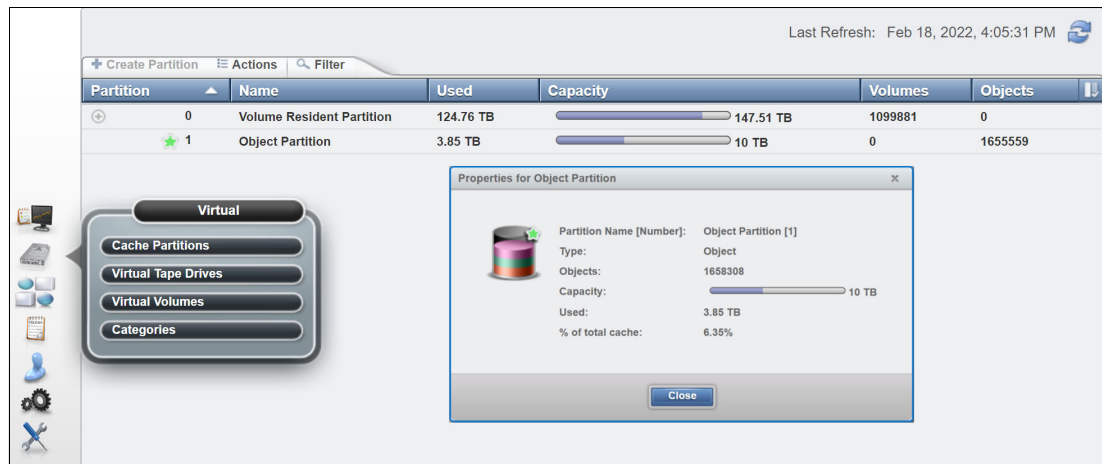


Figure 9-38 Advanced Object Store Partition example in a cloud-attach VED

For more information about Advanced Object Store, see *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583

9.8.2 Incoming Copy Queue window

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.

The user can specify which clusters (if any) copies reside on 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 9-7 lists the values that are displayed in the copy queue table.

Table 9-7 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 can be copied based on their creation time or last access time. ▶ Copy-refresh: Volumes can be in this queue if the MC that is assigned to the volumes changed and a LI REQ command was sent from the host to start 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 cluster 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 shown only when View by Last TVC is selected.</p>
Size	<p>Total size of the queue, which is displayed in GiB.</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 the queue for each type.</p>

Figure 9-39 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.

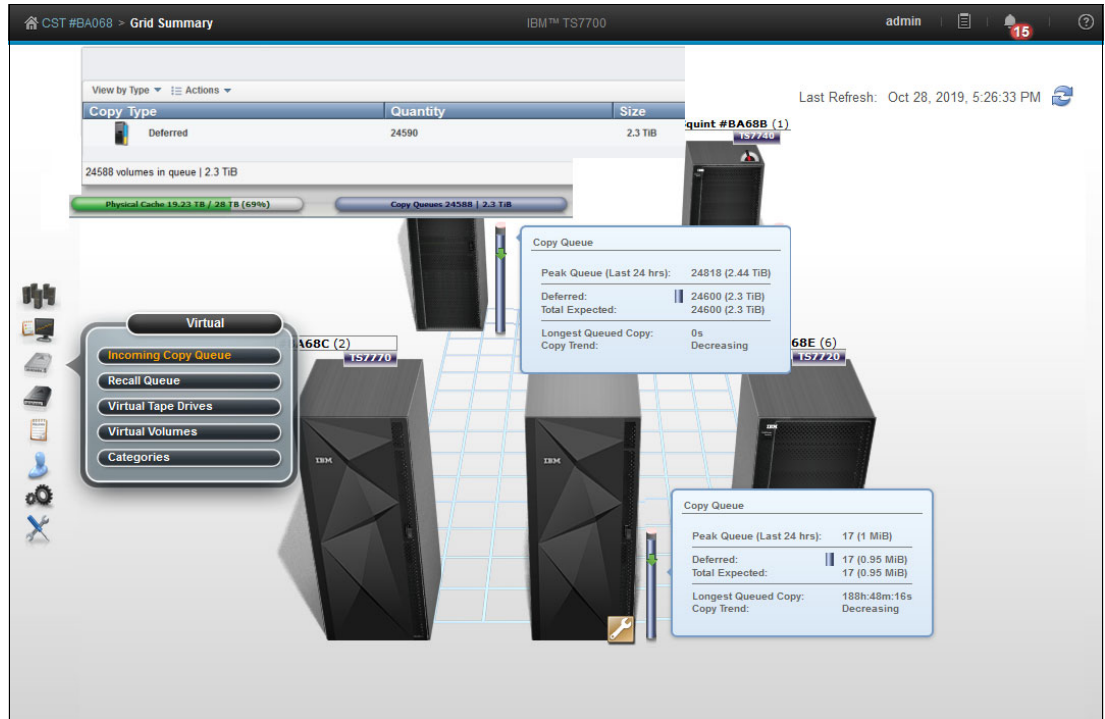


Figure 9-39 Incoming copy queue information in the Grid Summary window

Using the upper-left option, choose between **View by Type** and **View by Last TVC Cluster**. Click **Actions** to download the Incoming Queued Volumes list.

Figure 9-40 shows the Incoming Copy queue information.

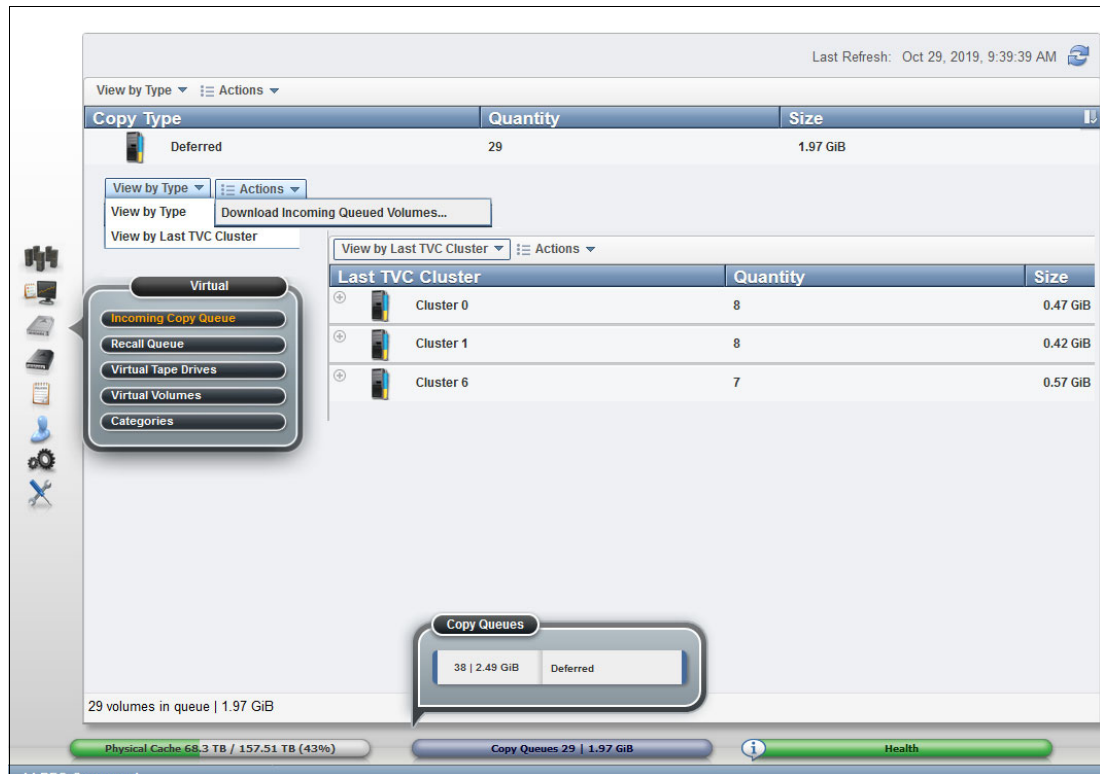


Figure 9-40 Incoming Copy queue window

9.8.3 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 no physical tape attachment to the selected cluster exists, but at least one tape-attached cluster is 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 no TS7700 tape-attached cluster is in the 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 the following groups:

- ▶ In Progress
- ▶ Scheduled
- ▶ Unscheduled

Figure 9-41 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

Total 61 | Selecting 1 rows

Figure 9-41 Recall Queue window

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 option that is 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 option is available by checking an unassigned volume in the table and clicking **Actions → Promote Volume**.

9.8.4 Virtual tape drives

The Virtual Tape Drives window of the MI shows the status of all virtual tape drives in a cluster. Use this window to check the status of a virtual mount, perform a stand-alone mount or unmount, or assign host device numbers to a specific virtual drive.

The Cache Mount Cluster field in the virtual tape drives page identifies to which cluster TVC the volume is mounted. The user can recognize remote (crossed) or synchronous mounts by reviewing in this field.

Remote mounts show other clusters that are used by a mounted volume instead of a local cluster (the cluster to which the Virtual Tape Drives belong to in the page that is displayed), whereas synchronous mounts show both clusters that are used by the mounted volume.

The page contents can be customized by selecting specific items to display. Figure 9-42 shows the Virtual tape drives window and the available items under Actions menu.

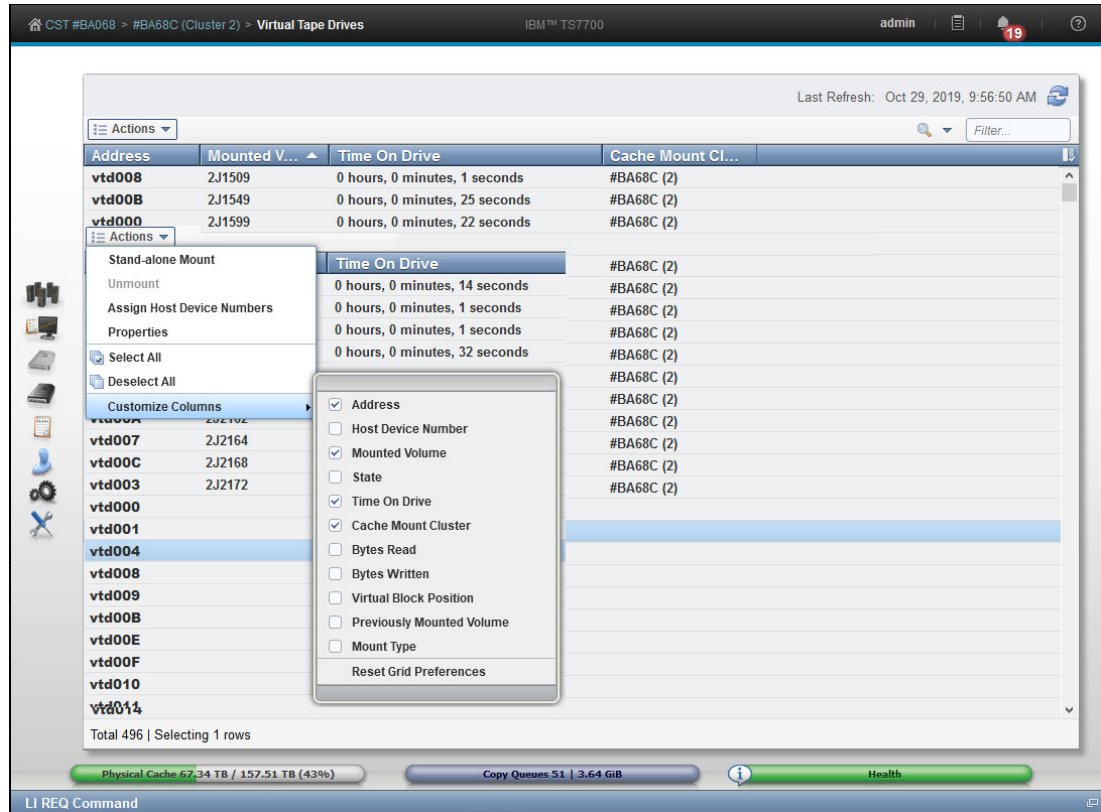


Figure 9-42 Virtual Tape Drives window

Assign Unit Host Device numbers

The user can assign the MI virtual tape drive identifier to the defined unit address in the attached host. The value in this field does not affect drive operations, but if the host device number is set, then it is simpler to compare the virtual tape drives to their associated host devices. This might be useful for troubleshooting and makes it simpler to use a drive as an IPL address with Stand-alone Services. To assign host device numbers:

1. Select one or more virtual tape drives
2. Select Assign Host Device Numbers from the Action menu
3. Enter the host device address that is assigned to the first virtual tape drive. If a range of devices is selected, host device numbers are added to subsequent virtual tape drives incrementally.

Figure 9-43 show an example of assigning a host device number.

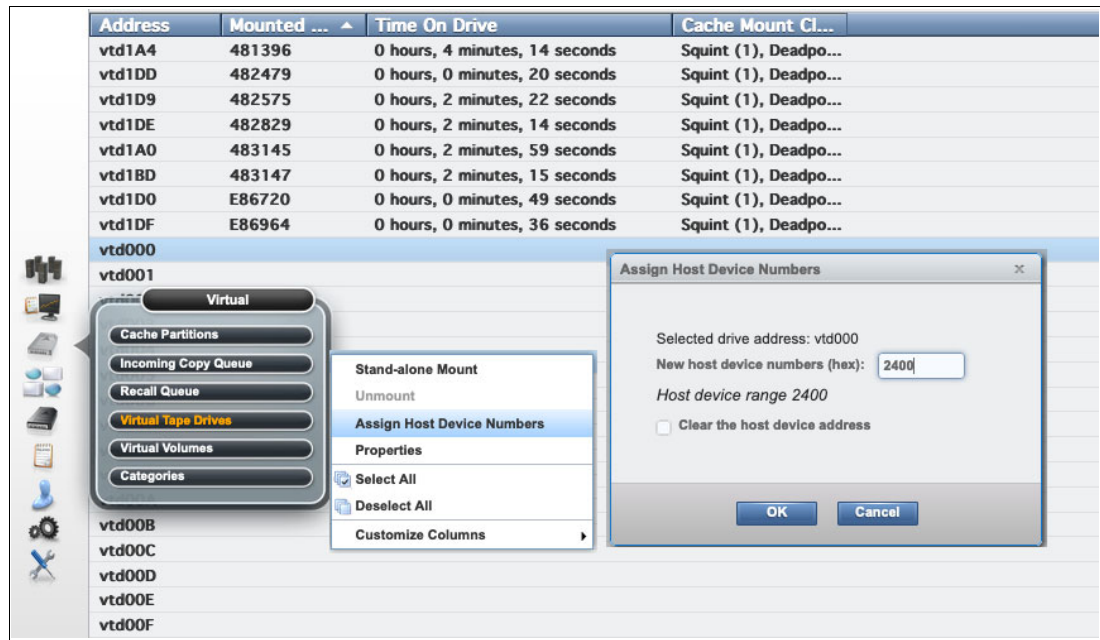


Figure 9-43 Assign Unit Host Device number sequence

Notice that when assigning the host device numbers in the MI Virtual Tape Drives page, the *VTDnnn* units should correspond to the Control Unit addresses defined by *CUADD* parameter in the IOCP/HCD, and the unit addresses 0–F within that control unit. For example, *VTD000-VTD00F* corresponds to *CUADD=0* and *VTD010-VTD01F* to *CUADD=1*, and so on. In the example shown in Figure 9-43, the *IODEV 2400* is the first address of the range that is defined for the *CUADD=0* in the HCD/IOCP, thus assigned to the *vtd000*.

Note: The *VTDnnn* units should relate to the corresponding Control Unit address defined by the *CUADD* parameter in the IOCP/HCD definitions and the unit address 0–F within that control unit.

Stand-alone Mount and Unmount

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 for volumes that were manually mounted, and for occasions when a logical volume was left mounted on a virtual drive by an incomplete operation or some test rehearsal, which creates the need to unmount it through MI operations.

To perform a stand-alone mount, complete the steps as shown in Figure 9-44:

1. Go to the Virtual Tape Drive page.
2. Select an empty virtual drive on which to mount a logical volume.
3. Click **Actions** → **Stand-alone Mount**.
4. Enter the logical volume to be mounted in the dialog box and then click **OK**.

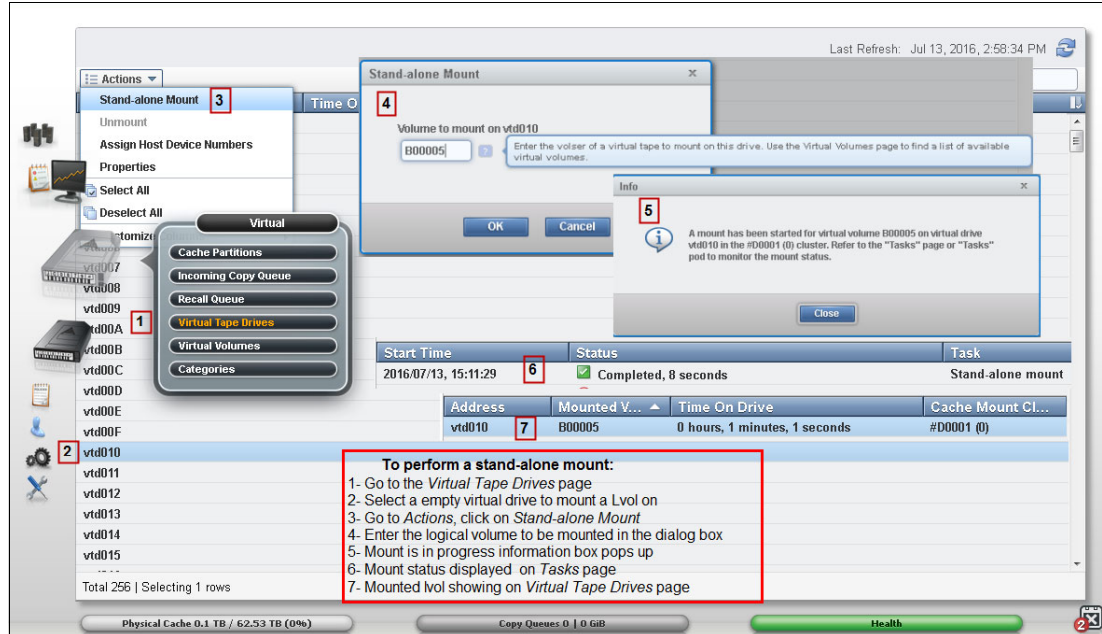


Figure 9-44 Stand-alone mount procedure

The user can mount only virtual volumes that are not already mounted on a virtual drive that is online.

If a logical volume that is mounted to a virtual tape drive must be unmounted, follow the steps that are shown in Figure 9-45 on page 416:

1. Go to the Virtual Tape Drives page.
2. Select the virtual tape drive where the lvol to be unmounted is on. The virtual drive must be in an idle state to be manually unmounted.
3. Click **Actions** → **Unmount**.
4. Click **Yes** to confirm the unmount action in the dialog box.

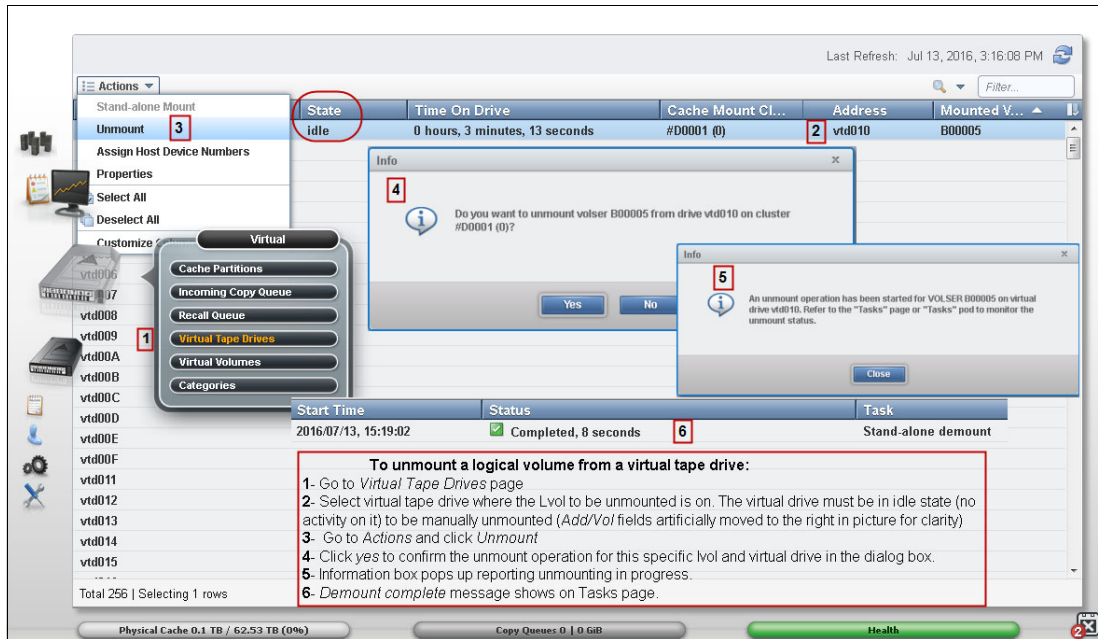


Figure 9-45 Demounting a logical volume by using MI

The user can unmount only those virtual volumes that are mounted and have a status of Idle.

For more information, see this IBM Documentation [web page](#).

9.8.5 Virtual volumes

In this section, we describe 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 9-46 on page 417 and Figure 9-47 on page 420 show an example of the resulting windows for a Virtual Volume query. The window can be subdivided in the following parts:

- ▶ Virtual volume summary
- ▶ Virtual volume details
- ▶ Cluster-specific virtual volume properties

A tutorial available is available about the virtual volume display and how to interpret the windows that are accessible directly from the MI window. To watch it, click the **View Tutorial** link on the Virtual Volume Detail window.

Figure 9-46 shows a composite Virtual Volume Details window in the MI, which is a graphical summary of the status of the virtual volume being displayed throughout the grid. The example shows at center left a virtual volume that was offloaded to cloud. Also shown is the adjacent cluster being aware of the logical volume in the cloud. The cloud awareness is a new function with R5.1.

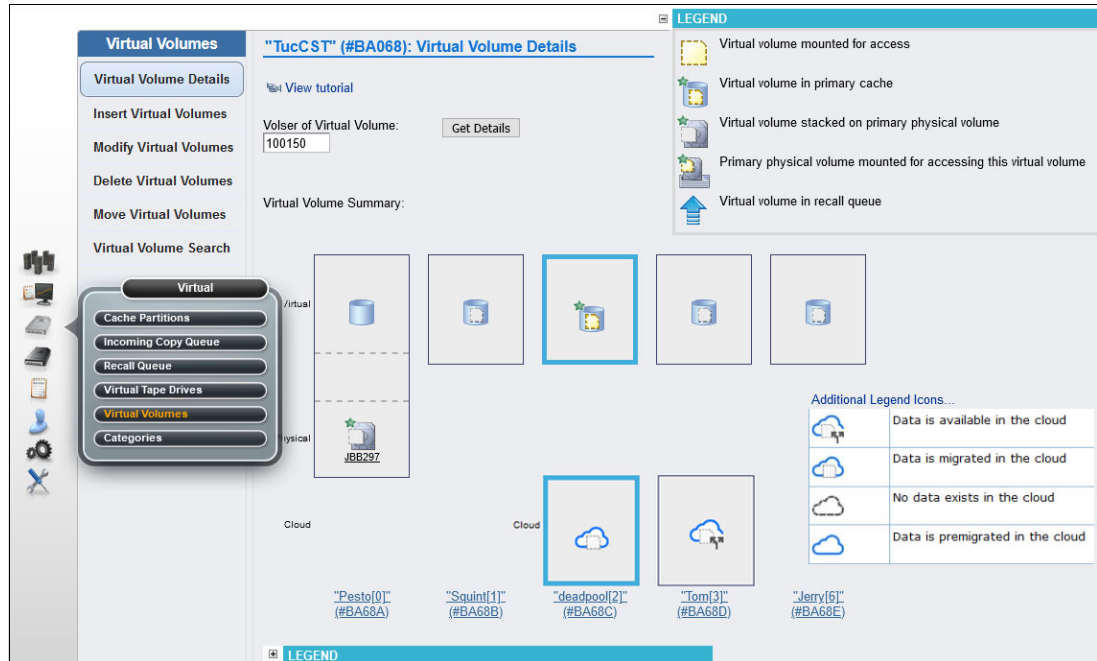


Figure 9-46 Virtual Volume Details: Graphical summary and legend icon descriptions

This graphical summary brings details of the present status of the virtual volume within the grid, plus the current operations that are occurring 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 is being recalled from tape in a remote location. For more information about examples of usage for this MI page, see Chapter 12, “Monitoring the IBM TS7700C”, of *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.

Hovering over the graphical representation of the virtual volume shows the meaning of the icon in the page, as shown in Figure 9-46.

For more information about the Virtual volume legend, see TS7700 Customer Documentation 5.2.2, under Virtual volume details, or click **Additional Legend Icons** at the bottom of the display (see Figure 9-46).









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 TS7700T.


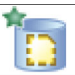









The Virtual Volume Details window shows all clusters where the selected virtual volume is within the grid. The icon that represents each cluster is divided in the following areas by broken lines:

- ▶ The top area relates to the logical volume status.
- ▶ The intermediate area shows actions that are 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. Table 9-8 shows the icons that are used in the legend in the virtual volume window with representation for the pending actions.

Table 9-8 The Virtual Volume Details legend icons

Icon	Meaning
	Data is available in the cloud.
	Data is migrated in the cloud.
	Data is pre-migrated in the cloud.
	No data exists in the cloud.
	Primary partition for tape or cloud partitions.
	Primary Object Partition (only when DS8K Object Store is implemented and we have object partitions).
	Virtual volume not in primary cache. The volume was in the primary cache, but was migrated. The current cluster might not be the owner. This volume must not ever be mounted.
	Virtual volume not in cache on the owning cluster. When the green star is applied to the cache icon, it signifies the owning cluster.

Icon	Meaning
	Virtual volume in cache.
	Virtual volume in primary cache. The current cluster might not be the owner. This volume can be mounted. When the green star is applied to the cache icon, it signifies the owning cluster.
	Virtual volume mounted for access.
	Virtual volume mount in progress.
	Virtual volume stacked on a primary physical volume. When applied to a physical volume, the green star signifies a primary physical volume, or the first physical volume that is written to when a virtual volume is migrated to a physical tape
	Virtual volume stacked on secondary physical volume.
	Primary physical volume that is mounted for accessing this logical volume. When applied to a physical volume, the green star signifies a primary physical volume, or the first physical volume that is written to, when a virtual volume is migrated to a physical tape.
	Primary physical volume that is mounted for accessing another logical volume. When applied to a physical volume, the green star signifies a primary physical volume, or the first physical volume that is written to, when a virtual volume is migrated to a physical tape.
	A secondary physical volume that is mounted for accessing this virtual volume.
	A secondary physical volume that is mounted for accessing another virtual volume.
	Pending action: Virtual volume in incoming copy queue (icon is above Virtual Action areas).




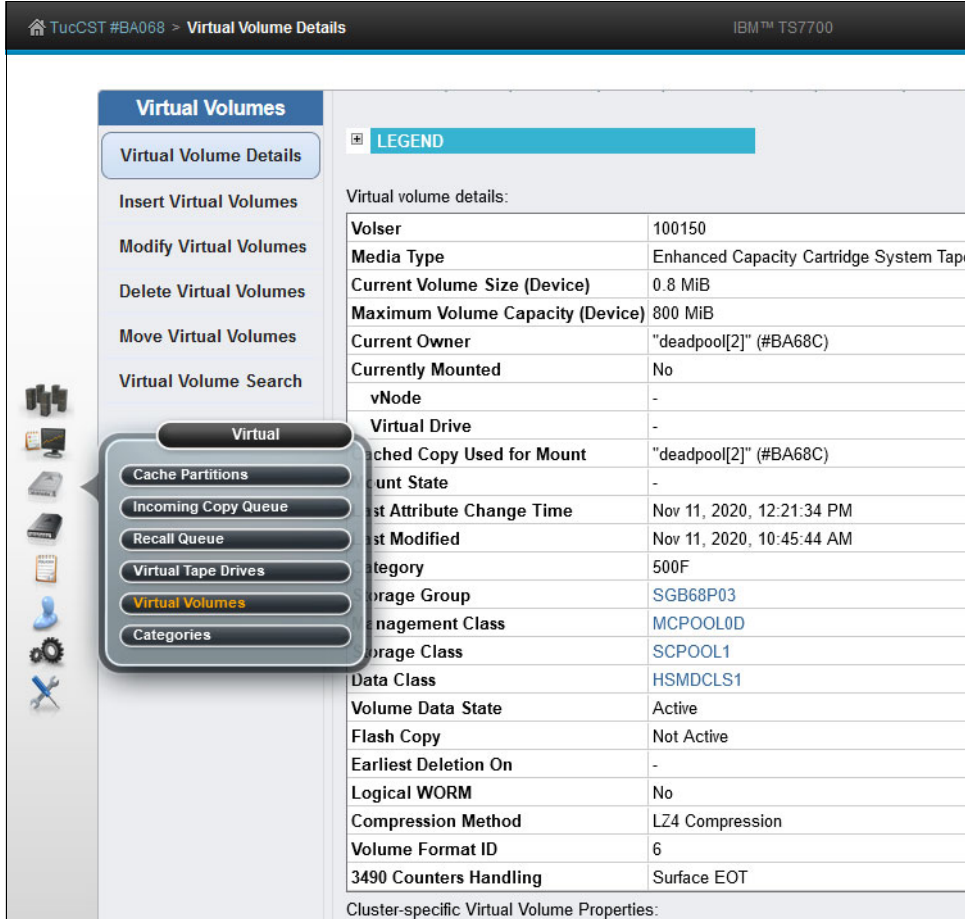
Icon	Meaning
	Pending action: Virtual volume deferred in the incoming copy queue.
	Pending action: Virtual volume in recall queue.
	Pending action: Virtual volume migrating (icon is between Virtual and Physical Action areas).

Figure 9-47 shows an example of the text section of the Virtual Volume Details window.



Virtual volume details:	
Volser	100150
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	0.8 MiB
Maximum Volume Capacity (Device)	800 MiB
Current Owner	"deadpool[2]" (#BA68C)
Currently Mounted	No
vNode	-
Virtual Drive	-
Attached Copy Used for Mount	"deadpool[2]" (#BA68C)
Mount State	-
Last Attribute Change Time	Nov 11, 2020, 12:21:34 PM
Last Modified	Nov 11, 2020, 10:45:44 AM
Category	500F
Storage Group	SGB68P03
Management Class	MCP00L0D
Storage Class	SCPOOL1
Data Class	HSMDCLS1
Volume Data State	Active
Flash Copy	Not Active
Earliest Deletion On	-
Logical WORM	No
Compression Method	LZ4 Compression
Volume Format ID	6
3490 Counters Handling	Surface EOT

Cluster-specific Virtual Volume Properties:

Figure 9-47 Virtual volume details: Text section

Virtual volume details: Text

The virtual volume details and status are displayed in the Virtual volume details table. These virtual volume properties are consistent across the grid. Table 9-9 lists the volumes details that are found on this page.

Table 9-9 Virtual Volume details and descriptions

Volume detail	Description
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 following possible values are available: <ul style="list-style-type: none"> ▶ Cartridge System Tape ▶ Enhanced Capacity Cartridge System Tape
Maximum Volume Capacity	The maximum size (MiB) of the virtual volume. This capacity is set on insert by the Data Class of the volume. <p>Note: Consider the following points:</p> <ul style="list-style-type: none"> ▶ If an override is configured for the Data Class's maximum size, it is applied only when a volume is mounted and a load-point write (scratch or fast ready mount) occurs. During the volume close operation, the new override value is bound to the volume and cannot change until the volume is reused. Any other changes to a Data Class override are not inherited by a volume until it is written again during a fast ready mount and then closed. ▶ When the host mounts a scratch volume and unmounts it without completing any writes, the system might report that the virtual volume's current size is larger than its maximum size. This result can be disregarded.
Current Volume Size	Size of the data (MiB) for this unit type (channel or device).
Current Owner	The name of the cluster that owns the latest version of the virtual volume.
Currently Mounted	Indicates whether the virtual volume is mounted in a virtual drive. If this value is Yes, the FOLLOWING qualifiers are also shown: <ul style="list-style-type: none"> ▶ 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 name of the cluster that is associated to the TVC that is selected for mount and I/O operations to the virtual volume. 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 possible values are available: <ul style="list-style-type: none"> ▶ Mounted: The volume is mounted. ▶ Mount Pending: A mount request was received and is in progress. ▶ Recall Queued/Requested: A mount request was received and a recall request was queued. ▶ Recalling: A mount request was received and the virtual volume is being staged into a tape volume cache from physical tape.
Cache Management Preference Group	The preference level for the Storage Group. This setting determines how soon volumes are removed from cache following their copy to tape. This information is displayed only if the owner cluster is a TS7740 or if the owner cluster is a TS7700T and the volume is assigned to a cache partition that is greater than 0. The following values are possible: <ul style="list-style-type: none"> ▶ 0: Volumes in this group have a preference to be removed from cache over other volumes. ▶ 1: Volumes in this group have a preference to be retained in cache over other volumes. A "least recently used" algorithm is used to select volumes for removal from cache if no volumes exist to remove in preference group 0. ▶ Unknown: The preference group cannot be determined.

Volume detail	Description
Last Accessed by a Host	The most recent date and time that a host accessed a live copy of the virtual volume. The recorded time reflects the time zone in which the user's browser is located.
Last Modified	The date and time that the virtual volume was last modified by a host mount or demount. The recorded time reflects the time zone in which the user's browser is located.
Category	The number of the scratch category to which the virtual volume belongs. A scratch category groups virtual volumes for quick, nonspecific use.
Storage Group	The name of the Storage Group that defines the primary pool for the premigration of the virtual volume.
Management Class	The name of the Management Class applied to the volume. This policy defines the copy process for volume redundancy.
Storage Class	The name of the Storage Class applied to the volume. This policy classifies virtual volumes to automate storage management.
Data Class	The name of the Data Class applied to the volume. This policy classifies virtual volumes to automate storage management.
Volume Data State	<p>The state of the data on the virtual volume. The following values are possible:</p> <ul style="list-style-type: none"> ▶ New: The virtual volume is in the insert category or a nonscratch category and data was never written to it. ▶ Active: The virtual volume is within a private category and contains data. ▶ Scratched: The virtual volume is within a scratch category and its data is not scheduled to be automatically deleted. ▶ Pending Deletion: The volume is within a scratch category and its contents are a candidate for automatic deletion when the earliest deletion time passes. Automatic deletion then occurs sometime thereafter. <p>Note: The volume can be accessed for mount or category change before the automatic deletion; therefore, the deletion might be incomplete.</p> <ul style="list-style-type: none"> ▶ Pending Deletion with Hold: The volume is within a scratch category that is configured with hold and the earliest deletion time is not yet passed. The volume is not accessible by any host operation until the volume leaves 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 currently is within a scratch category or was in a scratch category where it became a candidate for automatic deletion and was deleted. Any mount operation to this volume is treated as a scratch mount because no data is present.
FlashCopy	<p>Details of any existing flash copies. The following values are possible, among others:</p> <ul style="list-style-type: none"> ▶ Not Active: No FlashCopy is active. No FlashCopy was enabled at the host by an LI REQ operation. ▶ Active: A FlashCopy that affects this volume was enabled at the host by an LI REQ operation. Volume properties did not change since FlashCopy time zero. ▶ Created: A FlashCopy that affects this volume was enabled at the host by an LI REQ operation. Volume properties between the live copy and the FlashCopy changed. Click this value to open the FlashCopy Details page.
Earliest Deletion On	<p>The date and time when the virtual volume is deleted. The time that is recorded reflects the time zone in which your browser is located.</p> <p>This value displays as “—” if no expiration date is set.</p>

Volume detail	Description
Logical WORM	Whether the virtual volume is formatted as a WORM volume. Possible values are Yes and No.
Compression Method	The compression method that is applied to the volume.
Volume Format ID	Volume format ID that belongs to the volume. The following values are possible: <ul style="list-style-type: none"> ▶ -2: Data is not yet written ▶ 5: Logical Volume old format ▶ 6: Logical Volume new format
3490 Counters Handling	3490 Counters Handling value that belongs to the volume.

Cluster-specific Virtual Volume Properties

The Cluster-specific virtual volume properties table displays information about requesting virtual volumes on each cluster. These properties are specific to a cluster. Figure 9-48 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:					
	"Pesto[0]" (#BA68A)	"Squint[1]" (#BA68B)	"deadpool[2]" (#BA68...)	"Tom[3]" (#BA68D)	"Jerry[6]" (#BA68E)
In Cache	No	Yes	Yes	Yes	Yes
Device Bytes Stored	0.8 MiB (Device)	0.8 MiB (Device)	0.8 MiB (Device)	0.8 MiB (Device)	0.8 MiB (Device)
Primary Physical Volume	JBB297	-	-	-	-
Secondary Physical Volume	None	-	-	-	-
Copy Activity	Complete	Complete	Complete	Complete	Complete
Queue Type	-	-	-	-	-
Copy Mode	Deferred	Deferred	Deferred	Deferred	Deferred
Deleted	-	-	-	-	-
Removal Residency	-	No Removal Attempted	No Removal Attempted	No Removal Attempted	No Removal Attempted
Removal Time	-	-	-	-	-
Partition Number	-	-	0	-	-
Storage Preference	-	Prefer Keep	Prefer Keep	Prefer Keep	Prefer Keep
Cloud Data Status	No data exists in the cloud	No data exists in the cloud	Data is migrated in the cloud	Data is available in the cloud	No data exists in the cloud
Cloud Pool	-	-	MYPOOL	MYPOOL	-
Cloud Account	-	-	MYACCT	MYACCT	-
Cluster to Premigrate to Cloud	-	-	2	2	-
Cloud Consistency	Unaware	Unaware	Checked.	Checked	Aware

Figure 9-48 Cluster-specific Virtual Volume Properties

The Cluster-specific Virtual Volume Properties table displays information about requesting virtual volumes on each cluster. Properties are specific to a cluster. Virtual volume details and the status that is displayed include the properties that are listed in Table 9-10.

Table 9-10 Cluster-specific virtual volume properties

Property	Description
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.
Device Bytes Stored	The number of bytes used (MiB) by each cluster to store the volume. Actual bytes can vary between clusters, based on settings and configuration.

Property	Description
Primary Physical Volume	The physical volume that contains the specified virtual volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details page for this physical volume. A value of None means that no primary physical copy is to be made. This column is only visible if a physical library is present in the grid. If at least one physical library exists in the grid, the value in this column is shown as “—” for those clusters that are not attached to a physical library.
Secondary Physical Volume	Secondary physical volume that contains the specified virtual volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details page for this physical volume. A value of None means that no secondary physical copy is to be made. This column is visible only if a physical library is present in the grid. If at least one physical library exists in the grid, the value in this column is shown as “—” for those clusters that are not attached to a physical library.
Copy Activity	<p>Status information about the copy activity of the virtual volume copy. The following values are possible:</p> <ul style="list-style-type: none"> ▶ 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 is 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 are 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	<p>The type of queue as reported by the cluster. The following values are possible:</p> <ul style="list-style-type: none"> ▶ Rewind Unload (RUN): The copy occurs before the rewind-unload operation completes at the host. ▶ Deferred: The copy occurs some time after the rewind-unload operation completes at the host. ▶ 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. For more information about synchronous mode copy settings and considerations, see 3.3.5, “Copy Consistency policy” on page 130. ▶ Immediate Deferred: A RUN copy that was moved to the deferred state because of copy timeouts or TS7700 Grid states. ▶ Time Delayed: The copy occurs sometime after the delay period is exceeded.
Copy Mode	<p>The copy behavior of the virtual volume copy. The following values are possible:</p> <ul style="list-style-type: none"> ▶ Rewind Unload (RUN): The copy occurs before the rewind-unload operation completes at the host. ▶ Deferred: The copy occurs some time after the rewind-unload operation at the host. ▶ No Copy: No copy is made. ▶ Sync: The copy occurs on any synchronization operation. For more information about synchronous mode copy settings and considerations, see 3.3.6, “Synchronous mode copy” on page 131. ▶ Time Delayed: The copy occurs sometime after the delay period was exceeded. ▶ 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 is not deleted, this value displays as “—”.

Property	Description
Removal Residency	<p>The automatic removal residency state of the virtual volume. In a TS7700 Tape Attach configuration, this field is displayed only when the volume is in the disk partition. This field is not displayed for TS7740 Clusters. The following values are possible:</p> <ul style="list-style-type: none"> ▶ —: Removal Residency does not apply to the cluster. This value is displayed if the cluster attaches to a physical tape library, or inconsistent data exists on the volume. ▶ Removed: The virtual volume is removed from the cluster. ▶ No Removal Attempted: The virtual volume is a candidate for removal, but the removal did not yet occur. ▶ 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 that can replicate. 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 when no outage exists and replication is allowed to continue. ▶ Pinned: The virtual volume is pinned by the virtual volume Storage Class. 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 passes. When the Removal Time passes, 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 passes, it is moved back to the Held state.
Removal Time	<p>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 field.</p> <p>Possible values include:</p> <ul style="list-style-type: none"> ▶ —: Removal Time does not apply to the cluster. This value is displayed if the cluster attaches to a physical tape library. ▶ 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: Removal Time is not applicable. <p>The time that is recorded reflects the time zone in which the user's browser is located.</p>

Property	Description
Volume Copy Retention Group	<p>The name of the group that defines the preferred auto removal policy applicable to the virtual volume. This field is displayed only for a TS7700 Cluster or for partition 0 (CP0) on a TS7700 Tape Attach Cluster, and only when the grid contains a disk-only cluster.</p> <p>The Volume Copy Retention Group provides more options to remove data from a cluster as the active data reaches full capacity. Volumes become candidates for removal if an appropriate number of copies exist on peer clusters <i>and</i> the volume copy retention time elapses.</p> <p>The Retention Time parameter describes the number of hours a volume remains in the cache before becoming a candidate for removal. Retention time is measured beginning at the volume's creation time (when a write operation was performed from the beginning of a tape) or at the time the volume was most recently accessed. Volumes in each group are removed in order based on their least recently used access times.</p> <p>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 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.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> ▶ —: Volume Copy Retention Group does not apply to the cluster. This value is displayed if the cluster attaches to a physical tape library. ▶ 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 are removed. ▶ 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 then moved to scratch become priority candidates for removal. <p>Caution: Care must be taken when assigning volumes to this group to avoid cache overruns.</p>
Retention Time	<p>The length of time in hours that must elapse before the virtual volume that is named in the Volume Copy Retention Group can be removed. Depending on the Volume Copy Retention Reference setting, this time period can be measured from the time when the virtual volume was created or when it was most recently accessed.</p>
Volume Copy Retention Reference	<p>The basis for calculating the period is defined in Retention Time. The following are the possible values:</p> <ul style="list-style-type: none"> ▶ Volume Creation: Calculate retention time starting with the time when the virtual volume was created. This value refers to the time that a write operation was performed from the beginning of a tape. It can be either a scratch mount or a private mount where writing began at record 0. ▶ Volume Last Accessed: Calculate retention time starting with the time when the virtual volume was most recently accessed.
TVC Preference	<p>The group containing virtual volumes that have a preference for premigration. This field is displayed only for a TS7740 Cluster or for partitions 1–7 (CP1 through CP7) in a TS7720/TS7760 Tape Attach cluster. The following values are possible:</p> <ul style="list-style-type: none"> ▶ PG0: Volumes in this group have a preference to be premigrated over other volumes. ▶ PG1: Volumes in this group have a preference to be premigrated over other volumes. A “least recently used” algorithm is used to select volumes for premigration if no volumes exist to premigrate in preference group 0.

Property	Description
Time-Delayed Premigration Delay	The length of time (in hours) that must elapse before a delayed premigration operation can begin for the virtual volumes that are designated by the TVC Preference parameter. Depending on the Time-Delayed Premigration Reference setting, this period can be measured from the time when the virtual volume was created or when it was most recently accessed.
Time-Delayed Premigration Reference	The basis for calculating the time period defined in Time-Delayed Premigration Delay. The following values are possible: <ul style="list-style-type: none"> ▶ Volume Creation: Calculate the premigration delay starting with the time when the virtual volume was created. This value refers to the time that a write operation was performed from the beginning of a tape. It can be a scratch mount or private mount where writing began at record 0. ▶ Volume Last Accessed: Calculate the premigration delay starting with the time when the virtual volume was most recently accessed.
Storage Preference	The priority for removing a virtual volume is when the cache reaches a predetermined capacity. The following values are possible: <ul style="list-style-type: none"> ▶ Prefer Remove: These virtual volumes are removed first when the cache reaches a predetermined capacity and no scratch-ready volumes exist. ▶ Prefer Keep: These virtual volumes are the last to be removed when the cache reaches a predetermined capacity and no scratch-ready volumes exist. ▶ Pinned: These virtual volumes are never removed from the TS7700 Cluster.
Partition Number	The partition number for a TS7700 Tape Attach Cluster. Possible values are C0–C7. “Inserted” logical volumes are volumes with a -1 partition, meaning that no consistent data exists yet.
Cloud Pool	The nickname of the pool.
Cloud Account	The name of the cloud account
Cluster to Premigrate to Cloud	The ID of the cluster pre-migrating the volume to the cloud
Cloud Consistency	Status information regarding cloud storage operations: <ul style="list-style-type: none"> ▶ Unaware This cluster has no visibility to the container where this volume is stored or is not cloud attached. ▶ Aware This cluster is aware of the existence of this object but access and premigration has not been attempted. ▶ Checked This cluster is aware of the existence of this object through a previous check (for instance, a premigration attempt).

FlashCopy

This section provides more 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 was modified after FlashCopy time zero. For the volumes with a FlashCopy *active* (meaning no difference between the FlashCopy and live volume) as shown in Figure 9-47 on page 420, only the Virtual Volume details window is available (FlashCopy and live volume are identical).

Figure 9-49 shows a FlashCopy details page in the MI compared with the output of `LI REQ Lvol` flash command.

Virtual volume details:		
Volser	A08472	<p style="text-align: center;">The same volume seen through LI REQ LVOL FLASH command</p> <pre> > 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: CLUSTER2 " MOUNTED VNODE: " MOUNTED DEVICE: " TVC LIBRARY: CLUSTER2 " MOUNT STATE: " CACHE PREFERENCE: --- " CATEGORY: 002F </pre>
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	

Figure 9-49 FlashCopy details window

The following virtual volume details and statuses 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:** The size of the data in MiB for this virtual volume.
- ▶ **Current Owner:** The name of the cluster that owns the latest version of the virtual volume.
- ▶ **Currently Mounted:** Indicates whether the virtual volume is mounted in a virtual drive. If this value is Yes, the following qualifiers are also displayed:
 - 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 name of the cluster that owns the cache that is 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 was 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 storage class that is 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 within a private category and contains data.
 - **Scratched:** The virtual volume is within a scratch category and its data is not scheduled to be automatically deleted.
 - **Pending Deletion:** The volume is within a scratch category and its contents are a candidate for automatic deletion when the earliest deletion time 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 within a scratch category that is configured with hold and the earliest deletion time did not yet pass. The volume is not accessible by any host operation until the volume leaves 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 WORML** Not applicable to FlashCopy copies (-).

Cluster-specific FlashCopy volume properties

The Cluster-specific FlashCopy Properties window displays the following 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:** Indicates whether the virtual volume is in cache for this cluster.
- ▶ **Device Bytes Stored:** The number of actual bytes (MiB) that are 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 did not start or complete.
 - **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 are 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. 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.
 - 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 was moved to the deferred state because copy timeouts or TS7700 Grid states occurred.
 - Time Delayed: The copy occurs sometime after the delay period was exceeded.
- ▶ 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 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 is exceeded.
 - Exists: A consistent copy exists at this location although No Copy is intended. This issue occurs 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 was not 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 9-50.

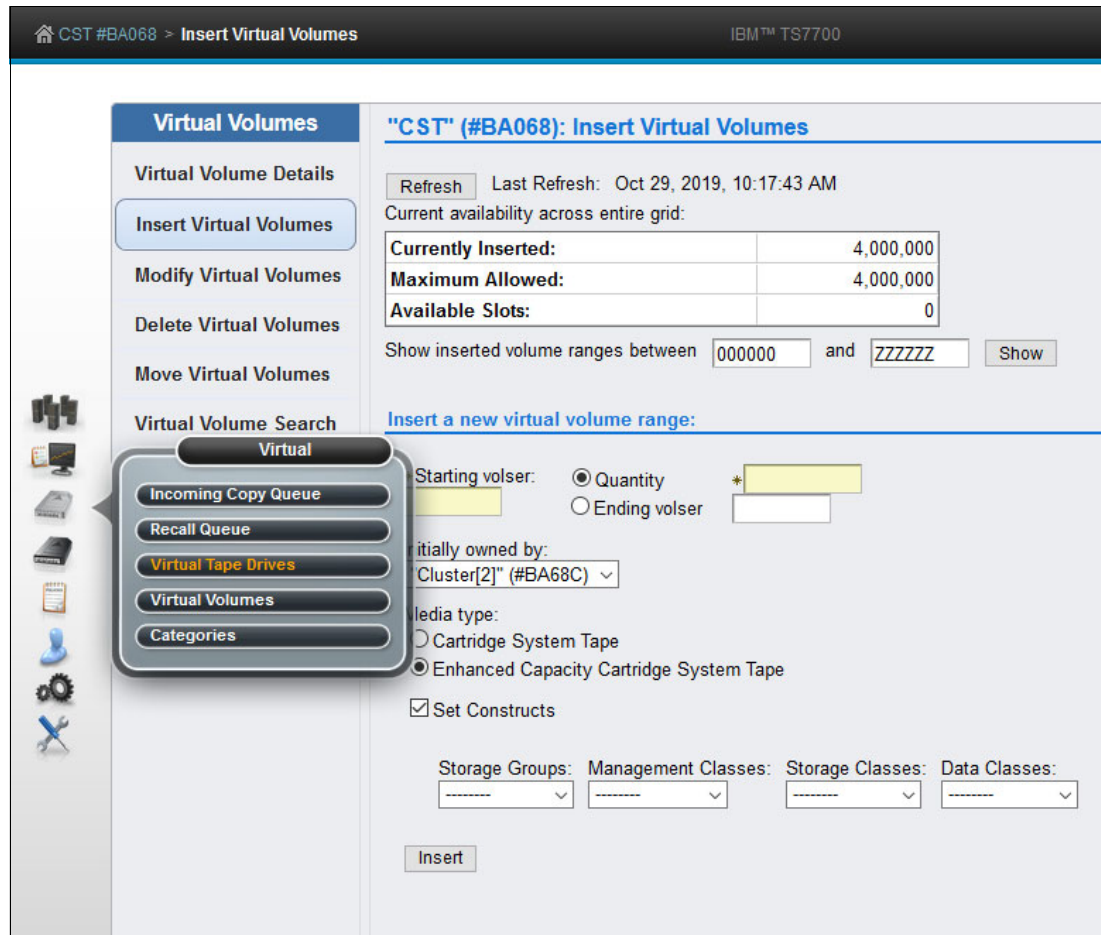


Figure 9-50 Insert Virtual Volumes window

The Insert Virtual Volume window shows the Currently availability across the entire grid table. This table shows the total of the inserted volumes, the maximum number of volumes that is 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 (see Figure 9-51) with the inserted volume ranges, quantities, media type, and capacity.

From	To	Quantity	Media Type	Capacity
111111	111115	5	ECCST	800 MiB
222222	223221	1,000	ECCST	800 MiB
Total: 2				

Figure 9-51 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.
For code levels before 8.50.1.25 a user could insert up to 10,000 virtual volumes at a time. From code level 8.50.1.25 and higher, the maximum quantity to insert virtual volumes was increased from 10,000 to 100,000 volumes when all clusters in the Grid are at 8.50.1.25 or at a later level of code, as in the case of the level R5.1 of code.
- ▶ **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 owns 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)
- ▶ **Set Constructs:** Select this option to specify constructs for the new virtual volumes. Then, use the menu that is under each construct to select a predefined construct name.
Set constructs only for virtual volumes that are used by hosts that are *not* z/OS hosts. z/OS 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 to modify the constructs that are associated with existing virtual volumes in the TS7700 composite library.

Note: You can use the Modify Virtual Volume function to manage virtual volumes that belong to a *non z/OS* host that is not aware of constructs.

z/OS hosts automatically assign constructs for virtual volumes, and manual changes are not recommended. The Modify Virtual Volumes window acts on any logical volume that belongs 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 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 menu feature the following 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 to delete *unused* virtual volumes from the TS7700 that are in the Insert Category (FF00).

Note: Only the unused logical volumes can be deleted through this window; that is, volumes in the insert category FF00 that were never mounted or had their category, constructs, or attributes modified by a host. Otherwise, those logical volumes can be deleted from the host only.

Complete the following steps to delete unused virtual volumes:

1. Select one of the following options and click **Delete Volumes**:
 - Delete ALL unused virtual volumes: Deletes all unused virtual volumes across all VOLSER ranges.
 - Delete a 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 the Delete specific range of unused virtual volumes option is selected.
 - To: The end of the VOLSER range to be deleted if the Delete specific range of unused virtual volumes option is selected.

A confirmation window is displayed.

2. Click **OK** to delete or **Cancel** to cancel.
3. 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.

Scratch volumes that have been used or touched by the host should be deleted by ejecting the volume from the host. With Data Facility Storage Management Subsystem (DFSMS)/Removable Media Management (RMM) as the tape management system (TMS), the process is done by using RMM commands.

Move Virtual Volumes window

Use the Move Virtual Volumes window to move a range of virtual volumes that are used by the TS7700T from one physical volume or physical volume range to a new target pool, or to cancel a move request that is in progress. If a move operation is in progress, a warning message opens. The user can view move operations in progress from the Events window.

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 a 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. 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 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 occurs 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

By using the window that is shown in Figure 9-52 on page 435, you can search for virtual volumes in a specific TS7700 cluster by way of one of the following parameters:

- ▶ VOLSER
- ▶ Category
- ▶ Media type
- ▶ Expiration date
- ▶ Inclusion in a group or class,

With the TS7700T, a search option is available to search by Partition Number. The user provides a search name to every query so that they can be recalled as necessary.

Virtual Volumes "[2] (#BA68C): Virtual Volume Search

Virtual Volume Details
 Insert Virtual Volumes
 Modify Virtual Volumes
 Delete Virtual Volumes
 Move Virtual Volumes
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: Media Type: Current Owner: [Clear all options](#)

Equal to Expire Time: Time Units:

Removal Residency: Removal Time:

Storage Preference:

Storage Group: or --Storage Groups in --

Management Class: or --Management Classes in --

Storage Class: or --Storage Classes in --

Data Class: or --Data Classes in --

Compression Method:

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

Select the following properties to show on the Search Results page.

Select	Properties
<input checked="" type="checkbox"/>	Category
<input checked="" type="checkbox"/>	Current Owner
<input checked="" type="checkbox"/>	Media Type
<input checked="" type="checkbox"/>	Expire Time
<input checked="" type="checkbox"/>	Storage Group
<input checked="" type="checkbox"/>	Management Class
<input checked="" type="checkbox"/>	Storage Class
<input checked="" type="checkbox"/>	Data Class
<input checked="" type="checkbox"/>	Compression Method
<input checked="" type="checkbox"/>	Mounted Tape Drive
<input checked="" type="checkbox"/>	Removal Residency
<input checked="" type="checkbox"/>	Removal Time
<input checked="" type="checkbox"/>	Storage Preference
<input checked="" type="checkbox"/>	Logical WORM

Search Results Options

Figure 9-52 MI Virtual Volume Search entry window

A maximum of 10 search queries results or 2 GB of search data can be stored at one time. If either limit is reached, the user deletes one or more stored queries from the Previous Virtual Volume Searches window before creating a search.

To view the results of a previous search query, select **Previous Searches** to see a table that contains a list of previous queries. Click a query name to display a list of virtual volumes that match the search criteria.

To clear the list of saved queries, select the checkbox next to one or more queries to be removed, select **Clear** from the menu, and then 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 search query, enter a name for the new query. Enter a value for any of the fields and select **Search** to start 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**.

When looking for the results of earlier searches, click **Previous Searches** in the Virtual Volume Search window (see Figure 9-52).

Search Options

Use this table to define the parameters for a new search query. Only one search can be run at a time. Define one or more of the following search parameters:

- ▶ **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 example, 0002/0102 (scratch MEDIA2), 000E (error), 000F/001F (private), and FF00 (insert) are possible values for Scratch and Specific categories. Wildcard characters can also 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.
- ▶ **Current Owner:** The cluster owner is the name of the cluster where the logical volume is stored. Use the drop-down menu to select from a list of available clusters. This field is available in a grid environment only and 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 TS7700T (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 selection is the default selection.
 - Removed: The search includes only virtual volumes that were 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 also be completed.
 - 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 also be complete.
 - 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 were not 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 also be complete.
- Removable After: The search includes only virtual volumes that are candidates for removal after a specific date and time. If this value is selected, the Removal Time field must also be complete.
- ▶ Removal Time: This field is not available for the TS7740. Values that are displayed in this field depend on the values that are shown in the Removal Residency field. The following values reflect the time zone in which the browser is used:
 - Date: The calendar date according to month (M), day (D), and year (Y). The format MM/DD/YYYY is used. This field includes a date chooser calendar icon. The user can enter the month, day, and year manually, or 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 a suitable number of copies exist on peer clusters *and* the volume copy retention time 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 the cache before becoming a candidate for removal.

This field is visible only if the selected cluster does not attach to a physical library. 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 selection 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 TS7700T (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 stored. 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.
- ▶ **Compression Method:** The name of the compression method to which the virtual volume belongs. You can select a name from the adjacent drop-down menu, or the 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 selection is the default selection.
 - Yes: Includes only mounted virtual volumes.
 - No: Includes only unmounted virtual volumes.
- ▶ **Storage Preference:** Removal policy for a virtual volume, which determines when the volume is removed from the cache of a TS7700 cluster in a grid configuration. The following values are possible:
 - Prefer Remove: These virtual volumes are removed first when cache reaches a predetermined capacity and no scratch-ready volumes exist.
 - Prefer Keep: These virtual volumes are the last to be removed when the cache reaches a predetermined capacity and no scratch-ready volumes exist.
 - Pinned: These virtual volumes are never removed from the TS7700 Cluster.
- ▶ **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 (default selection).
 - Yes: Includes only WORM logical volumes.
 - No: Does not include any WORM logical volumes.
- ▶ **Partition Number:** The partition number (0–7) for a TS7700 Tape Attach volume. *Inserted* logical volumes are those volumes with a -1 partition; that is, no consistent data exists yet.

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.

Search Results Option

Use this table to select the properties that are displayed on the Virtual volume search results window. Complete the following steps:

1. Click the down arrow that is next to the **Search Results Option** (see Figure 9-52 on page 435) to open the Search Results Options table.
2. Select the checkbox that is next to each property that must be included in the Virtual Volume Search Results window. The following properties can be selected for display:
 - Category
 - Current Owner (Grid only)
 - Media Type
 - Expire Time
 - Storage group

- Management Class
- Storage Class
- Data Class
- Compression Method
- Mounted Tape Drive
- Removal Residency
- Removal Time
- Volume Copy Retention Group
- Storage Preference
- Logical WORM

3. Click **Search** to start a new virtual volume search.

When the search is complete, the results are displayed in the Virtual 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. The following subwindows are available:

- ▶ Previous virtual volume searches: Use this window to view previous searches of virtual volumes in the MI currently accessed cluster.
- ▶ Virtual volume search results: Use this window to view a list of virtual volumes on this cluster that meet the criteria of a search query that was run.

9.8.6 Categories

Use this page to add, modify, or delete a scratch category of virtual volumes.

This page can also be used to view the total number of logical volumes that are classified in each category, which are grouped under Damaged, Scratch, and Private groups.

Clicking in the “+” that is next to the category expands the information about that category and shows how many volumes in that category exist in each cluster of the grid (as shown in Figure 9-54 on page 440).

A *category* is a grouping of virtual volumes for a predefined use. A scratch category groups virtual volumes for nonspecific use. This grouping enables faster mount times because the TS7700 can order category mounts without recalling data from a stacked volume (fast ready).

Figure 9-53 shows the Category window in the TS7700 MI.

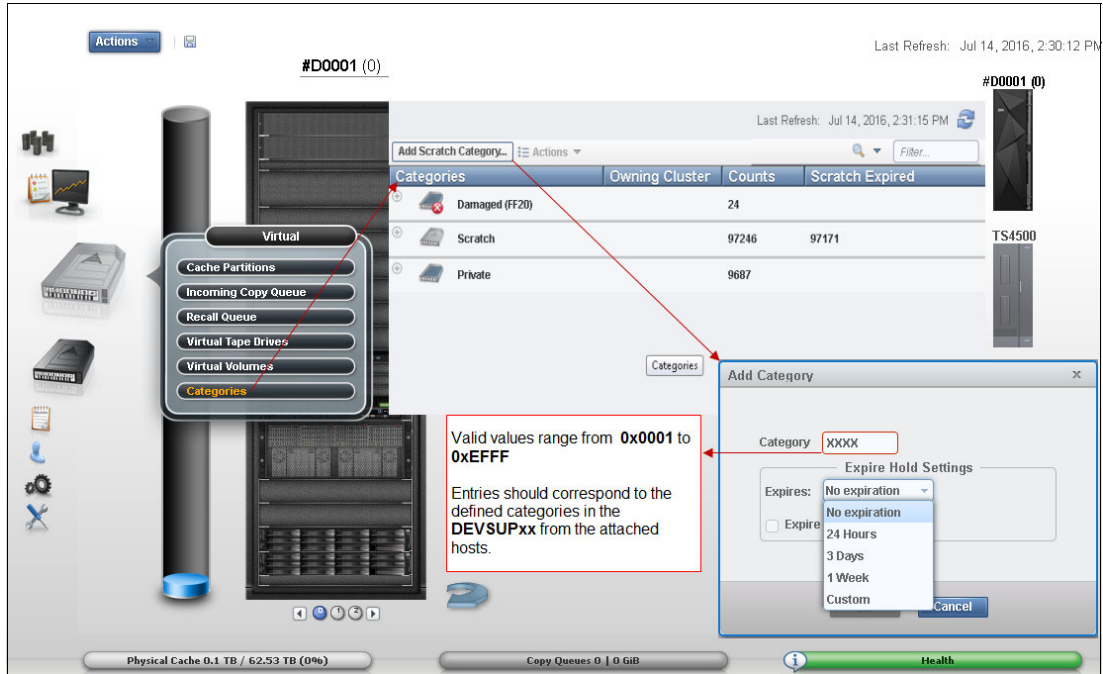


Figure 9-53 Categories window

You can display the defined categories, as shown in the Figure 9-54.

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 9-54 Displaying existing categories

Table 9-11 lists the values that are displayed on the Categories table, as shown in Figure 9-54 on page 440.

Table 9-11 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. 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. For more information about these values, see Table 9-11. ▶ 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 specific type, that type is not displayed in 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 9-12 lists the actions that can be performed on the Categories window.

Table 9-12 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 category before its data content is automatically delete-expired¹. <p>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:</p> <ul style="list-style-type: none"> • Time: Enter a number in the field according to the following restrictions: <ul style="list-style-type: none"> 1–2,147,483,647 if unit is hours 1–89,478,485 if unit is days 1–244983 if unit is years • Time Unit: Select a corresponding unit from the menu. – Set Expire Hold: Select this option to prevent the virtual volume from being mounted or having its category and attributes changed before the expire time elapses.² Selecting this option activates the hold state for any volumes that are in the scratch category and for which the expire time has not yet elapsed. Clearing this option removes the access restrictions on all volumes that are in the hold state within this scratch category.
Modify a scratch category	<p>The user can modify a scratch category in the following 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 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 by using one of the following methods:</p> <ul style="list-style-type: none"> ▶ Select a category on the table, and then select Actions → Delete Scratch Category. ▶ 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 checkbox 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>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>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 by which to filter. 3. Refine the selection: <ul style="list-style-type: none"> – Categories: Enter a whole or partial category number and press Enter. – Owning 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.

1. A volume becomes a candidate for delete-expire when *all* of 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 data recorded in that volume was created or last modified is greater than 12 hours.
 - At least 12 hours passed since the volume was migrated out or recalled back into the disk cache.

Up to 1,000 delete-expire candidate volumes can be deleted per hour. The volumes that were within the scratch category the longest are deleted first. If a volume is selected for a scratch mount before it delete-expired, the previous data in that volume is deleted immediately at first write.
2. 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 host. The volume category can be changed, but only to a private category, which allows accidental scratch occurrences to be recovered to private. If EXPIRE HOLD is not set, 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 host. Starting with R5.0, the maximum Expire hold time allowed was increased from 10 to 2000 years (730,000 days or 17,520,000 hours) to accommodate users requests.

Note: No cross-check is conducted between defined categories in the z/OS systems and the definitions in the TS7700.

9.9 Object Store icon

By using the TS7700 Management Interface pages that are collected under the Object Store icon, a user can create, modify, and delete object replication policies. It also helps the user set up an object store on TS7700 and link the object policies to an object store.

Note: Object Store and Object Policy pages are protected by **Dual Control** when mode is enabled (changes needs approval of a checker).

Figure 9-55 shows this aspect of the Object Store icon, and options.

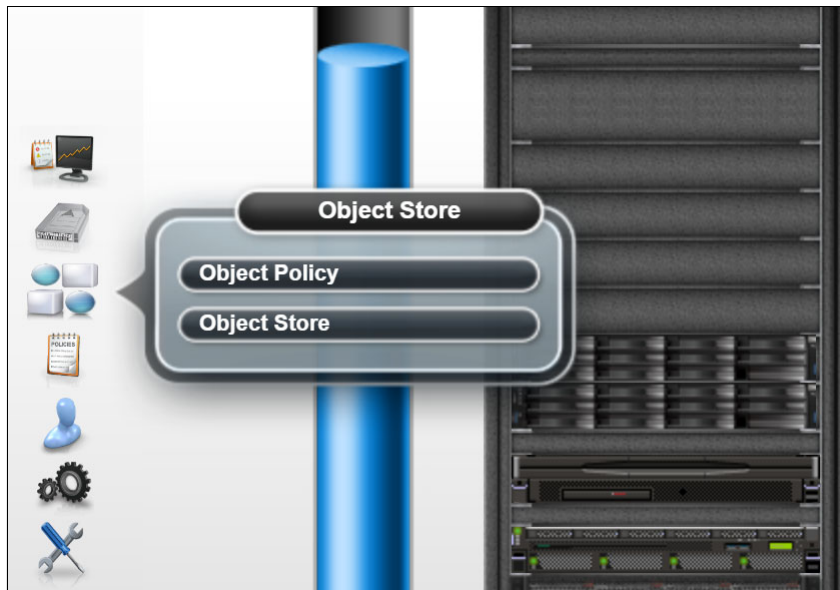


Figure 9-55 The Object Store icon and options

9.9.1 Object Policy

Use the Object Policy page of the Management Interface to add, modify, and delete object policies.

Object policies are used to determine how to manage objects that are directed to the TS7770 from an external host, such as a DS8900 as part of transparent cloud tiering (TCT).

TS7700 currently allows the capability to set up copy policies for object redundancy across an object-enabled TS7700 grid.

Object policies are grid scope and pertain to how an object is managed for each cluster in the grid. Object policies can be set up for every object-enabled cluster in the grid by accessing this page from one of those object-enabled clusters.

The maximum number of object policies that can exist on a TS7700 grid is 256. If a policy does not exist in a TS7700 that is object-enabled, the default object policy is used. The default object policy can be modified, but cannot be deleted.

Figure 9-56 on page 445 shows the Object Policy page, and Create Object Policy detail.

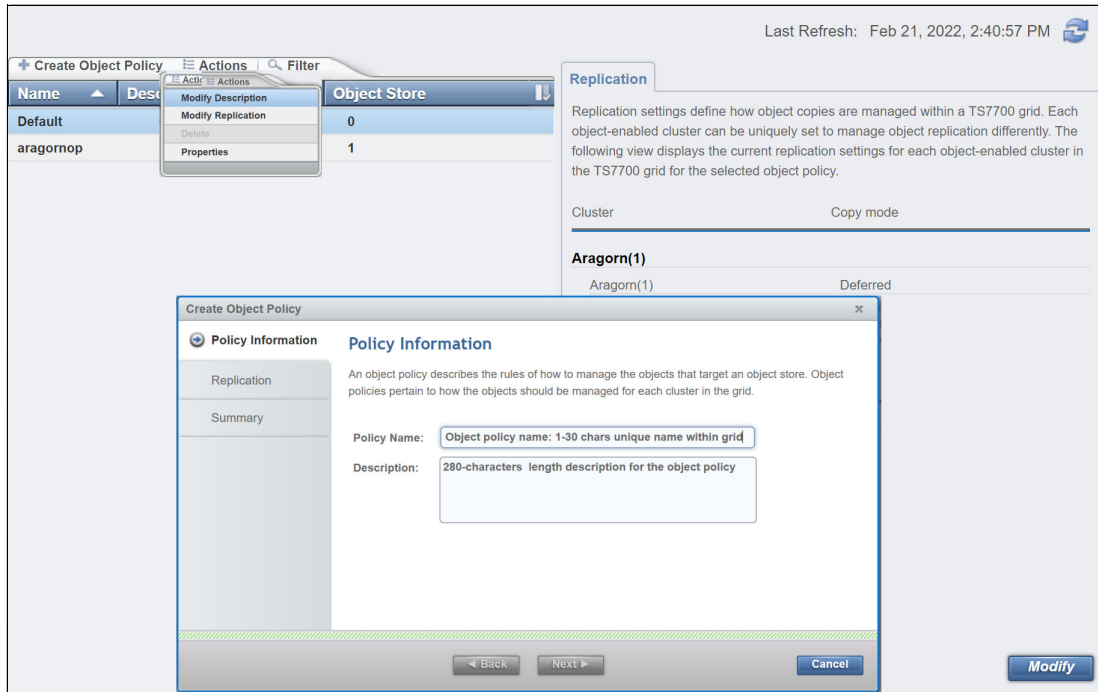


Figure 9-56 Object Policy page and Create Object Policy detail.

Note: Only the object-enabled clusters in the grid are selectable for the object copy policies. FC 5283 Advanced Object Store is required to be installed on a TS7700 VED to show this icon for access.

To create an object policy in the object policy page, complete the following steps:

1. Click **Create Object Policy** (see Figure 9-56).
2. Enter the Policy Name and description in the suitable boxes.
3. Click **Next**. The Replication tab opens.
4. For each object-enabled TS7700 cluster, the user selects the wanted copy mode:
 - a. Sync: This copy mode provides the ability to write two copies of the same object synchronously. If synchronous copy mode is requested, two clusters must be selected. If one of the clusters is not available, or fails to write one object, TS7700 continues the write operation to the available cluster, but the TS7700 grid enters a synchronous-deferred state.
 After the cluster is available and completes the synchronous-deferred copies, the grid removes the synchronous-deferred state. Synchronous-deferred copies are prioritized over deferred copy queue entries.
 - b. Deferred: In this copy mode, the object duplication occurs later based on the internal schedule of the copy engine.
 - c. No copy: If this copy mode is selected, no object duplication occurs in the cluster.
5. If the information that is provided is correct, click **Finish**. Otherwise, click **Back** to modify the policy information. To exit without saving the policy, click **Cancel**.

Policies can be viewed, modified, or deleted in this same Object Policy page in the Management Interface.

Modify, delete, and view policy

Existing object policies can be viewed, modified, or deleted from the object policy page, which is shown in Figure 9-57:

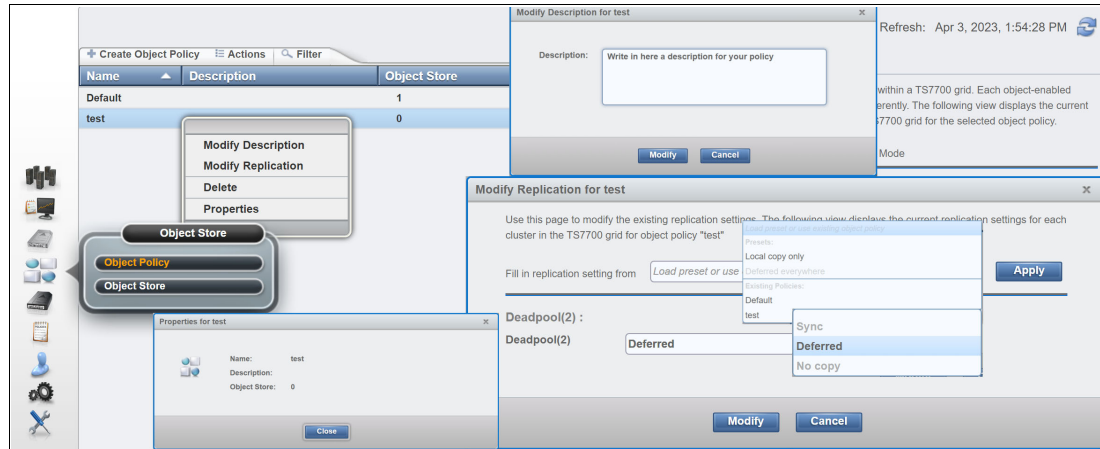


Figure 9-57 Options are available in the object policy page

1. To modify an existing policy description:
 - Right-click in an existing policy to be modified, and select Modify Description (or alternatively, select Actions → Modify Description).
 - Update the description in the Description field and click Modify. If no update is required, click Cancel.
2. To modify an existing replication policy:
 - Right-click the policy that replication settings needs to be modified and select Modify Replication (alternatively, select Actions → Modify Replication)
 - Make the required changes and click Modify. If no update is required, click Cancel.
3. To delete an existing policy:
 - Right-click the policy to be deleted and select Delete (alternatively, select Actions → Delete)
 - The Delete Object Policy window appears asking for confirmation: **Are you sure you want to remove “Sync Policy”?**
 - Click Delete to perform the action, or click Cancel to abandon the delete action.

Note: An object policy cannot be deleted if it is already assigned to a cloud name.

4. To view the policy information:
 - Right-click on the policy and select Properties (alternatively select Actions → Properties)

9.9.2 Object Store

Use the Object Store page of the MI to create an object store, modify an object store description and object policy, and view object store information.

To use the DS8K transparent cloud tiering (TCT), a cloud name must be created in the z/OS DFSMS Cloud Network Connection Construct name in the ISMF cloud panel.

Also, a cloud name must be created on the TS7700 before DS8K configuration and z/OS host transactions for that cloud name. If the cloud name does not exist on TS7700, the DS8K configuration process or host TCT transactions fails.

The object store page is used to add cloud names to the TS7700 grid. These cloud names are used as object store targets for TCT. Each cloud name can be assigned to an object policy. The object store page also lists the containers that are created by the z/OS host application. Figure 9-58 shows the Object store page.



Cloud Name	Description	Container Name	Object Policy	Containers	Used Capacity	Objects
hydratctcloud1	DSS full volume dump		Deferred Policy	7	2.03 TB	220313
hydratctcloud2	HSM full deferred copies		Deferred Policy	178	18.01 TB	1673770
hydratctcloud3	HSM full synchronous copies		Sync Policy	31	3.20 TB	500211

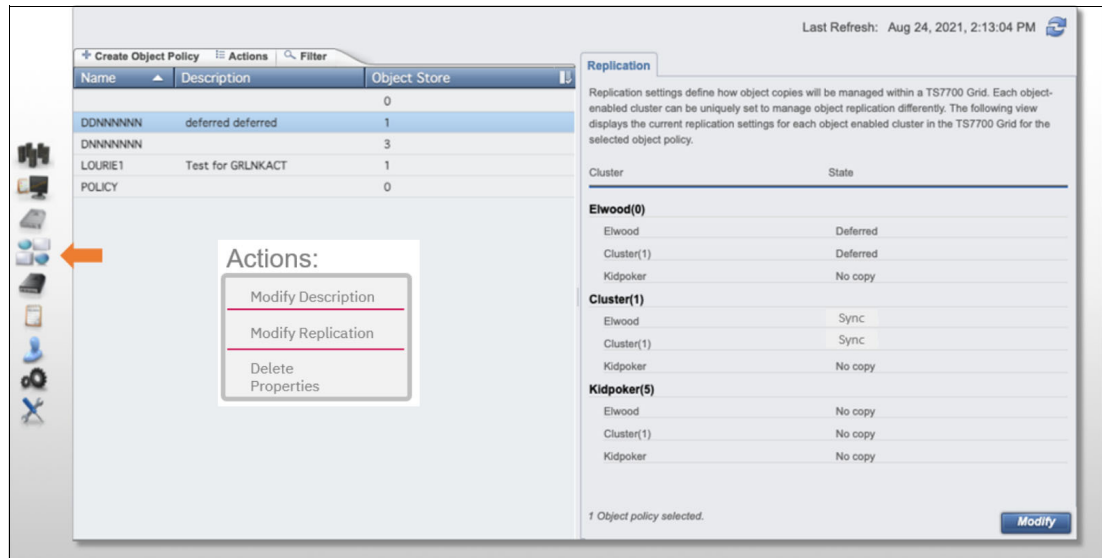
Figure 9-58 Object Store Page

To create an object store in the Object Store page, complete the following steps:

1. Click **Create Object Store** (see Figure 9-58).
2. Enter the following information:
 - Name of the object store in the Cloud Name field. The name of the object store should match the z/OS DFSMS Cloud Network Connection Construct name that is used to target TS7700.
 - Description of the object store in the Description field.
 - The object policy by selecting from the Object Policy drop-down list.

Click **Create**. A window confirms that Object Store was successfully created. Otherwise, return and correct the information, or click **Cancel** to exit.

Figure 9-59 shows the new grid scope page (under the Object Icon when a grid exists) where a user can update object policy changes for the entire grid.



Name	Description	Object Store
DDNNNNNN	deferred deferred	0
DDNNNNNN	deferred deferred	1
DNNNNNNN		3
LOURIE1	Test for GRLNKACT	1
POLICY		0

Figure 9-59 The "Object" page when a grid exists

User can create or alter policies that are to be used for object management. The following fields are available:

- ▶ Name: User-defined name for the policy
- ▶ Description: User description for the policy
- ▶ Object Store: Count of the number of object stores associated with that policy
- ▶ Replication: Select copy mode type of *Synchronous*, *Deferred* or *No copy* for each cluster in the grid

Note: Changes to policies affect future PUT operations.

For more information, see the following resources:

- ▶ *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583
- ▶ [Object store](#)



IBM TS7700 Management Interface operations: Part 2

This chapter and Chapter 9, “IBM TS7700 Management Interface operations: Part 1” on page 351 provide information about how to configure and operate the IBM TS7700 by using the Management Interface (MI).

For more information about TS3500 or TS4500 tape libraries, see *IBM TS4500 R8 Tape Library Guide*, SG24-8235.

Note: Support for TS1160 (Jaguar 6) and JE/JM media is included in the R5.3 and higher level of code.

This chapter includes the following topics:

- ▶ 10.1, “Physical icon” on page 449
- ▶ 10.2, “Constructs icon” on page 476
- ▶ 10.3, “Access icon” on page 492
- ▶ 10.4, “Settings icon” on page 512
- ▶ 10.5, “Service icon” on page 543

10.1 Physical icon

This section includes a description of monitoring and handling physical volumes in the TS7700T 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 10-1 on page 450.

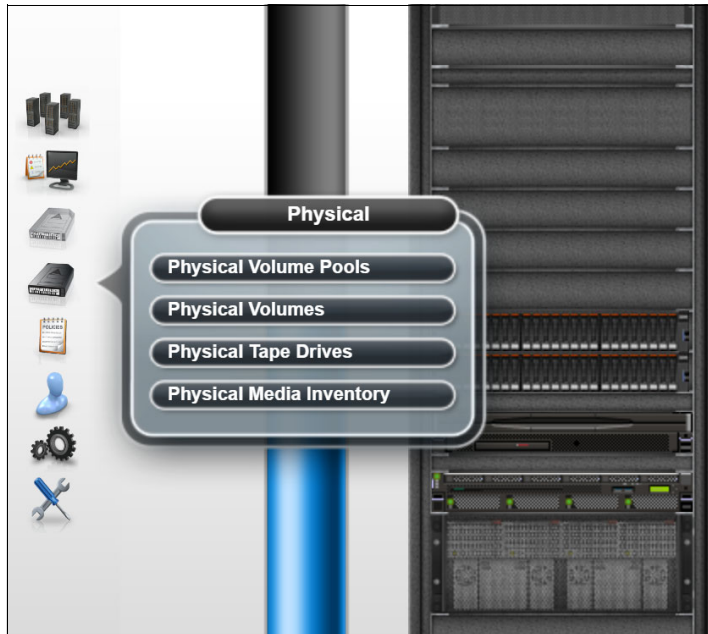


Figure 10-1 Physical icon

10.1.1 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 the following 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 10-2 show an example of the Physical Volume Pools window. You can use this window to view or modify settings for physical volume pools.

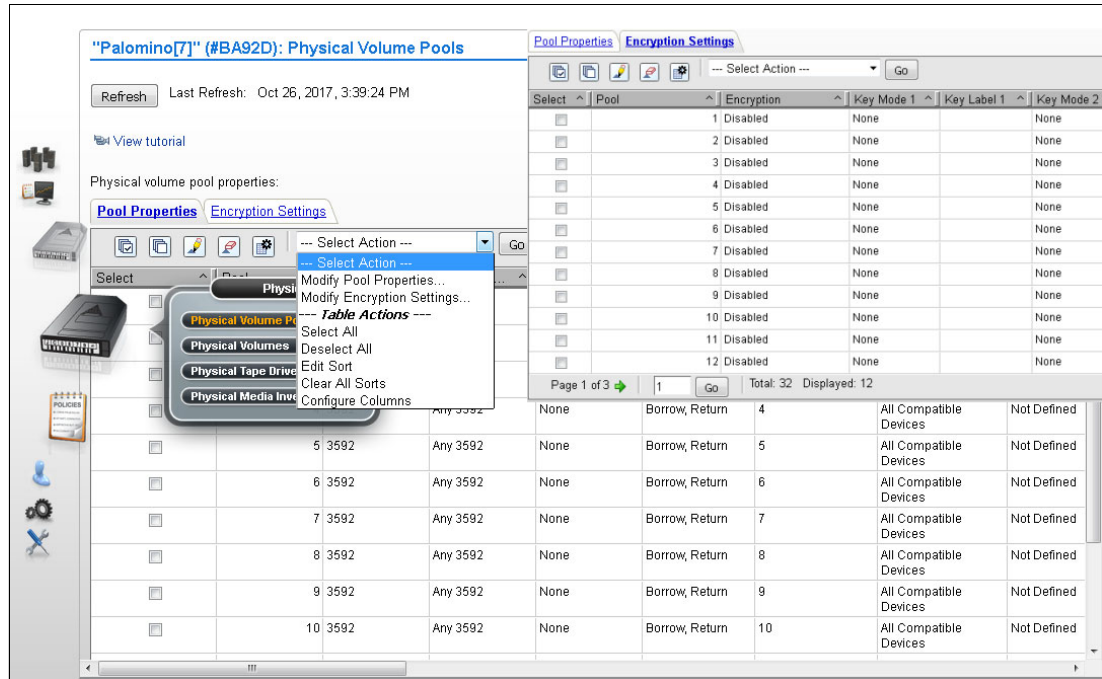


Figure 10-2 Physical Volume Pools Properties table

A link is available in the window for a tutorial that shows 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. The following message is displayed:

The cluster is not attached to a physical tape library.

The Physical Volume Pool Properties table displays the encryption settings and media properties for every physical volume pool that is defined in a TS7700T cluster. 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.

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 Figure 4-7 on page 207 for First and Second Media values that are based on drive configuration.

Table 10-1 shows the possible values for media type in a Tape Attach cluster (only the media type that is supported by the tape drives in the current configuration shows up).

Table 10-1 Possible values for media type field

Value	Explanation
Any 3592	Any media with a 3592 format. The only option available if the Primary Media type is any 3592.
None	This option is valid only when the Borrow Indicator value is <i>No Borrow, Return</i> or <i>No Borrow, or Keep</i> .
JA	Enterprise Tape Cartridge (ETC)
JB	Enterprise Extended-Length Tape Cartridge (ETCL)
JC	Advanced Type C Data (ATCD)
JD	Advanced Type D Data (ATDD)
JE	Advanced Type E Data (ATED)
JJ	Enterprise Economy Tape Cartridge (EETC)
JK	Advanced Type K Economy (ATKE)
JL	Advanced Type L Economy (ATLE)
JM	Advanced Type M Economy (ATLM)

- The primary media type can have the values that are shown in Table 10-1.
- 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 possible values are shown on Table 10-1.
- 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 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 must be set to the same value as the Copy Export pool. If the reclaim pool is modified, Copy Export disaster recovery (DR) capabilities can be compromised.

If the reclaim pool that is designated for the Copy Export pool must be modified, 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, 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 that is 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 exporting the volumes. This field is enabled only 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 value is also the default value for the Export Format field.

E06 Format of a 3592-E06 Tape Drive.

E07 Format of a 3592-E07/EH7 Tape Drive.

E08 Format of a 3592-E08/EH8 Tape Drive.

E09 Format of a 3592-60G/60F 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 checkbox 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 checkbox deactivates this function.
- **Age of Last Data Written:** The number of days the pool persisted without write access to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume features 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 checkbox deactivates this function.
- **Days Without Data Inactivation:** The number of sequential days that the data ratio of the pool was higher than the Maximum Active Data that is used to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume features 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 checkbox 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 option is new in Release 3.3. To modify pool properties, select the checkbox next to one or more pools that are shown on the Pool Properties tab, select **Modify Pool Properties** from the menu, and then 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 the 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 that corresponds 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 the default key.

To modify encryption settings, complete the following 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 then click **Go**.

Modify pool properties

Use this page to view or modify settings for physical volume pools, which manage the physical volumes that are used by the IBM TS7700T cluster. To modify properties for one or more physical volume pools, complete the following steps:

1. From the Physical Volume Pools window, click the **Pool Properties** tab.
2. Select the checkbox next to each pool to be modified.
3. Select **Modify Pool Properties** from the Physical volume pools drop-down menu.
4. Click **Go** to open the Modify Pool Properties window.

Note: Pools 1–32 are preinstalled. Pool 1 functions as the default pool and is used if no other pool is selected. All other pools must be defined before they can be selected.

5. You can modify values for any of the following fields:
 - **Media Class:** The supported media class of the storage pool. The value that is used is 3592.
 - **First Media (Primary):** The primary media type that the pool can borrow 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. The possible values are shown in Table 10-1 on page 452.
 - **Second Media (Secondary):** The second choice of media type that the pool can borrow from. The options that are shown exclude the media type chosen for First Media. The
 - **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 when emptied.
 - **Borrow, Keep:** A cartridge is borrowed from the CSP and retained, 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.
 - **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 should be set to the same value as the copy export pool. If the reclaim pool is modified, copy export DR capabilities can be compromised.

If the reclaim pool that is designated for the copy export pool must be modified, 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, 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 your 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 supported export 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 that is 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 enabled only if the value in the Export Pool field is Copy Export. The following values are possible:
 - Default: The highest common format that is supported across all drives in the library. This value is also the default value for the Export Format field.
 - E06: Format of a 3592 E06 Tape Drive.
 - E07: Format of a 3592 E07/EH7 Tape Drive.
 - E08: Format of a 3592 E08/EH8 Tape Drive.
 - E09: Format of a 3592 60G/60F Tape Drive.
- Days Before Secure Data Erase: The number of days that 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. This timer 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 checkbox deactivates this function.
- Days Without Access: The number of days that a determined pool can stay without being accessed for a recall. When a physical volume reaches that number of days without being accessed, it is set 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 checkbox deactivates this function.

Note: This control is applied to the reclamation of both sunset and R/W media.

- Age of Last Data Written: The number of days the pool persisted without write access to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume features 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 checkbox deactivates this function.

Note: This control is applied to the reclamation of both sunset and R/W media.

- Days Without Data Inactivation: The number of sequential days that the data ratio of the pool was higher than the Maximum Active Data that was used to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume features 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 checkbox deactivates this function. If deactivated, this field is not used as a criteria for reclaim.

Note: This control is applied to the reclamation of both sunset and R/W media.

- 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 selected.
- Reclaim Threshold: The percentage used to determine when to perform reclamation of free storage on a stacked volume. When the amount of active data on a physical stacked volume drops below this percentage, a reclaim operation is performed on the stacked volume.

Physical volumes hold between the threshold value and 100% of data. For example, if the threshold value is 35% (the default), the percentage of active data on the physical volumes is $(100\% - 35\%) / 2$ or 15%. Setting the threshold too low results in more physical volumes being needed. Setting the threshold too high might affect the ability of the TS7700 Tape Attach to perform host workload because it is using its resources to perform reclamation. Experiment to find a threshold that matches your needs.

The valid range of possible values is 0–95(%) and can be entered in 1% increments. The default value is 35%. If the system is in a heterogeneous tape drive environment, this threshold is for R/W media.

- Sunset Media Reclaim Threshold: This field is always available, but it affects only sunset media. The percentage used to determine when to perform reclamation of free storage on a stacked volume. When the amount of active data on a physical stacked volume drops below this percentage, a reclaim operation is performed on the stacked volume.

Physical volumes hold between the threshold value and 100% of data. For example, if the threshold value is 35% (the default), the percentage of active data on the physical volumes is $(100\% - 35\%) / 2$ or 15%. Setting the threshold too low results in more physical volumes being needed. Setting the threshold too high might affect the ability of the TS7700 Tape Attach to perform host workload because it is using its resources to perform reclamation. Experiment to find a threshold that matches your needs.

The valid range of possible values is 0–95(%) and can be entered in 1% increments. The default value is 35%. If the system is in a heterogeneous tape drive environment, this threshold is for sunset media.

Note: If the system contains TS1140, TS1150, or TS1160 tape drives, the system requires at least 15 scratch physical volumes to run reclamation for sunset media.

6. To complete the operation, click **OK**. To abandon the operation and return to the Physical Volume Pools window, click **Cancel**.

Modify encryption settings

Use this page to modify encryption settings for the physical volume pools that manage the physical volumes that are used by the IBM TS7700T cluster.

To watch a tutorial that shows how to modify pool encryption settings, click **View tutorial** on the Physical Volume Pools page.

To modify encryption settings for one or more physical volume pools, complete the following steps:

1. From the Physical Volume Pools page, click the **Encryption Settings** tab.
2. Select each pool to be modified.
3. Click **Select Action** → **Modify Encryption Settings**.
4. Click **Go** to open the Modify Encryption Settings window.
5. Modify values for any of the following fields:
 - Encryption: The encryption state of the pool. The following values are available:
 - 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 the encryption key server default key: Select this option to populate the Key Label field by using a default key provided by the encryption key server.

Note: Your encryption key server software must support default keys to use this option.

This option is available before Key Label 1 and Key Label 2 fields. Select this option for each label to be defined by using the default key. If this option is selected, the following fields are disabled:

- Key Mode 1
 - Key Label 1
 - Key Mode 2
 - Key Label 2
- Key Mode 1: Encryption Mode 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 that corresponds to its associated public key.
 - None: Key Label 1 is disabled.
 - -: The default key is in use.
 - Key Label 1: The current encryption key Label 1 for the pool. The label must consist of ASCII characters and cannot exceed 64 characters. Leading and trailing blanks are removed, but an internal space is allowed. Lowercase characters are internally converted to uppercase upon storage. 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.

- Key Mode 2: Encryption Mode used with Key Label 2. The following values are possible 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 that corresponds to its associated public key.
 - None: Key Label 2 is disabled.
 - -: The default key is in use.
 - Key Label 2: The current encryption key Label 2 for the pool. The label must consist of ASCII characters and cannot exceed 64 characters. Leading and trailing blanks are removed, but an internal space is allowed. Lowercase characters are internally converted to uppercase upon storage. Therefore, 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**.

For more information about a practical example of how to configure the TS7700 Virtual Engine to apply external key management, see Appendix J, “Configuring externally managed encryption” on page 1013.

10.1.2 Physical volumes

This section provides a description of monitoring and manipulating physical volumes in the TS7700T clusters.

Relabeling cartridges with the TS7700T

With R5.0 or higher level of code, the user can change the external label of a cartridge and have the TS7700 to initialize it again, writing a new internal label to match the new bar code label when the physical cartridges are inserted into a TS7700T cluster. This capability is useful when tape cartridges are reused from decommissioned equipment, or scratch cartridges are reused to predefined volume ranges in the TS7700T environment.

The TS7700T cluster always checks the internal label of a cartridge and compares it with the external bar code label before using a physical tape. If a mismatch exists between the internal and the external labels, that cartridge is marked in error and rejected by the TS7700T, being ejected from the physical tape library.

The relabeling function allows the user to rewrite the internal label to match the new bar code of such cartridge, which makes it a good candidate for use in the TS7700T by way of a **LI REQ** command.

The **LI REQ** command must be issued before the TS7700T uses that cartridge (otherwise, an internal versus external label mismatch occurs and the cartridge is rejected).

Use **LI REQ PVOL, zzzzzz, RELABEL, YES/NO** to relabel physical tape cartridges. **Yes** queues a physical volume for relabeling, and **NO** removes the corresponding entry from the queue.

Suggestion: Assign the cartridges to be relabeled to an unused pool that is defined as no-borrow/no-return. After inserting the cartridges in the tape library, issue the **LI REQ** command to relabel the cartridges. After all **LI REQ** commands are issued, move the cartridges to the CSP, which makes them available as empty scratch to the TS7700T.

The Physical volumes 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 10-3 shows the navigation and the Physical Volumes window. The Physical Volumes page is available for all TS7700 tape-attached models.

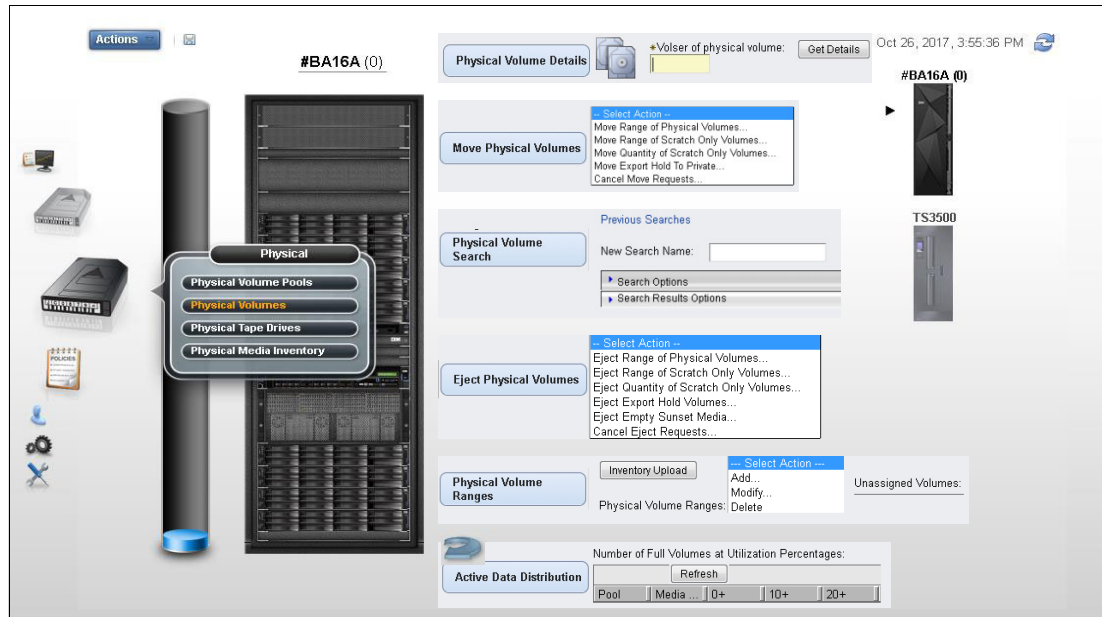


Figure 10-3 Physical Volumes navigation and options

The following options are available selections under Physical Volumes:

- ▶ Physical Volume Details window
- ▶ Move Physical Volumes window
- ▶ Eject Physical Volumes window
- ▶ Physical Volume Ranges window
- ▶ Physical Volume Search window
- ▶ Active Data Distribution

These options are described next.

Physical Volume Details window

Use this window to obtain more information about a physical stacked volume in the TS7700T cluster. You can download the list of virtual volumes in the physical stacked volume that is displayed by clicking **Download List of Virtual Volumes** under the table.

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 possible values are shown on Table 10-1 on page 452.

Note: JD (ATDD) and JL (ATLE) media types are available only if the highest common format (HCF) is set to E08 or later. JE (ATED) and JM (ATME) media types are available only if the highest common format (HCF) is set to E09 or later.

Table 10-2 contains the possible values for the recording format that can be found in the **Physical Volume Details** page.

Table 10-2 Possible values for recording format

Recording Format	The format used to create the media
Undefined	The recording format that is used by the volume is not recognized as a supported format.
J1A	Written by a J1A drive
E05	Written by an E05 drive
E05E	Written by an Encryption-capable E05 drive
E06	Written by an E06 drive
E06E	Written by an Encryption-capable E06 drive
E07	Written by an E07 drive
E07E	Written by an Encryption-capable E07 drive
E08	Written by an E08 drive
E08E	Written by an Encryption-capable E08 drive
55F	Written by an 55F drive
55FE	Written by an Encryption-capable 55F drive
60F	Written by a 60F drive
60FE	Written by an Encryption-capable 60F drive

- ▶ **Recording Format:** The format that is used to write the media. Possible values are shown on Table 10-2.
- ▶ **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 was 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 that is marked Copy Exported. The data mover then reclaims the virtual volumes from the primary copies.
 - Copy Export No Files Good: The physical volume was 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 present in the library inventory, but now cannot 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 no labels are present, the volume is not encrypted. If the EK used is the default key, the value in this field is the default key.
 - ▶ Encrypted Time: The date the physical volume was first encrypted by 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 exists.
 - ▶ Mount Count: The number of times the physical volume was mounted since being inserted into the library.
 - ▶ Virtual Volumes Contained: The 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 timestamp portion of the database backup name.

Move Physical Volumes window

Use this option to move a range or quantity of physical volumes that is used by the TS7700T to a target pool, or cancel a previous move request.

The Select Move Action menu provides the following options for moving 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 that Number of Volumes, Source Pool, and Target Pool are selected. 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 that a Source Pool is selected.
- ▶ **Cancel Move Requests:** Cancels any previous move request.

Note: The Move physical volume option applies to private media only, not scratch tapes. When moving out of a SDE-enabled pool, the function also causes a physical volume to be erased, although the number of days that are specified did not yet elapse. This process includes returning borrowed volumes.

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 this option or the Existing Ranges option to define the range of volumes to move, but not both. Specify the following range:
 - To: VOLSER of the first physical volume in the range to move.
 - FromL VOLSER of the last physical volume in the range to move.
- ▶ **Existing Ranges:** The list of physical volume ranges. The user can use 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 **Select Move Action** → **Cancel Move Requests**. 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

Use this page to eject a range or quantity of physical volumes that are used by the TS7700T, or to cancel a previous eject request.

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 in the same pool. The stacked volume that is selected for ejection does not go through Secure Data Erase before ejection.

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 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**. The eject operation parameters include the following parameters:

- ▶ **Volume Range:** The range of physical volumes to eject. The user can use this option or the Existing Ranges option to define the range of volumes to eject, but not both. Specify the following 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 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 occurs. 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 is 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 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 **Select Eject Action** → **Cancel Eject Requests**. 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

Use this window to view physical volume ranges or unassigned physical volumes in a library that is attached to a TS7700T cluster. Figure 10-3 on page 460 shows the options that are available on this page.

When working with volumes that were recently added to the attached TS3500 tape library that are not shown in the Physical Volume Ranges window, click **Inventory Upload**. The physical inventory from the defined logical library in the tape library is then uploaded to the TS7700T, which repopulates the Physical Volume Ranges window.

Important: When inserting a VOLSER that belongs to a defined tape attach TS7700 range, it is presented and inventoried according to the setup that is in place. If the newly inserted VOLSER does not belong to any defined range in the TS7700T, an intervention-required message is generated, which requires the user to correct the assignment for this VOLSER.

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. Table 10-1 on page 452 shows the possible values.

Note: JA and JJ media are supported for only read-only operations with 3592 E07/EH7 tape drives. 3592-E08/EH8 does not support JA, JJ, or JB media. 3592-60G/60F support the 4 TB format with JC media read-only.

- ▶ 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 that is associated with a specific physical volume.
 - Media Type: The media type for all volumes in a VOLSER range. Table 10-1 on page 452 shows the possible values.

Note: JA and JJ media are supported for only read-only operations with 3592 E07 tape drives. 3592-E08 does not support JA, JJ, or JB media.

- ▶ Pending Eject: Whether the physical volume that is 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 TS7700T.

Physical Volume Search window

Use this window to search for physical volumes in a TS7700T cluster according to one or more identifying features. Figure 10-3 on page 460 shows the options that are available on this page. Click **Previous Searches** to view the results of a previous query on the Previous Physical Volumes 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 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. Table 10-1 on page 452 shows the possible values.
- ▶ **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 possible values are shown on Table 10-2 on page 461.
- ▶ **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. Sometimes, 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): 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 possible values are shown on Table 10-1 on page 452.
- ▶ **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 possible values are shown at Table 10-2 on page 461.
- ▶ **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. Sometimes, 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 field is a text field. The following values are valid:
 - A name that is identical to the first or second key label on a physical volume.
 - Any physical volume that was encrypted by using the designated key label is included in the search.
 - Search for the default key. Select this option to search for all physical volumes that were encrypted by 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 flags 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 checkbox 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 start a new physical volume search. After the search is started but before it completes, the Physical Volume Search window displays the following message:

The search is 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 in the Physical Volume Search window until the in-progress search completes or is canceled.

To check the progress of the search being run, click **Previous Search Results** 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

Use this window to view the distribution of data on physical volumes that are marked full on a TS7700T cluster. The distribution can be used to select an appropriate reclaim threshold. The Active Data Distribution window displays the usage percentages of physical volumes in increments of 10%.

Number of Full Volumes at Utilization Percentages window

The tables in this page show the number of physical volumes that are marked as full in each physical volume pool, according to the percentage of volume that is used. The following fields are displayed:

- ▶ **Pool:** The physical volume pool number. Click this number to display a graphical representation of the number of physical volumes per utilization increment in a pool. If the user clicks the pool number, the Active Data Distribution subwindow opens.

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 available 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: Click this percentage 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, which are 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 possible values are shown on Table 10-1 on page 452.
- ▶ **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 example, 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.

10.1.3 Physical Tape Drives window

Use this window to view a summary of the state of all physical drives that are accessible to the TS7700T cluster.

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 10-4 shows the Physical Tape Drives window.

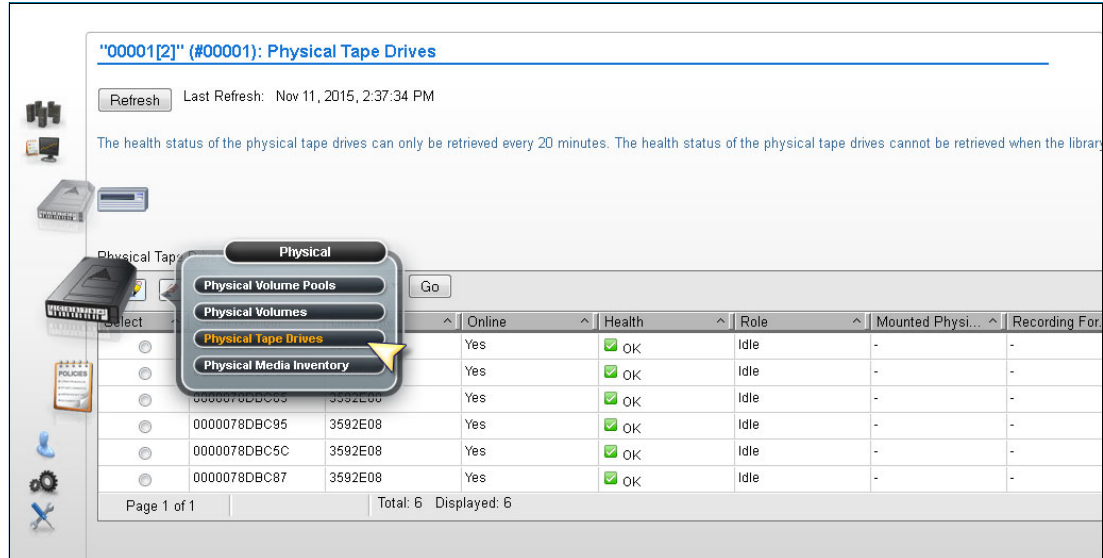


Figure 10-4 Physical Tape Drives window

The Physical Tape Drives table displays status information for all physical drives that are 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 possible values are shown on Table 10-3:

Table 10-3 Drive type values.

Drive Type	Description	Name
3592-J1A	3592 model J1A (Gen. 1)	3592
3592-E05	3592 model E05 (Gen. 2)	TS1120
3592-E05E	3592 model E05, encryption capable	TS1120
3592-E06	3592 model E06 (Gen. 3)	TS1130
3952-E07	3592 model E07 (Gen. 4)	TS1140
3952-E08	3592 model E08 (Gen. 5)	TS1150
3592-60F	3592 model 60F (Gen. 6)	TS1160
3592-60G	3592 model 60G (Gen. 6)	TS1160

- ▶ **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 must be taken to correct the errors.
 - **DEGRADED:** The drive is operational, but 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 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 that is mounted by the drive.

- ▶ **Recording Format:** The format in which the drive operates. The possible values are shown in Table 10-2 on page 461.
 - **Not Available:** The format cannot be determined because no physical media exists in the drive or the media is being erased.
 - **Unavailable:** The format cannot be determined because the Health and Monitoring checks were yet completed. Refresh the current window to determine whether the format state 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 more information for a specific selected drive, see the Physical Drives Details table on the Physical Tape Drive Details window. Complete the following steps:

1. Select the radio button next to the serial number of the physical drive in question.
2. Click **Select Action** → **Details**.
3. Click **Go** to open the Physical Tape Drives Details window.

The Physical Drives Details table displays the following 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 possible values are shown on Table 10-3 on page 472.
- **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 TS7700T. The following values are possible:
 - **OK:** The drive is fully functioning.
 - **WARNING:** The drive is functioning, but reporting errors. Action must 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:** The VOLSER of the physical volume mounted by the drive.

- Recording Format: The format in which the drive operates. The possible values are shown in Table 10-2 on page 461.
- 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 installed.
- 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 can encrypt.
 - 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.

10.1.4 Physical Media Inventory window

Use this window to view physical media counts for media types in storage pools in the TS7700.

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 10-5 shows the Physical Media Inventory window.

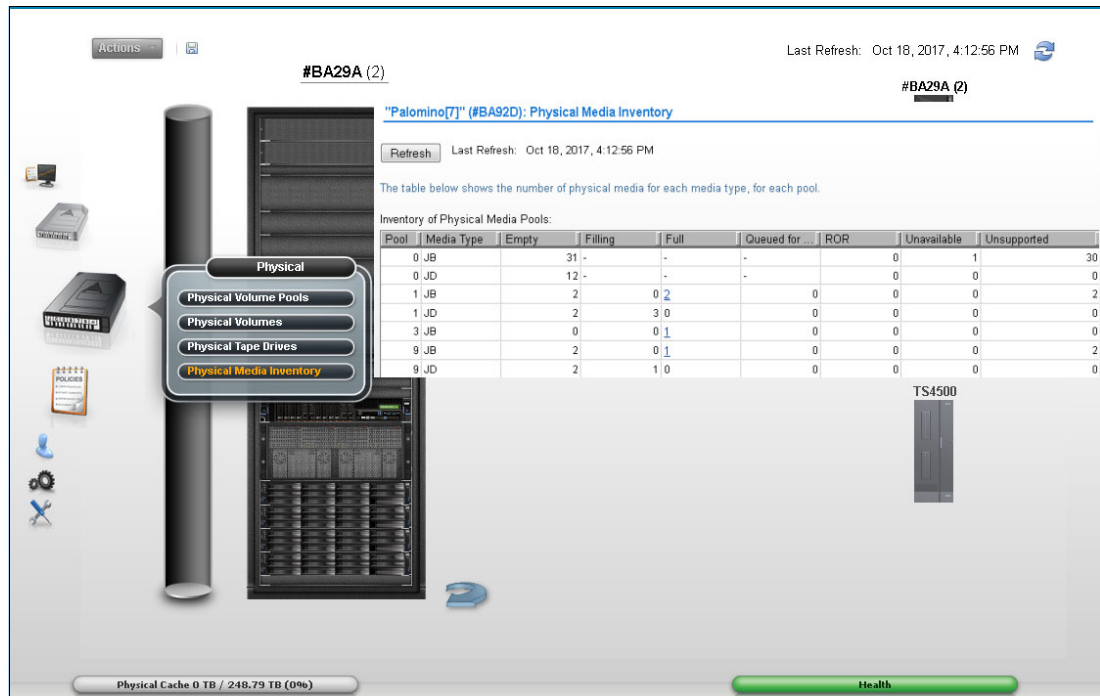


Figure 10-5 Physical Media Inventory window

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 that is defined for the pool. A storage pool can feature multiple media types and each media type is displayed separately. The possible values are shown at Table 10-1 on page 452.
- ▶ 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: Click the Full field 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 must 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 TS7700T. 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.

10.2 Constructs icon

TS7700 storage constructs are described in this section. Figure 10-6 shows the Constructs icon and the options that are available under it.

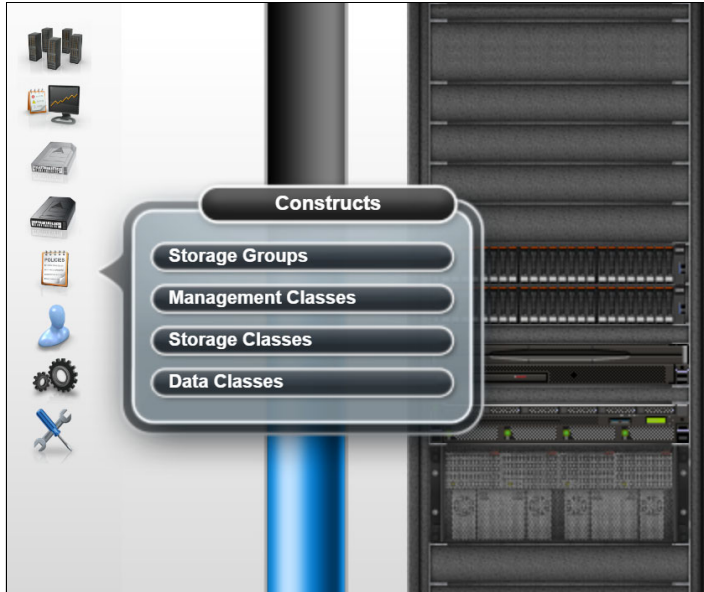


Figure 10-6 Constructs icon

10.2.1 Storage Groups window

Use the window that is shown in Figure 10-7 on page 477 to add, modify, or delete an SG. The figure shows a TS7760C (cloud attach) example.

The SGs table displays all SGs that are 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.

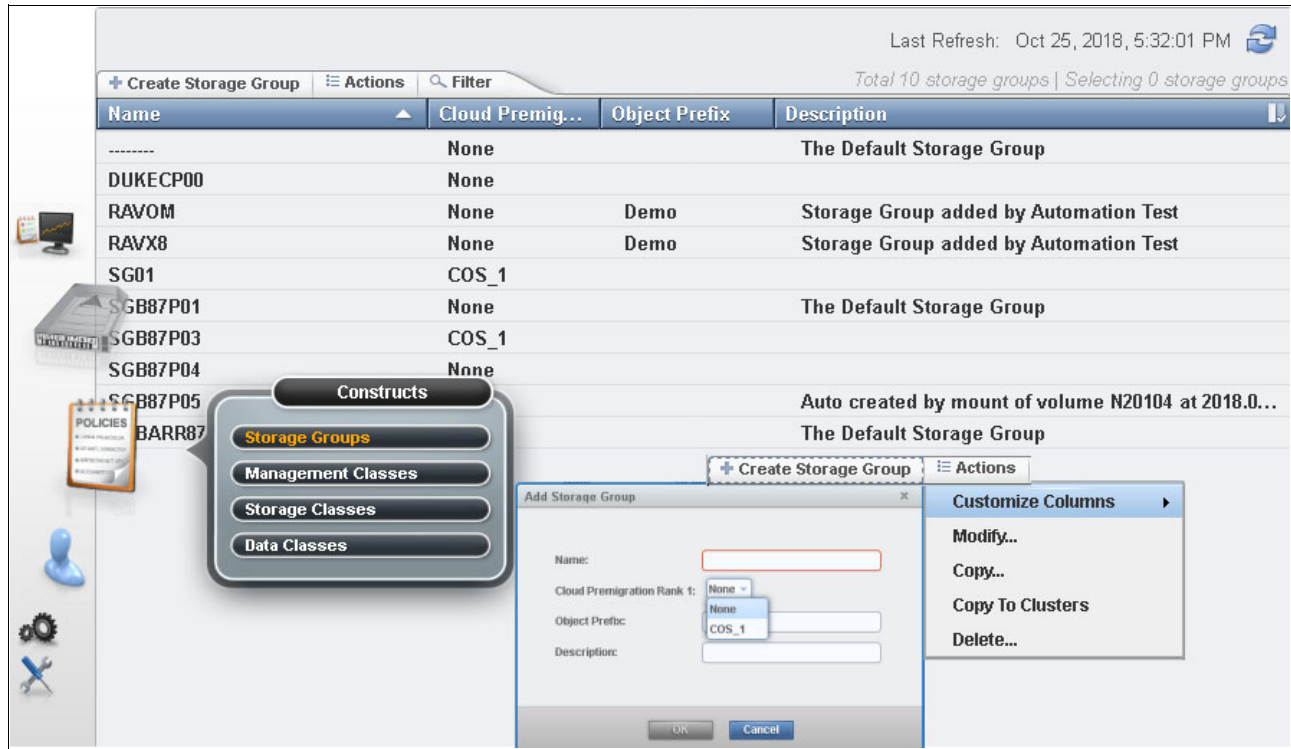


Figure 10-7 MI Storage Groups window with Cloud Tier

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 %. The first character of this field cannot be a number. This is the only field that cannot be modified once added.
- ▶ **Cloud Premigration Rank**
- ▶ This field represents which Cloud Pool is used for logical volumes that offload to an object store. The logical volume's corresponding storage class determines which disk partition is used and this field determines which object store pool is used. Logical volumes which exist in the resident only partition ignore this setting, unless they are moved to a partition through a mount/demount or LIB REQ PARTRFSH operation.
- ▶ **Object Prefix**
- ▶ This field represents an optional prefix to add to the target object store key name or object name. This allows you to separate objects by workload within the target object store.
- ▶ **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 is 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 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 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.

Important Note: The cloud premigration rank and object prefix values should be setup identically within all clusters in a grid when cloud support is enabled.

When a host creates or accesses a logical volume, the mount point cluster's storage group's cloud settings are used to determine what cloud premigration rank and object prefix name will be used across the entire grid. Object stores can potentially be accessed from all clusters in a grid, so where content resides must be in complete agreement when storing data in an object store. This is why the mount point cluster's cloud settings are used for all clusters for a given logical volume. If two clusters don't agree, it can result in different behaviors for common workloads or excessive movement of data between pools when data is accessed from a cluster other than the one that created it.

10.2.2 Management Classes window

Use this window to define, modify, copy, or delete the MC that defines the TS7700 copy policy for volume redundancy (see Figure 10-8). The table displays the copy policy that is in effect for each component of the grid.

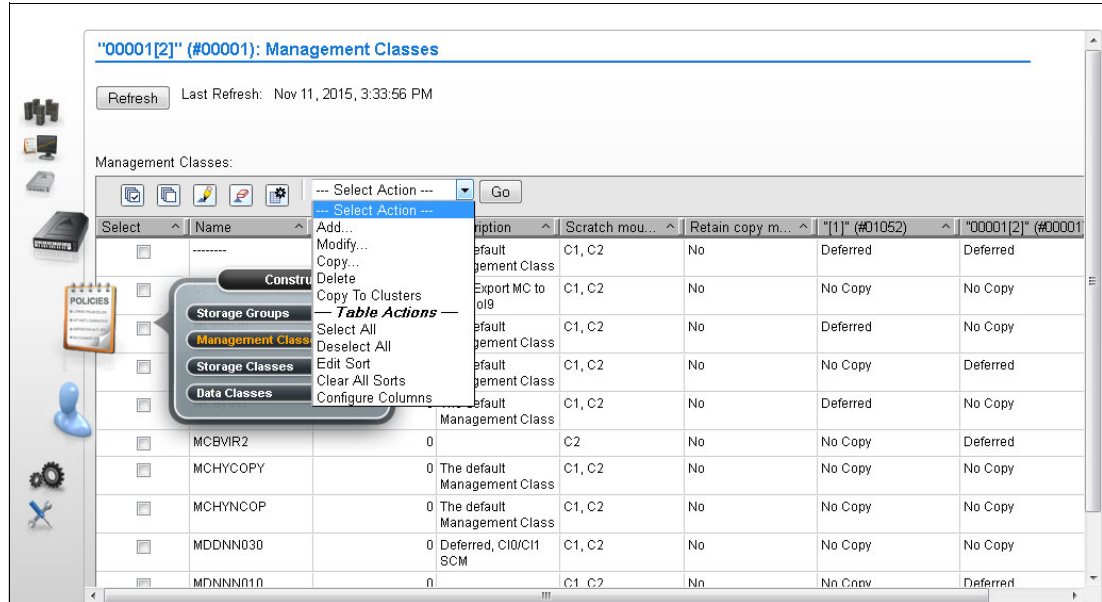


Figure 10-8 MI Management Classes window on a grid

Use the Select Action drop-down menu on the Management Classes table to

- ▶ Create a management class
- ▶ Modify an existing management class
- ▶ Copy an existing management class

- ▶ Copy one or more existing management classes from the accessing cluster to another cluster in the grid
- ▶ Delete one or more existing management classes.

The default management class can be modified, but cannot be deleted. The default management class has a name that is “-----”.

The status information that is displayed in the Management Classes table includes:

- ▶ **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 field 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.

Note: Refer to the Storage Groups page to define the primary pool. If the Copy Export function is used, set the secondary pool to a pool other than the primary pool of a defined storage group. For the TS7700 Tape Attach, the secondary pool is applied only to logical volumes that do not reside in cache partition 0.

- ▶ **Description:** A description of the MC definition. The value in this field must be 1–70 characters in length.
- ▶ **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.

Note: this function is available if the Setting Assist Scratch (SAA) library request option is enabled across the grid. **Disabled** is the default setting for SAA.

When using **SAA**, enough devices that are connected to the scratch mount candidate clusters should be online at the host (otherwise, tape jobs might go into allocation recovery at the host).

- ▶ **Retain Copy Mode:** Whether previous copy policy settings on non-fast-ready virtual volume mounts are retained. Possible values are:
 - Yes: Previous copy policy settings on non-fast-ready virtual volume mounts are retained
 - No: Previous copy policy settings on non-fast-ready virtual volume mounts are not retained.

Figure 10-9 shows the MCs options, including the Time Delayed option.

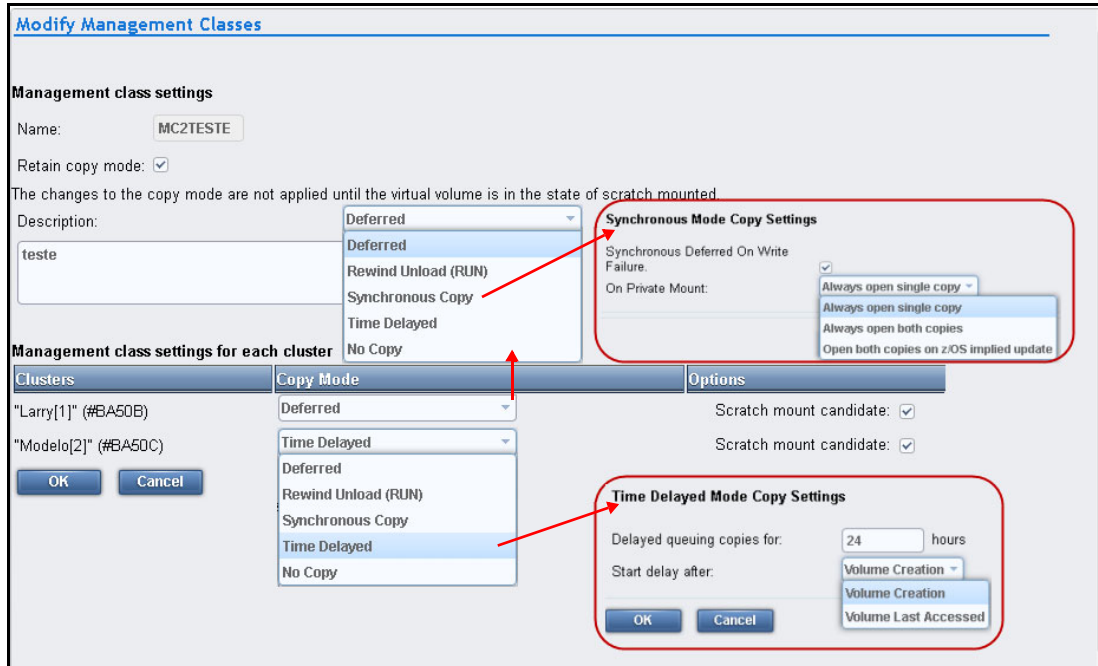


Figure 10-9 Modify Management Classes options

The Management class settings that are shown in Figure 10-9 for each cluster are:

- ▶ Clusters: The nickname and ID of the cluster to which the settings apply.
 - Copy Mode: The copy mode that is used in the volume duplication.
 - Deferred: Volume duplication occurs later based on the internal schedule of the copy engine.
 - Rewind Unload (RUN): Volume duplication occurs when the RUN command is received. The command only returns after the volume duplication completes successfully.
 - Synchronous Copy: Volume duplication occurs upon any synchronous mode copy operation.
- ▶ Time Delayed: Volume duplication occurs after the user-specified delay period passes.
- ▶ Copy mode:
 - Delay Queuing Copies for [X] Hours: The number of hours copy queuing is delayed if Time Delayed Mode Copy is selected. This value is a number in the range of 1 to 65,535.
 - Start Delay After: The trigger that initiates the delayed start time if Time Delayed Mode Copy is selected. Possible values include
 - Volume Create: The time-delayed start occurs when the volume is created.
 - Volume Last Accessed: The time-delayed start occurs when the volume was last accessed.
 - No Copy: No volume duplication occurs if this action is selected.

Synchronous Mode Copy

Synchronous Mode Copy settings specify how the library or virtual tape drive behaves when the Copy Mode value is Synchronous Copy and synchronization between the two synchronized (“S”) locations is not satisfied. Synchronous Mode Copy settings determine whether the library opens the volume on both TVC clusters when a private mount occurs. These options are only visible when synchronous mode copy is enabled.

Default settings

By default, the synchronous-mode-copy clusters fail mount and tape operations if two copies of a volume cannot be maintained during an update (synchronous failure). When the synchronous failure setting is used, a zero RPO is provided for the target workload, independent of failures. Consider the following circumstances when using the default strict synchronization behavior include:

- ▶ If the failure to synchronize is detected after the mount has already occurred, then tape operations fail to the targeted volume until a RUN occurs and a demount command is issued.
- ▶ If content was written before the synchronization failure, then previous content on the emulated volume up to the last successful tape synchronization operation point is considered persistently synchronized and can be accessed later from either “S” consistency point.
- ▶ If either “S” consistency point is unavailable, then scratch mount operations fail.

Synchronous Deferred On Write Failure

Enable this 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 Synchronous Deferred On Write Failure option is not checked and a write failure occurs at either of the “S” locations, then host operations fail.

Note: An “R”, “D,” or “T” site is chosen as the primary consistency point only when both “S” locations are unavailable.

On Private Mount: Always open single copy

By default, synchronous mode copy opens only one TVC during a private mount. The best TVC choice is used to satisfy the mount. The best TVC choice selection is made with location preferences in this order: synchronized [“S”], RUN [“R”], deferred [“D”], and time-delayed [“T”]. If a write operation occurs, the job enters the synchronous-deferred state regardless of whether the Synchronous Deferred On Write Failure option is enabled.

On Private Mount: Always open both copies

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 the disk cache within those locations. The Always open both copies option is useful for applications that require synchronous updates during appends. Private mounts can be affected by cache misses when this option is used. Other circumstances to consider include:

- ▶ If a private mount on both locations is successfully opened, then all read operations use the primary location. If any read fails, then the host read also fails and no failover to the secondary source occurs unless a z/OS DDR swap is initiated.
- ▶ If a write operation occurs, both locations receive write data and 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.

On Private Mount: Open both copies on z/OS implied update

Open both previously written “S” locations only when requested by the host to do so. This takes place when the mount request from the host has either write from BOT or update intent specified.

Synchronous Mode Copy includes the following settings:

- ▶ Synchronous Deferred on Write Failure:
 - The copy that is deferred on the unavailable cluster in a synchronous copy pair. By default, synchronous deferred on write failure is disabled and only visible when synchronous copy mode is used. When a private mount occurs, these options determine whether one or both tape volume cache (TVC) synchronized ('S') locations are opened.
 - **Always open single copies**
Only one TVC location is opened during a private mount. The best TVC choice is used to satisfy the mount in the following order: 'S', RUN ('R'), and deferred ('D').
 - **Always open both copies**
Both previously written “S” locations are opened when a private mount occurs.
- ▶ On Private Mount
 - Open both copies on z/OS implied update:** Both previously written 'S' locations are opened when a private mount occurs unless the host implies an update to the last updated mount. An update is implied when the host sets either the intend-to-update flag or the write-from-beginning-of-tape flag on the last updated mount.
 - If an update is implied but does not occur, then the mount either enters the synchronous-deferred state, or the mount fails.

Use the drop-down menu on the Management Classes table to add a management class, modify or copy an existing management class, or delete one or more existing management classes.

To add an MC, complete the following steps:

1. Select **Add** from the Management Class menu that is shown in Figure 10-8 on page 478 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 MC to each cluster in the TS7700 Grid.

To modify an MC, complete the following steps:

1. Select the checkbox 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**. Any of the fields that are listed in the MCs' table can be changed by the user *except* the MC name.

To delete one or more MCs, complete the following steps:

1. Select the checkbox 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.

10.2.3 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 10-10. Also, you can use this window to copy an SC to the same cluster that is 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 10-10 shows the window in a TS7700T model.




Figure 10-10 MI Storage Classes window on a TS7700T

The SC 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 or cache partition that is selected by the MI. Figure 10-11 shows appearance for the Create Storage Class box in different TS7700 partitions.

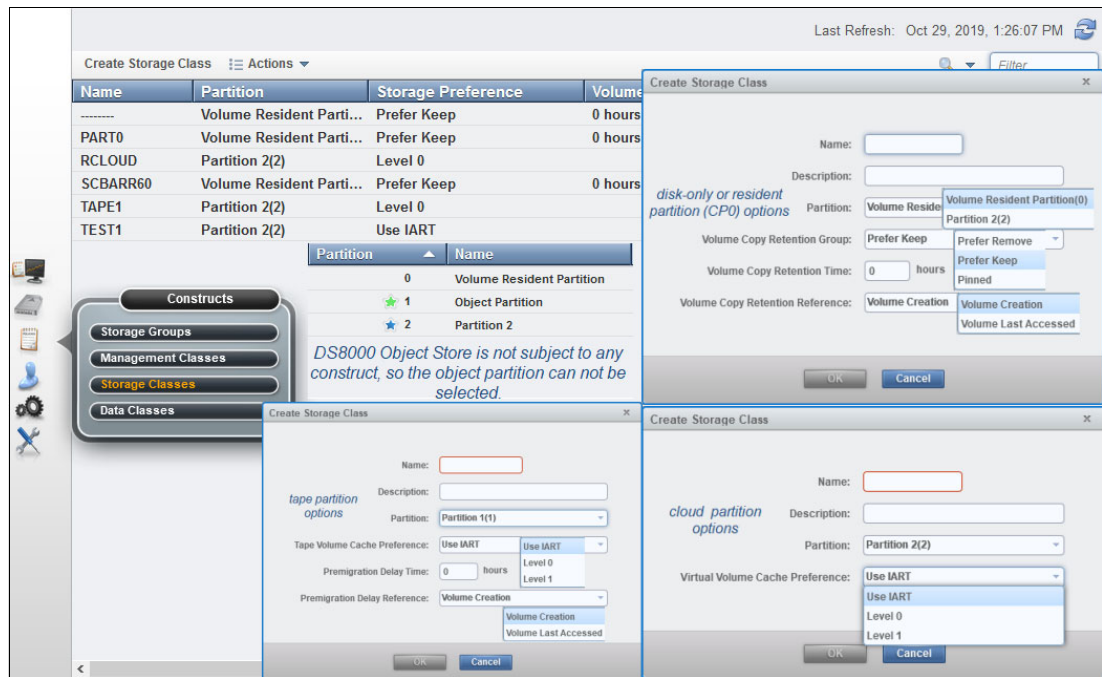


Figure 10-11 The Create Storage Class window and options

The default SC can be modified, but cannot be deleted. The default SC includes dashes (-----) as the symbolic name.

The SC 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 field 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 TS7700T or TS77C (tape or cloud attach) cluster. 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.

Note: A DS8000 object partition cannot be selected.

A dash (-) indicates that the Storage Class contains a partition that was deleted. Any volumes that are assigned to go to the deleted partition are redirected to the primary 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 when 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 level 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 timestamp 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 if the TS7700 cluster attaches to a physical library.
 - ▶ Premigration Delay Reference: The volume operation that establishes the timestamp 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 that is 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 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 the cache before becoming a candidate for removal. This field is displayed only 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 changing 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 visible only if the selected cluster is a TS7700T. 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 timestamp from which Volume Copy Retention Time is calculated. 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.

This field is disabled if the Volume Copy Retention Group displays a value of Pinned.

10.2.4 Data Classes window

This page is used to define, modify, copy, or delete a TS7700 Data Class for volume sizes and LWORM policy assignment. Data classes are used to automate storage management through the classification of data sets (see Figure 10-21 on page 491). On this page, the user can choose between the traditional compression (FICON) or one of the compression algorithms that optimize the host data compression by selecting the compression method.

The introduction of the LWORM Retention added many more settings and complexity to the original Data Classes window. Therefore, with R5.4 code, the Data Classes window has been redesigned. The new page presentation is split in two sections: a data grid section on the left that shows the existing Data Classes and an inspector area to the right.

Note: The new Data Classes window is only accessible when all clusters in the grid are at R5.4 or later level of code. In a mixed code grid configuration, the R5.4 cluster provides the pre-R5.4 version of the data class menu.

The new MI Data Class page includes the support for LWORM Retention settings and support for 65000 MiB logical volumes. No feature code is required to use these changes, LWORM Retention settings and support for 65000 MiB logical volumes are available when all clusters in the grid are at R5.4. Any other configuration can support the function by using an RPQ.

R5.4 implements Dual control for Data Class settings to ensure the LWORM Retention settings are safeguarded from a single user malicious or accidental change to these sensitive settings. The existing Data Class settings are also included in that protection.

Figure 10-12 on page 487 shows the Data Classes page in R5.4.

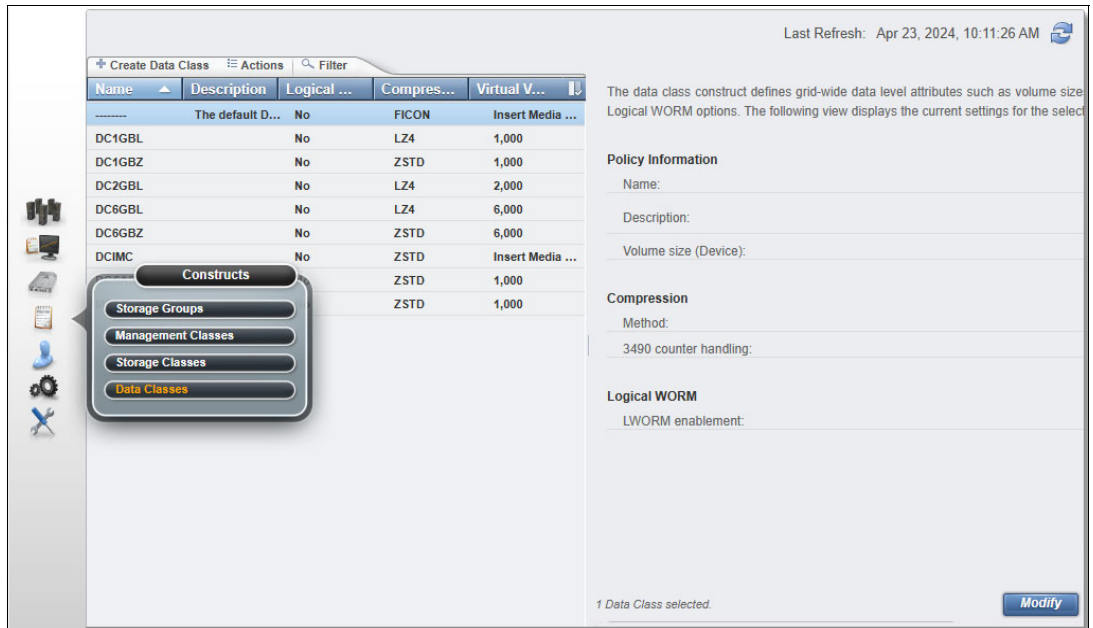


Figure 10-12 TS7700 MI Data Classes page

The following actions are supported in the Data Classes page:

- ▶ Modify description: Provides the ability to quickly change the description of a data class
- ▶ Modify properties: Allows the user to change the remaining properties of a data class
- ▶ Copy: Used to copy an existing data class record into a new record for duplication. The data class name is required to be changed in the operation (no duplicate names allowed).
- ▶ Delete: Delete a data class record
- ▶ Properties: Displays the properties of a selected data class

Supported actions and modify a selected data class button are shown in Figure 10-13:

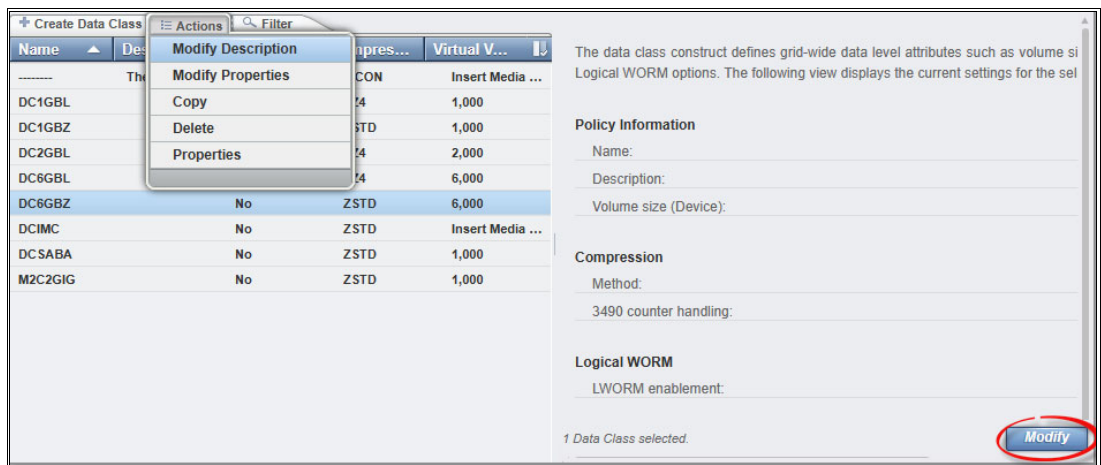


Figure 10-13 Supported actions and Modify button at Inspector Area

Selecting one of the possible actions or selecting one existing data class at the left of the page and clicking modify at the right area starts a wizard for setup assistance. For instance, the wizard invoked to modify the selected data class is shown in Figure 10-14 on page 488.

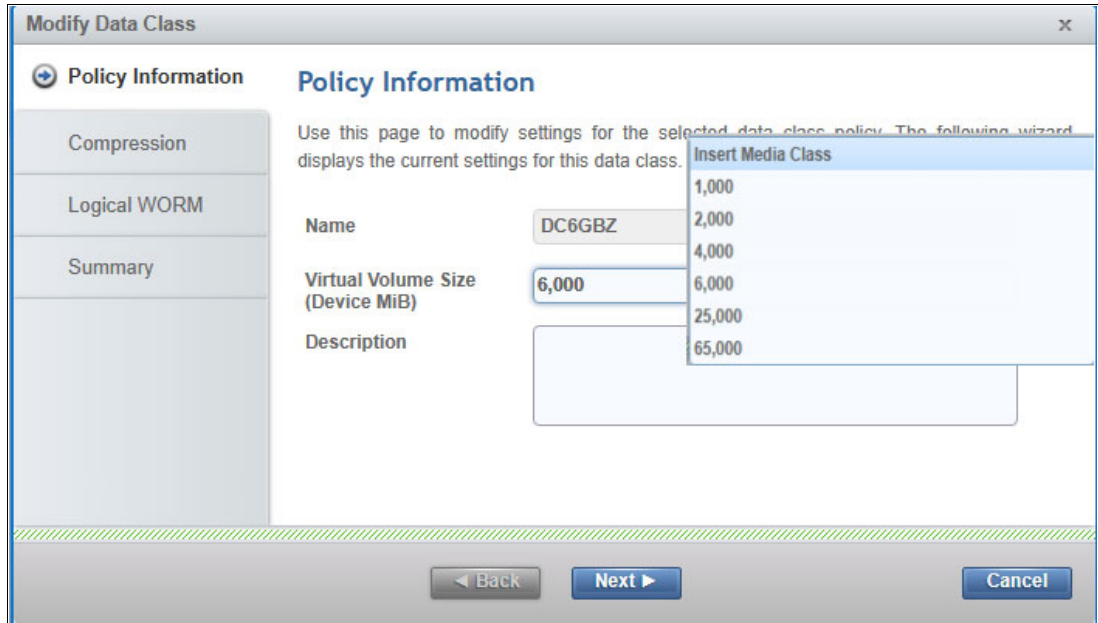


Figure 10-14 Policy Information wizard

Clicking next opens the Compression wizard, which is shown in Figure 10-15.

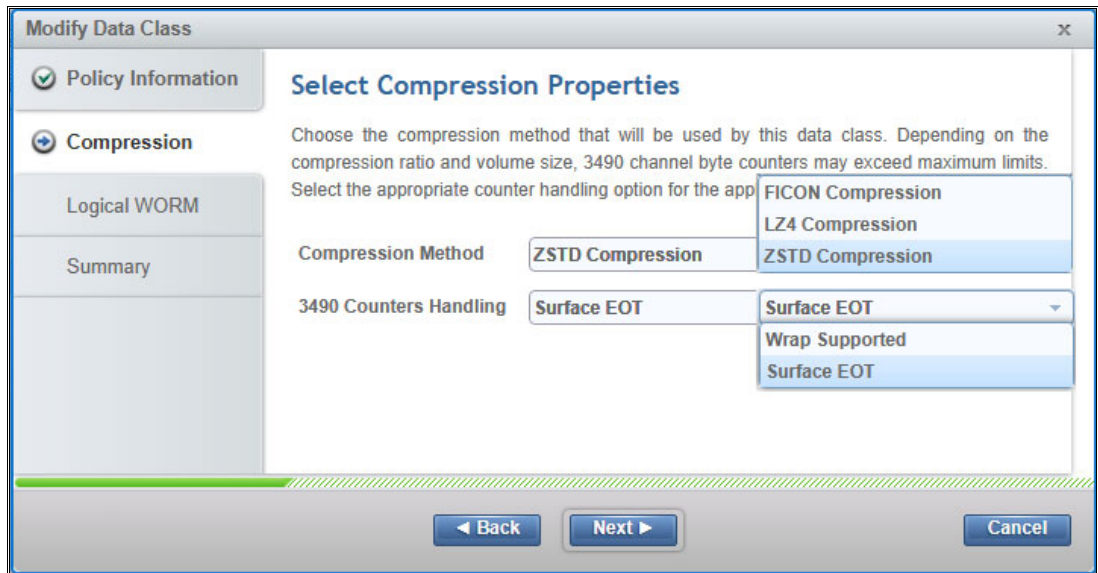


Figure 10-15 Selection Compression Properties and possible options.

Clicking **Next** button opens the Logical WORM wizard panel, which is shown in Figure 10-16 on page 489.

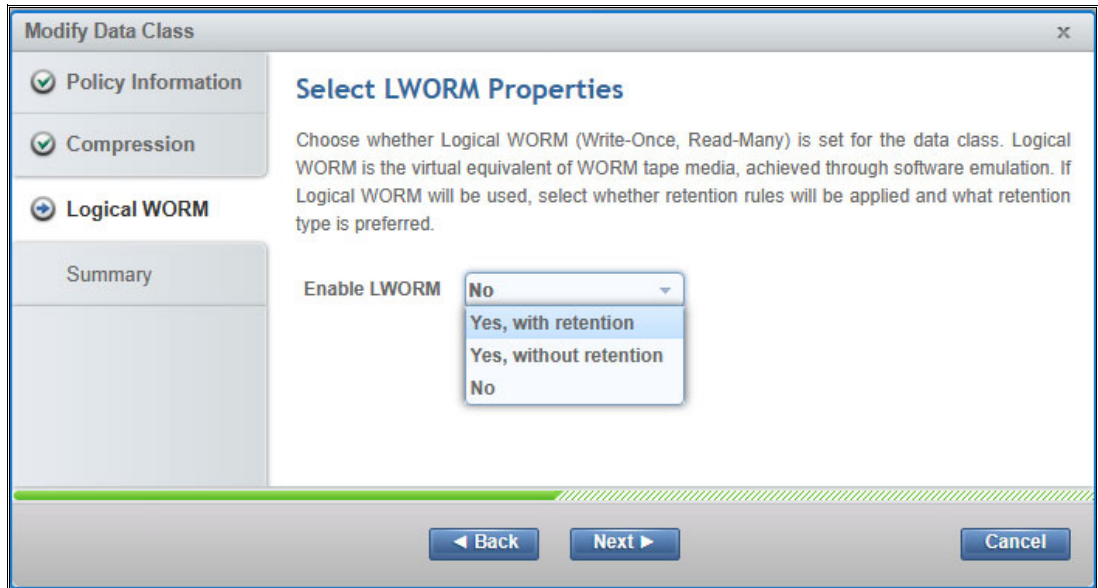


Figure 10-16 LWORM Properties wizard

The **Next** button opens the Summary page. This provides the user with an opportunity to review the choices before submitting the change. The aspect of the Summary section is shown in Figure 10-17.

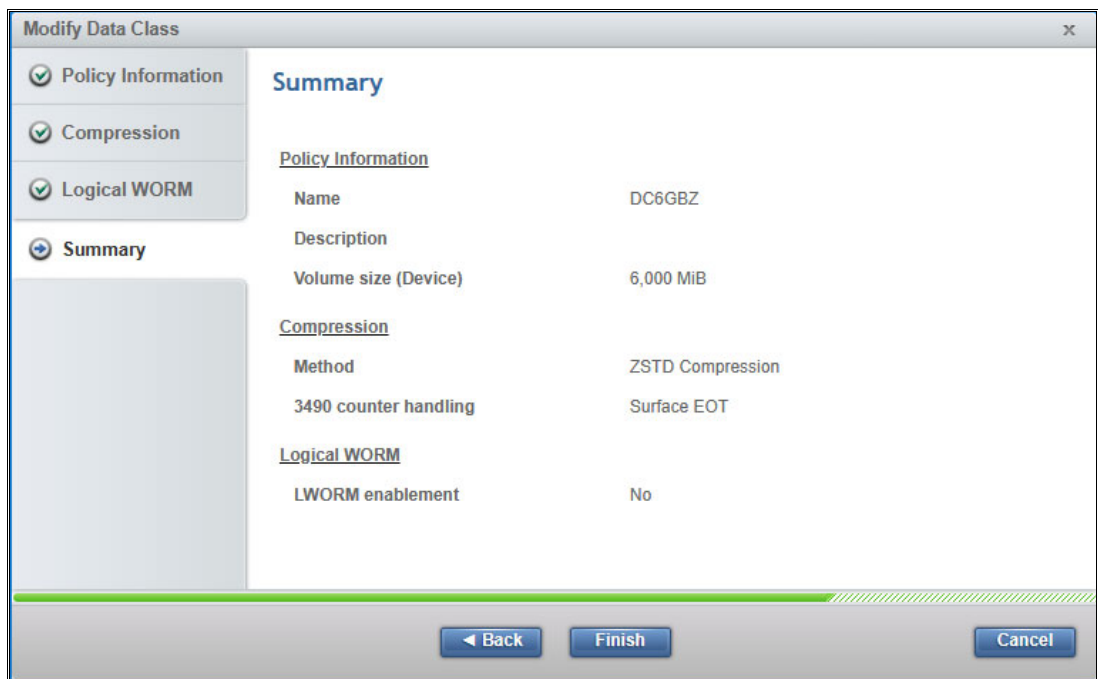


Figure 10-17 Summary section of the wizard.

Note: Users of the 65000 MiB logical volumes might have block size requirements to avoid out-of-block ID occurrences. Current design supports maximum of 4.1 million blocks (22-bits of the read block ID CCW response). This effect existed for the 25000 MiB logical volume sizes but it is expected that this situation is encountered more frequently when using 65000 MiB logical volume sizes. A 256k block size might be needed in this use case.

Figure 10-18 shows the wizard for LWORM properties.

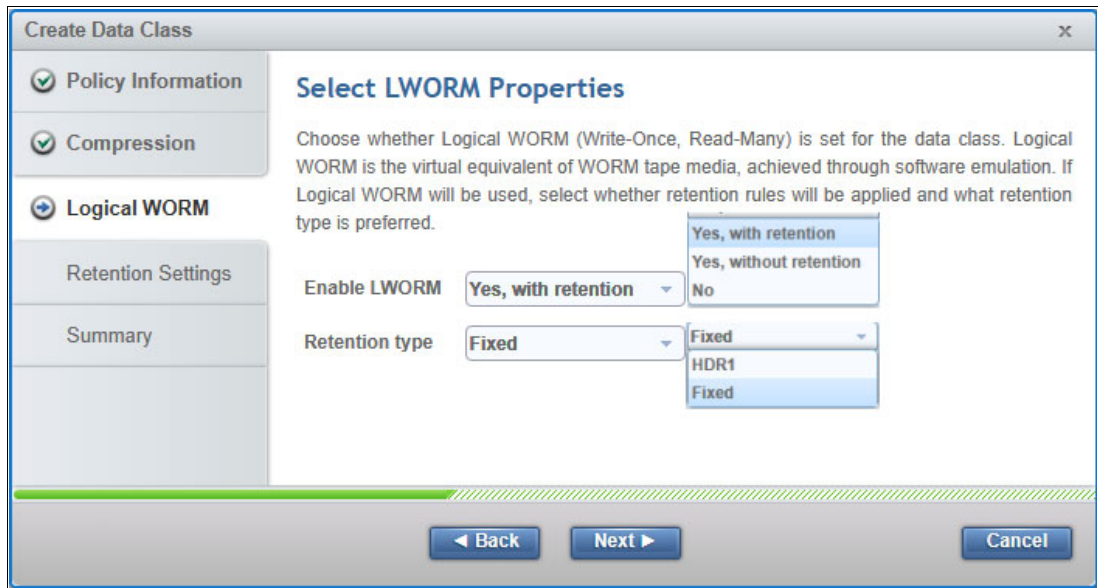


Figure 10-18 Data Class wizard for LWORM properties

Figure 10-19 shows the Retention Settings wizard for LWORM volumes.

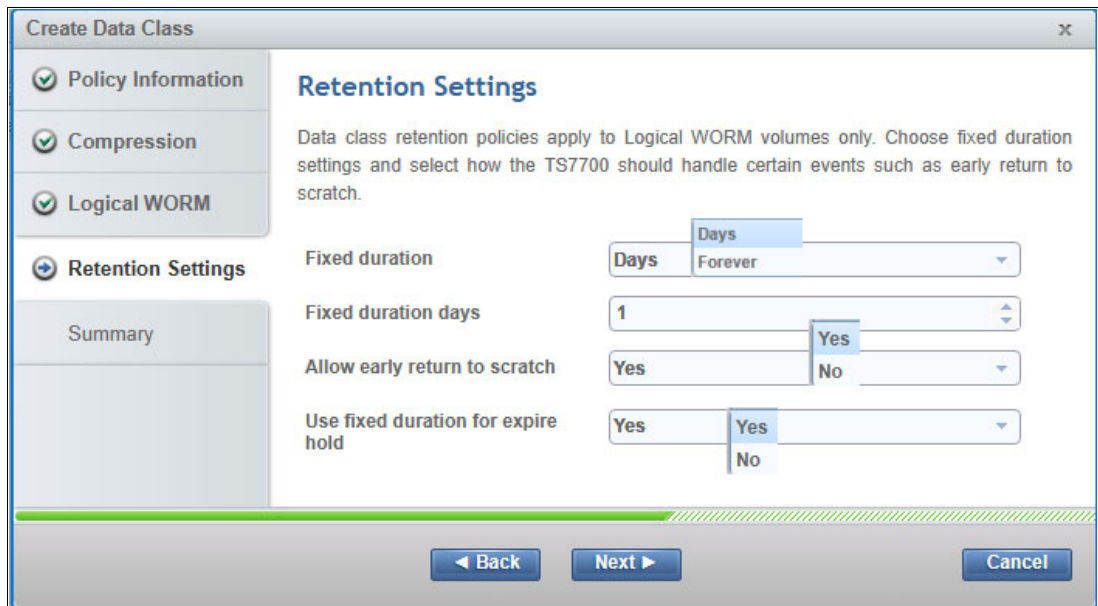


Figure 10-19 Retention Settings wizard

Figure 10-20 shows the summary section of an LWORM Data Class as an example.

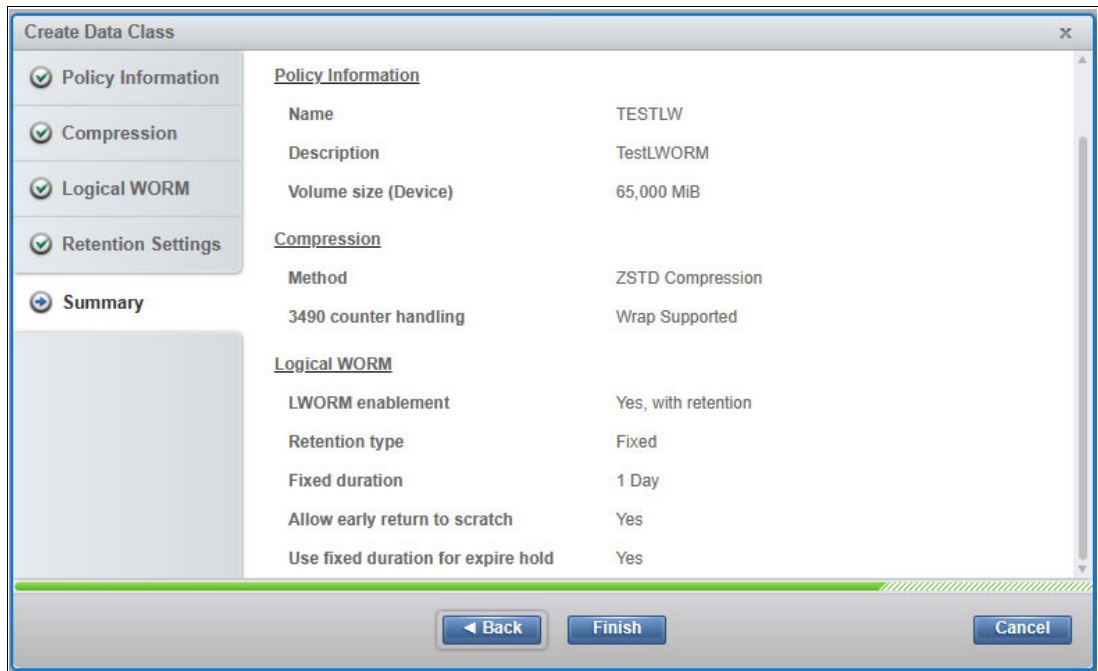


Figure 10-20 Example of the Summary section of an LWORM data class.

Figure 10-21 shows the pre-R5.4 Data Classes page for reference.

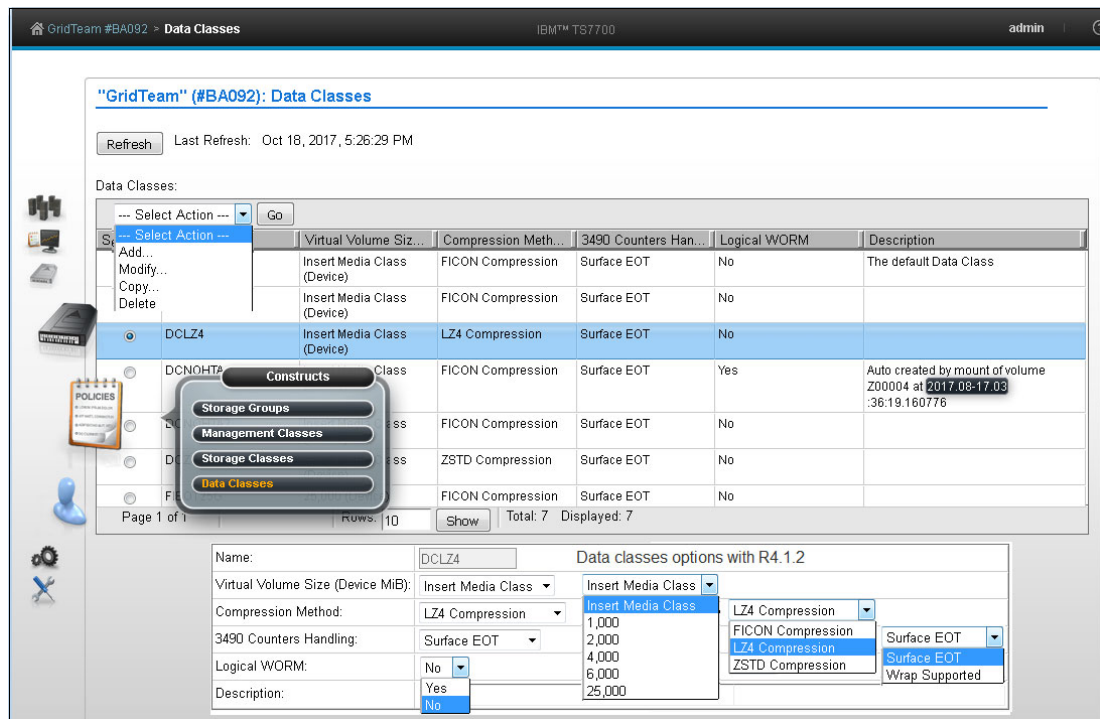


Figure 10-21 MI Data Classes window

For more information, such as the complete description of the fields and information available in the Data Classes page, see the locally available TS7700 5.4 Documentation available locally at *TS7700 MI Help* or [Data Classes](#).

For more information about this function, see Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Important: Scratch 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 user can create up to 256 DCs per TS7700 grid.

10.3 Access icon

The topics in this section present information that is related to managing user access in a TS7700 subsystem.

The user can access the following options through the User Access (blue person icon) link, as shown in Figure 10-22.



Figure 10-22 Access Icon and options

The TS7700 management interface (MI) pages that are collected under the Access icon help the user to view or change security settings, roles and permissions, passwords, and certifications. The user can also update the IBM Documentation (InfoCenter) files from this menu. The following windows are available:

- ▶ **Security Settings:** Use this window to view security settings for a TS7700 grid. From this window, the user can also access other pages 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.

- ▶ Customer Documentation Settings: Use this window to upload a new TS7700 IBM Documentation to the cluster's MI.
- ▶ Dual Control: Use this window to enable and configure Dual Control.

10.3.1 Dual Control window

Dual control adds an extra layer of protection against an accidental or malicious change in certain MI settings that might affect the clients data protection.

Note: Before activating Dual Control, set up users or groups to be administrators or checkers as described in 10.3.2, "Security Settings window" on page 495.

The following settings are protected by dual control (when enabled) as implemented in R5.1:

- ▶ Modify category
- ▶ Delete category
- ▶ Modify logical volume version retention for a cloud pool

When dual control is active, it also protects its own control settings against a single user (even an administrator) circumventing the protection. The following settings are also protected when dual control is enabled:

- ▶ Disable Dual Control
- ▶ Modify User Password
- ▶ Modify Local User
- ▶ Enable Local or Remote Security Policy
- ▶ Modify Remote Security Policy User and Group Mapping

MI Dual Control support for the Object Store and Object Policy pages includes:

- ▶ Add Category
- ▶ Modify object policy replication
- ▶ Delete object policy
- ▶ Modify object Store
- ▶ Delete object Store

Regarding dual control, the specific terms are used in the following contexts:

- ▶ Maker: The person that performs a request to run an operation.
- ▶ Checker: The person who authorizes the request.

Dual control is a grid-wide setting, affecting all clusters in the grid when activated. The sequence of actions necessary to activate the dual control and select the users with checker authority are shown in Figure 10-23 on page 494. The dialog shows all the local and remote user mappings from all active security policies, and from that list the users with checker authority are selected.

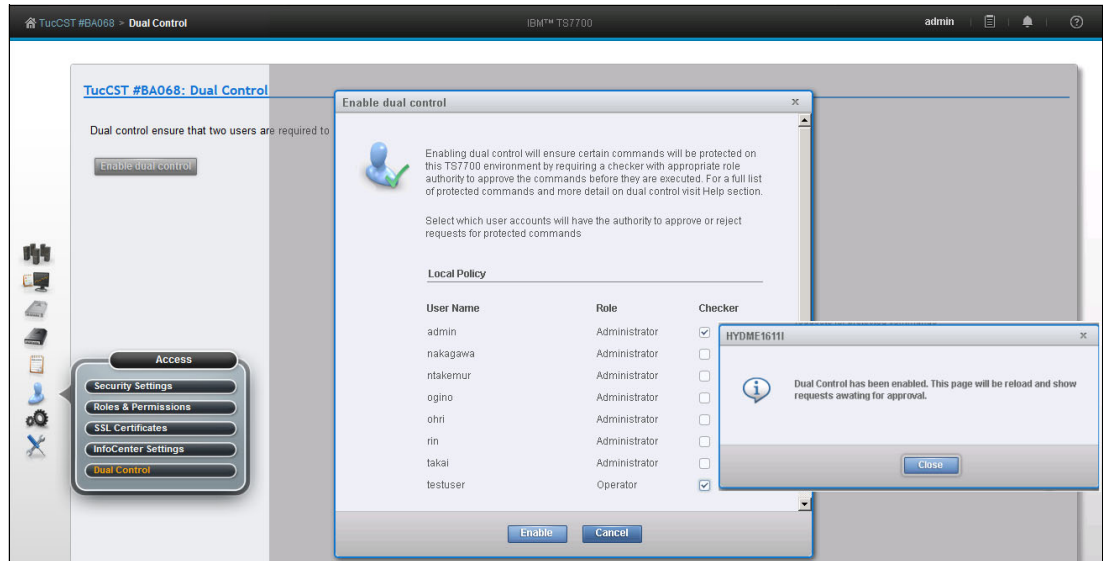


Figure 10-23 Enabling Dual Control and selecting users with checker authority

To enable dual control, at least two of the existing users must be defined as Checker and at least two Administrators must exist. If there are only two Administrator users defined, they cannot also be the only users who are assigned as checkers (this configures a deadlock situation and the Enable button grays out). After the deadlock situation is resolved, the administrator can then see the Enable dual control button available again.

After Dual Control is enabled, the function itself becomes protected, requiring the approval from a checker to be disabled again.

The following Dual Control rules apply:

- ▶ At least two users with Checker authority and two Administrators must exist in all currently assigned policies, counting three or more users minimum.
- ▶ If there are only two administrator users who are defined in the policies, they cannot be the only checkers also.
- ▶ It takes at minimum three different users (between the two Admin and two users who are defined as Checker) to avoid deadlock at the activation time.
- ▶ Only a checker can approve a request. No checker user can approve their own submitted requests.
- ▶ A checker or a maker can reject a request.
- ▶ If conditions are not met (the check does not pass), the request stays in the request table.

For more information about the Dual Control at R5.4, see the IBM Documentation locally at the TS7700 MI by clicking the question mark at the upper right of the window, or see [Dual Control](#).

10.3.2 Security Settings window

Figure 10-24 shows the Security Settings window, which is the entry point to enabling security policies.

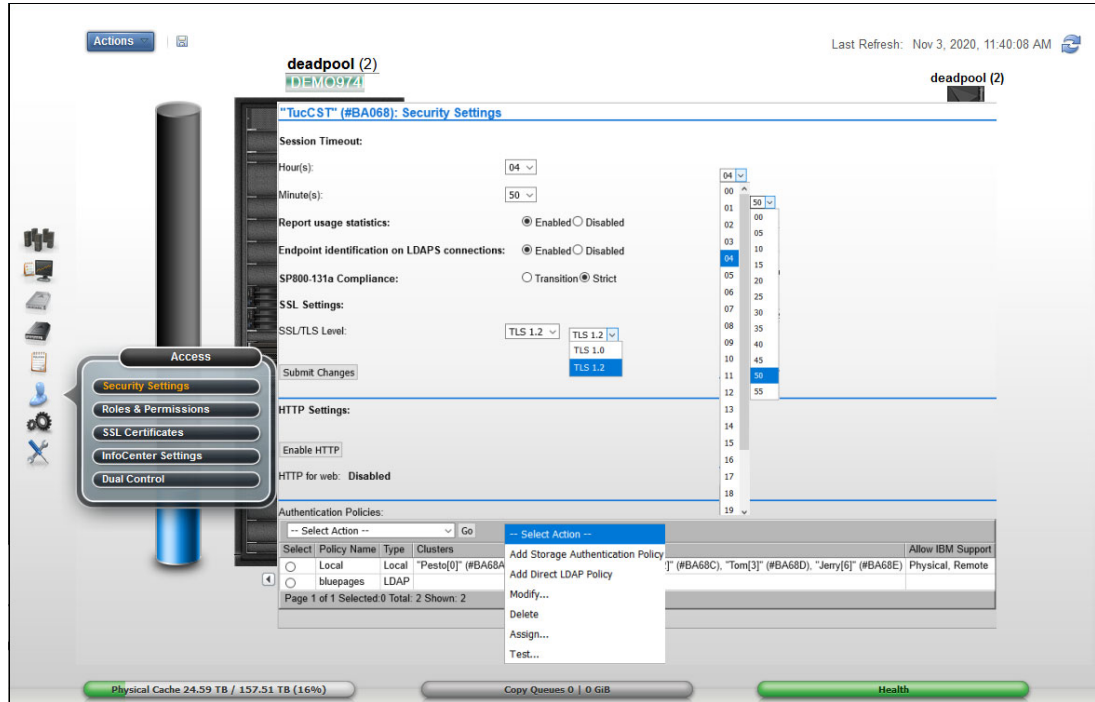


Figure 10-24 TS7700 Security Settings

This page allows you to change the following settings:

- ▶ **Session Timeout:** This setting is defined from the Security Settings window. The user can specify the number of hours and minutes that the management interface can be idle before the current session expires and the user is redirected to the login page.

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.

- ▶ **Endpoint identification on LDAPS connections:** Use Endpoint identification on LDAPS connections to enable or disable the endpoint identification algorithms for secure LDAP over TLS to improve the robustness of LDDAPS connections. If enabled, the user must ensure that the following prerequisites are met:

- The LDAP server certificate includes Subject Alternative Names (SAN)
- The address from LDAP authentication policy must match one of those SANs

If the server's SSL certificate does not have SANs, the LDAP policy address must match the Common Name (CN) of the certificate.

- ▶ **SP800-131A Compliance:** Support the requirements that are defined by the National Institute of Standards and Technology (NIST) Special Publications 800-131a SP 800-131a. SP 800-131a strengthens security by defining which algorithms can be used,

and minimum strengths. Set the property to transition to specify that SP800-131a transition compliance is requested.

Set the property to strict to allow only strict adherence to the SP800-131a recommendation. When this option is selected, the SSL/TLS level under the SSL Settings section is defined as TLS1.2.

- ▶ **SSL settings:** Use this section to set the SSL/TLS level. Two choices are available: TLS 1.0 (transition) and TLS 1.2 (strict). The default setting is TLS 1.0.

If the browser that is used to access TS7700 MI does not support TLS 1.2 and HTTPS-only is enabled, a warning message is displayed. The message states that access to MI might be lost if you proceed.

- ▶ **HTTP Settings:** HTTP enablement can be changed in this session. The HTTP for web setting can be Enabled or Disabled from here. If the HTTP settings change, the TS7700 MI restarts. All users that are logged on at the time lose connection to the Management Interface and must log in again.
- ▶ **Usage Reporting:** Use this section to enable or disable usage reporting. Usage reporting is enabled or disabled for all users on the client and server sides. By default, usage reporting is disabled.

Authentication policies

The user can add, modify, assign, test, and delete the authentication policies that determine how users are authenticated to the TS7700 Management Interface. Each cluster is assigned a single authentication policy. The user must be authorized to modify security settings before changing authentication policies.

Two categories of authentication policies are available: Local, which replicates users and their assigned roles across a grid, and External, which stores user and group data on a separate server and maps relationships between users, groups, and authorization roles when a user logs in to a cluster. External policies include Storage Authentication Service policies and Direct LDAP (lightweight directory access protocol) policies.

Note: A restore of cluster settings (from a previously taken backup) does not restore or otherwise modify any user, role, or password settings that are defined by a security policy.

Policies can be assigned on a per cluster basis. One cluster can employ local authentication, while a different cluster within the same grid domain can employ an external policy. Also, each cluster in a grid can operate its own external policy. However, only one policy can be enabled on a cluster at a time.

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 an 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 cannot be deleted.
 - External: Policies that map user, group, and role relationships upon user login. External policies can be modified. However, they cannot be deleted if in use on any cluster. The following external policies are available:
 - Storage Authentication Service: A centrally managed, role-based access control (RBAC) policy that authenticates and authorizes users by using the System Storage Productivity Center to authenticate users to an LDAP server.
 - LDAP: An RBAC policy that authenticates and authorizes users through direct communication with an LDAP server.
 - Clusters: The clusters for which the authentication policy is in force. Cluster names are displayed only for policies that are enabled and assigned. Only one policy can be assigned to a cluster at a time.
 - Allow IBM Support: The type of access that is granted to IBM service representatives for service support. This access is most often used to reset the cluster authentication policy to Local during an LDAP authentication issue.

Note: If an IBM service representative resets a cluster authentication policy to Local, the Local authentication policy is enabled on all clusters in the grid, regardless of previous LDAP policy setting. Previously enabled LDAP policies are disabled and should be re-enabled following resolution of any LDAP authentication issue.

The following are values (among others) are possible:

- Physical: IBM service representatives can log in physically without LDAP credentials to connect to the cluster. At least one IBM representative must have physical access to the cluster. An onsite IBM representative can grant temporary remote access to an off-site IBM representative. This option is recommended.
- Remote: IBM service representatives can log in remotely without LDAP credentials to connect to the cluster.

Important: If this field is blank for an enabled policy, IBM service representatives must log in by using LDAP login credentials that are obtained from the system administrator. If LDAP server is inaccessible, IBM service representatives cannot access the cluster.

Adding a user to the Local Authentication Policy

A Local Authentication Policy replicates authorization that is based on user accounts and assigned roles. It is the default authentication policy. This section provides a description of 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 in the Authentication Policies table.

3. Select **Modify** from the **Select Action** menu and click **Go**.
4. In the Local Accounts table, select **Add** from the **Select Action** menu and click **Go**.
 - 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. Table 10-4 displays the list of user roles and a brief description of each possible value.

Table 10-4 User roles

User role	Description
Operator	The operator can access 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 can access monitoring information and 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 adding or removing user accounts. The administrator can access all service functions and TS7700 resources.
Manager	The manager can access monitoring information and performance data and functions. The manager can also perform actions for users, including adding, modifying, and deleting user accounts. The manager is restricted from changing most other settings, including those settings for logical volume management, network configuration, feature licenses, and custom roles.
Read Only	The read only role can view all pages but cannot perform any actions. You can make any custom role read-only by applying the Read Only role template to the custom role.
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 access multiple clusters.
5. To complete the operation, click **OK**. To stop the operation and return to the Modify Local Accounts window, click **Cancel**.

Modifying the user or group of the Local Authentication Policy

Use this window modify a user or group property for a TS7700 grid.

Tip: Passwords for the users are changed from this window also.

To modify a user account belonging to the Local Authentication Policy, complete the following steps:

1. In the TS7700 MI, click **Access** (blue person 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 username of the policy to modify.
5. Select **Modify** from the **Select Action** menu and click **Go**.
6. Modify the values for any of the following fields:
 - Role: The role that is assigned to the user account. The possible values are shown on Table 10-4 on page 498.
 - Cluster Access: The clusters to which the user can access. A user can access multiple clusters.
7. To complete the operation, click **OK**. To stop the operation and return to the Modify Local Accounts window, click **Cancel**.

Note: The user cannot modify the username or Group Name. Only the role and the clusters to which it is applied can be modified.

In the Cluster Access table, select the **Select** checkbox 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 person 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. The following fields are available:
 - 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`

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 username 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.
4. To complete the operation, click **OK**. To abandon the operation and return to the Security Settings window, click **Cancel**.

Note: Generally, select the **Allow IBM Support** option to grant access to the IBM services representative to the TS7700.

Click **OK** to confirm the creation of the Storage Authentication Policy. In the Authentication Policies table, no clusters are assigned to the newly created policy; therefore, 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.

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:
 - **Username:** 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 possible values are shown on Table 10-4 on page 498.
 - **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.

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 username 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. In 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.
3. Click **Go** to open the Assign Authentication Policy window.
4. To apply the authentication policy to a cluster, select the checkbox next to the cluster's name.

Enter values for the following fields:

- Username: Username 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. Ensure that no clusters are assigned to the policy so that it can be deleted. If clusters are assigned to the policy, use **Modify** from the **Select Action** menu to remove the assigned clusters.

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. 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: Click **OK** to delete the policy.

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 that are 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 must be configured to an LDAP server before adding it to 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 username by an existing policy:

1. 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 to test.
 - b. Select **Test** from the **Select Action** menu.
 - c. Click **Go** to open the Test Authentication Policy window.
2. Select the checkbox next to the name of each cluster on which to conduct the policy test.
3. Enter values for the following fields:
 - Username: The username 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 username that is entered belongs to a username 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. The results include a statement that indicates whether the test succeeded or failed, and the reason for the failure if it failed. The Test Authentication Policy results window also displays the Policy Users table. The 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.

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. In the TS7700 MI, click **Access** → **Security Settings** from the left navigation window.
2. From the menu, select **Add Direct LDAP Policy** and click **GO**.
3. Select the check boxes to grant local or remote access to the TS7700 to the IBM SSR so that the SSR can perform service support.

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 following LDAP attributes fields can also 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.
- ▶ **Username Attribute:** The attribute name that is used for the username 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.
- ▶ **Username filter:** Used to filter and verify the validity of an entered username. 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

This process is similar to the process that is described “Adding a Direct LDAP policy” on page 502. Some configurations are required on the host side regarding the RACF, SDBM, and IBM Security Directory Server that must be performed in advance before this capability can be enabled.

For more information about the parameters and configurations, see Chapter 12, “IBM z/OS host console operations” on page 625. When those configurations are ready, the RACF-based LDAP Policy can be created and activated.

Adding a RACF policy is shown in Figure 10-25.

Figure 10-25 Adding a RACF based Direct LDAP Policy

As shown in Figure 10-25, 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.

As shown in Figure 10-25, the Group Member Attribute was set to racfgroupuserids (it shows truncated in MI’s text box).

The User Distinguished Name time should be specified with all of the following parameters:

- ▶ racfid
- ▶ profiletype
- ▶ cn

When the previous setup is complete, users can be added to the policy, or clusters can be assigned to it, as described next. No specific restrictions exist for these RACF/LDAP user IDs. They can be used to secure the MI, or the IBM service login (for the IBM SSR) as with any other LDAP user ID.

For more information, see this IBM Documentation [web page](#), which is also available locally in the MI window by clicking the question mark in the upper right bar and selecting **Help**.

Adding users to a Direct LDAP Policy

For more information, see the process that is described in “Adding a user to a Storage Authentication Policy” on page 500. The same steps apply when adding users to a Direct LDAP Policy.

Assigning a Direct LDAP Policy to a cluster or clusters

For more information, see the procedure that is described in “Assigning clusters to a Storage Authentication Policy” on page 501. The same steps apply when working with a Direct LDAP Policy.

Deleting a Direct LDAP Policy

For more information, see the procedure that is described in “Deleting a Storage Authentication Policy” on page 501. The same steps apply when deleting a Direct LDAP Policy.

10.3.3 Roles and Permissions window

You can use the window that is shown in Figure 10-26 to set and control user roles and permissions for a TS7700 Grid.

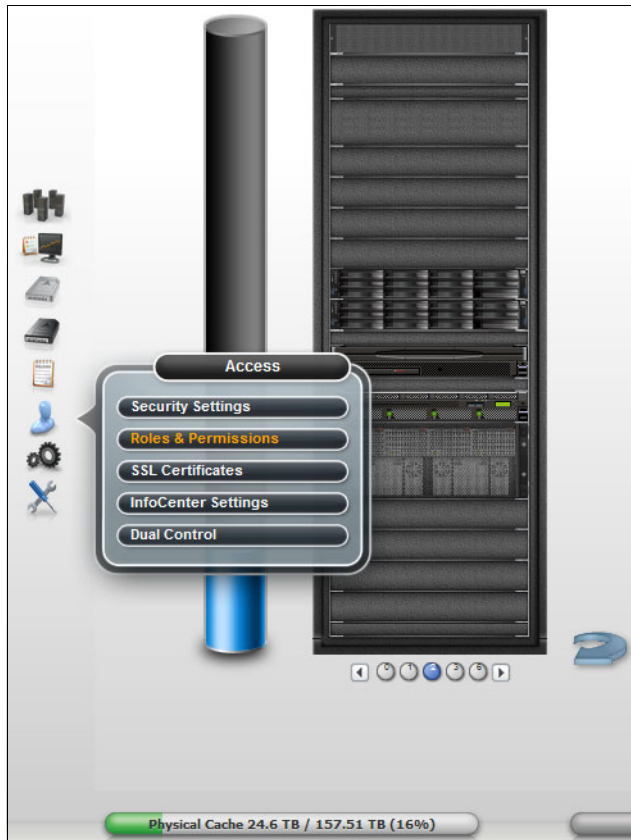


Figure 10-26 TS7700 MI Roles and Permissions window

Clicking **Manage users and assign user roles** opens the Security Settings window in which the user can add, modify, assign, test, and delete the authentication policies that determine how users are authenticated to the TS7700 Management Interface.

The available roles are shown on Table 10-4 on page 498.

Note: Valid characters for the Name of Custom Role field are A-Z, 0-9, \$, @, *, #, and%. The first character of this field *cannot* be a number.

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 checkbox to the left of the role to be displayed. The user can select more than one role to display a comparison of permissions.
2. Click **Select Action** → **Properties**.
3. Click **Go**.

The first column of the Roles and Assigned Permissions table lists all the tasks that are available to users of the TS7700. Subsequent columns show the assigned permissions for selected role (or roles). A checkmark denotes permitted tasks for a user role. A null dash (-) denotes prohibited tasks for a user role.

Permissions for predefined user roles 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.

Note: Valid characters for this field are A-Z, 0-9, \$, @, *, #, and%. The first character of this field *cannot* be a number.

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.

10.3.4 SSL Certificates window

Use the window that is shown in Figure 10-27 to view, import, or delete SSL certificates to support secure connections to a Storage Authentication Service server from a TS7700 cluster. This page also allows the user to replace the MI HTTPS SSL certificate with a custom one.

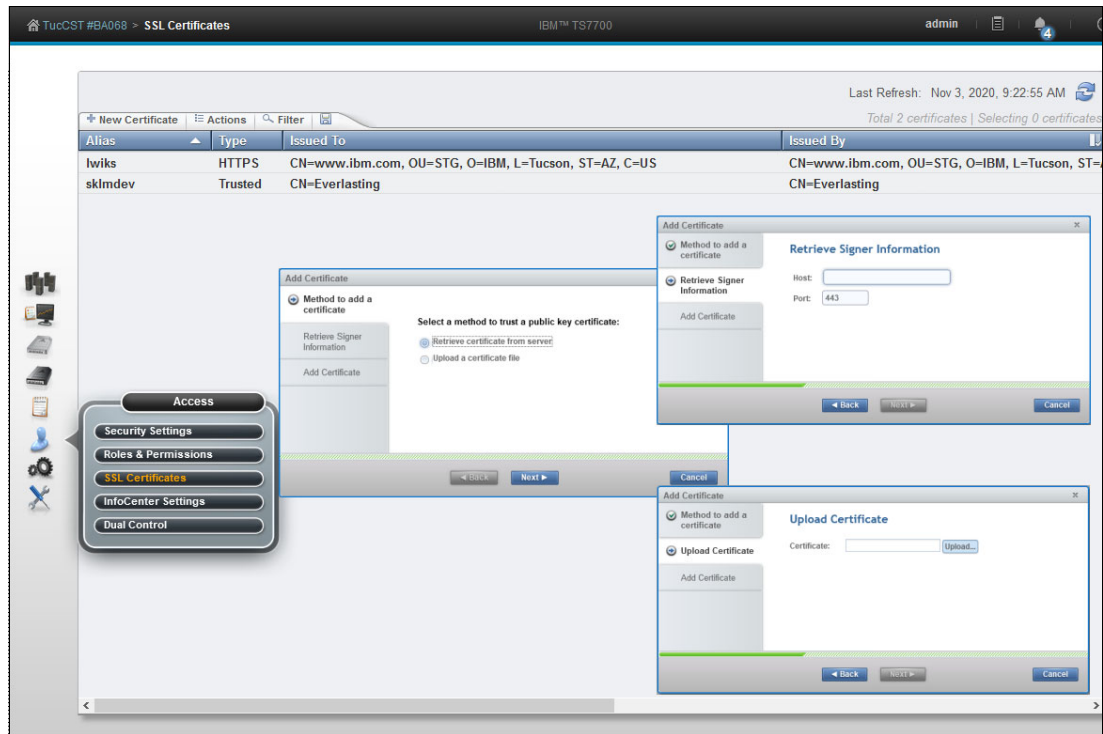


Figure 10-27 SSL Certificates window

If HTTPS is used, the TLS version must be the same at the TS7700 and the Cloud Server to avoid problem with the SSL Certificate Retrieve. If a Primary or alternative Server URL (which is defined by a Storage Authentication Service Policy) uses the HTTPS protocol, a certificate for that address must be defined in this window. The same is true for Direct LDAP policies if the Primary or alternative server uses LDAPs. If the policy uses LDAP, a certificate is not required. The Certificates table displays identifying information for SSL certificates on the cluster.

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) 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.
- ▶ **Issued By:** The issuer of the certificate.
- ▶ **Type:** Shows the type of the SSL certificate. Can be a trusted certificate that is installed from a remote server, or HTTPS for a certificate that is used in https connections to the local MI.

To import a new SSL certificate, complete the following steps:

1. Click **Select Action** → **Retrieve from port** 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 import a new SSL certificate, select **New Certificate** from the top of the table, which displays a wizard dialog. Consider the following points:

- ▶ To retrieve a certificate from the server, select **Retrieve certificate from server** and click **Next**. Enter the host and port from which the certificate is retrieved and click **Next**. The certificate information is retrieved. The user must set a unique alias in this window. To import the certificate, click **Finish**. To abandon the operation and close the window, click **Cancel**. The user can also return to the Retrieve Signer Information window.
- ▶ To upload a certificate, select **Upload a certificate** file and click **Next**. Click the **Upload** button, select a valid certificate file, and click **Next**. Verify that the certificate information (serial number, issued to, issued by, fingerprint, and expiration) is displayed on the wizard. Complete the alias field with valid characters.

When the Finish button is enabled, click **Finish**. Verify that the trusted certificate was successfully added in the SSL Certificates table. To stop the operation and close the window, click **Cancel**. The user can also return to the Retrieve Signer Information window.

To delete an SSL certificate, complete the following steps:

1. Select the radio button that is next to the certificate to delete, select **Select Action** → **Delete**, and then click **Go**. The Confirm Delete SSL Certificate window opens and prompts you to confirm the decision to delete the SSL certificate.
2. Click **Yes** 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.



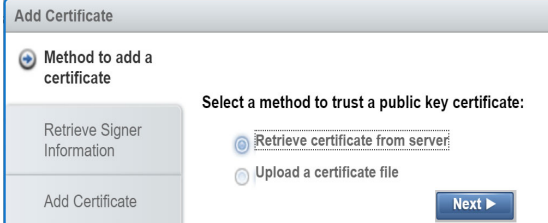
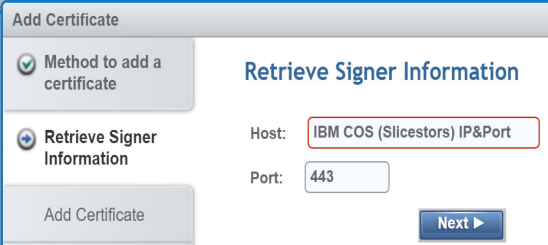
Adding SSL Certificates for the Cloud.

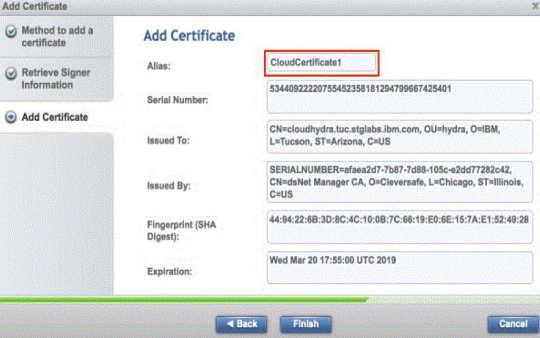
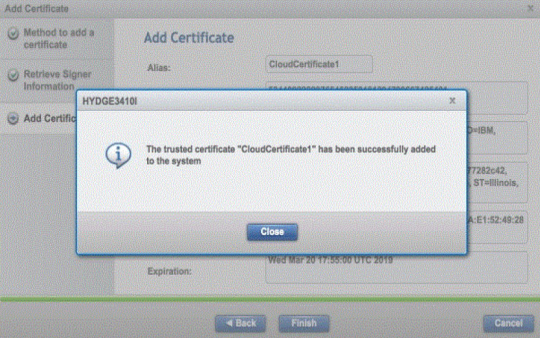
Follow the procedure that is presented in this section to add SSL certificates for the cloud.

Note: If HTTPS is used to communicate with IBM Cloud Object Storage, the ENDPOINT connection is the IBM Cloud Object Storage Accesser® (from where certificates are to be retrieved). Then, the cloud URL can be created by setting the Certificate Alias that was just created.

The steps necessary to add SSL certificates for the Cloud are shown in Table 10-5 on page 509.

Table 10-5 Procedure to add SSL certificates for the Cloud

Step	Activity
<p>1</p> 	<p>Log in to the TS7700 MI. From the Access icon, go to SSL certificates.</p>
<p>2</p> 	<p>Click New Certificate.</p>
<p>3</p> 	<p>In the Add Certificate tab, select Retrieve certificate from server. Click Next.</p>
<p>4</p> 	<p>In the next tab, enter the Cloud IP and Port. IBM Cloud Object Storage is composed for many modules (and all of them have a different associated IP).</p> <p>Retrieve SSL from the IBM Accesser® nodes.</p> <p>Enter the Cloud Object Storage IP and port and then click Next.</p>

Step	Activity
<p>5</p> 	<p>In the Add Certificate tab, enter an alias name (for example, CloudCertificate1).</p> <p>Click Finish.</p>
<p>6</p> 	<p>A window opens that indicates that the certificate was added successfully.</p> <p>Click Close.</p>

Some cases exist in which the SSL certificate must be manually uploaded to the SSL store in the TS7700. For more information, see *Updating the TS7700 CA trust list* in the *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.

10.3.5 Customer Documentation Settings window

To upload a new TS7700 IBM Documentation to the cluster's MI, use the window that is shown in Figure 10-28.

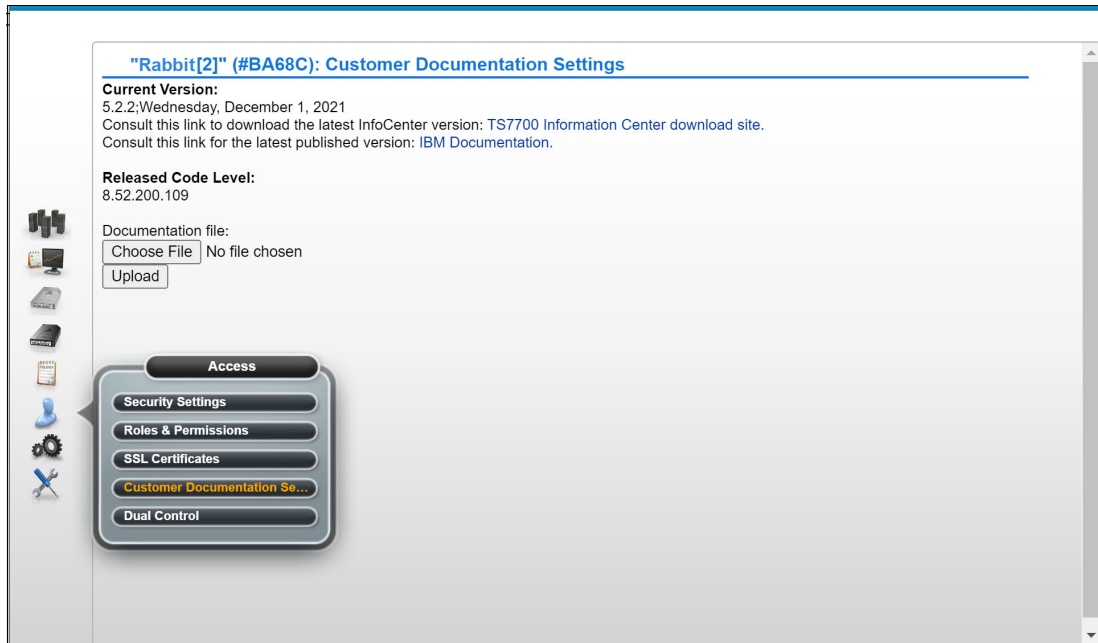


Figure 10-28 Customer Documentation Settings window

This window features the following items:

- ▶ Current Version section, where the following items can be identified or accessed:

- Identify the version level and date of IBM Documentation that is installed on the cluster.
- Access a product database to download a JAR file that contains a newer version of IBM Documentation.
- Access an external site displaying the most recently published version of IBM Documentation.

- ▶ The TS7700 IBM Documentation 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 Documentation as a .jar file (if available). Complete the following steps:

- Select **System Storage** from the **Product Group** menu.
- Select **Tape Systems** from the **Product Family** menu.
- Select **TS7700** from the **Product** menu.
- Click **Continue**.
- In the **Select Fixes** window, select the box that is next to the wanted InfoCenter Update file (if available).
- Click **Continue**.
- In the Download Options window, select **Download using Download Director**.
- Select the checkbox that is next to **Include prerequisites and co-requisite fixes**.

- i. Click **Continue**.
 - j. In the Download files using Download Director window, ensure that the checkbox that is next to the correct InfoCenter Update version is selected 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 Documentation (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 Documentation, complete the following steps:
- a. Click **Browse** to open the File Upload window.
 - b. Browse 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 Documentation on the cluster's MI.

10.4 Settings icon

The TS7700 management interface pages that are listed under the Settings icon help the user to view or change cluster network settings, install or remove feature licenses, and configure SNMP, and several other functions that are relative to the cluster functioning. Figure 10-29 shows the Settings Icon and the options under it.

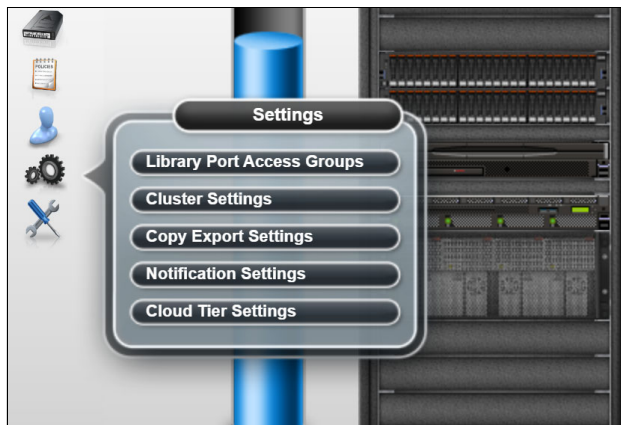


Figure 10-29 Settings icon and options

The Library Port Access Groups and Cloud Tier Settings items can appear under the Settings option, depending on whether FC 5278 (Enable Cloud Storage Tier) and FC 5271 (Selective Device Access Control, SDAC) are present at the cluster configuration.

The following items are included under Settings:

- ▶ Cluster network settings: Used to set or modify IP addresses for the selected cluster. Modifications of the assigned IP addresses interrupt the active MI session.
- ▶ Notification settings: Set or modify notification settings for the selected cluster.
- ▶ Cloud Tier settings: Used to configure the cloud tier settings on the cluster.
- ▶ Feature Licenses: View, activate, or remove feature licenses on a selected cluster.
- ▶ SNMP: View or modify the simple network management protocols on the cluster.
- ▶ Library port access groups: View information about library port access groups that are used by the selected cluster.

- ▶ Cluster settings: View or change settings that determine how a cluster performs copy policy overrides, applies Inhibit Reclaim schedules, uses an encryption key server, implements write protect mode, and performs backup or restore operations.
- ▶ Copy export settings: Change the maximum number of physical volumes that can be exported by the TS7700T.
- ▶ Secure Data transfer: (SDT) provides a way to securely transfer data (logical volumes or objects) among clusters and DS8900s by way of the network links infrastructure of a TS7700 grid.

Note: SDT replaces the IPsec method to encrypt objects or logical volume data that is being transferred through the grid links with no significant effect on the performance of the TS7700. SDT requires FC 5281 enabled on both clusters that are performing a copy transaction. If one of the clusters does not have encryption that is enabled, the data exchange is not encrypted.

10.4.1 Cloud Tier Settings

Figure 10-30 shows the Setting Icon appearance with the Cloud Tier Settings and options.

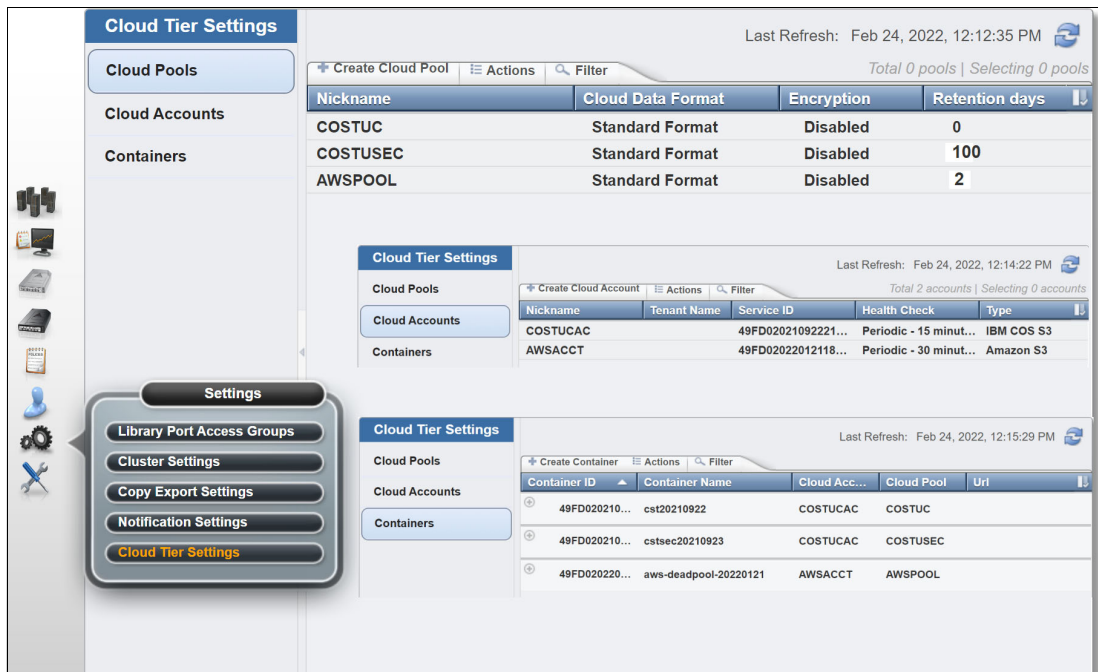


Figure 10-30 The Cloud Tier Settings and options

The Cloud Tier Settings page is used to set or modify the cloud tier settings on the TS7700 cluster. With the R5.1 level of code, the TS7700C now supports:

- ▶ Up to 256 cloud pools, which is excellent for separation of data and multi-tenancy usage. Each pool can be configured differently, with unique bucket, vault, provider, or location.
- ▶ The Volume Version Retention in the cloud for later recovery.

With R5.1 or higher, users can configure how long previous stale versions of a logical volume are kept by using a new Cloud Pool setting. The retention period can be configured as Retention Days by Cloud Pool on the TS7700 MI Cloud Pools page (see Figure 10-30).

- ▶ TS7700 Cloud Export and Recovery suite of features, providing capability to export a backup into an attached cloud. This backup can then later be used to restore an empty TS7700C cluster. This suite of features also includes the ability to perform DR testing to ensure disaster preparedness and covers the concept of retaining older versions of logical volumes in the cloud.

These features require that all VEC and VED models within the grid to be at the R5.1 or higher code level before they are enabled.

For more information about the IBM TS7700 Cloud Export, Cloud Export Recovery, Cloud Recovery testing, logical volume version retention, and restore concepts for the TS7700 Cloud Storage Tier feature (Feature Code 5278), see the white paper [TS7700 Cloud Storage Tier Export Recovery and Testing Guide](#).

The following items are included in the Cloud Tier Settings page, as shown in Figure 10-31:




Figure 10-31 Cloud Pools and parameters

- ▶ Cloud Pools: Cloud Pool parameters

Each cloud pool contains the following components:

- Nicknames: Specify a name that identifies the pool (up to eight characters). This nickname identifies the pool in other cloud setting panels.
- Cloud Data Format: Select the format for the data (the default is Standard Format).
- Encryption: Specifies whether the data in the pool is to be encrypted.
- Retention days: Volume version retention defines how many days data is retained before deletion by cloud pool.0 (zero) means do not retain any version.

Cloud pool Volume version retention allows previous versions of logical volumes to be retained in the cloud for the specified period. The retention period needs to be set according to the recovery requirements for the data

Note: Changes to the volume version retention settings are retroactive. Expiry dates for existing data awaiting expiration change according to the new settings.

- ▶ **Cloud Accounts:** Cloud account parameters. IBM Cloud Object Storage S3 and Amazon S3-type accounts are supported.

Each cloud account contains the following fields. The maximum number of cloud accounts is 256:

- **Nickname:** Cloud account nickname (up to eight characters). Nicknames can be changed without restriction to the Account Type.
 - **Tenant Name:** Not required for the type of Accounts that are supported (IBM Cloud Object Storage S3 and Amazon S3).
 - **Service ID:** Service ID is unique and cannot be modified. The value is auto-generated; no input is needed by the user.
 - **Health Check:** Periodic, Event, or Disabled. If Periodic is chosen, specify a number of minutes to perform health check.
 - **Type:** Amazon S3 or IBM Cloud Object Storage S3. After a Cloud account is created, this type cannot be modified.
 - **Access Key ID:** The username to access data in the cloud when the Cloud Account is also provided, and both settings can be modified later.
- ▶ **Cloud Container parameters:** A container is owned by the account that created it. The TS7700 cloud supports up to 256 containers in each of the accounts. Container ownership is not transferable, but can be deleted if empty.

After a container is deleted, the name becomes available for reuse, but it can take some time to the name becomes available for reuse. Also, another account can create a container with that name in the meantime; therefore, the name might not be available for reuse.

Tip: If you are planning to use the same container name, do not delete the original empty container.

An unlimited number of objects can be stored in a container and no difference exists in performance whether many or only a few containers are used. All objects can be stored in a single container, or they can be organized across several containers.

Note: A container cannot be created within another container

Each container features the following fields:

- **Container name:** After creation, an S3 container's name cannot be changed. The following rules apply to naming S3 container in all AWS Regions:
 - Container names can be repeated to support Cross-Region Replication on Amazon.
 - Container names must be 3–63 characters.
 - Container names cannot contain uppercase characters or underscores.
 - Container names must start with a lowercase letter or number.
 - Container names must be a series of one or more labels. Adjacent labels are separated by a single period (.). Container names can contain lowercase letters, numbers, and hyphens. Each label must start and end with a lowercase letter or a number.
 - The single period (.) is not allowed in AWS account container names but is allowed in IBM Cloud Object Storage account container names.

- Container names must not be formatted as an IP address (for example, 192.168.5.4).
 - When you use virtual hosted-style containers with Secure Sockets Layer (SSL), the SSL wildcard certificate matches only containers that do not contain periods. To work around this issue, use HTTP or write your own certificate verification logic. It is recommended that you do not use periods (“.”) in container names when virtual hosted-style containers are used.
- Cloud Account.
 - Cloud Pool.

For more information, see the following resources:

- ▶ [IBM TS7700 R5.3 Cloud Storage Tier Guide, REDP-55733](#)
- ▶ [Amazon updated naming conventions](#)
- ▶ [Virtual Hosting of Buckets](#)
- ▶ [Amazon S3 Transfer Acceleration](#)

After an Amazon container is created, the cloud URLs must be added. For Amazon S3, select one of the endpoint URLs from [AWS service endpoints](#).

For more information about containers, see the following resources:

- ▶ [IBM Cloud Object Storage System](#)
- ▶ [IBM Cloud Object Storage System Account Creation](#)

10.4.2 Library Port Access Groups window

Use this selection to view information about library port access groups that are used by the TS7700. Library port access groups enables the user to segment resources and authorization by controlling access to library data ports.

Tip: This window is visible only if at least one instance of the Selective Device Access Control (FC 5271, SDAC) is installed on all clusters in the grid.

Figure 10-32 on page 517 shows the Library Port Access Groups window and options.

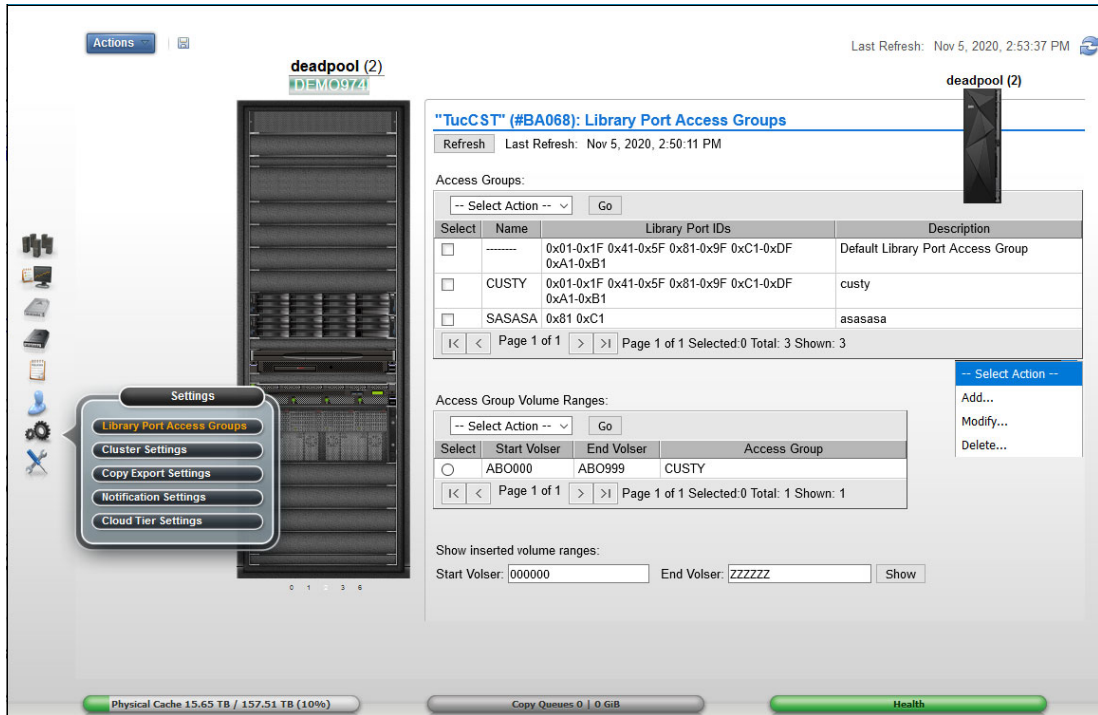


Figure 10-32 Library Port Access Groups page.

Access Groups table

This table displays information about library port access groups.

The user can use the Access Groups table to create a library port access group. The user can also modify or delete an existing 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 that are accessible by using the defined access group. This field contains a maximum of 750 characters, or 31 Library Port IDs that are 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 that are not in a defined range.

Click the **Select Action** menu in 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 library port access groups. The user can also click the **Select Action** menu in 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.

Click the **Select Action** menu in 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 more information, see “Backup settings” on page 531, and “Restore Settings window” on page 534.

10.4.3 Cluster Settings page

Use the Cluster Settings to view or change settings that determine how a cluster runs copy policy overrides, applies Inhibit Reclaim schedules, uses an encryption key server, implements write protect mode, and performs backup and restore operations.

Figure 10-33 shows the Cluster Settings icon and options.

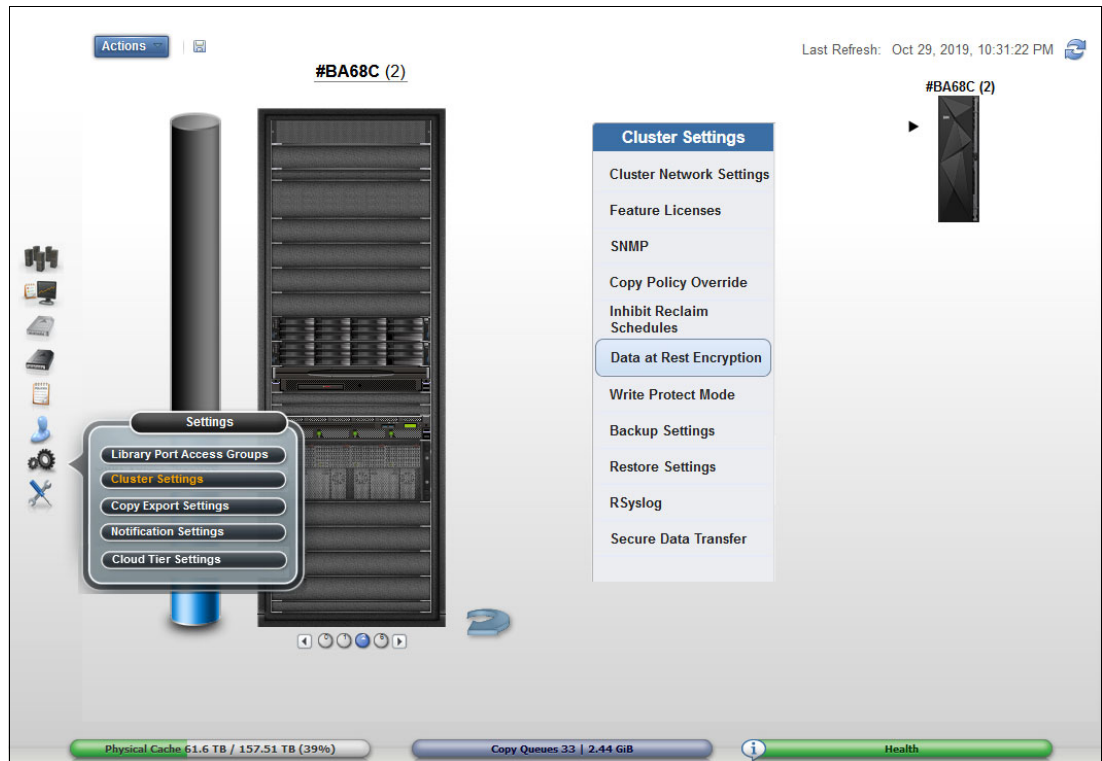


Figure 10-33 The Cluster Settings options

Cluster network settings

Use this page to set or modify IP addresses for the selected IBM TS7700 cluster.

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.

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 are displayed on this tab:

- ▶ IPv4: Select this option 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 option 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 other IPv6 information, the minimum required local addresses for each customer network interface automatically are 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
 - Rsyslog server address

Note: Amazon S3 cloud object storage requires a DNS server that can convert *amazonaws.com addresses into IP addresses.

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.

The Encrypt Grid Communication tab is no longer visible in the GUI TS7700 MI if the IPsec grid communication encryption is disabled in all clusters of the grid. Use the Secure Data Transfer option instead.

Important: Secure Data Transfer has no significant effect on the performance of the TS7700 when enabled. Secure Data Transfer requires FC 5281 installed.

Feature licenses

The user can view information about feature licenses, or activate or remove feature licenses from the TS7700 cluster from this TS7700 MI page.

The Feature Licenses window includes the following fields:

- ▶ 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 caches 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: These settings can be backed up by using the Backup Settings function under the Cluster Settings tab and restoring 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.

SNMP settings

Use this section to configure global settings that apply to Simple Network Management Protocol (SNMP) traps on an entire cluster. SNMP is a networking protocol that enables the TS7700 to gather and automatically transmit information about alerts and status to other entities in the network. SNMP traps events, such as logins, configuration changes, status changes (vary on or off, and service prep), shutdown, and code updates.

Consider the following points when adding or modifying SNMP destinations:

- ▶ 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.

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.
 - V2: An alternative trap version.
 - V3: An alternative trap version.
 - Enable SNMP Traps: Enables or disables SNMP traps on a cluster. The option 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 does 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 that are listed in the Destination Settings table by using the current SNMP trap values. The Enable SNMP Traps option does not need to be selected 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.

- ▶ IP address: The IP address of the SNMP server. This value can take any of the following formats: IPv4, IPv6, a *hostname* 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) when they are added, modified, or deleted. These updates do not depend on the user clicking **Submit Changes**.

Any change to SNMP settings is logged on the Tasks window.

Copy Policy Override

Use this page to override local copy and I/O policies for a specific TS7700 cluster.

Reminder: The items in this window can modify the cluster behavior regarding local copy and certain I/O operations. Some **LI REQUEST** commands can also perform this action.

For the selected cluster, the user can tailor copy policies to override certain copy or I/O operations. Select one or more of the following settings to specify a policy override:

► Prefer local cache for scratch mount requests

When this setting is selected, a scratch 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 Copy Mode field that is defined for the local cluster is not Deferred when one or more peer clusters are defined as Rewind Unload (RUN).
- 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 private 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 includes a valid copy of the data, even if that data is only on physical tape.

If the local cluster does not have a valid copy of the data, the default cluster selection criteria apply.

► Force volumes that are mounted on this cluster to be copied to the local cache

When this setting is selected for a private (non-scratch) mount, a copy operation is performed on the local cluster as part of the mount processing. When this setting is selected for a scratch 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 it does influence 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 setting in this window, complete the following steps:

- a. Select or clear the box next to the setting that must be changed. If the user enables the Enable fewer RUN consistent copies before reporting the RUN command complete option, the user can alter the value for the Number of required RUN consistent copies including the source copy.
- b. Click **Submit Changes**.

Inhibit Reclaim Schedules window

Use this window to add, modify, or delete Inhibit Reclaim Schedules that are used to postpone tape reclamation in a TS7700T cluster.

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.

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 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 is 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 00:00–23:59.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 the 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.

To modify an Inhibit Reclaim schedule, complete the following steps:

1. From the Inhibit Reclaim Schedules window, go to the Schedules table.
2. Select 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.

To delete an Inhibit Reclaim schedule, complete the following steps:

1. From the Inhibit Reclaim Schedules window, go to the Schedules table.
2. Select 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.

TS7770 Data at Rest Encryption

Use this page to set the EKS addresses in the TS7700 cluster. This selection is available only on the TS7700 Management Interface in any of the following circumstances:

- ▶ A physical tape library and the tape encryption enablement feature (FC 9900) are installed.
- ▶ The Disk Encryption with External Key Management feature (FC 5276) is installed.

Note: IBM Guardium Key Lifecycle Manager 4.1 and later is supported on TS7700 with R5.1 or higher level of code.

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.

Before the TS7700-VED server and 3956-CSB cache model, the cache models 3956-CC9, 3956-CS9, and 3956-CSA encrypted their data at rest by using the full data encryption (FDE) DDMs, with local or external management of the encryption key (only the IBM Security Key Lifecycle Manager, supported).

Both cache models 3956-CC9/CS9 and 3956-CSA required IPP key servers for encryption and the tape encryption. Figure 10-34 shows an example of the Data at rest Encryption window.

Cluster Settings	
Cluster Network Settings	"Venus[1]" (#BA12B): Data at Rest Encryption
Feature Licenses	Disk encryption: Disabled
SNMP	Tape encryption: License installed
Copy Policy Override	Key server type: IBM SKLM (IPP)
Inhibit Reclaim Schedules	Primary key server
Data at Rest Encryption	Address: 9.7.124.236 Port: 3801 Test connectivity
Write Protect Mode	IPP TLS 1.2 <input type="checkbox"/>
Backup Settings	Secondary key server
Restore Settings	Address: 9.7.124.204 Port: 3801 Test connectivity
RSyslog	IPP TLS 1.2 <input type="checkbox"/>
Secure Data Transfer	Submit Changes

Figure 10-34 Data at rest Encryption by using IBM Security Key Lifecycle Manager (IPP)

The TS7770-VED features a new type of cache, the 3956-CSB. In this cache model, the encryption is performed by the CSB processor at the nodes on the control enclosure, which is a different approach from previous models in which the encryption is performed on the FDE-DDMs. In the 3956-CSB, the encrypted data is written into regular DDMs (non-FDE) in the drawers. The DRAID arrays in the CSB cache must be encrypted at creation time at manufacturing.

Consider the following points regarding the CSB cache model:

- ▶ FC 5272 Disk Enabled Encryption is not available for field installations on the TS7770. It must be shipped from manufacturing for any encryption.
- ▶ FC 7405 must be ordered on every 3956-CSB in the TS7770 configuration.
- ▶ FC 7405 provides four USB sticks per 3956-CSB used to store the local encryption keys

Therefore, the External Key Encryption (FC 5276) requires that FC 5272 is installed on the TS7770 server before initial installation. All TS7770 configurations with 3956-CSB /XSB ordered for encryption cache data are always shipped from manufacturing with local key management enabled (FC 5272).

After a TS7770 with FC 5272 is configured in a client environment and can communicate with an external key server, FC 5276 can be activated to transition to external key management.

The CSB cache model requires KMIP key servers for encryption. Figure 10-35 shows an example of an IBM Security Key Lifecycle Manager key server configuration for CSB cache encryption (KMIP portion, port 5696) and tape encryption (IPP portion, port 441). The IPP portion uses port 441 because the IPP TLS 1.2 option is selected; otherwise, the default port is still 3801. With the TLS option, the option to import a Key Server certificate is displayed.



Figure 10-35 Data at Rest Encryption window with CSB (KMIP)

For more information about how to configure the TS7700 Virtual Engine to apply external key management, see Appendix J, “Configuring externally managed encryption” on page 1013.

For more information, see Chapter 2, “Architecture, components, and functional characteristics” on page 15, and [TS7700 5.4.1](#).

Note: Only IBM Guardium Key Lifecycle Manager 4.1 and later (formerly *IBM Security Key Lifecycle Manager*) supports external disk encryption and TS1140, TS1150, and TS1160 tape encryption. The settings for Encryption Server are shared for tape and external disk encryption. That is, a TS7700 supports only one dedicated key manager.

For 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 145, and Chapter 7, “Hardware configurations and upgrade considerations” on page 261.

A tutorial is available on MI that shows the properties of the EKS. To watch it, click **View tutorial** in the MI window. This window is visible on the TS7700 MI if the cluster includes the feature code for disk or tape encryption enabled.

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, the port settings are the same as the default values when this file is restored.

Write Protect Mode window

Use this page on the TS7700 MI to view Write Protect Mode settings in a TS7700 cluster. This window also is displayed if the Write Protect Mode is enabled because FlashCopy is enabled (the Current State field shows the Write protect for FlashCopy option is enabled).

With FlashCopy in progress, the Write Protect Mode window cannot be modified 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 by using the Library Request Host Console (**LI REQ**) command.

Meanwhile, host commands that are sent to virtual devices in peer clusters can continue with full read and write access to all volumes in the library. Write Protect Mode is used primarily for client-initiated DR 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. With R5.0, up to 128 categories can be identified and set to include or exclude from Write Protect Mode by using the Category Write Protect Properties table. Also, write-protected volumes in any scratch category can be mounted as private volumes if the Ignore Fast Ready characteristics option of the write-protected categories is selected.

Note: A maximum of 128 categories can be added per cluster when all clusters in the grid operate at R5.0 or later level of code.

The following settings are available:

- ▶ Write Protect Mode settings

Write Protect Mode does not affect standard grid functions, such as virtual volume replication or virtual volume ownership transfer. The settings that are available in the Write Protect Mode page are listed in Table 10-6 on page 529.

Table 10-6 Write Protect Mode settings for the active cluster

Setting	Description
Write Protect State	<p>Displays the status of Write Protect Mode on the active cluster. The following values are possible:</p> <ul style="list-style-type: none"> ▶ Disabled: Write Protect mode is disabled. No Write Protect settings are in effect. ▶ Enabled: Write Protect Mode is enabled. Any host command to modify volume data or attributes by using virtual devices in this cluster fails, subject to any defined category exclusions. ▶ Write protect for FlashCopy enabled: Write Protect Mode is enabled by the host. The Write Protect for FlashCopy function was enabled by using the LI REQ zOS command and a DR test is likely in progress. <p>Important: Write Protect Mode cannot be modified while LI-REQ-initiated write protection for FlashCopy is enabled. User must disable it first by using the LI REQ command before trying to change any Write Protect Mode settings on the TS7700 MI:</p> <ul style="list-style-type: none"> – DR family: The name of the DR family the Write Protect for FlashCopy was initiated against. – Flash time: The date and time at which the FlashCopy was enabled by the host. This mimics the time at which a real disaster occurs.
Disable Write Protect Mode	Select this to disable Write Protect Mode for the cluster.
Enable Write Protect Mode	<p>Select this option to enable Write Protect Mode for devices that are associated with this cluster. When enabled, any host command fails if it attempts to modify volume data or volume attributes through logical devices that are associated with this cluster, subject to any defined category exclusions. After Write Protect Mode is enabled through the TS7700 Management Interface, it persists through any outage and can be disabled only through the TS7700 Management Interface.</p> <p>Note: Write Protect Mode can be enabled only on a cluster if it is online and no Write Protect FlashCopy is in progress.</p>
Ignore fast ready characteristics of write protected categories	<p>Select this setting to permit write protected volumes that were returned to a scratch or a fast ready category to be recognized as private volumes.</p> <p>When this setting is selected, DR test hosts can mount production volumes as private volumes, even if the production environment returned them to scratch. However, peer clusters, such as production clusters, continue to view these volumes as scratch volumes. This setting does not override the Fast Ready characteristics of the excluded categories.</p> <p>Consider the following points:</p> <ul style="list-style-type: none"> ▶ When the production environment returns volumes to scratch that are still needed for read operations on a DR test host, the user must ensure that the data is available to be read by the DR hosts. To meet this requirement, scratch categories in the production environment can use, delete, expire, or hold for a period that is long enough to cover the DR test. Otherwise, even if selective write protect is being used on the DR clusters, the data is no longer available to the DR hosts if the production volume is returned to scratch and deleted or reused. One alternative is to avoid running return to scratch processing from the production environment during the DR test. Another alternative in a DR environment that is composed only of TS7700s is to use FlashCopy for DR testing, which is immune to production changes, including any scratch processing, deletion, and reuse. ▶ When this setting is selected, any volume that is deleted by the production host also is deleted on the DR cluster.

▶ 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 while 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.

Note: Categories that are configured and displayed in this table are not replicated to other clusters in the grid.

When Write Protect Mode is enabled, any categories that are added to this table must display a value of Yes in the Excluded from Write Protect field before the volumes in that category can be modified by an accessing host.

The following category fields are displayed in the Category Write Protect Properties table:

- Category Number: The identifier for a defined category. This alphanumeric hexadecimal value is 0x0001–0xFEFF (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.

The user can add up to 128 categories per cluster when all clusters in the grid operate at code level R5.0 or later.

The following windows can be used to manage categories:

- ▶ Add category: Use this window to add a new Write Protect Mode category in an IBM TS7700 cluster.
- ▶ Modify category: Use this window to modify a Write Protect Mode category in an IBM TS7700 cluster.
- ▶ Confirm Delete Category: Use this page to delete a Write Protect Mode category in an IBM TS7700 Cluster.

The following tips regarding Disaster Recovery (DR) tests from IBM Documentation can be valuable when planning for DR tests. For more information, see Chapter 16, “Disaster recovery testing in a grid configuration” on page 819:

- ▶ Use the Write Protect Mode during a DR test to prevent any accidental DR host-initiated changes to your production content.
- ▶ During a DR test, it is recommended to avoid housekeeping (or return to scratch processing) tasks within the DR test host configuration unless the process specifically targets volumes only within the DR host test range. Otherwise, even with the Selective Write Protect function enabled, the DR host can attempt to return production volumes to

scratch. This problem can occur because the tape management system snapshot that is used for the DR test can interpret the volumes as expired and ready for processing.

Never assume the return to scratch process acts only on DR test volumes. If Write Protect Mode is enabled before DR testing and return to scratch processing is run on the DR host, the Selective Write Protect function prevents the return to scratch from occurring on protected categories. Further, options in the tape management system can be used to limit which volumes within the DR host are returned to scratch.

For example, in the DFSMSrmm tape management system, the **VOLUMES** or **VOLUMERANGES** options can be used on the **EXPROC** command to limit volumes that are returned to scratch. When tape management and write protect safeguards are used, protection against data loss occurs both on the TS7700 and at the host.

- ▶ It is not necessary to disable return to scratch processing within the production environment if the following category settings are used:
 - Delete expire hold option is set to prevent the reuse of production volumes before they are verified within the DR test application.
 - Expire hold is not required when the FlashCopy for DR testing mechanism is used and no TS7700 Tape Attach is present within the DR location, since the flash snapshot is unchanged by any scratch and reuse of volumes in the production environment.
 - Ignore fast ready characteristics option is set to ensure that scratched volumes are viewed as private when accessed through the DR test devices.
- ▶ Write protection extends only to host commands that are issued to logical devices within a write-protected cluster. Volume ownership and data location are independent of write protection. For example, production access to non-write-protected devices within a production cluster can still alter content through a remote mount to a DR cluster in the Write Protect State.
- ▶ You can leave a cluster in the Write Protect State indefinitely. Doing so can prevent unexpected host access in a DR location from modifying production content. However, write protect must be disabled if a true failover to the secondary location occurs.

For more information about DR testing, see Chapter 5, “Disaster recovery” on page 213, and Chapter 16, “Disaster recovery testing in a grid configuration” on page 819.

Backup settings

This section describes backing up the settings from a TS7700 cluster.

Important: Backup and restore functions are not supported between clusters that are operating at different code levels. Only clusters that are 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 that are operating different code levels are visible, but the options are disabled.

The cluster settings that are available for backup (and restore) in a TS7700 cluster are listed in Table 10-7.

Table 10-7 Backup and restore settings reference

Setting	Can be backed up from	Can be restored to
Storage Classes Data Classes	Any TS7700 cluster	Any TS7700 Cluster

Setting	Can be backed up from	Can be restored to
Partitions	TS7700T TS7700C	TS7700T TS7700C
Inhibit Reclaim Schedule Physical Volume Pools Physical Volume Ranges	A TS7700 cluster that is attached to a tape library	A TS7700 cluster that is attached to a tape library
Library Port Access Groups Categories Storage Groups Management Classes Session Timeout Account Expirations Account Lock Roles and Permissions Encryption Key Server Addresses Cluster Network Settings Feature License Settings Copy Policy Override SNMP Write Protect Mode Categories	Any TS7700 Cluster	Any TS7700 Cluster

The Backup Settings table lists the following cluster settings that are available for backup:

- ▶ **Constructs:** Select this setting to include all of the following constructs for backup:
 - Storage Groups: Select this option to back up defined Storage Groups.
 - Management Classes: Select this option to back up defined Management Classes.
 - Storage Classes: Select this option to back up defined Storage Classes.
 - Data Classes: Select this option to back up defined Data Classes.
- ▶ **Partitions:** Select this setting to back up defined partitions.

Note: Resident and Object partitions are not considered in the options.

- ▶ **Inhibit Reclaim Schedule:** Select this option to back up the Inhibit Reclaim Schedules that are used to postpone tape reclamation. This option is not available if the cluster does not have an attached tape library.
- ▶ **Library Port Access Groups:** Select this option to back up defined library port access groups. Library port access groups and access group ranges are backed up together.
- ▶ **Categories:** Select this option to back up scratch categories that are used to group virtual volumes.
- ▶ **Physical Volume Ranges:** Select this option to back up defined physical volume ranges. This option is not available if the cluster does not have an attached tape library.
- ▶ **Physical Volume Pools:** Select this option to back up physical volume pool definitions. This option is not available if the cluster does not have an attached tape library.
- ▶ **Security Settings:** Select this option to back up defined security settings:
 - Session Timeout
 - Account Expiration
 - Account Lock
- ▶ **Roles & Permissions:** Select this option to back up defined custom user roles.

Important: A restore operation after cluster settings are backed up does *not* restore or otherwise modify any user, role, or password settings that are defined by a security policy.

- ▶ Encryption Key Server Addresses: Select this option to back up the information at the TS7700 Data at Rest Encryption page:
 - IPP only (3956-CSA and Tape encryption):
 - IP address (primary and secondary key server)
 - Ports IPP (primary and secondary key server)
 - TLS enabled
 - KMIP only (3956-CSB and CFC):
 - IP address (primary and secondary key server)
 - Ports KMIP (primary and secondary key server)
 - Certificate information
 - KMIP and IPP (3956-CSB/CFC and Tape):
 - IP address (primary and secondary key server)
 - Ports IPP (primary and secondary key server)
 - Ports KMIP (primary and secondary key server)
 - TLS enabled
 - Certificate information

Note: The certificate is not backed up or restored by this procedure. It must be trusted in advance by the TS7700 by using the SSL Certificates page in the TS7700 MI.

- ▶ Feature Licenses: Select this option 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.
- ▶ Copy Policy Override: Select this option to back up the settings to override local copy and I/O policies.

- ▶ **SNMP:** Select this option to back up the settings for SNMP.
- ▶ **Write Protect Mode Categories:** Select this option to back up the settings for write protect mode categories.

To back up cluster settings, select 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.

Note: 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 (SDT)
- ▶ SSL certificates

Record these settings in a safe place and recover them manually, if necessary.

For more information about backup settings, see this IBM Documentation [web page](#).

Restore Settings window

Use this window to restore the settings from a TS7700 cluster to a recovered or new cluster.

Note: Backup and restore functions are not supported between clusters that are operating at different code levels. Only clusters that are operating at the same code level as the current cluster can be selected from the Current Cluster Selected graphic. Clusters that are operating at different code levels are visible (but not available) in the graphic.

For more information about Backup and Restore settings, see Table 10-7 on page 531.

Complete the following steps to restore cluster settings:

1. Use the banner bread crumbs to navigate to the cluster where the restore operation is applied.
2. In the Restore Settings window, click **Browse** to open the File Upload window.
3. Go to the backup file that was used to restore the cluster settings. This file has a `.xmi` extension.

4. Add the file name to the **File name** field.
5. Click **Open** or press Enter from the keyboard.
6. 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 setting to restore scratch categories that are used to group virtual volumes.
- ▶ **Physical Volume Pools:** Select this setting 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.

- ▶ **Constructs:** Select this setting to restore all of the displayed constructs. When these settings are restored, new settings are added and existing settings are modified, but no settings are deleted:
 - **Storage Groups:** Select this setting to restore defined SGs.
 - **Management Classes:** Select this setting 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 that are present in the restore file, MC copy policies on the cluster are overwritten with those policies 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 setting to restore defined SCs.
 - **Data Classes:** Select this setting 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.

Media type volume sizes are restored based on the restrictions of the restoring cluster.
- ▶ **Inhibit Reclaim Schedule:** Select this setting 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 schedule settings are reset to default.

- ▶ Tape Partitions: Select this setting to restore defined tape or cloud partitions.

Note: If a partition exists with the same name or if the maximum number of partitions is reached, the restore cannot be performed.

- ▶ Library Port Access Groups: Select this setting to restore defined library port access groups.
Library port access groups and access group ranges are backed up and restored together.
- ▶ Physical Volume Ranges: Select this setting 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 setting 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 that are defined by a security policy.
- ▶ Security Settings: Select this setting to restore defined security settings, for example:
 - Session Timeout
 - Account Expiration
 - Account Lock
- ▶ Encryption Key Server Addresses: Select this setting 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 information at the TS7700 Data at Rest Encryption page can be restored:
 - IPP only (3956-CSA and Tape encryption):
 - IP address (primary and secondary key server)
 - Ports IPP (primary and secondary key server)
 - TLS enabled
 - KMIP only (3956-CSB/CFC):
 - IP address (primary and secondary key server)
 - Ports KMIP (primary and secondary key server)
 - Certificate information.
 - KMIP and IPP (3956-CSB/CFC and Tape):
 - IP address (primary and secondary key server)
 - Ports IPP (primary and secondary key server)
 - Ports KMIP (primary and secondary key server)
 - TLS enabled
 - Certificate information

Note: The certificate is not backed up or restored by this procedure. It must be trusted in advance by the TS7700 by using the SSL Certificates page in the TS7700 MI.

- ▶ Cluster Network Settings: Select this setting 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 setting 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 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 was yet used.
- ▶ **Copy Policy Override:** Select this setting to restore the settings to override local copy and I/O policies.
- ▶ **SNMP:** Select this setting to restore the settings for Simple Network Management Protocol (SNMP). When these settings are restored, new settings are added and existing settings are modified, but no settings are deleted.
- ▶ **Write Protect Mode Categories:** Select this setting to restore the settings for write protect mode categories. When these settings are restored, new settings are added and existing settings are modified, but no settings are deleted.

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 the Restore window while the operation is in progress, the current 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.

Restoring to or from a cluster without a physical library: If the cluster that created the backup file or the cluster that is performing the restore operation do not possess a physical library, upon completion of the restore operation all physical tape library settings are reset to default. One of the following warning messages is displayed on the confirmation page:

The file was backed up from a system with a physical tape library attached but this cluster does not have a physical tape library attached. If you restore the file to this cluster, all the settings for physical tape library have default values.

The file was backed up from a cluster without a physical tape library attached but this cluster has a physical tape library attached. If you restore the file to this cluster, all the settings for physical tape library have default values.

The following settings are affected:

- ▶ Inhibit Reclaim Schedule
- ▶ Physical Volume Pools
- ▶ Physical Volume Ranges

Confirm restore settings: This page confirms that a restore settings operation is in progress on the IBM TS7700 cluster.

For more information about Restore settings, see this IBM Documentation [Restore settings](#).

RSyslog

Use this page to set or modify RSyslog settings on the IBM TS7700 cluster. The user can add or modify settings of a remote target (the Rsyslog server) to which the system logs are sent. All TS7700 cluster models can have the Rsyslog function configured and operational. For more information, see Chapter 2, “Architecture, components, and functional characteristics” on page 15.

The following information is displayed on this page:

- ▶ Target Number: The number of the remote target (server) to which to send the system logs.
- ▶ IP address or hostname: IP address or hostname of the remote target to receive the system logs.

Note: If this value is a Domain Name Server (DNS) address, you must activate and configure a DNS on the Cluster Network Settings window.

- ▶ Port: Port number of the remote target.
- ▶ Status: Status of the remote target, which can be Active or Inactive.

Note: To send logs to a syslog server, RSyslog must be enabled and the status of the remote target must be Active. Otherwise, no logs are sent.

The available actions that are on this page are listed in Table 10-8.

Table 10-8 Actions available from the RSyslog table

Action	Description
Enable or disable RSyslog	To change the state, select Enable RSyslog or Disable RSyslog .

Action	Description
Create a remote target	<ol style="list-style-type: none"> 1. Select Actions → Create Remote Target. 2. Add the remote target settings: <ul style="list-style-type: none"> ▶ IP address or hostname ▶ Port ▶ Status <p>The user can have a total of two remote targets.</p> 3. Click OK to add the remote target, or Cancel to quit the operation.
Modify a remote target	<ol style="list-style-type: none"> 1. Highlight a remote target. 2. Select Actions → Modify Remote Target. 3. Modify the remote target settings: <ul style="list-style-type: none"> ▶ IP address or hostname ▶ Port ▶ Status 4. Click OK to modify the remote target, or Cancel to quit the operation.
Delete a remote target	<ol style="list-style-type: none"> 1. Highlight a remote target. 2. Select Actions → Delete Remote Target. 3. Click OK to delete the remote target, or Cancel to quit the operation.
Change the order of the remote targets	Click the Target Number column heading.
Hide or show columns on the table	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Select the checkbox next to a column heading to hide or show that column in the table. Column headings that are selected display on the table.
Reset the table to its default view	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Click Reset Grid Preferences.

For more information, see [Rsyslog](#).

Secure Data Transfer

Secure data transfer (SDT) provides a way to secure the data that is transferred between clusters within a grid. SDT encrypts the data that is in transit by way of grid links, including objects that are exchanged between the DS8900 and the TS7700 in both directions, remote reads, writes, and replication data.

To enable SDT on a cluster in a grid, install FC 5281. The SDT encryption occurs only between clusters that include SDT enabled on both ends of the data transfer transaction. If a member of a grid has the SDT encryption disabled, then data sent to the cluster is not encrypted.

Note: If DS8000 object encryption is required, FC 5281 Secure Data Transfer must be activated on any TS7700 that is to be a DS8000 TCT target. No extra features are required to support SDT on the DS8000.

Secure Data Erase (SDE) explores the in-core encryption acceleration, which is available on VEC and VED (IBM Power Systems P8 and P9) cluster models, with no significant effect on the global performance of the cluster. SDE encryption can be enabled or disabled concurrently with cluster operation by way of TS7700 MI.

SDT uses OpenSSL software libraries with the TLS1.2 protocol that follow AES standards. AES-256 and AES-128 bit keys are supported. Logical volume data is encrypted within the TS7700 before transport, so no special network requirements are needed.

Note: SDT cannot be enabled on a stand-alone cluster. FC 5281 is required to enable Secure Data Transfer.

Figure 10-36 shows an example of an SDT configuration window.

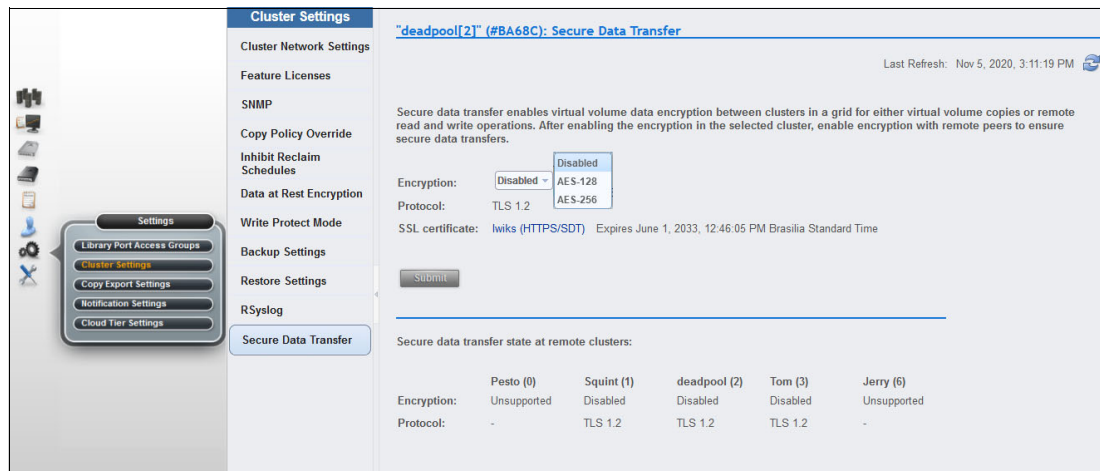


Figure 10-36 Secure Data Transfer page

SDT features the following options:

- ▶ **Encryption:** This option allows customers to enable and disable encryption and select an AES encryption key size by way of the MI as a concurrent activity. Selecting a key size automatically enables encryption for this TS7700. AES-128 Encryption is enabled with 128-bit or 256-bit encryption.
- ▶ **Protocol:** TS7700 supports TLS1.2.
- ▶ **SSL Certificate:** The TS7700 uses a default certificate that is known as “lwiks” for SDT server authentication. Optionally, users can upload their own trusted certificates from the MI SSL Certificate window at **Access** → **SSL Certificate**.

For more information about Secure Data Transfer, see 4.3.9, “Secure Data Transfer” on page 189.

10.4.4 Copy Export Settings window

Use this window to change the maximum number of physical volumes that can be exported by the TS7700.

Note: The *Copy Export Settings* option is only available for tape-attached stand-alone clusters or when there is at least one tape-attached member in the Grid.

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 15, “Copy Export” on page 783.

10.4.5 Notification settings

Use this page to set or modify notification settings on the IBM TS7700 Cluster.

This page displays information for an entire grid if the accessing cluster is part of a grid, or for only the accessing cluster if that cluster is a stand-alone machine. Use this page to modify settings of the notifications that are generated by the system, such as Event, Host Message, and Call Home Microcode Detected Error (MDE).

The following types of notifications are generated by the system:

- ▶ Events (OPxxxx messages in a CBR3750I message on the host)
- ▶ Host Message (Gxxxx, ALxxxx, EXXXX, Rxxxx in a CBR3750I message on the host)
- ▶ Call Home MDE in SIM (in a IEA480E message on the host)

During normal and exception processing, intervention, or action by an operator or storage administrator is sometimes required. As these conditions or events are encountered, a message is logged to the host console. Those messages can be viewed through the TS7700 MI by using the Events window as shown in 9.7.1, “**Events**” on page 388.

Notification Settings allows the user to adjust different characteristics for all CBR3750I messages. Also, the user can add personalized text to the messages, which makes it simpler to implement or manage the automation-based monitoring in the IBM z/OS.

The Notification Settings window also works as a catalog that can be used to search for descriptions of the Event, Host Message, or Call Home MDE. From this window, the user can send the text messages to all hosts that are attached to the subsystem by enabling the Host Notification option.

Figure 10-37 on page 542 shows the Notification Settings window and options.

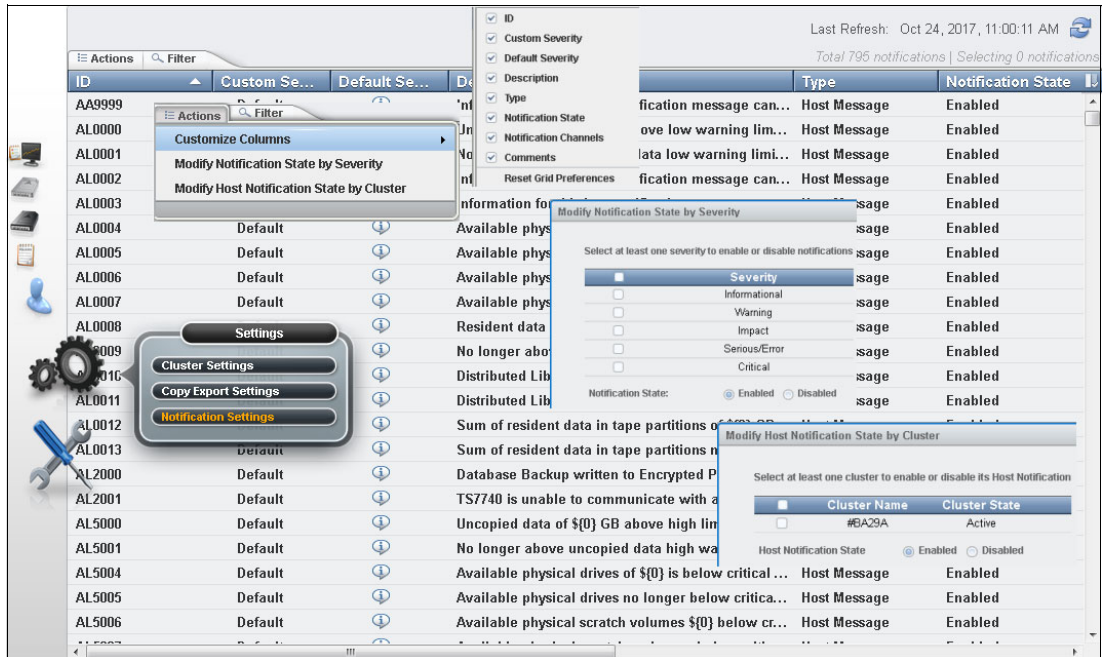


Figure 10-37 Notification Settings and options

Information that can be viewed on the Notification settings is listed in Table 10-9.

Table 10-9 Notification settings

ID	ID of the Event, Host Message, or Call Home MDE
Custom Severity	Custom severity of the notification. The following values are possible: <ul style="list-style-type: none"> ▶ Default ▶ Informational ▶ Warning ▶ Impact ▶ Serious/Error ▶ Critical
Description	Description of the notification.
Type	The type of notification that is listed: Event, Host Message, or Call HOME MDE.
Notification State	Whether the notification is active or inactive. If inactive, it is not sent to any notification channel.
Notification Channels	How the notification is sent. The following selections are possible: <ul style="list-style-type: none"> ▶ Host ▶ Management Interface ▶ SNMP ▶ RSyslog <p>Note: Host Messages cannot be sent to the Management Interface.</p>
Comments	Field that is available to add user comments. The comments are sent with the message through the notification channels when the message is triggered by the system. <p>Note: Comments for Call Home MDEs are not sent to Host. Only the MDE code is sent to the Host.</p>

The actions that are available from the Notifications table are listed in Table 10-10.

Table 10-10 Notifications actions

Action	Description
Enable or disable host notifications for alerts	<ol style="list-style-type: none"> 1. Select Actions → Modify Host Notification State by Cluster. 2. Select the cluster in the Modify Host Notification State by Cluster box. <ul style="list-style-type: none"> ▶ Select Active; then, OK to enable notifications. ▶ Select Inactive; then, OK to disable notifications.
Set notifications to Active or Inactive	<ol style="list-style-type: none"> 1. Select at least one notification. 2. Select Actions → Modify Notifications Settings. 3. Select Active; then, OK to enable notifications. 4. Select Inactive; then, OK to disable notifications.
Set a custom severity level for notifications	<ol style="list-style-type: none"> 1. Select at least one notification. 2. Select Actions → Modify Notifications Settings. 3. From the Custom Severity drop-down box, select the severity level. 4. Verify that the Notification State is Active; then click OK. <p>The option to set a custom severity can be performed only if all clusters in the grid are at microcode level 8.41.2 or later.</p>
Modify notification state by severity	<ol style="list-style-type: none"> 1. Select at least one notification. 2. Select Actions → Modify Notifications State by Severity. <ul style="list-style-type: none"> ▶ Select Active; then click OK to enable notifications. ▶ Select Inactive; then click OK to disable notifications. <p>This action works based on the current value of the severity, whether is the default severity or custom severity.</p>
Download a CSV file of the notification settings	<p>Select the File icon (export to CSV).</p> <p>The time that is reported in the CSV file is shown in Coordinated Universal Time.</p>
Hide or show columns on the table	<ol style="list-style-type: none"> 1. Right-click the table header. 2. Click the checkbox that is 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 by using a string of text	<ol style="list-style-type: none"> 1. Click the Filter field. 2. Enter a search string. 3. Press Enter.
Filter the table data by using a column heading	<ol style="list-style-type: none"> 1. Click the down arrow next to the Filter field. 2. Select the column heading by which to filter. 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 Grid Preferences.

10.5 Service icon

This section provides a description of the running service operations and troubleshooting problems for the TS7700. Figure 10-38 on page 544 shows the Service icon options for a stand-alone and grid configuration of TS7700 clusters.

The Ownership Takeover Mode option shows only when a cluster is a member of a grid, whereas Copy Export Recover and Copy Export Recovery Status options appear for a single TS7700T configuration (that is connected to a physical library).

Note: The Copy Export Recover and Copy Export Recovery Status options are available only in a single cluster configuration for a tape-attached cluster.

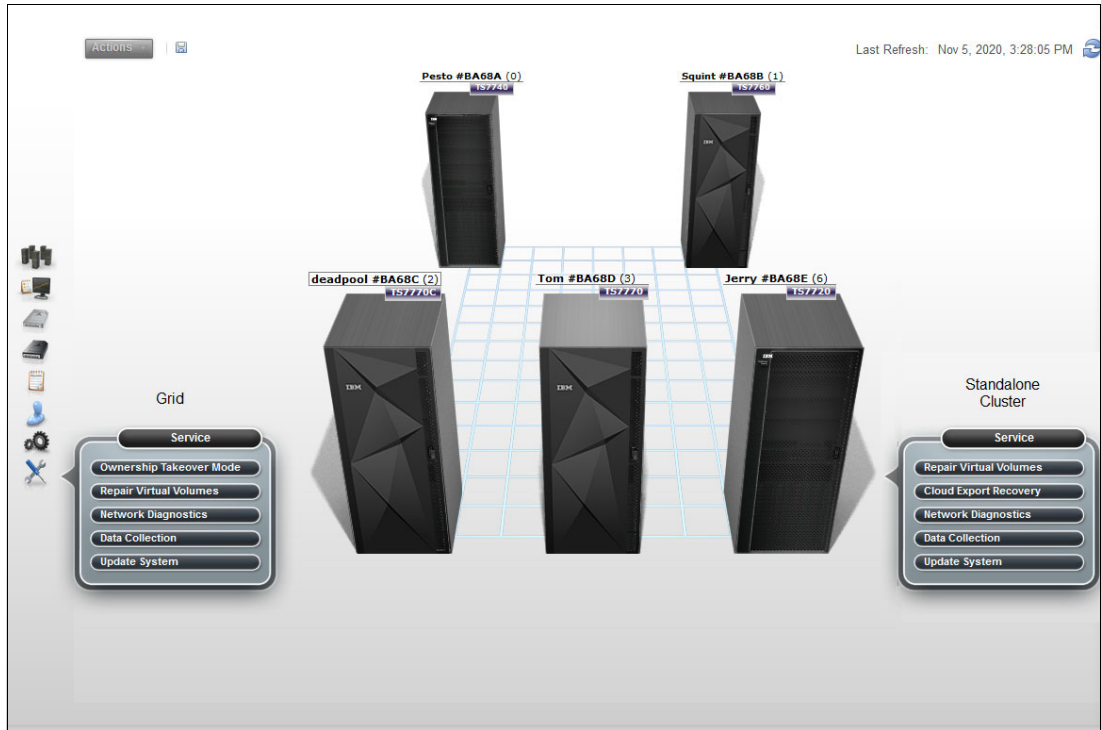


Figure 10-38 Service Icon options for a grid or stand-alone cluster

10.5.1 Ownership Takeover Mode

To select an Ownership Takeover Mode for a failed cluster in a TS7700, use the Ownership takeover mode page. 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: Have the IP addresses for all clusters in the grid configuration available if a cluster failure occurs. Therefore, the MI can be accessed from a surviving cluster to start the ownership takeover actions.

Normally, ownership is transferred from one cluster to another through communication between the clusters. When a cluster enters a failed state or the communication links between clusters fail, the other clusters in the grid cannot obtain the ownership of the logical volumes that belong to the cluster that stopped answering the requests.

When this occurs any host mount for that logical volume fails, possibly with **CBR4174I** message:

Cannot obtain ownership of volume volser in library library-name.

If one or more clusters become isolated from one or more peers in the grid, and a host mount is issued to a volume owned by the cluster that is not responding, this volume cannot be mounted without first enabling an ownership takeover mode. Volumes that are owned by one of the accessible clusters in the grid 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.

Read/Write takeover (WOT) allows the local cluster to read or write to any volumes it takes over from the remote cluster that is not responding. Nonetheless, read/write takeover should not be used if only the communication path between the clusters has failed and the isolated cluster remains operational and accessible to a host. In this scenario, read-only takeover mode (ROT) is advisable.

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.

When Read Only takeover mode is enabled, those volumes that require 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.

When available and configured, AOTM 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 that is associated with each cluster to determine whether the cluster is alive or failed, which enables the ownership takeover only in case the unresponsive cluster did fail. If the cluster is still alive, AOTM does not start a takeover, and the decision is up to the human operator.

Note: AOTM does *not* start a takeover if an isolated cluster is up and operational; for example, assume all grid links to a cluster were lost, which isolates this cluster from the grid. In such a case, manual takeover is still possible.

Figure 10-39 on page 546 shows the Ownership Takeover Mode window.

Figure 10-39 on page 546 also shows the local cluster summary, the list of available clusters in the grid, and 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 10-39, cluster zero features 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 take over the ownership must be made by a human operator or AOTM.

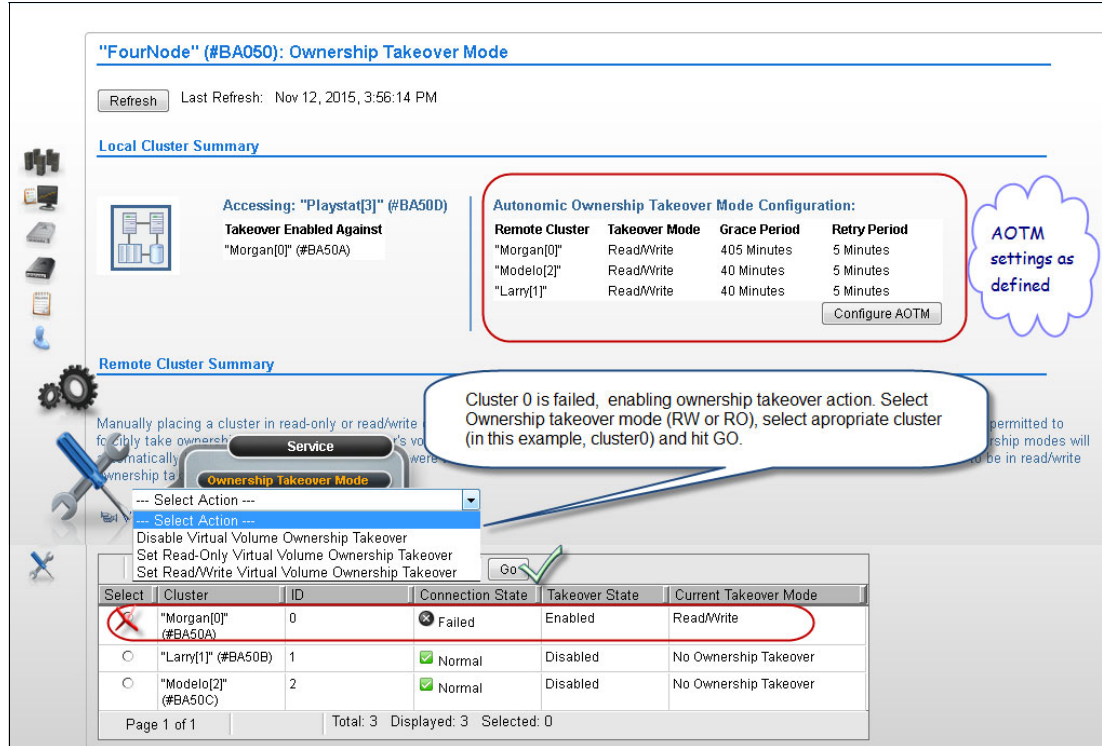


Figure 10-39 Ownership Takeover Mode

Complete the following steps to start an ownership takeover against a failed cluster:

1. Authenticate the MI to the surviving cluster showing 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 then retry the host operation that failed.

Figure 10-39 also shows that AOTM was configured in this grid (for read/write, with a grace period of 405 minutes for Cluster 0). In this case, the automatic ownership takeover occurs 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, or by changing the AOTM settings to more suitable values. Click **Configure AOTM** 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 (as part of the installation process, or later) *before* AOTM can be enabled and configured for any cluster in the grid.

The operation of read/write and read-only ownership takeover modes are compared in Table 10-11.

Table 10-11 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 the following 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 perform read operations on the virtual volumes that are owned by the failed cluster.</p> <p>Operational clusters in the grid cannot run the following 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>

For more information, see local by clicking the question mark icon in the upper right of the MI window, or see [TS7700 5.4.1](#).

10.5.2 Repair Virtual Volumes window

Damaged volumes typically occur because of 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, error messages are displayed, as shown in the following example:

```
CBR4125I Valid copy of volume volser in library library-name inaccessible
```

Use the Repair virtual volumes page to repair virtual volumes in the damaged category for the TS7700 Grid.

The user can print the table data by clicking **Print report**. A comma-separated value (.csv) file of the table data can be downloaded by clicking Download spreadsheet. The following information is displayed in this window:

- ▶ **Repair policy:** The Repair policy section defines the repair policy criteria for damaged virtual volumes in a cluster. The following criteria are 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 was 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 rendered the volume obsolete. If the volume was returned to scratch, the data on the volume is no longer needed. If data loss 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 Virtual Volume information is shown, which is 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.

Complete the following 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. Select **Move to insert category keeping all data** or **Move to insert category deleting all data**.
4. In the Damaged Virtual Volumes table, select the checkbox that is next to one or more (up to 10) damaged virtual volumes to be repaired by using the repair policy criteria.
5. Click **Select Action** → **Repair**.
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.

10.5.3 Network Diagnostics window

The Network Diagnostics window can be used to start ping or trace route commands to any IP address or hostname from this TS7700 cluster. The user can use these commands to test the efficiency of grid links and the network system.

Figure 10-40 shows the navigation to the Network Diagnostics window and a ping test example.

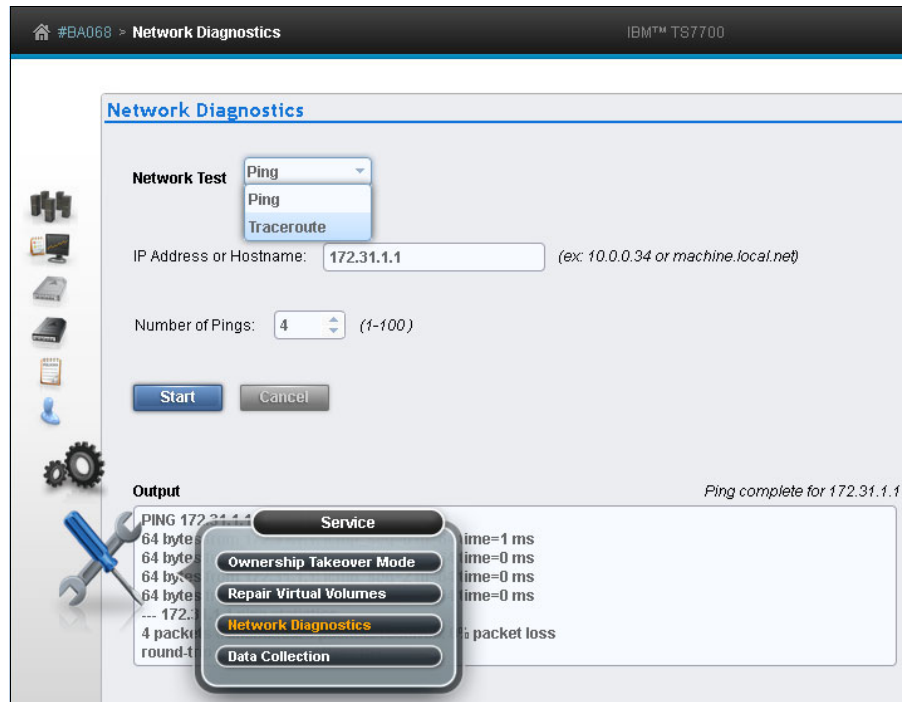


Figure 10-40 Network Diagnostics window

The following information is shown in 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 start a ping test against the IP address or hostname that is entered in the IP Address/Hostname, and back again. This option can test whether a connection to the target IP address or hostname is enabled, the speed of the connection, and the distance to the target.
 - **Traceroute:** Select this option to start a trace route test against the IP address or hostname 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 hostname for the selected network test. The value in this field can be an IP address in IPv4 or IPv6 format or a fully qualified hostname.
- **Number of Pings:** Use this field to select the number of pings that are sent by using the **Ping** command. The range of available pings is 1–100. The default value is 4. This field is displayed only 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 that is provided.

The status of the network command is displayed in line with the Output field label and right-aligned over the Output field. The information is displayed in the following format:

```
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
```

10.5.4 Data Collection window

Use this 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.

If the user is experiencing a performance issue on a TS7700, the following options are available to the user to collect system data for later troubleshooting:

- ▶ **System Snapshot**, which 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.
- ▶ **TS7700 Log Collection**, which enables you to collect historical system information for a 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.

The following information is shown on the Data Collection window:

- ▶ **System Snapshot:** Select this option 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:** Select this option 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 timestamp next to the hours field displays the earliest timestamp 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 a log file, a message prompts the user to select a different value for the hours field.

- ▶ **Continue:** Click this button to start the data collection operation. This operation cannot be canceled after the data collection process 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 the data collection process is started, a message is displayed that contains a button that links 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; therefore, no link to it is included in the message.

- ▶ **Data Collection Limit Reached:** This window 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 in this field a detailed description of the conditions or problem that was experienced before any data collection was started. 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.

10.5.5 Copy Export Recovery window

Use this window to test a Copy Export recovery, or to run an actual Copy Export recovery on the TS7700 cluster.

Tip: This window is visible only in a single tape-attached cluster.

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 DR. The user can also use this function to test system recovery.

For more information about the Copy Export function, see Chapter 15, “Copy Export” on page 783. Also, see [IBM TS7700 Series Copy Export Function User's Guide](#).

Reminder: The recovery cluster needs tape drives that are compatible with the exported media. Also, access to the EKs must be provided if encrypted tapes are used for export.

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 secondary copy, the original secondary copy is deleted because it becomes inactive data. For example, if the user modifies constructs for virtual volumes that 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 suffered a disaster. If this option is selected (default status), the Copy Export Recovery runs as a test. If the option is not selected, 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. Also, 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 option 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 (option is not selected) rewrites virtual volumes to physical storage if the constructs change so that the virtual volume's data can be placed in the correct pools. Also, in this type of recovery, reclaim processing remains enabled and primary physical copies are restored, which requires 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 FC 5270 licenses that are installed on the TS7700 that is used for recovery.

- ▶ Erase all existing virtual volumes during recovery: This option is shown if virtual volume or physical volume data is present in the database. A Copy Export Recovery operation erases any data. No option exists to retain data while running the recovery. The user can select this option to proceed with the Copy Export Recovery operation.
- ▶ Submit: Click this button to start the Copy Export Recovery operation.
- ▶ Confirm Submission of Copy Export Recovery: The user is prompted to confirm the decision to start 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** option, 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.

10.5.6 Copy Export Recovery Status window

Use this window to view information about or to cancel a running Copy Export recovery operation on a TS7700 cluster.

Important: The Copy Export recovery status is available for a stand-alone TS7700T cluster only.

The table in 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 completed steps. 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 timestamp for the start of the operation.**
- ▶ **Duration:** The amount of time the operation is 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 was 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 **Cancel Recovery** 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.

10.5.7 Cloud Export Recovery

Introduced in R5.1 level of code, the TS7700 Cloud Export and Recovery suite of features provides the capability to export a backup into an attached cloud. This backup can be used later to restore an empty TS7700C cluster. Cloud Export Recovery is started on the TS7700 Management Interface Cloud Export and Recovery page, under the Service Icon, as shown in Figure 10-41 on page 554.



Figure 10-41 Cloud Export Recovery page

Use this page to recover stored virtual volumes for restoration of data.

Before initiating a cloud export recovery process, ensure the following:

- ▶ All VECs/VEDs are at R5.1 or higher code level.
- ▶ A copy of the database is required for recovery purposes.
- ▶ All devices are offline to the host throughout the process.
- ▶ Gather information about the last backup needed for recovery.

This suite of features also includes the ability to perform DR testing to ensure disaster preparedness. It covers the concept of retaining older versions of logical volumes in the Cloud.

Note: Cloud Export Recovery option is only visible in a stand-alone cloud-attached TS7700, and the user was given permission to access that page.

Similar to the traditional Copy Export function, which exports secondary copies of selected logical volumes that use tape cartridges, Cloud Export is started by using the z/OS **Library Export <VOLSER>** command, which captures a database snapshot that can be used as a recovery restore point. A practical example is periodic backups that are taken for air-gap recovery purposes. For more information, see [Cloud export directory](#).

The TS7700 Cloud Export Recovery Test Mode feature allows the unique ability to perform full read/write testing against data that is stored in the production cloud vault without manipulating the production data. In test mode, the TS7700 recovery box places all production cloud pools into a READ-ONLY state and provide a temporary DRTEMP cloud pool that is used to redirect any write operations to the test data.

The DRTEMP cloud pool must be configured to an empty cloud container that is specified by the user. The test recovery box reads the data from the production cloud vault but automatically redirects any write operations to the DRTEMP. This function provides a safe environment to do conduct DR testing.

Important: Never clear the Test Mode Enable option, except in a real DR in which all production TS7700 clusters are lost. Clearing the Test Mode Enable option allows this recovery cluster to modify data in the production cloud vaults.

Cloud Export Recovery is a user-initiated operation that restores a stand-alone TS7700C to a fully functional state where it can read and write data that was preserved in the cloud.

The recovery TS7700 cluster includes the following requirements:

- ▶ Must not be part of a grid configuration.
- ▶ Must be configured as cluster 0 and cloud-attached system.
- ▶ It and the source cluster must be at level R5.1 or higher of code.
- ▶ Must be fully operational and in an online state to start recovery.
- ▶ Must be empty of data. If data is detected in the recovery cluster, it must be removed as part of the recovery process.

To start Cloud Export Recovery, the new cluster must be connected to the cloud infrastructure. The user must create a cloud pool, cloud account, and cloud container that points to the cloud vault that contains the database backup (which were created by a previous Cloud Export operation, as described at 10.5.7, “Cloud Export Recovery” on page 553).

After the recovery TS7700C is online, cloud credentials for the cloud vault that contains the database backup that is used in the Cloud Export Recovery process must be entered. This process is done in the Cloud Tier Settings MI window that is beneath the Settings icon (see 10.4.1, “Cloud Tier Settings” on page 513).

The following information must be entered to start the Cloud Export Recovery operation:

- ▶ The container ID created on this TS7700 that can access the cloud vault where the TS7700 database backup is stored that is to be used for this recovery.
- ▶ The serial number of the original TS7700 cluster that wrote the database backup.
- ▶ The cloud backup ID that is to be restored.

Refer to [IBM TS7700 Virtualization Engine Cloud Storage Tier Export, Recovery, and Testing Guide](#) for complete information on preparation, use case examples and guidance on the topic.

10.5.8 Update System

Starting with R5.1 PGA1 (8.50.1.25) and higher level of code, the new Update System item was added under the Service Icon. This item is also referred to as a *web updater*.

Note: The Update System option is available for the Administrator role only.

When an MI update becomes available for installation, the admin user is notified by a Suggested Tasks that is displayed in the Banner section of the TS7700 MI. Some of the updates can be identified as critical; for example, when a security fix is implemented and the newer version is available for installation.

The update packages are provided and distributed to the TS7700 clusters by the IBM TS3000 Total System Storage Console (TSSC). The TSSC needs an active internet connection (which is usually the case for the Call Home function) to query the IBM product support pages for available updates. A settings window is available on the TSSC to define such interval; for example, once a day (see Figure 10-42 on page 556). R5.1 and R5.2 need specific vtd_execs to be installed (for example, vtd_exec.900 or vtd_exec.901).

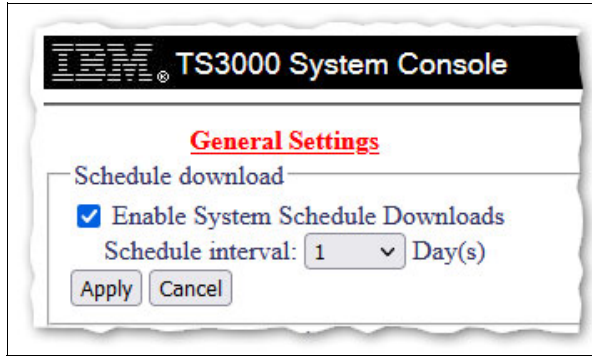


Figure 10-42 TS3000 System Console General Settings window for schedule download

Note: Consult with your IBM service representative (SSR) to activate, configure, and use this feature.

Figure 10-43 shows the Update System page, and the available options.

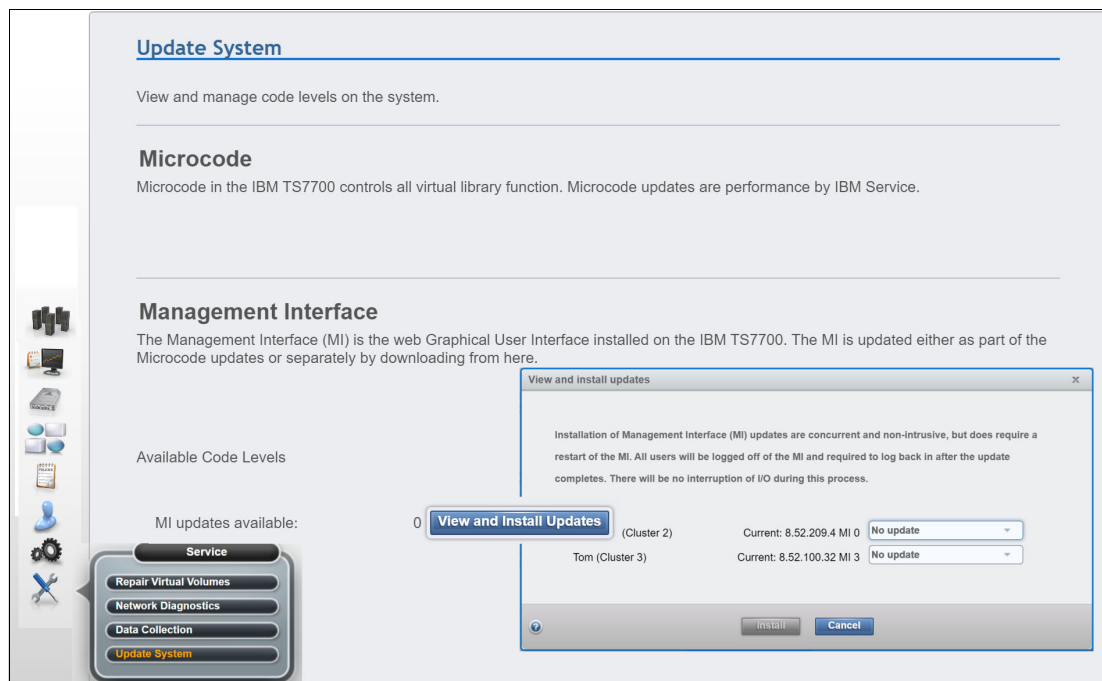


Figure 10-43 The Update System window and options

The Update System page shows the microcode and management interface levels, which are installed at all clusters of the grid. Clicking the **View and Install Updates** button opens the correspondent window. From this window, the admin user can select what MI update to install on what cluster or to decide to update one or more clusters at the same time.

The installation process creates a task for each chosen update that can be monitored individually. The installation process is *not* synchronous. At the end of the update, the TS7700 MI that was updated restarts, and the new level is displayed in the Update System page.

For more information, see the documentation locally at the TS7700 MI by clicking the question mark at the upper right of the window, or see [Update system](#).



IBM TS7700 common operations and procedures

This chapter provides information about IBM TS7700 common operations and procedures.

This chapter includes the following topics:

- ▶ 11.1, “Call Home and Electronic Customer Care” on page 558
- ▶ 11.2, “Common procedures” on page 561
- ▶ 11.3, “Basic operations” on page 608
- ▶ 11.4, “Cluster intervention scenarios” on page 615

11.1 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, 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.

The Total Storage System Console (TSSC) model TS3000 plays a central role in the serviceability of the TS7700 and associated tape libraries or Tape Virtual Cache (TVC) that are present in the configuration. The IBM service representative performs all required tasks to service the TS7700 at the IT center or by accessing the TSSC remotely by way of the Remote Support Facility interface, which is called *Assist On-site (AOS)*. These activities include installing the equipment, upgrading microcode level, installing new hardware features in the subsystem, or repairing a malfunctioning component.

The TS3000 System Console (FC 2725) with Optical Drive (FC 2748) and KVM (FC 5512) and Redundant Power (FC 1904) is required for all configurations of the TS7700-VED.

Moreover, the TSSC also centralizes the Call Home and Electronic Customer Care functions for all attached systems (the collection of subsystems under the domain of the TSSC).

The Call Home function generates a service alert automatically when a problem occurs within one of the following subsystems that are supported by the TSSC:

- ▶ TS7760
- ▶ TS7770
- ▶ 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 or work remotely to diagnosis and resolve the issue. 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.

Note: The Attached System provides only service-related trace log, configuration, and dump files, which contain information that is specific to machine functions. No user data or content is included in the Call Home information.

A high-level view of call home and remote support capabilities is shown in Figure 11-1.

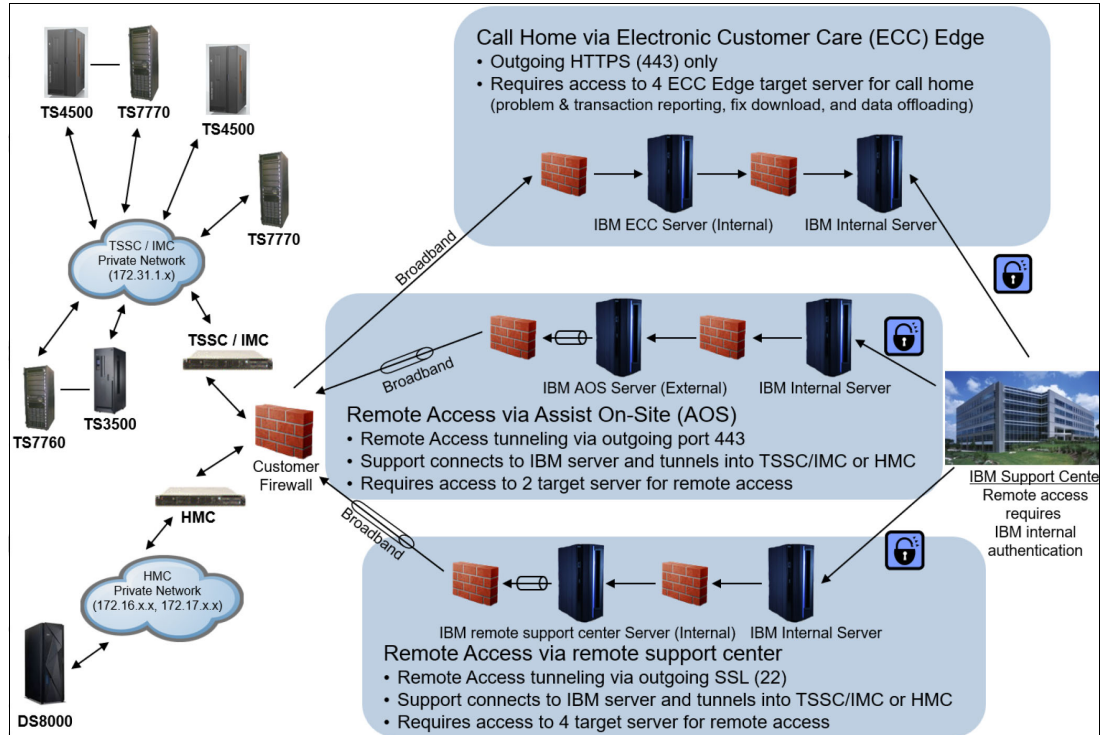


Figure 11-1 Call home and remote support functions

11.1.1 Electronic Customer Care

Electronic Customer Care (ECC) provides a method to connect IBM Storage Systems with IBM remote support. The package supports out-bound communication for broadband Call Home. All information that is sent back to IBM is Secure Sockets Layer (SSL) or Transport Layer Security (TLS) encrypted.

ECC is a family of services that feature 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 feature 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 an HTTP proxy does not exist between the configured TS3000 and the outside network to IBM. Selecting this method requires no further setup. ECC supports customer-provided HTTP proxy.

Also, a customer might require all traffic to go through a proxy server. In this case, the TSSC connects directly to the proxy server, which intermediates 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, in addition to the normal Security provided by IBM.

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 out-bound communication that is associated with ECC Call Home is through an Ethernet connection. A modem connection is not supported by the TS3000 TSSC. The local subnet LAN connection between the TSSC and the attached subsystems is isolated without any outside access. ECC adds another Ethernet connection to the TSSC, which brings the total number to three. The following 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 setup avoids both communications from using the same network connection.

All of these connections are set up by using the Console Configuration Utility User Interface that is on the TSSC. TS7700 events that start a Call Home are displayed in the Events window under the Monitor icon.

11.1.2 Assist On-site

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 the IBM Support team to quickly identify and fix issues with the system.

AOS uses the same network as broadband Call Home, and works on 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 157. The AOS function is disabled by default.

When enabled, the AOS can be configured to run in attended or unattended modes:

- ▶ Attended mode requires that the AOS session is started at the TSSC that is 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 to a specific relay server, so no connection is started from the outside to the TSSC. It is always the TSSC that starts the connection.

In unattended mode, the TSSC periodically checks whether a request was made for a session when it connects to the regional AOS relay servers. When a session request exists, the AOS authenticates and establishes the connection, which allows 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 10.3, “Access icon” on page 492.

For more information about security aspects relative to the Call Home and Remote Access capabilities that are provided by the TSSC that uses the network infrastructure on-premises, see [IBM TS7700, TS4500, Diamondback and DS8000 System Connectivity and Security v.5.04](#).

11.2 Common procedures

This section describes some procedures that are necessary during the implementation stage of the TS7700 (in stand-alone or grid mode). The information 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 the following criteria:

- ▶ Procedures that are related to the tape library that is connected to a TS7700 tape-attached model
- ▶ Procedures that are used with all TS7700 cluster models

11.2.1 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 tasks that are described next. Therefore, ensure that it is set up correctly and working. For access by using 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 TS7700T cluster requires its own logical library in the tape library.

Advanced Library Management System (ALMS) virtualizes the locations of cartridges in the TS3500 and the TS4500 Tape Library.

Ensuring that ALMS is enabled in the tape library

Check the status of ALMS with the TS3500 tape library GUI by clicking **Library** → **ALMS**, as shown in Figure 11-2.

The screenshot displays the IBM System Storage TS3500 Tape Library GUI. The left sidebar shows a navigation tree with 'Library' expanded to 'ALMS'. A red arrow points from 'ALMS' in the sidebar to a pop-up window. The pop-up window contains the following information:

Library Status:

Accessor	OK
3592 Capacity Utilization	64%
Active Errors	None

All Frames | Inventory/Audit

Total storage slots	2037
3592 Licensed Capacity	2037
3592 Unlicensed Capacity	0
Total empty storage slots	336
Offline storage slots	1
Accessors	1
Total I/O slots	16
Empty I/O slots	0
Total 3592 data cartridges	1665
3592	1665
3592 Not Labeled	0
Cleaning cartridges	51
Drives	14
Node cards	9
Total frames	6
Active frames	6
Service bays	0

Figure 11-2 TS3500 tape library GUI Summary and ALMS window

Figure 11-3 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.

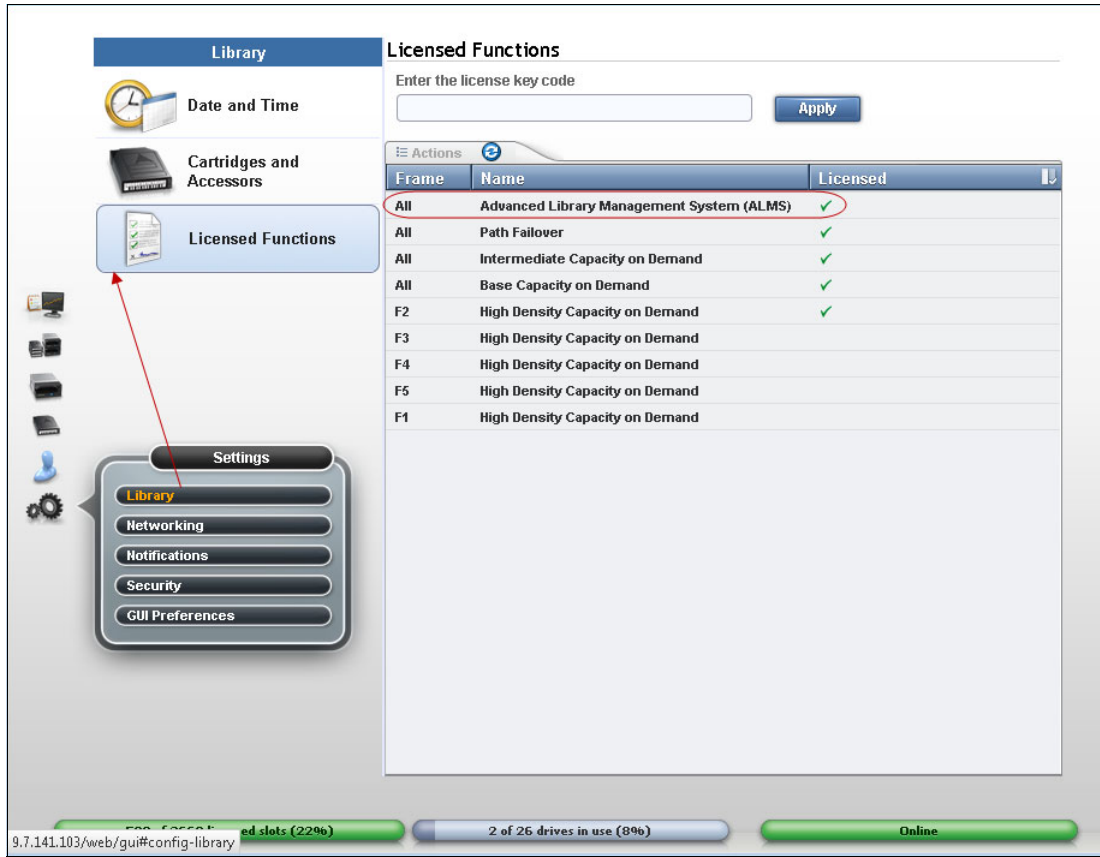


Figure 11-3 ALMS installed and enabled on TS4500

Creating a logical library with TS4500

Complete the following steps for a TS4500 Tape Library:

1. From the initial page of the TS4500 GUI, select **Library Icon** and click **logical library** (see Figure 11-4).

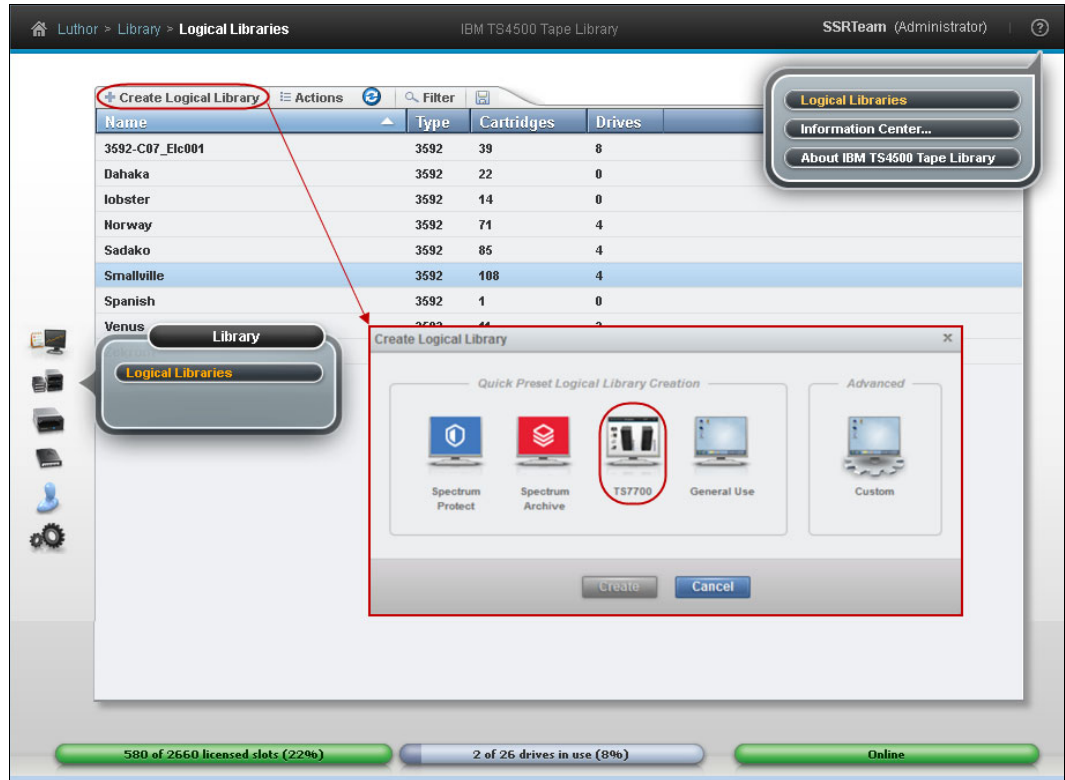


Figure 11-4 TS4500 create logical library page

2. Use the **Create Logical Library** option (see Figure 11-4) to complete the task.

Notice that the TS4500 GUI features selected presets, which help in setting up the new logical library. For the TS7700 library, use the TS7700 option that is highlighted in Figure 11-4. 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, and distributes them in two separate frames, when possible.

Note: The TS7700 preset is disabled when less than four unassigned tape drives are available to create a logical library.

Figure 11-5 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 changes in the environment. Those drives must be correctly cabled to the fiber switches that are dedicated to the TS7700. Also, the new back-end resources must be configured (or reconfigured) within the cluster for proper operation.




Figure 11-5 Display available tapes

The preset also indicates the System Managed encryption method for the new TS7700 logical library.

Note: After the tape-attached cluster has been configured with a Logical Library, this Logical Library should not be changed from the TS4500 MI. The TS7700T physical tape functions can be impacted or broken by the changes. Work with your IBM SSR if changes are necessary.

An example of a logical library definition is shown in Figure 11-6.

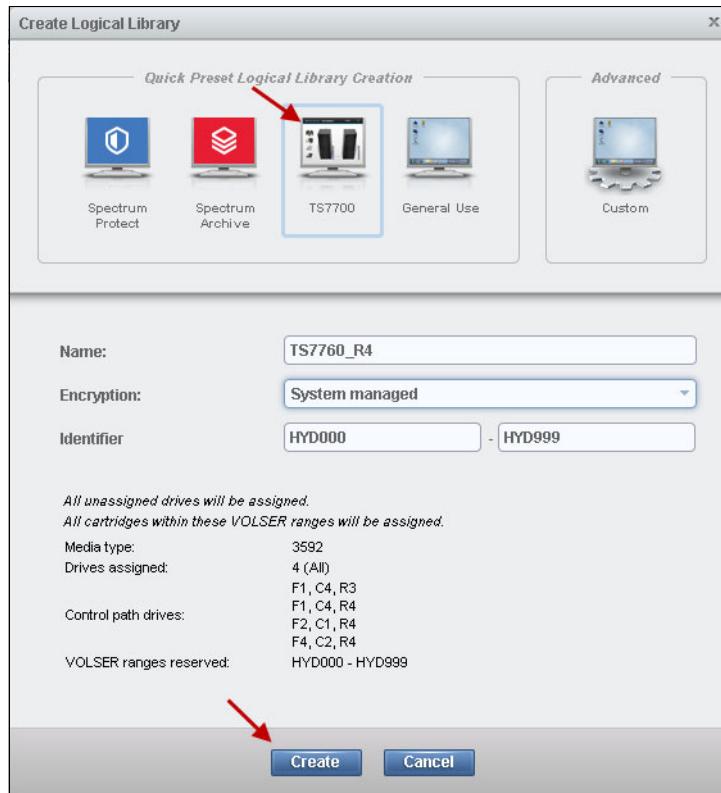


Figure 11-6 Defining the new Logical Library for the TS7700T

After the configuration of the logical library is completed, your IBM service representative can complete the TS7700 tape-attached cluster installation, and the tape cartridges can be inserted in the TS4500 Tape Library.

For more information about TS4500 operations and configuration, see locally available documentation, which is available locally by clicking the question mark icon in the TS4500 GUI, or see [Logical Libraries](#).

Creating a logical library with ALMS on the TS3500 tape library

Complete the following to create a logical library with ALMS on the 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 then click **Go**.

An extra window opens that is named Create Logical Library.

3. Enter the logical library name (up to 15 characters), select the media type (**3592** for TS7700T), and then click **Apply**. The new logical library is created and is displayed in the logical library list when the window is refreshed.

- 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 11-7 shows a summary of the windows in the Create logical library sequence.

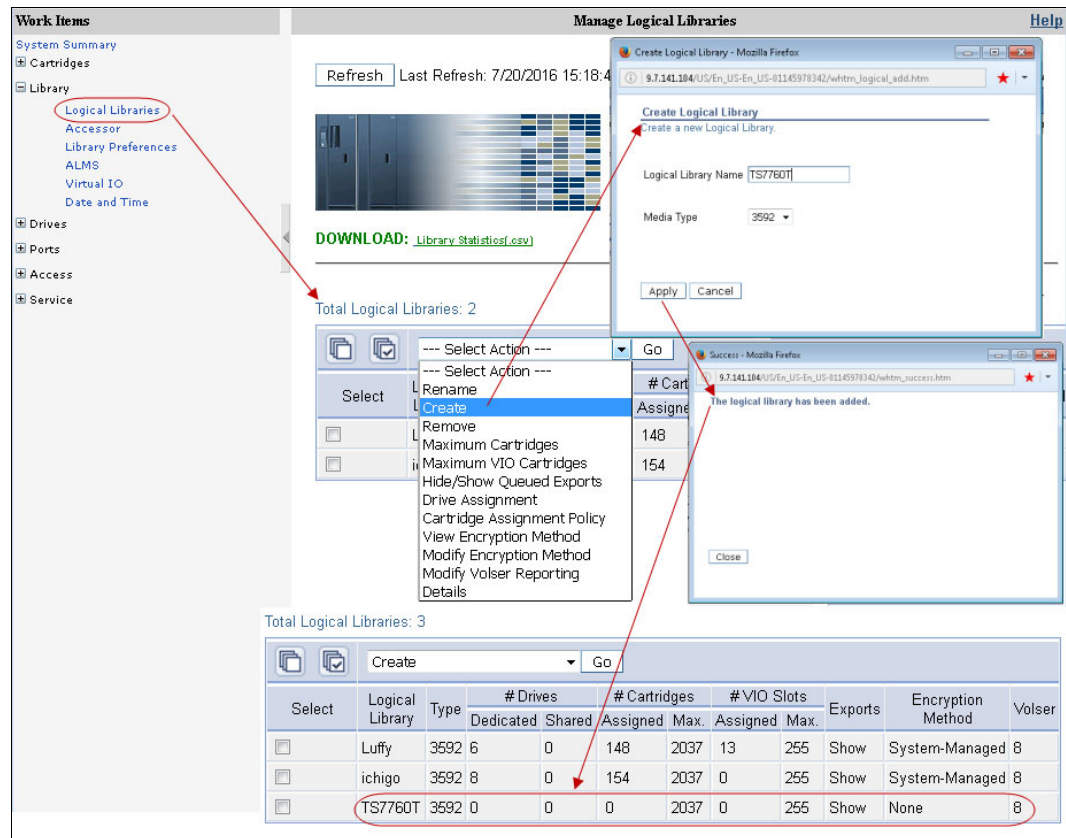


Figure 11-7 Creating a Logical Library with the TS3500 tape library

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 feature enables a logical library to grow without changing the configuration each time you want to add empty slots.

Ensure that the new logical library has the eight-character Volser reporting option set. Another item to consider is the virtual I/O usage; that is, if VIO is enabled and, if so, how many cells should be defined.

For more information, see the documentation regarding the TS3500 Tape Library that is available on virtual I/O slots and applicability at [this web page](#).

Assigning drives

Now, the TS7700T tape drives must be added to the logical library.

From the Logical Libraries window that is shown in Figure 11-8 on page 568, use the work items on the left side of the window to go to the requested web window by clicking **Drives** → **Drive Assignment**. Clicking this link takes you to a filtering window in which you can select to display the drives by drive element or logical library.

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 11-8 shows the drive assignment window of a logical library that has all drives that are assigned.

Unassigned drives appear in the Unassigned column with the box selected. 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 of the window that is shown in Figure 11-8 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.

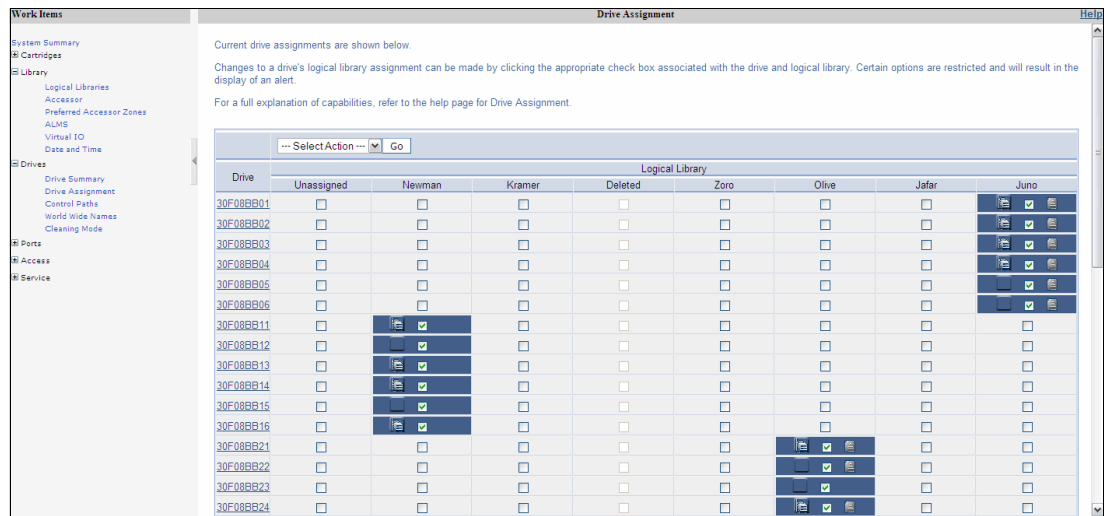


Figure 11-8 Drive Assignment window

TS7700T works with the TS1150 tape drives in a homogeneous or heterogeneous configuration. A 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 older media. Tape drives from previous generation are used only to read older media (JA/JB) while the TS1150 is read/write to the newer media types. Because no writes are made to the older media type, the support for heterogeneous configuration of the tape drives is deemed limited.

For more information about heterogeneous drive support, see Chapter 2, “Architecture, components, and functional characteristics” on page 15, and Chapter 8, “Migration” on page 303.

In a multi-platform environment, logical libraries appear as shown in Figure 11-8. Physical tape drives can be reassigned from one logical library to another. This physical tape drive reassignment 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: Work with your IBM SSR when changing the logical library configuration associated to a TS7700T. The TS7700T needs to be reconfigured to assume the new configuration.

Defining control paths

Each TS7700T requires four defined control path drives. If possible, distribute the control path drives over more than one tape library frame to avoid single points of failure.

Defining the encryption method for the new logical library

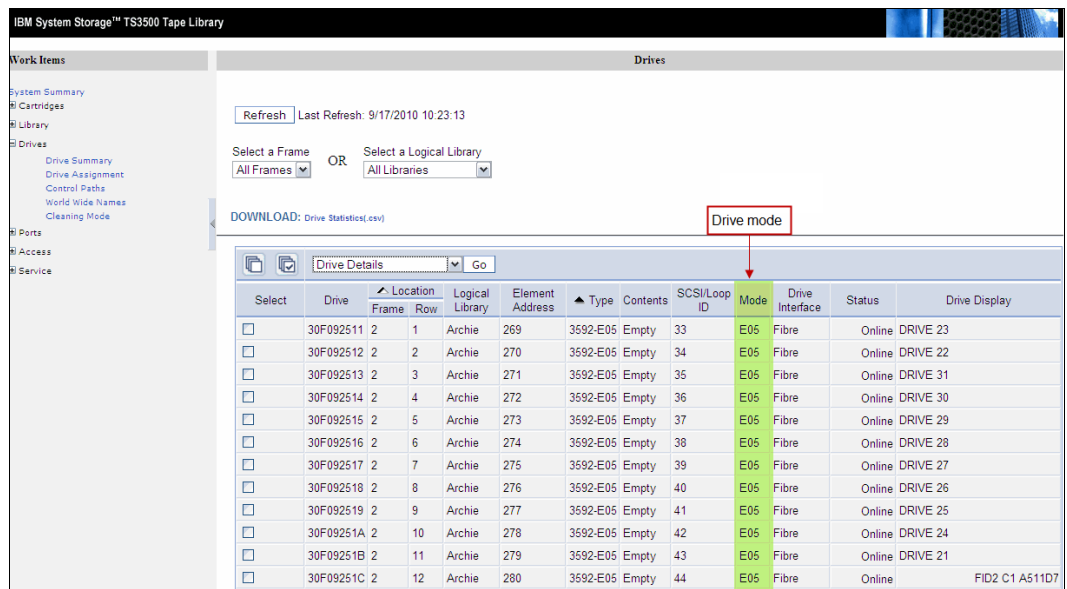
After adding tape drives to the new logical library, the encryption method for the new logical library (if applicable) must be defined.

Reminders: Tape drives must be set to Native mode when encryption is used.

To activate encryption, FC 9900 must be ordered for the TS7700T, 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 11-9. Review the Mode column. This column is displayed only if drives in the tape library are emulation-capable.



IBM System Storage™ TS3500 Tape Library

Work Items

System Summary

Cartridges

Library

Drives

Drive Summary

Drive Assignment

Control Paths

World Wide Names

Cleaning Mode

Ports

Access

Service

Drives

Refresh Last Refresh: 9/17/2010 10:23:13

Select a Frame: All Frames OR Select a Logical Library: All Libraries

DOWNLOAD: Drive Statistics.csv

Drive mode

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 11-9 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 11-10.

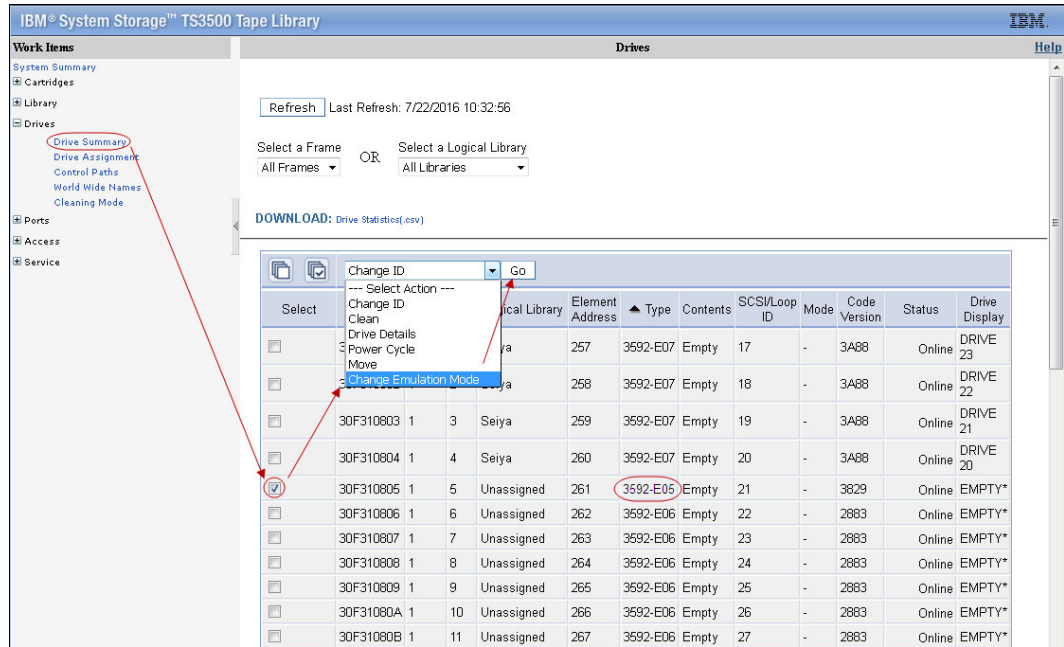


Figure 11-10 Change the drive emulation

- 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 11-11).

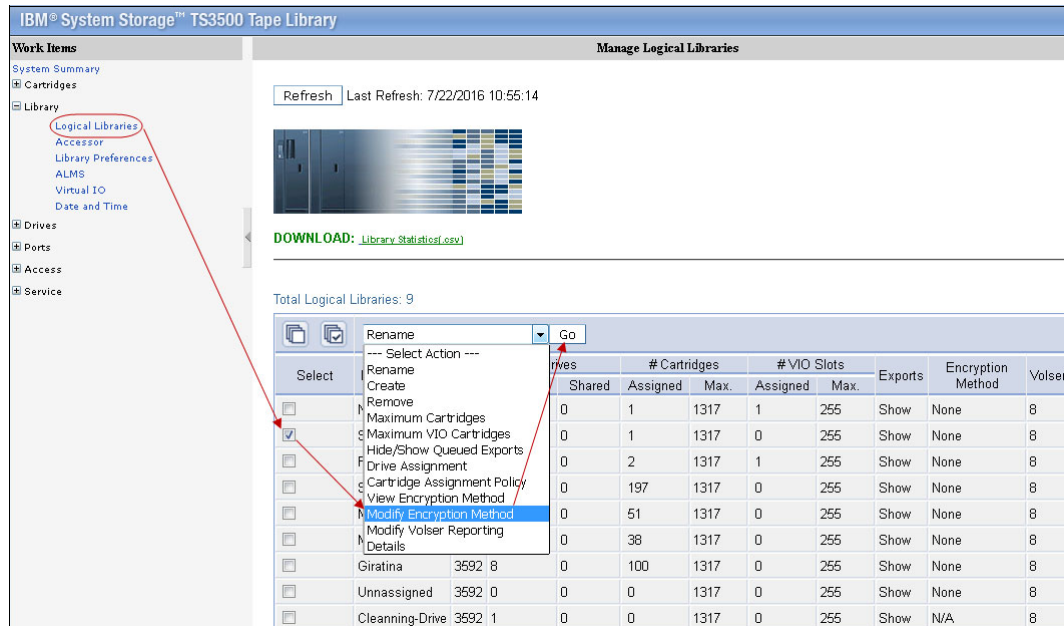


Figure 11-11 Select the encryption method

- In the window that opens, select **System-Managed** for the chosen method; then, select all drives for this partition (see Figure 11-12).

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 11-12 Set the encryption method

Keep the Advanced Encryption Settings as NO ADVANCED SETTING as a rule.

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 11-13 shows the VOLSER ranges that are defined for logical libraries.

The screenshot shows a web-based interface for 'Cartridge Assignment Policy'. On the left is a navigation tree with categories like 'Work Items', 'System Summary', 'Cartridges', 'Library', 'Drives', 'Ports', 'Access', and 'Service'. The main area displays a 'Refresh' button, a 'Last Refresh' timestamp, and a 'View' dropdown set to 'All Libraries'. Below this is a table with a 'Select Action' dropdown and a 'Go' button. The table lists logical libraries and their corresponding volume serial number ranges.

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"/>	Zoro	310650 - 310999
<input type="radio"/>	Zoro	J1M600 - J1M699
<input type="radio"/>	Zoro	J1M750 - J1M799

Figure 11-13 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 enables 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. Also, follow the same rules across host platforms whether a TS3500 tape library is shared between IBM Z and Open Systems hosts.

Tip: Creating or changing a CAP definition does not change the existing assignment of a tape cartridge. If needed, you must first unassign the cartridges, and then manually assign it to the correct logical library.

Inserting TS7700T physical volumes

The tape-attached TS7700 subsystem manages 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 TS7700T 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 helps ensure 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-attached TS7700 logical library partitions according to the definitions.

Important: Before inserting physical volumes that belong 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 578.

These procedures ensure that TS7700 back-end cartridges are never assigned to a host by accident. Figure 11-14 shows the flow of physical cartridge insertion and assignment to logical libraries for TS7700T.

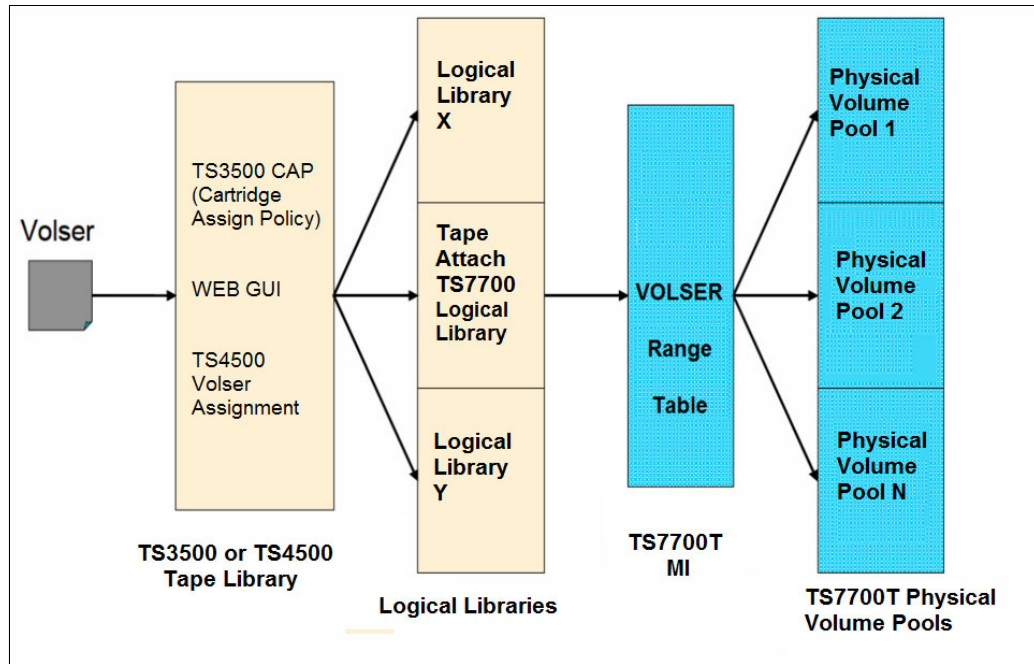


Figure 11-14 Physical volume assignment

Inserting physical volumes into the tape library

The following 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 are not inventoried.

After the new media is inserted, 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-attached TS7700 cluster.

Tip: When populating the tape library, place cartridges only in a frame whose front door is open. Do *not* add or remove cartridges to/from an adjacent frame.

Insertion by using the I/O station

The tape library can be operating with or without virtual I/O (VIO) enabled.

With VIO enabled, the tape library moves the cartridges from the physical I/O station into the physical library alone. In the first moment, the cartridge leaves the physical I/O station and moves 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 includes 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 moves 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 VIOS 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 that are inserted should belong to the range of volumes that are defined in the tape library (CAP or volser range) for the TS7700 logical library. Those ranges should also be defined in the TS7700 Physical Volume Range, as described in "Defining VOLSER ranges for physical volumes" on page 578. Both conditions must be met for a physical cartridge to be successfully inserted to the TS7700T.

If any VOLSER is not in the range that is defined by the policies, the cartridges must be assigned to the correct logical library manually by the operator.

Note: Ensure that CAP ranges are correctly defined. Insert Notification is not supported on a high-density library. If a cartridge that is 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. Figure 11-15 shows the MI page, with TS4500 and TS3500.

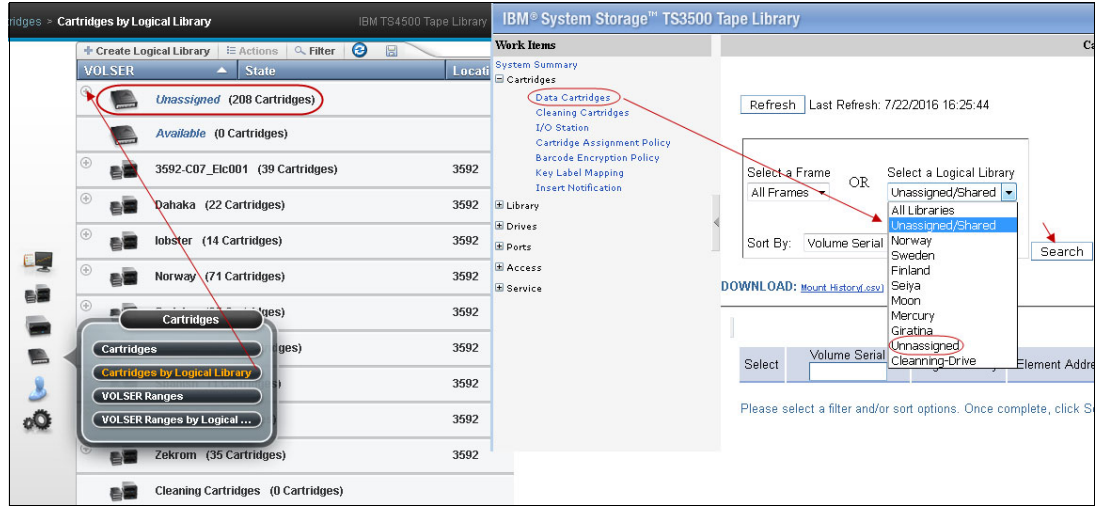


Figure 11-15 Check the volume assignment

When volumes that belong to a logical library are found unassigned, correct the CAP or volser assignment definitions and reinsert them again. Optionally, cartridges can be manually assigned to the correct logical library by the operator by using the GUI.

It is advised that correctly defined CAP or volser assignment policies in the tape library be in place 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 TS7700T if it does not fit in any previously 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 11-16.

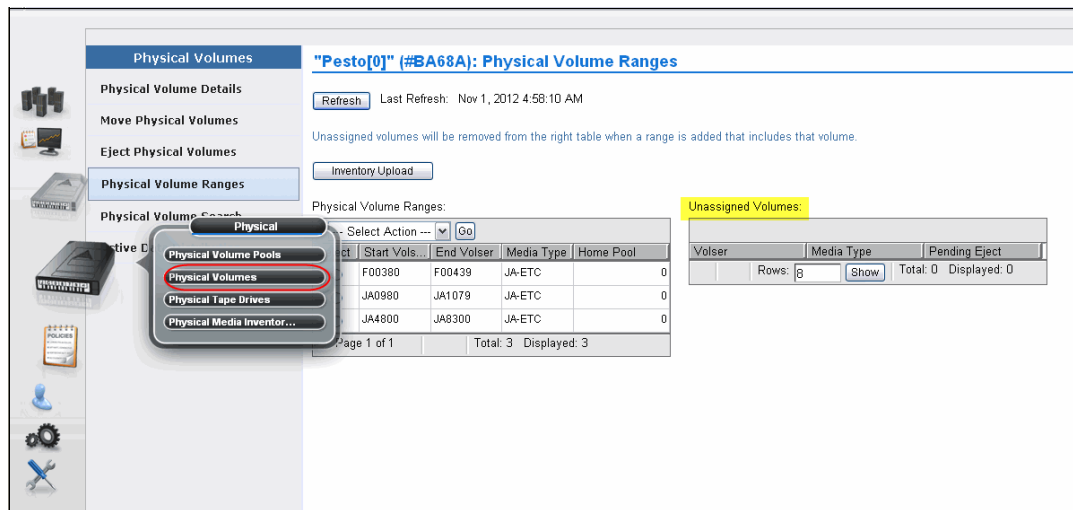


Figure 11-16 TS7700T unassigned physical volumes

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 578. 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, ensure that CAP or volser assignments are correct in the tape library.

Assigning cartridges in the tape library to the logical library partition

This procedure must be done 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 was 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 11-16 on page 575.

Cartridge assignment to a logical library is available only by using the tape library GUI.

Assigning a data cartridge

To assign a data cartridge to a logical library in the TS3500 tape library, complete the following steps:

1. Click the **Cartridge** icon on the TS3500 GUI, and select **Cartridges**.
2. Find the cartridge that you want to assign (it shows as unassigned at that point), and select it by clicking the line.
3. From the list of the menu options select **Assign**. Choose the correct logical library in the list available.
4. Complete the assignment insertion by clicking the **Assign** button.
5. For a TS7700T cluster, click **Physical** → **Physical Volumes** → **Physical Volume Ranges** and then click **Inventory Upload**, as shown in Figure 11-19 on page 579.

To assign a data cartridge to a logical library in the TS3500 tape library, complete the following 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.
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 11-19 on page 579.

Inserting a cleaning cartridge

Each drive in the tape library requires cleaning. Tape drives that are used by the TS7700 subsystem can request a cleaning action when necessary. This cleaning is completed by the tape library automatically. However, the necessary cleaning cartridges must be provided.

Remember: A 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. Figure 11-17 shows how to set the TS4500 to move expired cleaning cartridges to the I/O station automatically.

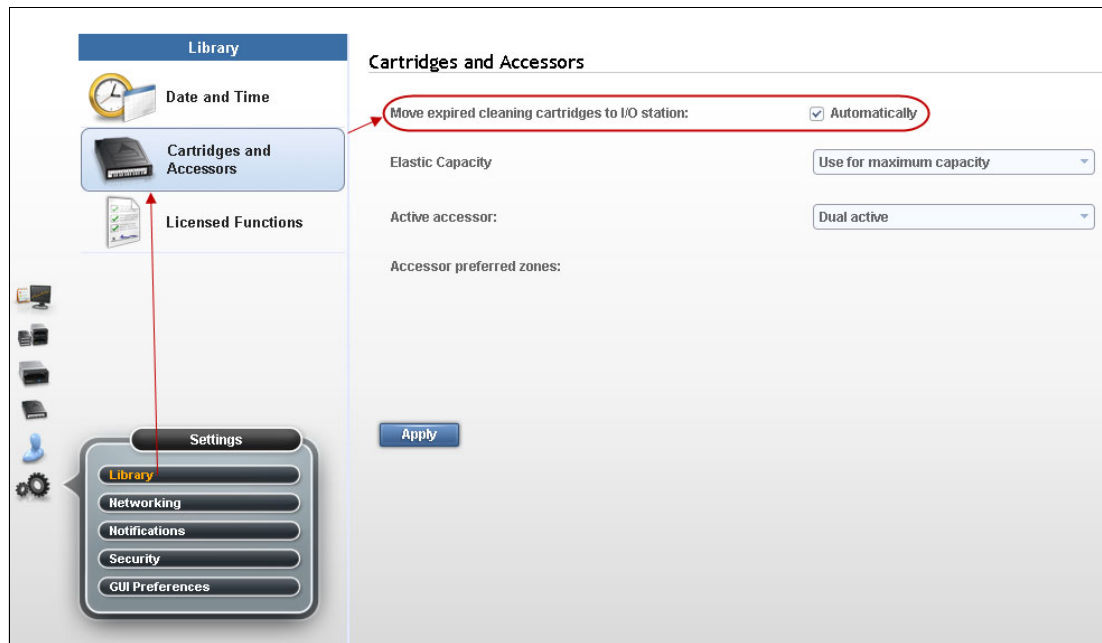


Figure 11-17 TS4500 Tape Library moves expired cleaning cartridge to I/O station automatically

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 11-18.

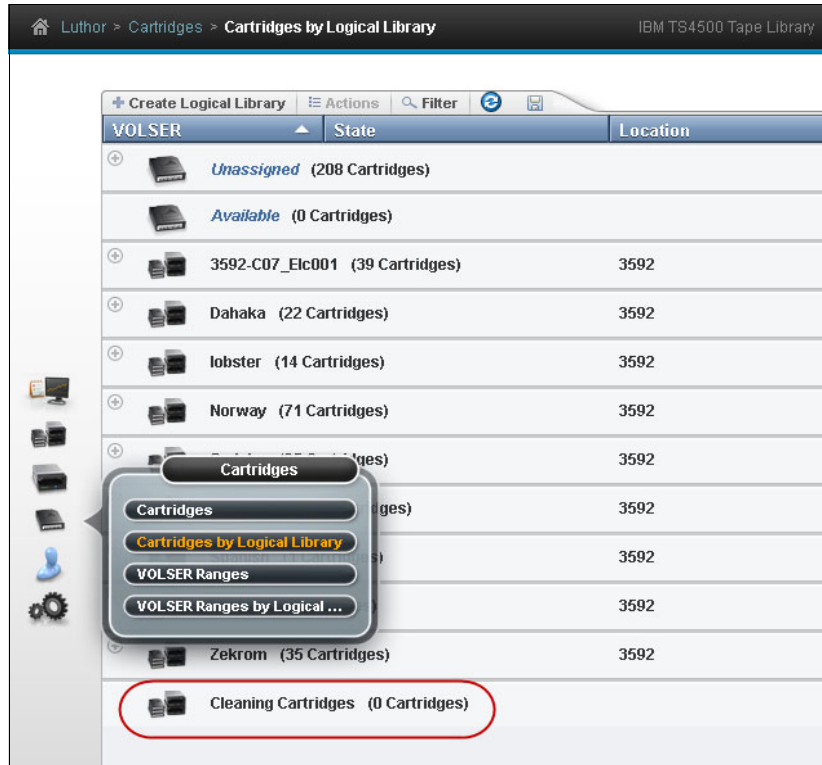


Figure 11-18 Displaying cleaning cartridges with the TS4500.

Also, TS4500 Tape Library command-line interface commands are available that can be used to check the status of the cleaning cartridges or alter settings in the tape library. For more information, see the documentation for TS4500 that is available locally by clicking the question mark icon at the top bar in the GUI, or see [TS4500 tape library CLI commands](#).

11.2.2 TS7700 definitions

This section provides information about some of the definitions and settings that are used in the TS7700 clusters. Not all definitions or settings applies to all clusters and models, and the examples that are described next are not in a specific order.

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 correct category. Define the correct policies in the VOLSER Ranges table *before* inserting the cartridges into the tape library.

Note: VOLSER Ranges (or CAP) should be correctly assigned at the tape library before the tape library is used with IBM Z hosts. Native physical volume ranges must fall within ranges that are assigned to IBM Z host logical libraries.

Use the window that is shown in Figure 11-19 to add, modify, and delete physical volume ranges. Unassigned physical volumes are listed in this window. If a volume is listed as an unassigned volume, and this volume belongs to this TS7700, a new range should be added and include 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.

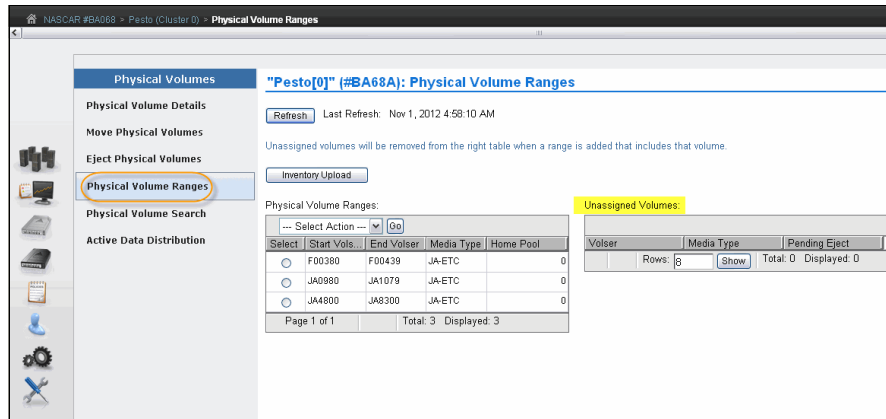


Figure 11-19 Physical Volume Ranges window

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 VOLSER range, or to modify or delete a predefined VOLSER range.

Important: Operator intervention is required to resolve unassigned volumes.

For more information about how to insert a new range of physical volumes by using the TS7700 Management Interface, see “Physical Volume Ranges window” on page 465.

Defining physical volume pools in the TS7700T

Pooling physical volume allows you to enable data to be placed into separate sets of physical media, treating each media group in a specific way. For example, production data might need to be kept separated from test data, or part of the data encrypted. These goals 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 to only tape attach clusters. Attempting to access those windows from a non-tape attach cluster results in the following HYDME0995E message:

This cluster is not attached to a physical tape library.

Use the window that is shown in Figure 11-20 to view or modify settings for physical volume pools, which manage the physical volumes that are used by the TS7700.

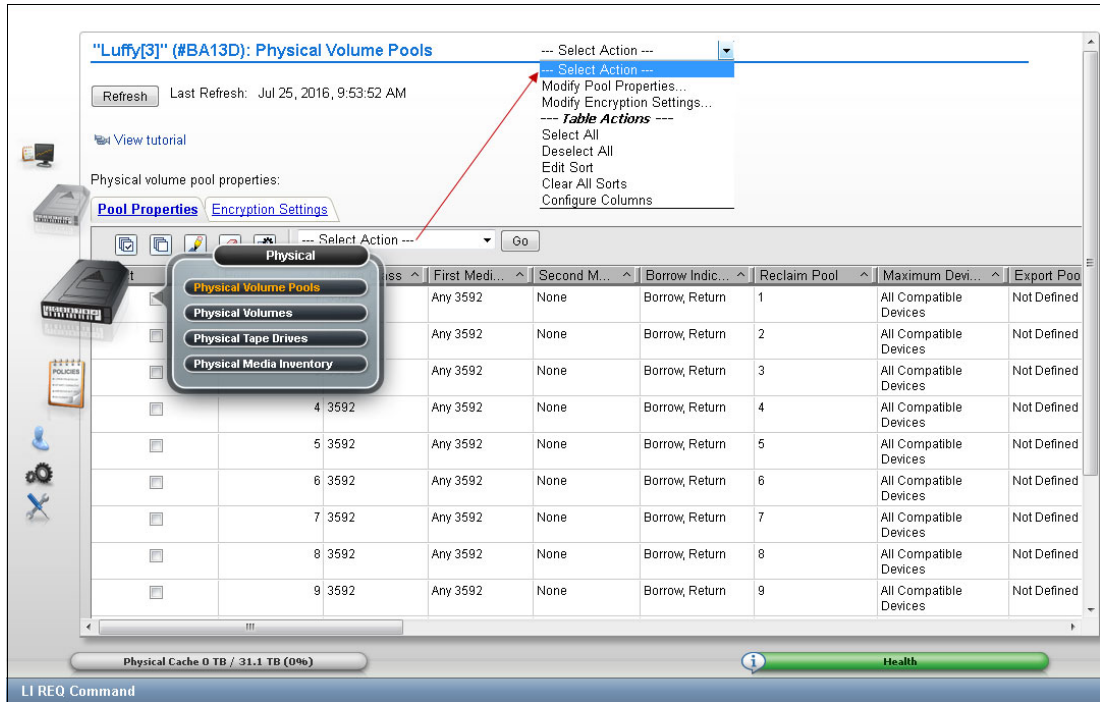


Figure 11-20 Physical Volume Pools

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.

For more information about how to view, create, or modify Physical volume tape pools by using the TS7700 management interface, see 10.1.1, “Physical Volume Pools” on page 450.

To modify encryption settings for one or more physical volume pools, complete the following steps (Figure 11-21 on page 581 shows this sequence). For more information, see 10.1.1, “Physical Volume Pools” on page 450:

1. Open the Physical Volume Pools window.

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 checkbox next to each pool to be modified.
4. Click **Select Action** → **Modify Encryption Settings**.

5. Click **Go** to open the Modify Encryption Settings window (see Figure 11-21).

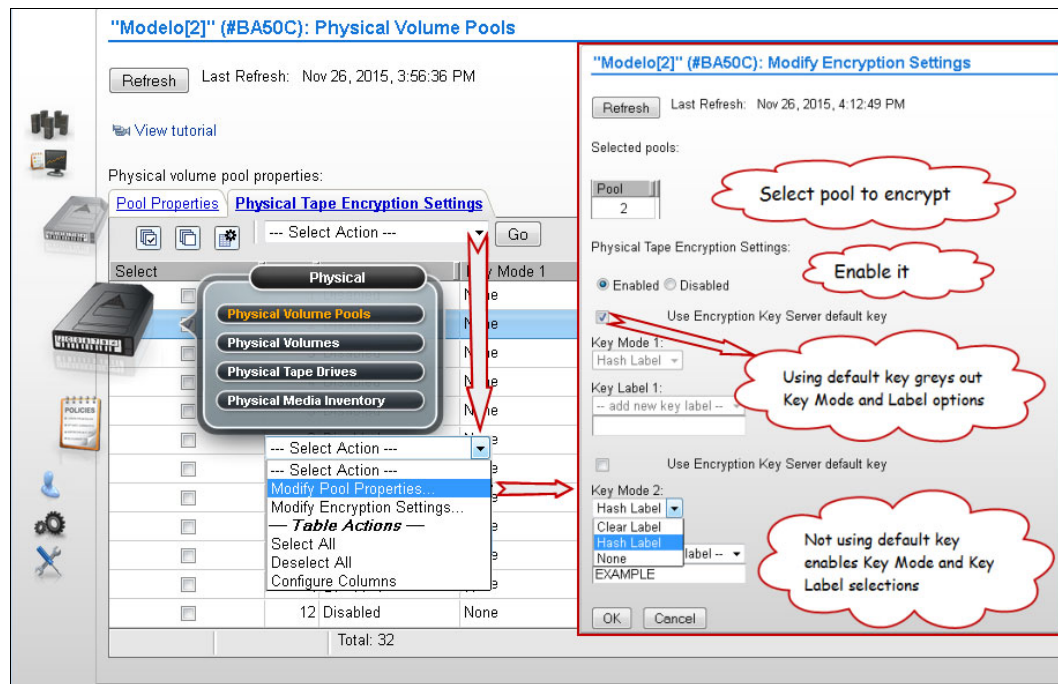


Figure 11-21 Modify encryption settings parameters

For more information about parameters and settings that are in the Modify Encryption settings window, see 10.1.1, “Physical Volume Pools” on page 450.

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 by using 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 physically 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 listed in Table 11-1.

Table 11-1 Installed drives versus maximum reclaim tasks

Number of available drives	Maximum number of reclaims
3	1
4	1
5	1
6	2
7	2
8	3
9	3
10	4
11	4
12	5
13	5
14	6
15	6
16	7

You might want to have fewer reclaims running, which spares the resources for other activities in the cluster. The user can set a maximum number of drives that are used for reclaim in a pool base. Also, reclaim settings can be changed by using **LI REQ** commands.

The reclamation level for the physical volumes must be set by using the Physical Volume Pools window in the TS7700 MI. Select a pool and click **Modify Pool Properties** in the menu to set the reclamation level and other policies for that pool.

For example, Figure 11-22 shows the borrow-return policy in effect for Pool 3; that is, cartridges can be borrowed from the common scratch pool. In addition, those cartridges are returned to the CSP upon reclamation. Also, the user has defined that volumes that belong to pool 3 should be reclaimed into pool 13.

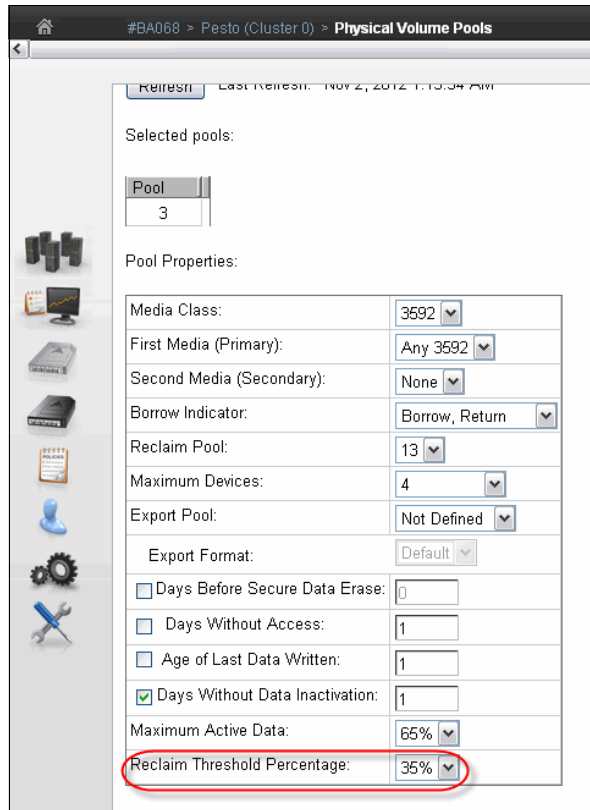


Figure 11-22 Pool properties

No more than four drives can be used for premigration in pool 3. The reclaiming threshold percentage was 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 a 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.

For more information about parameters and settings, see 10.1.1, “Physical Volume Pools” on page 450.

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 11-23 on page 586 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 falls below a minimum number of scratch cartridges that are available. Now, reclamation is enforced by the tape-attached TS7700 cluster.

Tip: The maximum number of Inhibit Reclaim Schedules is 14.

By using the Bulk Volume Information Retrieval (BVIR) process, the amount of active data on stacked volumes can be monitored on PHYSICAL MEDIA POOLS, which helps to plan for a reasonable and effective reclaim threshold percentage. Also, the 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 of the time. Other conditions must be met, such as stacked volumes that meet one of the reclaim policies and drives that are 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 drives 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 issue occurs, 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.

Individual pools can have separate reclaim policies set. The number of pools can also influence the reclamation process because the TS7700 tape attach 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 taking the following into account:

- ▶ A pool enters a Low scratch state when it can access less than 50 and two or more empty cartridges (scratch tape volumes).
- ▶ A pool enters a Panic scratch state when it has access to fewer than two empty cartridges (scratch tape volumes).

Panic Reclamation mode is entered when a pool features fewer than two scratch cartridges and no more scratch cartridges can be borrowed from any other pool that is defined for borrowing. For more information about borrowing, see “Using physical volume pools” on page 53.

Important: A physical volume pool that is running out of scratch cartridges might stop mounts in the TS7700T tape attach partitions, which affect host tape operations. Mistakes in pool configuration (media type, borrow and return, home pool, and others) 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, the reclaim threshold level should be 10%–35%. For more information about how to fine-tune this function and about the available host functions, see in 4.3.15, “Physical volumes for TS7700T” on page 194. Pools in either scratch state (Low or Panic state) get priority for reclamation.

The thresholds are listed in Table 11-2.

Table 11-2 Reclamation priority table

Priority	Condition	Reclaim schedule acknowledged	Active data threshold% acknowledged	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.	
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 inhibited 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, select 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 occurs for 10 minutes. The Inhibit Reclaim schedule is not accepted by the Secure Data Erase function for a volume that includes 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 11-23.

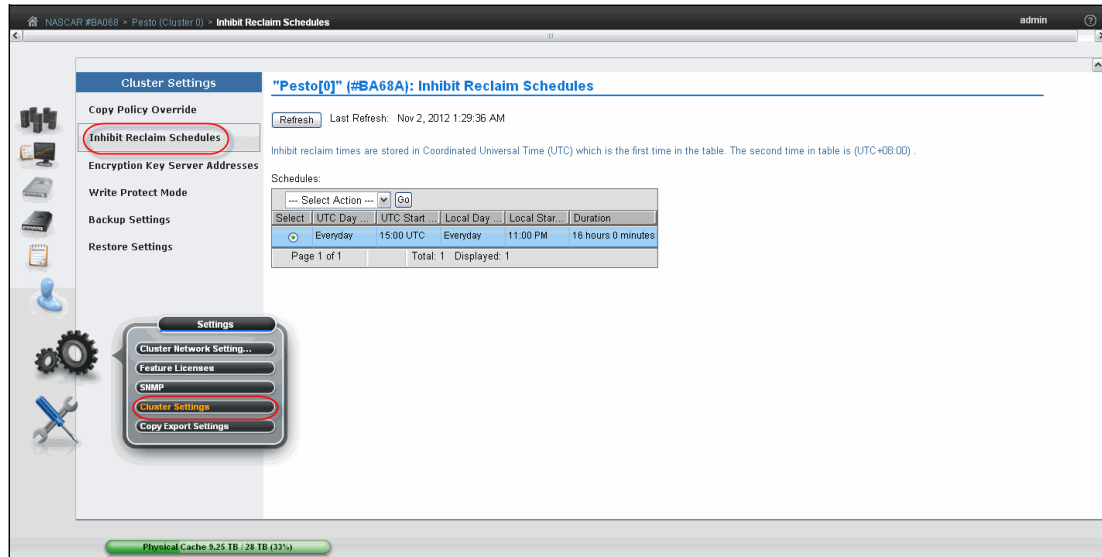


Figure 11-23 Inhibit Reclaim schedules

The Schedules table (see Figure 11-24) 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 Reclaim Inhibit Schedule, or modify or delete an existing schedule, as shown in Figure 11-24.

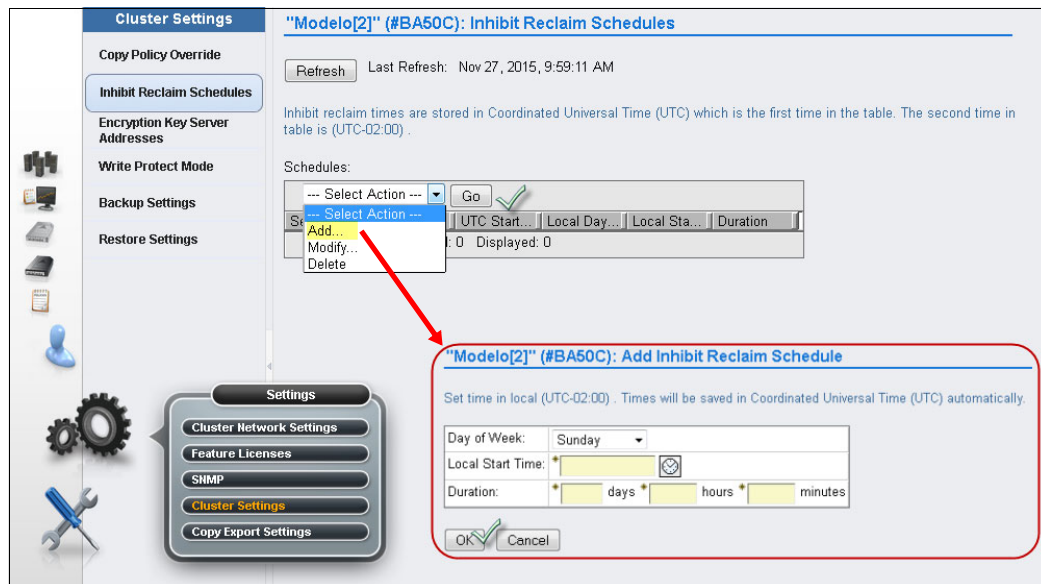


Figure 11-24 Add Inhibit Reclaim schedule

For more information, see “Inhibit Reclaim Schedules window” on page 524.

Defining Encryption Key Server addresses

Note: With Version 4.1.0, IBM Security Key Lifecycle Manager is renamed as IBM Security Guardium Key Lifecycle Manager (GKLM).

IBM Security Guardium Key Lifecycle Manager is supported for the TS7770 Data at Rest Encryption (FC 5276). IBM GKLM also supports the TS1150 and TS1140 tape drive encryption. For more information about EKM requirements and TS7700 feature codes, see Chapter 4, “Preinstallation planning and sizing” on page 145. Also, refer to Appendix J, “Configuring externally managed encryption” on page 1013.

Note: 3956-CSB external encryption requires TLS V1.2 to be configured on the Key Manager. Certificates must be exchanged and trusted among the Key Manager server and TS7700 before changing CSB encryption from local to external by way of TS7700MI.

Figure 11-25 shows the setup of the Data at rest Encryption when the TS7700-VED disk cache 3956-CSB is used with external encryption enabled. The 3956-CSB disk cache external encryption communicates with the key manager by using the Key Management interoperability Protocol (KMIP) for the encryption management exchanges with the key server. The communications between 3956-CSB (by way of the TS7700 as a proxy server) and the Encryption Key Server are protected by the TLS 1.2 encryption.

The screenshot displays the 'Data at Rest Encryption' configuration page for a cluster. The left sidebar contains navigation options: Cluster Settings, Cluster Network Settings, Feature Licenses, SNMP, Data at Rest Encryption (selected), Write Protect Mode, Backup Settings, Restore Settings, and RSyslog. A 'Settings' pop-up menu is visible over the sidebar, highlighting 'Cluster Settings' and 'Notification Settings'. The main content area is titled '[1] (#BA29B): Data at Rest Encryption' and shows the following configuration:

- Disk encryption:** AES256 Enabled external
- Tape encryption:** No license installed
- Key server type:** IBM SKLM (IPP and KMIP)
- Primary key server:**
 - Address: 9.7.124.236
 - KMIP port: 5696
 - IPP TLS 1.2:
 - IPP port: 441
 - Key server certificate: Import certificate
 - TS7700 certificate: lwiks (HTTPS)
 - Expires: June 1, 2033, 12:46:05 PM Brasilia Standard Time
- Secondary key server:**
 - Address: 9.7.124.204
 - KMIP port: 5696
 - IPP TLS 1.2:
 - IPP port: 441
 - Key server certificate: Import certificate
 - TS7700 certificate: lwiks (HTTPS)
 - Expires: June 1, 2033, 12:46:05 PM Brasilia Standard Time

A 'Submit Changes' button is located at the bottom of the page.

Figure 11-25 Data at Rest Encryption setup page for a 3956-CSB disk cache.

The user can use the default certificate with the TS7700 cluster, or add another secure socket layer (SSL) certificate or a certificate authority (CA) certificate for the cluster or grid.

Figure 11-26 shows how to add a certificate by using the Certificates page.

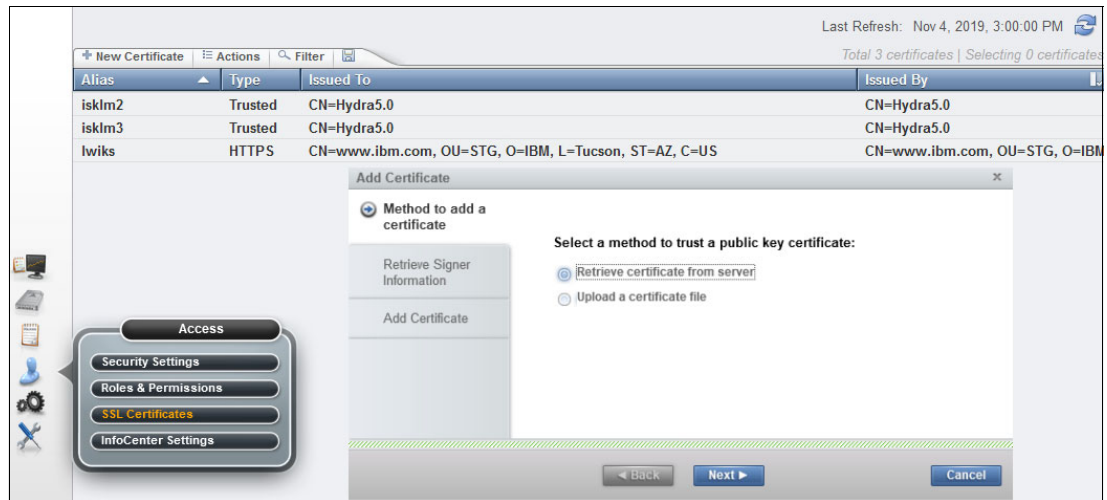


Figure 11-26 Installing a new customer provided certificate.

Figure 11-27 shows the setting of an Encryption Key Server as used for a 3956-CSA disk cache encryption and similar setup that is used with Tape Encryption. CSA disk cache and tape drive encryption use IBM Proprietary Protocol (IPP) to communicate with the IBM Security Guardium Key Lifecycle Manager (GKLM) for encryption management.



Figure 11-27 Encryption Key Server Addresses with 3956-CSA disk cache and tape-encryption.

This same setup (IPP GKLM) is also used with Tape Encryption. The EKS assists encryption-enabled tape drives in generating, protecting, storing, and maintaining EKs that are used to encrypt information that is written to and decrypt information that is read from tape media (tape and cartridge formats).

With the external key management disk encryption feature installed, the management of the key is removed from the TS7700 and disk subsystem controllers. With local encryption, the CSB disk cache uses four USB drives to store encryption keys; two drives are permanently attached to the controllers and the other two are used as backup for the keys, which should be kept at a safe location. CSA disk cache with local encryption keeps the encryption keys at NVSRAM into each controller.

Note: The settings for Encryption Server are shared for tape and external disk encryption.

For more information, see IBM Documentation, which is available locally at the TS7700 MI or see [Tape encryption](#).

Also, see the following white papers:

- ▶ [White Paper - TS7700 Encryption Support V2.0](#)
- ▶ [IBM Virtualization Engine TS7770 Series Disk Encryption Overview for 3956-CSB Version 1.0](#)
- ▶ [TS7700 Full Disk Encryption \(FDE\)](#)

Inserting virtual volumes

Use the Insert Virtual Volumes window (see Figure 11-28) 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.

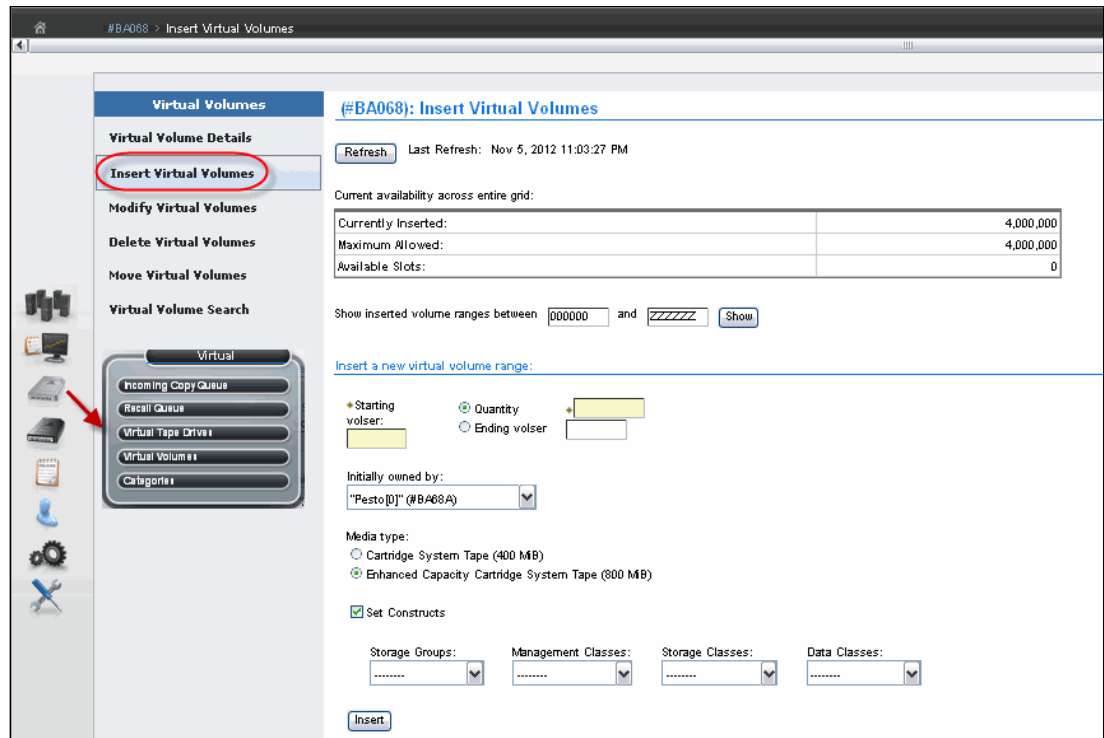


Figure 11-28 TS7700 MI Insert Virtual Volumes window

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 no online devices exist to a specific host, that host does not receive the attention interrupt from the library unless a device was varied online.

To work around this limitation, ensure that at least one device is online (or was online) to each host or use the **LIBRARY RESET, CBRUXENT** command to start 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, when OAM starts for the first time, and being the volumes in the Insert category, the entry processing starts, which do not allow 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.

For more information about this TS7700 MI page, see “Insert Virtual Volumes window” on page 430.

Note: Up to 100,000 logical volumes can be inserted at one time. This limit applies to 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. The TS7700 MI provides a way to define one or more scratch categories. 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 (for z/OS) and 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.

For more information about adding, modifying, or deleting a scratch category of virtual volumes, see “Insert a New Virtual Volume Range window” on page 431. 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.

– Private

Custom categories that are established by a user, within the range of 0x0001–0xEFFF.

– 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 because of some rare, internal categories not being displayed on the Categories table. Also, moving virtual volumes between scratch and private categories can be done multiple times per second. Any snapshot of volumes on all clusters in a grid is obsolete by the time that a total count completes.

The Categories table can be used to add, modify, and delete a scratch category. The table can also be used to change the way that information is displayed.

Figure 9-53 on page 440 shows the Category window and options. To add a scratch category, follow the instructions at Table 9-12 on page 442.

Tip: Add a comment to DEVSUPnn to ensure that the scratch categories are updated when the category values in DEVSUPnn are changed. The scratch categories and category values must always be in sync.

Defining the logical volume expiration time

The expiration time is defined from the MI window that is shown in Figure 11-29 on page 595. If the Delete Expired Volume Data setting is not used, logical volumes that were returned to scratch are still considered active data, which allocates physical space in tape cartridges on the tape attach TS7700. In that case, rewriting only this logical volume expires the old data, enabling the 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 were returned to scratch are expired after a specified 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 2,147,483,647 hours (approximately 244,983 years) can be specified.

Specifying a value of zero means that the data that is associated with the volume is to be managed as it was before the addition of this option, which means it is never deleted. In essence, specifying a value (other than zero) provides a “grace period” from when the virtual 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.

Remember: Consider the following points:

- ▶ 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 also applies to disk only clusters. If it is not used, logical volumes that were returned to scratch are still considered active data, which allocates 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.

- ▶ The value 0 is *not* a valid entry in the dialog box for Expire Time on the Add Category page. Use No Expiration instead.

Establishing the Expire Time for a volume occurs as a result of specific events or actions. The Expire Time of a volume can be affected by the following events or actions:

- ▶ A volume is mounted
The data that is associated with a logical volume is not deleted (even if it is eligible) if the volume is mounted. Its Expire Time is set to zero, which means it will not be deleted. It is reevaluated for deletion when its category is assigned.
- ▶ A volume’s category is changed
Whenever a volume is assigned to a category (including assignment to the same category in which it currently exists) it is reevaluated for deletion.
- ▶ Expiration
If the category has a nonzero Expire Time, the volume’s data is eligible for deletion after the specified time period, even if its previous category had a different nonzero Expire Time.

- ▶ No action

If the volume's previous category had a nonzero Expire Time or even if the volume was eligible for deletion (but was yet selected to be deleted) and the category to which it is assigned includes an Expire Time of zero, the volume's data is no longer eligible for deletion. Its Expire Time is set to zero.

- ▶ A category's Expire Time is changed

If a user changes the Expire Time value by using the scratch categories menu on the TS7700 MI, the volumes that are assigned to that category are reevaluated for deletion.

- ▶ Expire Time is changed from nonzero to zero

If the Expire Time is changed from a nonzero value to zero, volumes that are assigned to the category that include a nonzero Expire Time are reset to an Expire Time of zero. If a volume was eligible for deletion, but was not yet selected for deletion, the volume's data is no longer eligible for deletion.

- ▶ Expire Time is changed from zero to nonzero

Volumes that are assigned to the category continue to have an Expire Time of zero. Volumes that are assigned to the category later have the specified nonzero Expire Time.

- ▶ Expire Time is changed from nonzero to nonzero

Volumes that are assigned for that category are reevaluated for deletion. Volumes that are assigned to the category later have the updated nonzero Expire Time.

After a volume's Expire Time is reached, it is eligible for deletion. Not all data that is eligible for deletion is deleted in the hour that it is first eligible. Once an hour, the TS7700 selects up to 1,000 eligible volumes for data deletion. The volumes are selected based on the time that they became eligible, with the oldest ones being selected first. Up to 1,000 eligible volumes for the TS7700 in the library are selected first.

Defining TS7700 constructs

To use the Outboard Policy Management functions, the following 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, at least one SG is present. This SG 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.

Storage Groups also determines which cloud storage pool is used in a TS7700C cluster. The storage group that corresponds to a group of logical volumes defines which disk cache partition is used and which is the target cloud storage store pool for those volumes.

Figure 10-7 on page 477 shows the TS7700 MI Storage Groups page with Cloud Tier, which is used to add, modify, or delete an SG that is used to define a primary pool for logical volume premigration. The Storage Groups page for Cloud and Tape-attached clusters are shown there. For more information, see 10.2.1, “Storage Groups window” on page 476.

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 2.4.31, “Copy Export” on page 97.

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. In a Grid configuration, Management Classes defines the clusters that are candidates for Scratch Mount for volumes that are associated to the management class.

Also, MC defines the Copy Mode for those volumes in each cluster of that grid: If the volumes on that management class include deferred copies, or copies will be performed at rewind-unload time, or volumes be synchronously copied to other clusters or if the copies should occur after a time delay (and how this time delay is determined), or yet no copies should occur for determined cluster in the grid for logical volumes in that management class.

Figure 10-8 on page 478 shows the TS7700 MI Management Classes window on a grid. Use this page to work with the Management Classes. Refer to the 10.2.2, “Management Classes window” on page 478 for more information.

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.

Figure 10-7 on page 477 shows the TS7700 MI Storage Classes window, which is used 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.

The SC table displays defined SCs that are 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; therefore, the TVC preference for the disk-only clusters is always Preference Level 1.

For more information, see 10.2.3, “Storage Classes window” on page 483.

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 maximum size of a virtual volume is controlled through DC. If “insert media” is selected, it fills up to 400 MiB or 800 MiB, depending on the option that is chosen upon insertion of the virtual volume. Other maximum size options eligible through DC are 1000, 2000, 4000, 6000, 25000, and 65000 MiB.

Also, LWORM policy assignments are controlled from Data Classes.

The user can select the compression method for the logical volumes per Data Class policy. The compression method of the choice is applied to the z/OS host data that is written to logical volumes that belong to a specific Data Class.

Use Figure 11-29, the TS7700 MI Data Classes window to define, modify, or delete a TS7700 DC. The DC is used to automate storage management through the classification of data sets.

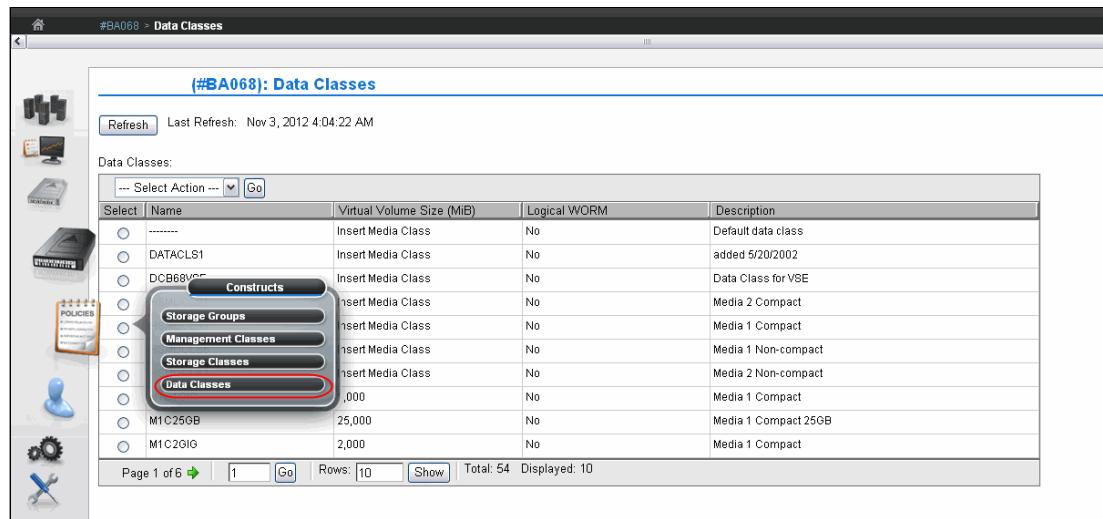


Figure 11-29 Data Classes window

For more information about how to create, modify, or delete a Data Class, see 10.2.4, “Data Classes window” on page 486.

Activating a TS7700 license key for a new Feature Code

The Feature Licenses page is used when you must view information about feature licenses that are installed on a TS7700 cluster, or to activate or remove some feature licenses from the IBM TS7700.

Figure 11-30 shows an example of the Feature Licenses window.

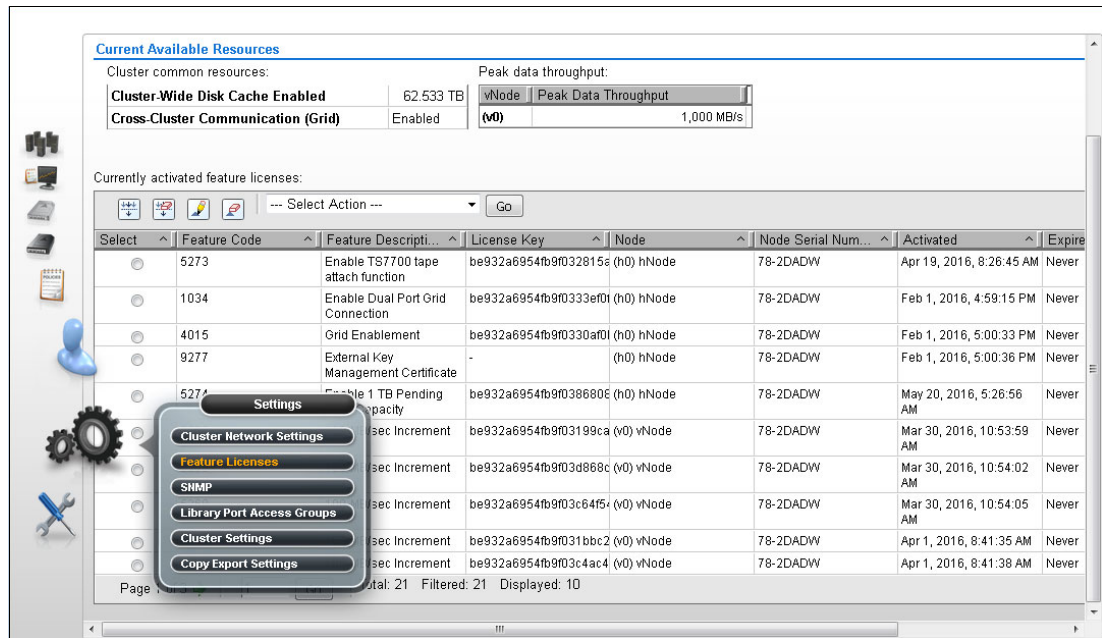


Figure 11-30 Feature Licenses window

For more information about how to use this window, see “Feature licenses” on page 520.

Defining Simple Network Management Protocol

SNMP is one of the notification channels that IBM TS7700 uses to inform the user that an unexpected occurrence, malfunction, or event happened. Use the window that is shown in Figure 11-31 to view or modify the SNMP that is configured on a TS7700 Cluster.

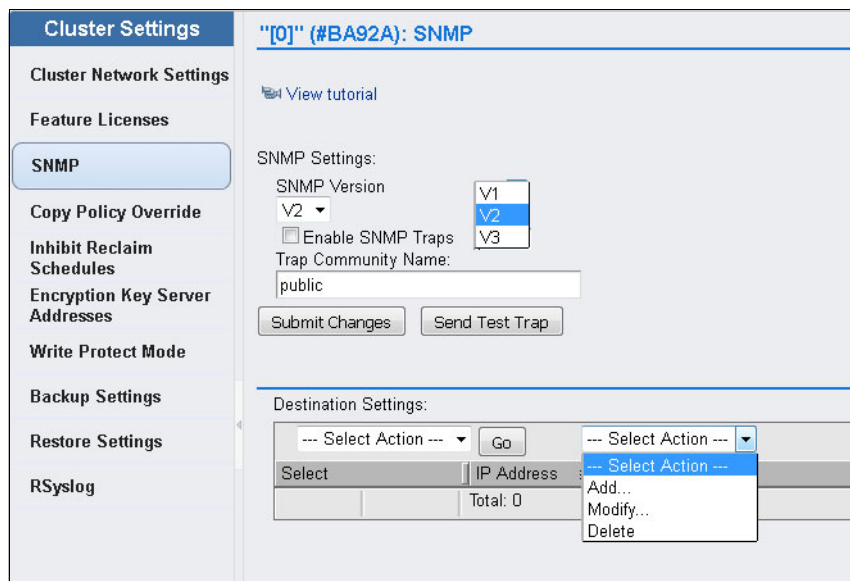


Figure 11-31 SNMP settings

For more information about how to use this page, see “Defining Simple Network Management Protocol” on page 596.

Enabling IPv6

IPv6 is supported by the TS7700 clusters.

Tip: The client network must choose 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 supported.

Figure 11-32 shows the Cluster Network Settings window from which you can enable IPv6.

The screenshot displays the "Cluster Network Settings" window for cluster "[0]" (#BA92A). The left sidebar lists various settings categories, with "Cluster Network Settings" selected. The main area has two tabs: "Customer IP Addresses" (active) and "Encrypt Grid Communication". Under "Customer IP Addresses", the "IPv6" radio button is selected. The IPv4 section includes fields for "IP Address" (9.11.219.80), "Primary Address" (9.11.219.73), "Secondary Address" (9.11.219.74), "Subnet Mask" (255.255.255.0), and "Gateway" (9.11.219.1). The IPv6 section includes fields for "Primary Address", "Secondary Address", "Prefix Length" (64), and "Gateway". A "DNS Server" field is also present. "Submit" and "Reset" buttons are at the bottom.

Figure 11-32 Cluster Network Settings

For more information about how to use Cluster Network Settings window, see “Cluster network settings” on page 519.

Defining security settings

Role-based access control (RBAC) is a general security model that simplifies administration by assigning roles to users and then assigning permissions to those roles. LDAP is a protocol that implements an RBAC methodology.

The TS7700 supports RBAC through the System Storage Productivity Center or by native LDAP by using Microsoft Active Directory (MSAD) or IBM Resource Access Control Facility (RACF).

For more information about setting up and checking the security settings for the TS7700 grid, see 10.3.2, “Security Settings window” on page 495.

11.2.3 TS7700 multi-cluster definitions

The following sections describe TS7700 multi-cluster definitions.

Defining the Secure Data Transfer

Secure Data Transfer (SDT) replaces the IPsec as an effective way to secure the logical volume data in transit through the grid links in a TS7700 grid. The SDT uses OpenSSL software libraries with the TLS1.2 protocol following AES standards, which support AES-256 and AES-128.

Note: Feature Code 5281 is required to enable SDT.

The SDT is negotiated between two clusters in a grid to start a data transfer between the clusters. The highest common key that is defined in the clusters is used in the exchange. If one of the clusters is found non-SDT-capable during this initial exchange, no encryption is performed in the data transfer between these two clusters.

Note: Enabling SDT does not affect the VEC or VED performance in a noticeable way. SDT can be enabled or disabled concurrently with cluster operation.

To enable (or disable) the Secure Data Transfer, see Figure 11-33.

The screenshot shows the 'Cluster Settings' page for '[2] (#BA68C): Secure Data Transfer'. The left sidebar contains navigation options: Cluster Network Settings, Feature Licenses, SNMP, Copy Policy Override, Inhibit Reclaim Schedules, Write Protect Mode, Backup Settings, Restore Settings, RSyslog, and Secure Data Transfer (highlighted). The main content area shows the configuration for Secure Data Transfer. It includes a 'Last Refresh' timestamp of Nov 4, 2019, 8:21:38 PM. A descriptive text states: 'Secure data transfer enables virtual volume data encryption between clusters in a grid for either virtual volume copies or remote read and write operations. After enabling the encryption in the selected cluster, enable encryption with remote peers to ensure secure data transfers.' The configuration fields are: Encryption: Disabled (with a dropdown menu showing AES-128 and AES-256 options), Protocol: TLS 1.2, and SSL certificate: lwiks (HTTPS/SDT) Expires June 1, 2033, 12:46:05 PM Brasilia Standard Time. A 'Submit' button is located below these fields. Below the configuration, a section titled 'Secure data transfer state at remote clusters:' contains a table with the following data:

	Pesto (0)	Squint (1)	Jerry (6)
Encryption:	Unsupported	Unsupported	Unsupported
Protocol:	-	-	-

At the bottom of the interface, there are status bars for Physical Cache (9.45 TB / 157.51 TB (6%)), Copy Queues (11 | 0.75 GiB), and Health.

Figure 11-33 Enable or Disable the Secure Data Transfer.

Adding an SSL or CA certificate

To import a new SSL certificate, select the **New Certificate** tab from the options in the SSL Certificates, then click **Access** → **SSL Certificate**. Select a method to add a certificate, as shown in Figure 11-34 on page 599.

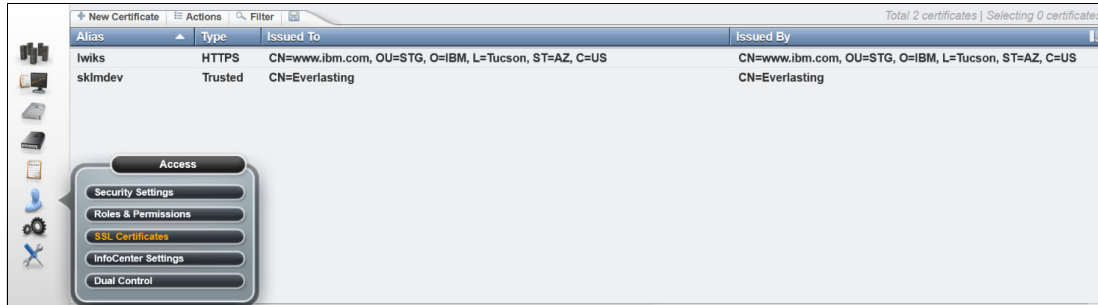


Figure 11-34 Adding a certificate

Proceed with the certificate installation process as shown in Figure 11-34. Figure 11-35 shows the certificates that are installed in a TS7770-VED.

Alias	Type	Issued To	Issued By
lwiks	HTTPS	CN=pigspearls.gdl.mex.ibm.com, OU=Systems, O=IBM, L=Gdl, ST=Jal, C=MX	
google	Trusted	CN=*.google.com, O=Google LLC, L=Mountain View, ST=California, C=US	
gdllinux	Trusted	CN=CP-7941G-SEP001D456151B2, OU=EVVBU, O=Cisco Systems Inc.	

Figure 11-35 Installed certificates

The following types of certificates are shown in Figure 11-35:

► **HTTPS (lwiks)**

The Secure Sockets Layer (SSL) certificate, which is used by the encrypted communications between the TS7700 MI web server and user web browsers. With R5.0 and later, the lwiks certificate is also used for Secure Data Transfer (SDT) and the secure encryption key exchanges. The lwiks (Liberty SWS) certificate is exclusive; each cluster features its own SSL certificate.

► **Trusted**

Client-provided certificates, which are required for Encryption Key Managers, LDAP servers, Cloud Tiering, or client proxies. The Trusted certificates are shared among the members of a grid. When a trusted certificate is installed in one cluster, it is propagated to all members of that grid.

For more information, see the following resources:

- 4.3.9, “Secure Data Transfer” on page 189
- 4.3.18, “Planning for tape encryption in a TS7700T” on page 198
- 4.3.19, “Planning for cache disk encryption in the TS7700” on page 200
- [SSL certificates](#)

Defining grid copy mode control

When upgrading a stand-alone cluster to a grid (FC 4015, 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 must be set.

The data consistency point is defined in the MC's construct definition through the MI. This task can be performed for a grid system. In a stand-alone cluster configuration, the Modify MC definition displays only the lone cluster.

For more information about how to modify the copy consistency by using the Copy Action table, see 10.2.2, “Management Classes window” on page 478.

Figure 11-36 shows an example of 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 extra menu is displayed only if a TS7700 is part of a grid environment (options are not available in a stand-alone cluster).

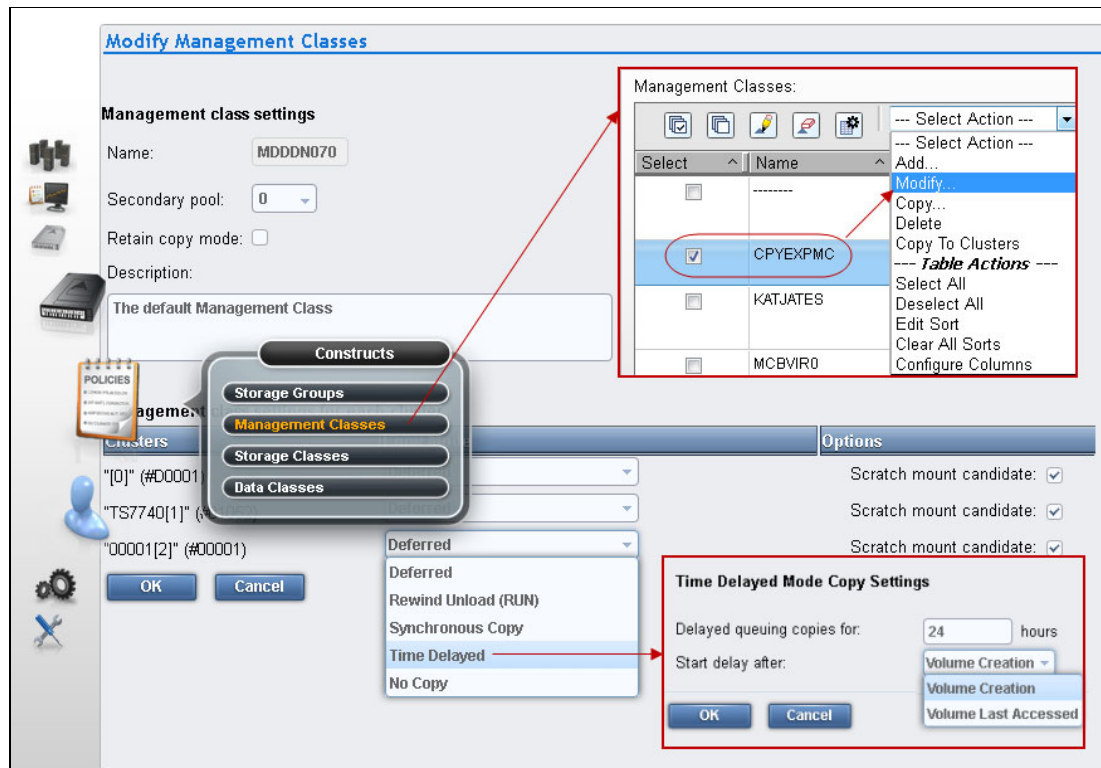


Figure 11-36 Modify Management Classes

For more information about how to modify the copy consistency by using the Copy Action table, see 10.2.2, “Management Classes window” on page 478.

For more information about this subject, see the following resources:

- ▶ [IBM Virtualization Engine TS7700 Series Best Practices - TS7700 Hybrid Grid Usage](#)
- ▶ [IBM TS7700 Series Best Practices - Copy Consistency Points](#)
- ▶ [IBM TS7700 Series Best Practices - Synchronous Mode Copy](#)

Defining Copy Policy Override settings

With the TS7700, the optional override settings that influence the selection of the tape volume cache and replication responses can be defined and set.

Reminder: The items in this window can modify the cluster behavior regarding local copy and certain I/O operations. Some **LI REQ** commands can also make these modifications.

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 that are 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

For more information about how to view or modify these settings, see “Copy Policy Override” on page 523.

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 TS7700D virtual tape drives for fast access. Other workloads can be directed to a TS7700T for archival purposes.

Clusters that are 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.

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, see Chapter 12, “IBM z/OS host console operations” on page 625.

Figure 11-37 shows an example of an MC. Select which clusters are candidates by an MC. If no clusters are selected, the TS7700 defaults to all clusters being candidates.

Modify Management Classes

Management class settings

Name:

Retain copy mode:

Description:

Management class settings for each cluster

Clusters	Copy Mode	Options
"[0]" (#BA92A)	Synchronous Copy	Scratch mount candidate: <input checked="" type="checkbox"/>
"Arabian[1]" (#BA92B)	No Copy	Scratch mount candidate: <input checked="" type="checkbox"/>
"Oak[2]" (#BA92C)	Synchronous Copy	Scratch mount candidate: <input checked="" type="checkbox"/>
"Palomino[7]" (#BA92D)	No Copy	Scratch mount candidate: <input checked="" type="checkbox"/>

Synchronous Mode Copy Settings

Synchronous Deferred On Write Failure:

On Private Mount:

Figure 11-37 Scratch mount candidate list in a Management Class

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, you usually want each cluster in the grid to provide the same list of candidates for load-balancing. For more information about how to create or change an MC see, 10.2.2, "Management Classes window" on page 478.

Note: The scratch mount candidate list as defined in MI (see Figure 11-37) is accepted only upon being enabled by using the **LI REQ** setting.

Retain Copy mode

Retain Copy mode is an optional setting in which Copy Consistency Points for a volume are accepted rather than applying the Copy Consistency Points that are defined at the mounting cluster. Retain Copy mode applies to private volume mounts for read/write appends. It is used to prevent more copies of a volume being created in the grid than wanted. Retain Copy mode is important in a grid with three or more clusters that has two or more clusters online to a host.

Figure 11-37 on page 601 also shows the Retain copy mode option in the TS7700 MI window.

Note: The Retain Copy mode option is effective on private (non-scratch) virtual volume mounts only.

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. By using cluster families, a common purpose or role can be assigned to a subset of clusters within a grid configuration. For example, the role that is assigned (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 tape volume cache selection, or replication can source volumes from other clusters within its family before using clusters outside of its family.

Use the **Cluster Families** option from the **Actions** menu of the Grid Summary window to add, modify, or delete a cluster family.

Figure 11-38 on page 603 shows an example of how to create a cluster family by using the TS7700 MI.

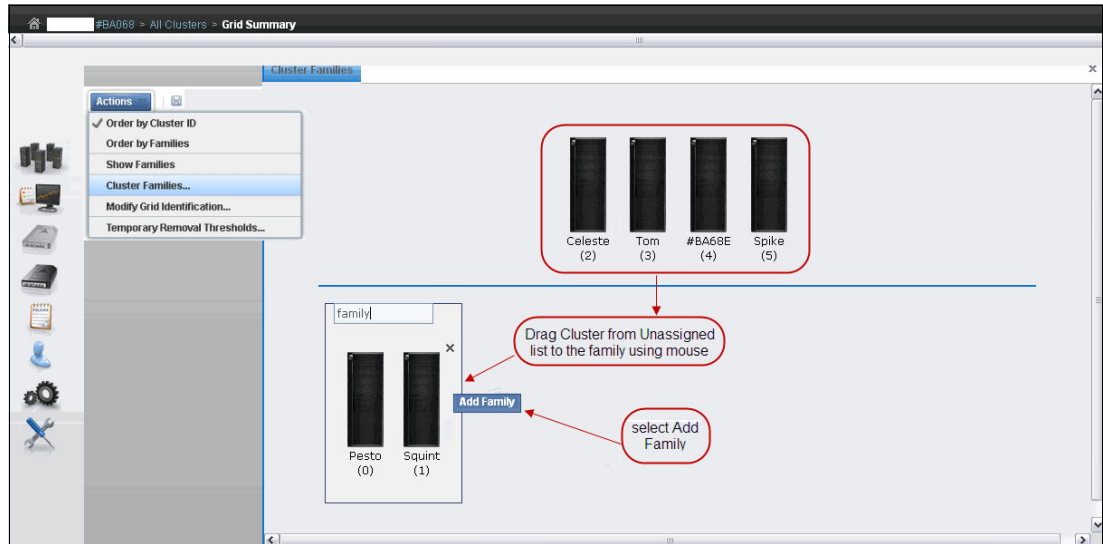


Figure 11-38 Create a cluster family

For more information about how to create, modify, or delete families on the TS7700 MI, see 9.6.3, “Cluster Families window” on page 371.

TS7700 cache thresholds and removal policies

This topic describe 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, the following thresholds are available:

- ▶ **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}
```

Automatic removal is temporarily disabled while disaster recovery write protect is enabled on a disk-only cluster so that a disaster recovery 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 less than 3 TB of free space remains 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 [TS7700 Cache thresholds and removal policies](#).

Clarification: Host writes to the disk-only TS7700 cluster and inbound copies continue during this state.

► Out of cache resources

This state occurs when less than 1 TB of free space remains 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 TS7700D 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 might appear:

- HYDME0997W
- HYDME1133W
- HYDME1201W

For more information, see [TS7700 Cache thresholds and removal policies](#).

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. Although read mounts can choose the disk-only TS7700 cluster in this state, modify and write operations fail. Copies inbound to this cluster are queued as Deferred until the disk only cluster exits this state.

The start and stop thresholds for each of the active cache capacity states that are defined are listed in Table 11-3.

Table 11-3 Active cache capacity state thresholds

State	Enter state (space free)	Exit state (space free)	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

State	Enter state (space free)	Exit state (space free)	Host message displayed
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 ¹	<(X + 1 TB) ²	>(X + 1.5 TB) ²	Console message

1. When enabled.
2. 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 11-39 shows several definitions in place.

Select	Name	Volume Copy Retention Group	Volume Copy Retention Time	Description
<input type="radio"/>	-----	Prefer Keep	0 hours	The default Storage Class
<input type="radio"/>	CACHEIN	Prefer Keep	0 hours	Default storage class
<input type="radio"/>	CACHEOUT	Prefer Keep	0 hours	Default storage class
<input type="radio"/>	CU270SHR	Prefer Keep	0 hours	Default storage class
<input type="radio"/>	SC\$00	Prefer Keep	0 hours	DEFINED
<input type="radio"/>	SC\$10	Prefer Keep	0 hours	Level 0 - Used for VM volumes
<input type="radio"/>	SC\$11	Prefer Keep	0 hours	Level 1 - Used for VM volumes
<input type="radio"/>	SCB68VSE	Prefer Keep	0 hours	VSE
<input type="radio"/>	SCBARR68	Prefer Remove	0 hours	The default Storage Class
<input type="radio"/>	SCPOOL1	Prefer Remove	0 hours	The default Storage Class

Figure 11-39 Storage Classes in TS7700 with removal policies

To add or change an SC, select the appropriate action in the menu, and click **Go** (see Figure 11-40).

Figure 11-40 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 are removed from a disk-only TS7700 cache.

Note: Logical volumes are removed only if another consistent copy exists within the grid.

Logical volumes are removed from a disk only TS7700 cache in the following order:

1. Volumes that are in scratch categories
2. Private volumes that are 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 is 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 visible only 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 that is 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 the cache for FlashCopy volumes that are 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 running out of cache resources, considering the worst possible scenario. The Temporary Removal Threshold resource is instrumental in helping to prevent this possibility.

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 clusters or CP0 partitions 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 was started enters service mode or the temporary removal process is canceled. The cluster that is starting the temporary removal process (a cluster within the grid that is not part of the DR testing or the cluster that is scheduled to go into Service) does not lower its own removal threshold during this process.

Note: The cluster that is elected to start Temporary Removal process is not selectable in the list of target clusters for the removal action.

Removal policy settings can be configured by using the Temporary Removal Threshold option in the Actions menu, which is available on the Grid Summary window of the TS7700 MI.

Figure 11-41 shows the Temporary Removal Threshold mode window.

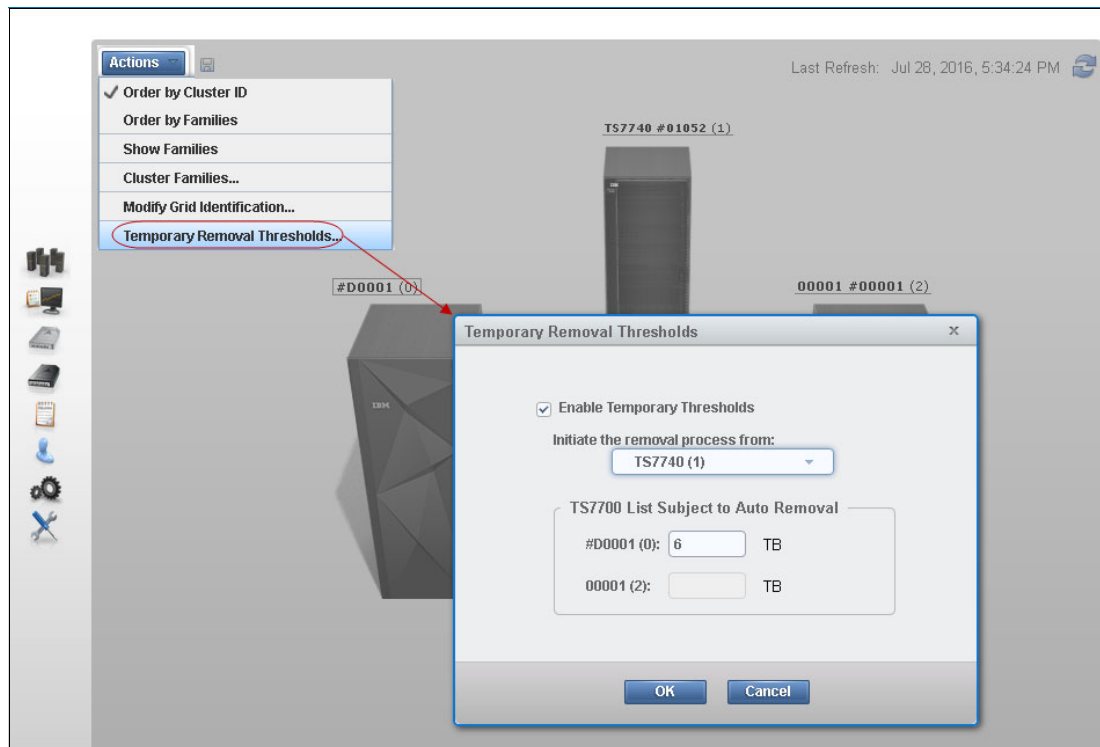


Figure 11-41 Selecting cluster to start removal process and temporary removal threshold levels

The Temporary Removal Threshold mode window includes the following options:

- ▶ Enable Temporary Thresholds
Select this option and click **OK** to start the pre-removal process. Clear this option and click **OK** to abandon a current pre-removal process.
- ▶ Initiate the removal process from (cluster to be serviced)
Select the cluster from this menu that will be put into service mode. The pre-removal process is started from this cluster.

Note: This process does not start Service Prep mode.

Even when the temporary removal action is started from a disk-only cluster, this cluster still is not selectable from the drop-down menu of the TS7700 List Subject to Auto Removal because the removal action does 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.

11.3 Basic operations

This section explains the tasks that might be needed during the operation of a TS7700.

11.3.1 Clock and time setting

In this section, we describe the TS7700 clock and time settings.

TS7700 time setting

The TS7700 time can be set from a Network Time Protocol (NTP) server or by the IBM SSR. It is set to Coordinated Universal Time. For more information about time coordination, see “Date and Time coordination” on page 66.

Note: Use Coordinated Universal Time in all TS7700 clusters whenever possible.

The NTP server can be configured simultaneously with cluster or grid operation at current code levels. Work with the IBM SSR to enable the NTP server in the TS7700 cluster or grid. Depending on the current cluster or grid time setting, the SSR IBM needs to correct the time so that the time is within a few minutes of the NTP server before activating the NTP server. The allowed time difference for an adjustment operation is two hours. If the difference between the cluster current timing setting and the NTP server is greater than two hours, the IBM SSR should run sequential adjustments until the Time Of Day (TOD) setting is close enough to enable the NTP server.

The grid time adjustment for the NTP server is made in cluster zero or the lowest numbered cluster in the grid by the IBM SSR.

For more information, see “Date and Time coordination” on page 66.

Tape Library time setting

The TS4500 Tape Library time can be set from management Interface, as shown in Figure 11-42. Notice that TS4500 can be synchronized with NTP server, when available.

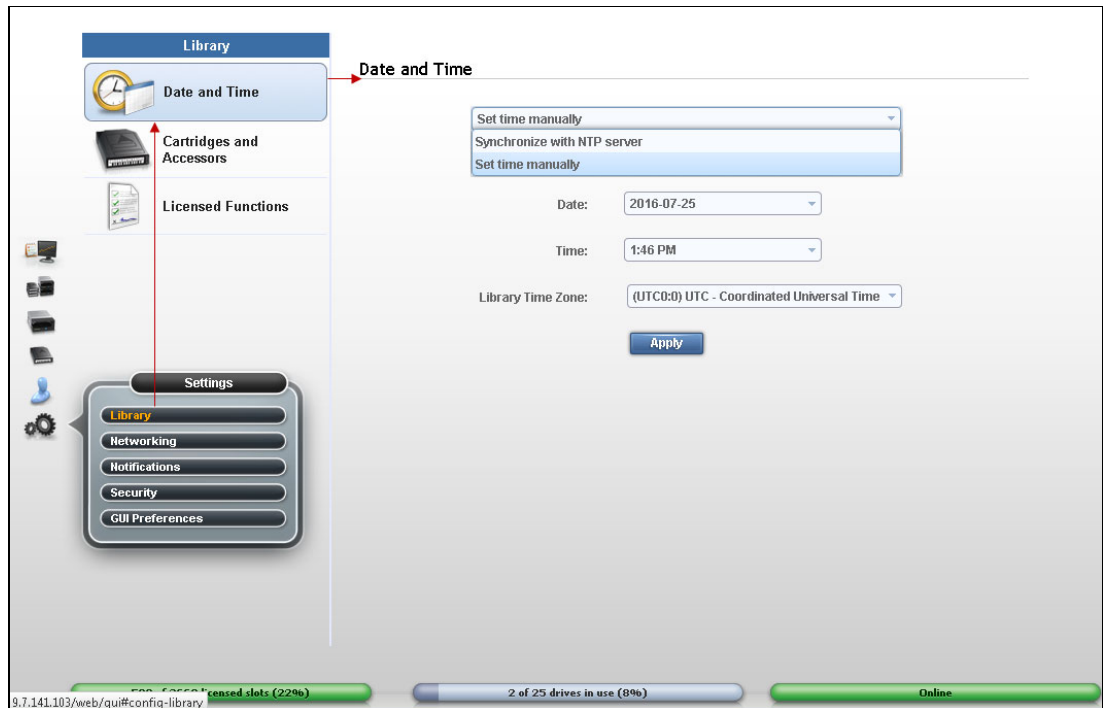


Figure 11-42 Adjusting date and time at TS4500 GUI

On the TS3500 tape library, the time can be set from IBM Ultra Scalable Specialist work items by clicking **Library** → **Date and Time**, as shown in Figure 11-43.

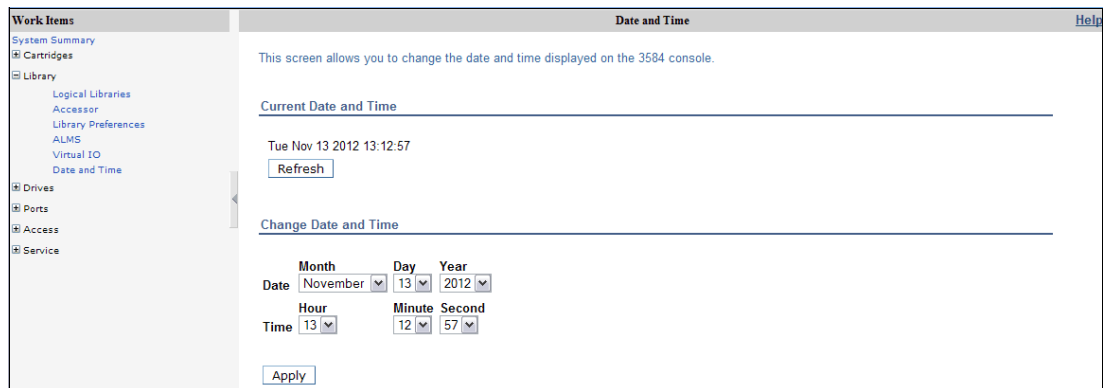


Figure 11-43 TS3500 tape library GUI Date and Time

For more information about TS4500 Tape Library see local documentation in the TS4500 GUI by clicking the question mark icon, or see [Set library time, date, and time zone](#).

11.3.2 Physical Tape Library paused or degraded

During the operation, the tape library can be paused, which might affect the related tape attach cluster, regardless of whether it is in a grid. 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 TS7700T for later processing when the library leaves the Pause mode. Because scratch mounts and private mounts with data in the cache are allowed to run (but not physical mounts), no other data can be moved out of the cache after the currently mounted stacked volumes are filled.

During an unusually elongated period of pause (such as during a physical tape library outage), the cache continues to fill up with data that is not migrated to physical tape volumes. This issue might lead in extreme cases to significant throttling and stopping of any mount activity in the TS7700T cluster.

The TS7700T cluster inherited many of the behaviors from the original TS7740 with a small cache size. Considering the current TS7700T with bigger cache partitions, it is possible to continue operating with physical tape library degraded or non-operational for an extended period.

The **LI REQ PHYSLIB SETTINGS** commands allow the user to keep the TS7700T tape partition operating normally in the case of the physical tape library extended outage (with exception of recalls). The **LI REQ PHYSLIB** parameters allow the TS7700T to go over the licensed amount of the pending premigration tape capacity (FC 5274) without throttling the tape attach partition.

For more information about this command, see [IBM TS7700 Library Request Command](#).

At previous levels of code, having the PRETHDEG setting disabled did help clients work around a physical tape library outage. However, after the tape library became operational again, the original values of the premigration throttling were turned back on immediately, which can slowdown the tape attach partition until the premigration queue deflated to a level that was below the original premigration threshold.

With R5.0 and later, the cluster waits for the premigration queue to deplete before reintroducing the original throttling values. Depleting the premigration queue allows for the tape partition to move out of the tape library outage back to the normal operation without affecting the performance of the tape attach partition in the cluster. For more information, see the “PHYSLIB,SLDPMPRI,ENABLE/DISABLE LI REQ command” topic in Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Similar behavior and defense mechanisms also occur in a Cloud Storage Tier enabled cluster, the TS7700C if the cloud repository becomes temporarily unavailable for some reason. For more information, see [IBM TS7700 R5.3 Cloud Storage Tier Guide](#), REDP-55733.

11.3.3 Preparing a TS7700 for service

When an operational TS7700 must be taken offline for service, the TS7700 Grid must first 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. For more information about this window, see 9.6.9, “Service mode window” on page 379.

The following message is posted to all hosts when the TS7700 Grid is in this state:

```
CBR3788E Service preparation occurring in library library-name
```

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 running and volumes to unload, which completes the transition to the offline state.

Virtual devices in other clusters are made online to provide a mount point to new jobs, which shifts workload to other clusters in the grid before start service preparation. After scheduled maintenance finishes and the TS7700 can be taken out of service, virtual devices can be varied back online for accessing hosts.

TS7700 supports Control Unit Initiated Reconfiguration (CUIR) at current levels of code. A library request command is performed so that notifications to the host from the TS7700 are enabled, which allows the automation in CUIR to function to minimize operator intervention during preparing TS7700 for service. When service-prep is started on the cluster, a Distributed Library Notification is surfaced from the cluster to prompt the attached host to vary off the devices automatically after the following conditions are met:

- ▶ All clusters in the grid are at microcode level 8.41.200.xx or later.
- ▶ The attached host logical partition supports the CUIR function.
- ▶ The CUIR function is enabled from the CLI by using the **LIBRARY REQUEST, library-name, CUIR, SETTING, SERVICE, ENABLE LI REQ** command.

The **LI REQ** command includes the options **SERVICE, FENCE, and ALL**. Enabling the use of **SERVICE** or **ALL** allows automatic notification during service prep.

When service is canceled and the local cluster comes online, a Distributed Library Notification is surfaced from the cluster to prompt the attached host to vary on the devices automatically after the following conditions are met:

- ▶ All clusters in the grid are at microcode level 8.41.200.xx or later.
- ▶ The attached host logical partition supports the CUIR function.
- ▶ The CUIR function is enabled from the CLI, by using the **LIBRARY REQUEST, library-name, CUIR, SETTING, AONLINE, ENABLE LI REQ** command.
- ▶ The AONLINE notification is enabled from the CLI, by using the following **LI REQ** command:

```
LIBRARY REQUEST, library-name, CUIR, AONLINE, SERVICE, ENABLE
```

If the AONLINE notification is disabled by using the **LIBRARY REQUEST, library-name, CUIR, AONLINE, SERVICE, DISABLE LI REQ** command, a blue informational icon is shown on the lower left of the cluster image to alert the user that the cluster's devices must be varied online.

In this case, devices can be varied online from the Actions menu by completing the following steps:

1. Select **Vary Devices Online** from the Actions menu.
2. Select the radio button next to the cluster that you want to vary devices online for and then click **OK**. Click **Cancel** to exit without varying devices online.

For more information about service preparation, see 12.2, “Messages from the library” on page 645.

If CUIR is not in place, all of the host actions, such as varying devices offline or online, must be performed manually across all LPARs and system plexes that are attached to the cluster.

With R5.0 and later, the user can set a cluster in service by using the **SERVICE,ENTER/CANCEL/SHOW, FORCE LI REQ** command, which starts, displays, or cancels service preparation in the cluster that is specified by the command.

With the R5.2.2 code, when a cluster is fenced as part of the Grid Resiliency mechanism, any cluster that is host-connected to the same hosts as the unhealthy cluster notifies the attached hosts about the sick cluster through CUIR. Hosts will then start the sequence of auto-vary offline for the devices to the unhealthy cluster. A single sick cluster can also notify its own attached hosts about its sick condition, also triggering CUIR and the auto-vary offline sequence.

Figure 11-44 shows CUIR being started due to a cluster fence action that occurred to a cluster that is attached to this z/OS host.

```
IEF880I 25FD NOW OFFLINE BY C.U.I.R.
IEF880I 25FE NOW OFFLINE BY C.U.I.R.
IEF880I 25FF NOW OFFLINE BY C.U.I.R.
IOS279I  C.U.I.R. QUIESCE REQUEST WAS ISSUED
          REQUEST REASON: UNHEALTHY CLUSTER
          THE FOLLOWING DEVICES ARE AFFECTED:
          2500-25FF
IOS281I  C.U.I.R. REQUEST SUCCESSFUL
CBR3750I Message from library HYDRAX: G0053 Library 00001 has applied
REBOOT local fence action. Reason: Manual local cluster fence has been
issued. Severity impact: INFORMATION.
CBR3750I Message from library HYDRAX: G0059 Library 00001 has
successfully surfaced required CUIR vary offline attentions. Severity
impact: INFORMATION.
CBR3750I Message from library HYDRAX: OP0208 The cluster 2 has
applied Reboot local cluster fence action. Reason: Manual local
cluster fence has been issued.. Severity impact: WARNING.
*CBR3762E Library HYDRAX intervention required.
```

Figure 11-44 Example of CUIR messages at z/OS console

After the outage, when the unhealthy cluster has completed restart and become online again (if this was the fence action that is designated to happen for the cluster), the cluster will then notify the host to auto vary the devices back online once it has recovered.

For more information, see 12.9, “CUIR for tape” on page 655 or see [IBM TS7700 Library Request Command](#).

Preparing for servicing the tape library

The user always had to be concerned about whether bringing the TS7700T cluster into service before taking down the associated tape library for service, considering the duration of the planned outage for the physical library, the role played by the tape attached cluster in the grid, available cache, and other factors.

Sometimes, the best option is to prepare TS7700T for service. Another option of keeping the cluster active was the best choice because it did not affect the operations.

R5.0 and later introduces a new **LI REQ** command that helps the user to prepare for a physical tape library outage. The **SETTING2, PHYSLIB, MAINT, ENABLE** command suspends all activities in the physical tape library and drives, such as premigration, recall, reclaim, copy export, secure data erase, read only recovery, and offsite reclaim in preparation for a tape library service window.

Activities that are in progress are not canceled; instead, they are allowed to finish gracefully before all physical tape activities cease. When the tape library is set in maintenance mode, the mechanism that is described in 11.3.2, “Physical Tape Library paused or degraded” on page 610 starts, which allows the tape attach partition to go beyond the licensed capacity for active premigration queue FC 5274.

After the tape library maintenance is complete, the maintenance mode is canceled by using the **SETTING2, PHYSLIB, MAINT, DISABLE LI REQ** command to restore the physical tape library operations. Again, the same mechanism allows the tape attach cluster to deplete the inflated premigration queue before enforcing the original throttling levels. For more information about how to set the TS7700 to service preparation mode by using the TS7700 MI, see 9.6.8, “Cluster Actions menu” on page 377.

11.3.4 Tape Library inventory

The inventory on the TS4500 Tape Library can be started from the Management Interface, as shown at Figure 11-45. A full tape library or frame inventory can be selected.




Figure 11-45 TS4500 inventory options

A partial inventory can be performed in any specific frame of the tape library by completing the following steps:

1. Left-click the wanted frame on the tape library image to select it (the frame changes colors).
2. Right-click to display the options.
3. Select **Inventory Frame** from the list.

A complete tape library inventory can be started from the Actions menu at the top of the page. Both options produce a dialog box, in which the user is prompted to scan tiers 0 and 1, or all tiers.

Selecting the Scan tier 0 and tier 1 option checks cartridges on the doors and the external layer of the cartridges on the walls of the library. This option scans only other tiers if a discrepancy is found. This option is the preferred option for normal tape library operations, and it can be performed concurrently.

Selecting the Scan all tiers option performs a 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 this option only when a full inventory of the tape library is required.

11.3.5 Inventory upload

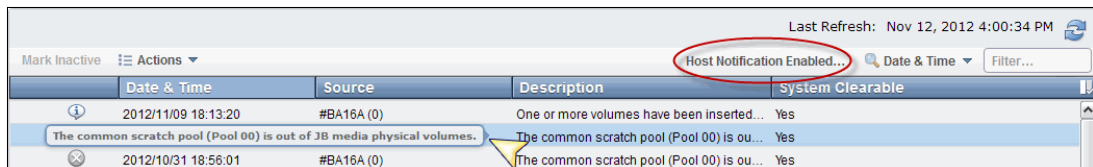
For more information about an inventory upload, see “Physical Volume Ranges window” on page 465.

Click **Inventory Upload** to synchronize the physical cartridge inventory from the attached tape library with the TS7700T database.

Note: An Inventory upload is performed automatically whenever the tape library completes an inventory action, or when the TS7700T cluster goes online.

11.4 Cluster intervention scenarios

This section describes some operator intervention scenarios that might occur. Most of the errors that require operator attention are reported on the MI or through a Host Notification, which is enabled from the Events window of the MI. A sample of one Event message that needs an operator intervention is shown in Figure 11-46.



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

Figure 11-46 Example of an operator intervention

11.4.1 Hardware conditions

Some potential hardware attention scenarios are described in the following sections. For more information about the operational or recovery procedures, see the IBM TS7700 R5.3 IBM Documentation. For more information, see the local TS7700 IBM Documentation from the TS7700 MI by clicking the question mark symbol in the upper right of the top bar, or see [TS7700 5.4.1](#).

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 feature 240 characters. This field includes a five-character message ID that might be examined by the message automation software to filter the events that gets operator attention.

The message ID classifies the event that is reported by its potential effect on the operations. The categories are critical, serious, impact, warning, and information. For more information, see [Notification settings](#).

For more information about these informational messages, see [IBM TS7700 Notification Messages v5.3](#). At the time of writing this book, no white paper is available for IBM TS7700 notification messages for R5.4.

IBM 3592 tape drive failure (TS7700T)

When the TS7700 determines that one of its tape drives is not operating correctly and requires service (because of 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 that belongs to the TS7700T exceeded the media error threshold, 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 includes 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 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, which ends 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 cartridge inventory.

Note: By design, the tape-attached 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, no other actions are needed, except inserting a new cartridge to replace the ejected cartridge.

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, which signals 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 because of 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, which completes the active data recovery.

If recovery fails because no other consistent copy is available within the grid, or this cluster is a stand-alone cluster, the media is not ejected and the 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 **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 no longer needed, they should be made into scratch volumes by using the TMS on the IBM Z 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 window, and through an MI function). Because these logical volumes that are not retrieved do not contain active data, the Automated ROR completes successfully this time, 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 that are reading or writing to virtual tapes fail as they fail with a real IBM 3490E. 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 following message:

```
HYDIN0571E Disk operation in the cache is degraded.
```

Also, the MI shows the Health Status bar (lower right of the Cluster Summary page in yellow warns the user 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 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.

Accessor failure and manual mode (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 feature 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 features 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 is dislodged or damaged, follow the instructions that are described 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 (for example, the tape is crushed or the media is physically broken), the TS7700T cannot recover the tape contents when configured as a stand-alone cluster without dual copy for volume redundancy. If this TS7700T cluster is part of a grid, the contents of the damaged tape (active logical volumes) will be retrieved from other clusters, and brought in automatically if there is another valid copy on the grid.

For all tape drive models or media types, if the damaged tape is the only data source available, check the list of logical volumes that are contained in the cartridge and work with your IBM SSR to attempt data recovery.

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 older generation of VTS, the VTS always indicated to the host that the mount was completed, even if a problem occurred. When the first I/O command is sent, the VTS fails that I/O because of the error. This occurrence results in a failure of the job without the opportunity to attempt 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, the TS7700 returns completion and reason codes to the host that indicate that a problem was encountered rather than indicating that the mount was successful. With DFSMS, the logical mount failure completion code results in the console messages that are shown in Example 11-1.

Example 11-1 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, review CBR4171I. Reason codes are documented in IBM Documentation. As an exercise, assume RSN=32. In the documented [Completion codes](#), the following reason code is shown:

Reason code x'32': Local cluster recall failed; the stacked volume is unavailable.

- ▶ CBR4196D: The 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 more information about CBR messages, the reason code, and specific actions that must be taken to correct the failure, see *z/OS MVS System Messages, Vol 4 (CBD-DMO)*, SA38-0671. For more information about OAM return and reason codes, see *z/OS DFSMSdfp Diagnosis*, SC23-6863.

Take the necessary corrective action and reply R to try again; otherwise, reply C to cancel.

Tip: Always review the appropriate documentation (TS7700 IBM Documentation and MVS System Messages) for the meaning of the messages and the applicable recovery actions.

Orphaned logical volume

A logical volume is orphaned when the TS7700 database includes a reference to a logical volume, but no reference to its physical location. This issue 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. The intervention-required condition is also posted at the TS7700T MI and sent to the host console (see Example 11-2).

Example 11-2 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.

Note: Starting with R5.0, the user can request the TS7700T to relabel a cartridge whose internal label is known to be different from the external bar code label. For example, the operator changed the bar code label to reuse cartridge in a different TS7700T cluster.

Use the `LIB REQ PVOL, zzzzzz, RELABEL, YES/NO` command to correct the internal label. For more information, see 10.1, “Physical icon” on page 449.

Failure during reclamation

If a failure occurs during the reclamation process, the process is managed by the TS7700T Licensed Internal Code. No user action is needed because recovery is managed internally.

Excessive temporary errors on a stacked volume

When a stacked volume is determined to include an excessive number of temporary data errors, the stacked volume is placed in read-only status to reduce the possibility of a permanent data error. 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.

11.4.2 Ownership takeover interventions

In a multi-cluster environment, the data integrity of the logical volumes is ensured by the ownership mechanism. The basic premise is that only one cluster at a time can modify volume data or attributes of a specific logical volume.

Within the composite library (the grid), all clusters can access all the logical volumes, regardless of where the data is stored. This feature makes possible that any logical volume in the composite library can be accessed from any virtual tape drive in any cluster that is participating in the grid, even when the local (accessing) cluster does not have a copy of that logical volume. The TS7700 volume ownership concept rules dictate that only one cluster owns a logical volume at any time.

If this logical volume must be mounted in a different cluster (not the cluster that owns that volume) in the grid, this cluster must obtain the ownership of that volume before running the mount. The ownership of the volumes is negotiated and transferred between the clusters by way of the grid links.

The ownership takeover intervention messages occur when the cluster that received the mount request from the host cannot get the ownership of that volume. This ownership failure can occur because the owning cluster cannot answer due to a malfunction or because the communication grid links between them are severed.

The following message can occur:

```
CBR4174I Cannot obtain ownership volume volser in library libname (Note: this message indicates that an operation was attempted that requires volume ownership and volume ownership could not be obtained)
```

The requesting cluster cannot automatically assume the ownership of the volume to prevent cases in which an invalid ownership acquisition might occur because only the communication of the grid links failed entirely, and both clusters are still operational.

In such cases, the cluster needs more direction to proceed, taking over the ownership of the volume and complete its mount, which makes the volume available to the requesting routine in host. If no directions are provided, the mount fails and the job abends.

For more information about directions that can be given manually by way of TS7700 MI, see 10.5.1, "Ownership Takeover Mode" on page 544.

Ownership takeover can occur in the following modes:

► Read-only Ownership Takeover

When Read-only ownership takeover (ROT) is enabled for a failed cluster, ownership of a volume can be taken from the failed TS7700 cluster when a volume is accessed by a host operation. Only read access to the volume is allowed through the other TS7700 clusters 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 is failed. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster is restored.

Any volumes that are accessed during the outage that were taken over in this mode are reconciled after the original owner returns and all clusters are made aware of the final owner. If a volume was accessed and modified during the outage by the original owner (network outage only), no error event occurs because the temporary owner only had read access.

► Write Ownership Takeover

When Write Ownership Takeover (WOT) is enabled for a failed cluster, ownership of a volume can be taken from the failed TS7700 cluster when a volume is accessed by a host operation. Full access is allowed by using other TS7700 clusters in the grid.

The mode for the failed cluster remains in place until a different mode is selected or the failed cluster is restored. Any volumes that are accessed during the outage that were taken over in this mode are reconciled after the original owner returns and all clusters are made aware of the final owner and the latest properties and volume data.

Replications are queued if data changed during the outage. If a volume was accessed and modified during the outage (network outage only) by the original owner and the temporary owner also modified the volume, the volume is moved into an error state where manual intervention is required to choose the most valid version.

Autonomic ownership takeover is designed to prevent such takeover enablement. Safety checks in manual enablement also prevent such a condition if the TS7700 and TSSC infrastructure determine that only a network outage exists.

► Service Ownership Takeover

When a TS7700 cluster is placed in service mode, the TS7700 Grid automatically enables Write Ownership Takeover mode against the serviced cluster. Although the result is identical to WOT, it is given a unique name to differentiate why it was enabled. This mode is not explicitly enabled through the management interface; rather, it is implicitly enabled by starting the service preparation process.

An automated optional method is provided to enable a user selected ownership takeover mode when all clusters are connected to a TSSC, by configuring the Autonomic Ownership Takeover Manager function (AOTM). AOTM is configured by the IBM service representative and by the user to take effect by taking a takeover action when the occasion presents.

Figure 11-47 shows an example of the AOTM configuration for a three-cluster grid in which a cluster failed.

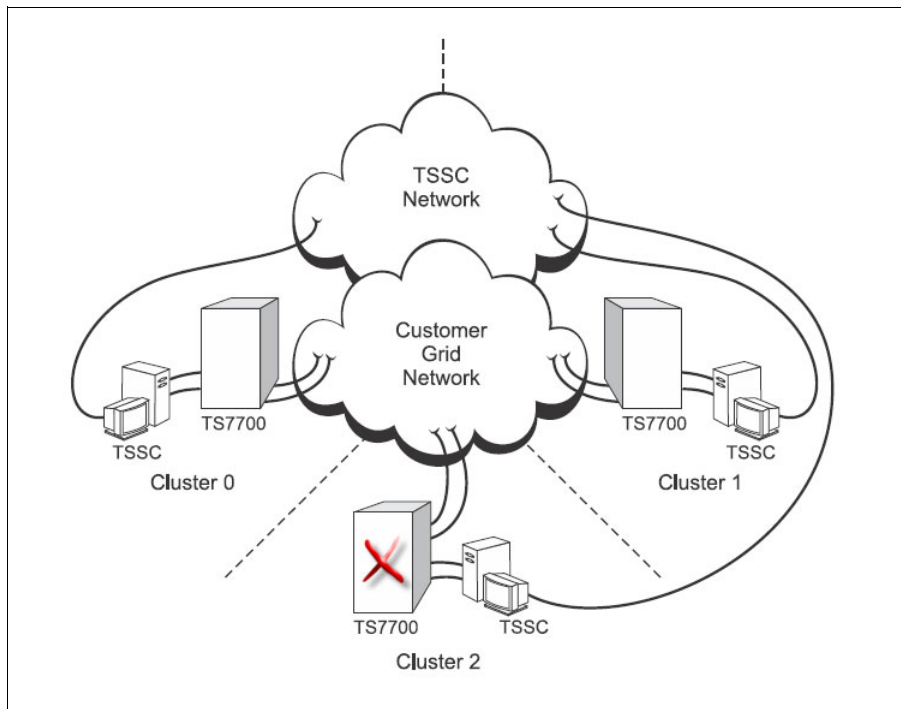


Figure 11-47 AOTM diagram for a three-way grid

Assume that Cluster 0 in the grid that is shown in Figure 11-47 cannot process transactions with Cluster 2 because communication with the remote cluster is lost. With AOTM configured and enabled, Cluster 0 starts a timer (named AOTM grace period, which is configured by the user).

When the grace period expires, the AOTM starts the following process to identify if remote cluster is down:

1. Cluster 0 communicates with its local TSSC.
2. TSSC 0 forwards a request to a TSSC that is attached to the Cluster 2 by way of TSSC Network.
3. TSSC 2 attempts to communicate with Cluster 2 (which is down).
4. If the remote TSSC 2 (connected to Cluster 2) request returns and agrees that the Cluster 2 failed, the takeover mode that was configured by the user is activated.

When takeover is enabled by AOTM, Cluster 0 features Read-Only or Read/Write access to the data that is owned by the failed cluster, according to the mode that is specified by the user in the AOTM configuration.

Note: AOTM needs to be configured during TSSC/Cluster installation by the IBM SSR (or later) before it can be used or configured by the user through TS7700 MI.

For more information about AOTM and how to configure it, see 2.4.33, “Autonomic Ownership Takeover Manager” on page 97 and [The Autonomic Ownership Takeover Manager](#).



IBM z/OS host console operations

This chapter describes commands and procedures that interact with the z/OS host operating system.

This chapter includes the following topics:

- ▶ 12.1, “System-managed tape” on page 626
- ▶ 12.2, “Messages from the library” on page 645
- ▶ 12.3, “Expire Hold and scratch processing considerations” on page 648
- ▶ 12.4, “Scratch count mismatch” on page 649
- ▶ 12.5, “Host cartridge entry processing” on page 650
- ▶ 12.6, “Effects of changing volume categories” on page 652
- ▶ 12.7, “Library messages and automation” on page 653
- ▶ 12.8, “Mount retry” on page 653
- ▶ 12.9, “CUIR for tape” on page 655
- ▶ 12.10, “Cloud Storage tier considerations (R4.2 enhancement)” on page 658
- ▶ 12.11, “Return-to-scratch enhancement (OA48240)” on page 658
- ▶ 12.12, “OAM Object SYSZTIOT enhancement” on page 659
- ▶ 12.13, “Enhanced SSMHONOR support” on page 659
- ▶ 12.14, “DFSMSHsm RECYCLE Enhancement for TS7700C” on page 660
- ▶ 12.15, “DFSMSHsm RECYCLE Considerations when using zEDC” on page 660
- ▶ 12.16, “LWORM retention changes” on page 661

12.1 System-managed tape

This section describes the commands that are used to operate a tape library in an IBM z/OS with 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.

12.1.1 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 be performed for SMS to recognize the library.

The following DFSMS operator commands support the tape library:

- ▶ `DISPLAY SMS,LIBRARY(libname|ALL),STATUS`

This command 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`

Although a single-system view, the `DETAIL` display gives more information. This method is suggested to display a high-level overview of each library that is defined to SMS in the configuration (see Example 12-1).

Example 12-1 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 operational state (auto, pause, or manual mode). The status lines at the end of the output are surfaced based on information that is obtained directly from the library (see Example 12-2).

Example 12-2 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. The information that is obtained directly from the library, such as the LIBRARY CATEGORY, LM constructs (SMS constructs that are stored in the library), and LM CATEGORY (see Example 12-3).

Example 12-3 D SMS,VOLUME(volser)

```
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 (as shown in Example 12-4) is useful for checking the status of the object access method (OAM) user exits and, if the CBROAMxx parmlib member was used during OAM startup and a SETTLIB parameter specified, the suffix of this member.

Example 12-4 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.
CBROAM: 00
```

- ▶ VARY SMS,LIBRARY(libname),ONLINE/OFFLINE

From the host perspective, 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. The use of this command on the composite library with the **OFFLINE** parameter stops tape library actions and gradually makes all of the tape units within this composite library unavailable to the MVS system the command was issued to. 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. However, varying a distributed library offline from the host really has no meaning and does not prevent outboard usage of the library. Message CBR3016I warns the user when a distributed library is offline during OAM initialization or is varied offline while OAM is active.

The use of this command with the **ONLINE** parameter is required to bring the SMS-defined library back to operation after it is offline.

- ▶ VARY SMS,LIBRARY(libname,sysname,...),ON/OFF and VARY SMS,LIBRARY(libname,ALL),ON/OFF

This extended versions of the VARY command can affect more than one system. The first version affects one or more named MVS systems. The second version 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}`

Consider the following points:

- When used against a real, physical volume, the **LIBRARY EJECT** command instructs the library the volume is in to EJECT the volume. In a TS7700, all the volumes are virtual, so no volume can be physically ejected.

When issued against a virtual volume in a TS7700, the **EJECT** command deletes that virtual volume from the TS7700. Virtual volumes can be ejected only if they are in a scratch category (and not being retained by the EXPIRE HOLD settings for that scratch category or for LWORM only the LWORM retention settings).

This command cannot be used to eject a physical backend volume from a TS3500 or TS4500 because the host has no knowledge of backend tape libraries. The following options are available for this command:

- 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 was 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 information is 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 delete a single virtual volume. Mass deletion of virtual volumes is performed through program interfaces, such as ISMF, a tape management system (TMS), or a batch job.

12.1.2 MVS system commands

For more information about the following commands, see *z/OS MVS System Commands*, SA38-0666:

- ▶ `VARY unit,ONLINE/OFFLINE`

The **VARY unit** command is used to vary a unit online or offline. However, when a unit belongs to a library, situations might arise that prevent a unit from being varied online.

When the library is offline, the tape units cannot be used. This state is internally indicated in a status (offline for library reasons), which is separate from the normal unit offline status. A unit can be offline for 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 bring it back online.

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. Although the order of these activities is not important, both are necessary.

- ▶ `LIBRARY DISPDRV,library_name`

The **LIBRARY DISPDRV (LI DD)** command (with output shown in Example 12-5) indicates whether a device is online or offline, and the reason if it is offline.

Example 12-5 LIBRARY DISPDRV command against a composite library

```
LIBRARY DISPDRV,COMPLIB1
CBR1220I Tape drive status: 914
DRIVE  DEVICE  LIBRARY  ON  OFFREASON  LM  ICL  ICL  MOUNT
NUM    TYPE    NAME      LI  OP  PT  CU  AV  CATEGRY  LOAD  VOLUME
1C00   3490    3484F    N   N  Y  Y  N   -  --N/A--  -
1C01   3490    3484F    Y   N  N  N  N   A  NONE     N
1C02   3490    3484F    N   N  Y  Y  N   -  --N/A--  -
1C03   3490    3484F    N   N  Y  Y  N   -  --N/A--  -
```

Another form of this command includes the **MOUNTED** keyword. This keyword specifies that status information is displayed for volumes that are mounted in the TS7700 for the specified library (composite or distributed). The status information includes 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).

An **ALL** parameter is available that can be passed with **MOUNTED**. **ALL** specifies that more drives can be displayed that are not owned by the distributed library that is targeted in the command. The extra drives are displayed if the distributed library that is specified is the primary or secondary TVC for the mounted volume.

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 or helping to identify which jobs must complete.

The **LIBRARY DISPDRV MOUNTED{,ALL}** command uses the following syntax:

```
LIBRARY DISPDRV,library_name,MOUNTED or
LIBRARY DISPDRV,library_name,MOUNTED,ALL

or

LI DD,library_name,M
LI DD,library-name,M,A
```

The *library_name* can be a composite or distributed library. The following examples illustrate the differences in the display output (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 12-6, all volumes that are mounted in the grid are displayed along with the distributed library in which they are mounted.

Example 12-6 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 12-7, 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 12-7 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 12-8, the ALL keyword is added 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 12-8 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 more information about this command and its output, see *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

- ▶ `DISPLAY M=DEV(xxxx)`

The **D M=DEV** command is useful for checking the operational status of the paths to the device (see Example 12-9).

Example 12-9 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 (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)`

You can use the **MOUNT** command to allocate an I/O device to all job steps that require a particular volume without intervening de-mountings and re-mountings of the volume.

- ▶ `UNLOAD devnum`

The **UNLOAD** command enables you to unload a drive if the Rewind Unload (RUN) process was not successful initially.

12.1.3 Host Console Request function

The **LIBRARY REQUEST** host console command (**LI REQ**) provides a way for an operator to determine the status of the TS7700 to obtain information about the resources of the TS7700, or 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.

Optionally, RACF (or an equivalent security product) can be used to control which users can issue the command. Access is controlled by way of the MVS.LIBRARY.REQUEST resource. For more information, see *DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, [SC23-6867](#).

Note: Be careful to not submit too many LI REQ commands on a host at one time. Some commands (such as those commands that display TS7700 data) run quickly, but others (those commands that instruct the TS7700 to take an action) can take longer.

LI REQ commands are processed in the M2 command class that is shared with other commands such as LOGON and START. A total of 50 commands can be run concurrently in the M2 class. Commands that are issued in excess of this value remain queued until 1 of those 50 running commands completes processing.

It is recommended to insert a delay between issuing LI REQ commands that are not strictly used to display TS7700 data. This recommendation is made to prevent flooding the queue with LI REQ commands. Flooding the queue can lead to the other commands in the queue not being processed in a timely manner.

If more LI REQ commands are submitted than can concurrently be processed, the following message is issued:

```
IEE806A COMMANDS EXCEED LIMIT IN COMMAND CLASS M2
```

You can remove all LI REQ queued commands from the queue. For example, if thousands of LI REQ commands are queued and are preventing other commands from running, such as LOGON or START, the following command can be issued:

```
CMDS REMOVE,CLASS=M2,CMD=LI
```

If issuing from a console on a z/OS host is not wanted or available, the TS7700 Management Interface (MI) enables an operator to issue a Library Request host console command as though 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 that is 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 in multiline write to operator (WTO) message CBR1280I.

Note: The information that is presented in the CBR1280I message is generated directly by the hardware as a response to the LI REQ command. Contact TS7700 support if you have a question about the information that is presented to the host in the CBR1280I messages that are generated as a response to an LI REQ command.

This section describes a few of the more useful LI REQ commands that a client can use from the host to obtain information about a grid or volumes in the grid. For more information about the vast amount of Host Console Request functions available, see *IBM TS7700 Series z/OS Host Command Line Request User's Guide*.

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 12-10.

Example 12-10 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 must be directed. This library name can be a composite or a distributed library, and which library is applicable depends on the other keywords specified.
keyword1	Specifies the operation is to be run on the TS7700.

The following parameters are optional. These parameters depend on the first keyword that is specified and based on that first keyword, more of these keywords may be needed:

- ▶ **keyword2**
Specifies more information in support of the operation that is specified with the first keyword.
- ▶ **keyword3**
Specifies more information in support of the operation that is specified with the first keyword.
- ▶ **keyword4**
Specifies more 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.

Consider the following points:

- ▶ 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 specified.
- ▶ 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. The response is then routed back through the requesting device address.

LI REQ changes for R5.1 and later

In this section, LI REQ changes for R5.1 and later are described.

Changes for version 5.1

Version 5.1 features the following additions:

- ▶ New format of LVOL, CLDINFO to show multiple cloud settings per cluster
- ▶ New keywords:
 - LVOL, CLDVERS
 - CLDBKUP, SHOW
 - CLDSET, PFRCCPG0
 - CLDSET, PFRCTDL
 - CLDSET, PFRCTRY
 - CLDSET, USEHIGH
 - CLDSET, USELOW
 - CLDSET, USESHOW

Changes for version 5.1a

Version 5.1a features the following changes:

- ▶ Restores changes of version 5.0a that were missing in version 5.1
- ▶ Adds GGM copy and DS8000 Transparent Cloud Tiering statistics to STATUS, GRLNKACT output
- ▶ Adds expanded explanation about SETTING2,VOLINVT

Changes for version 5.1b

Version 5.1b features the following additions:

- ▶ LWORM retention fields in LVOL, INFO output
- ▶ New keywords:
 - CLDVR
 - LWORMR

Changes for version 5.11b

Version 5.11b features an updated STATUS, GRIDLINK percent retransmitted description

Changes for version 5.22

Version 5.22 features the following changes:

- ▶ New OCOPY,SUMMARY keyword added
- ▶ New OBJSET1 keyword added
- ▶ Moved RSDOHIGHIRSDOLOW to OBJSET1 from SETTING2

Changes for version 5.22a

Version 5.22a features the following change:

- ▶ Add the missing state in PVOL response (VOLUME STATE)

Changes for version 5.3

Version 5.3 features the following changes:

- ▶ Change LI REQ keyword table format
- ▶ Add new drive format modes for TS1160 support
- ▶ Add a new CRCSET keyword

- ▶ Add a note about Time Delayed Premigration configured logical volumes to the description about RSDTHIGH/LOW keyword
- ▶ Add explanations to the description about CACHE2 response
- ▶ Correct CACHE2 response format information
- ▶ Add LI REQ for zTape Air-Gap (FC 5995)

For more information about these LI REQ changes and the available Host Console Request functions, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#) and [TS7700 zTape Air-Gap FC 5995 Users Guide](#).

LIBRARY REQUEST,composite_library,STATUS,GRID

In addition to surfacing other information about the status of a grid, this command is a simple way to determine what microcode level exists for each cluster in a grid. Sample output for this command is shown in Example 12-11.

Example 12-11 LIBRARY REQUEST,libname,STATUS,GRID

```
LIBRARY REQUEST,ATL3484F,STATUS,GRID
CBR1020I Processing LIBRARY command: REQUEST,ATL3484F,STATUS,GRID.
CBR1280I Library ATL3484F request. 467
Keywords: STATUS,GRID
-----
```

GRID STATUS V3 .0							
COMPOSITE LIBRARY VIEW							
LIBRARY	STATE	IMMED-DEFERRED		OWNERSHIP-T/O		RECONCILE HCOPY	
		NUM	MB	MODE	NUM	NUM	ENB
cluster0	ON	0	0	-	0	0	Y
cluster1	ON	0	0	-	0	0	Y
cluster2	ON	0	0	-	0	0	Y

```
COMPOSITE LIBRARY VIEW
      SYNC-DEFERRED
LIBRARY  NUM    MB
cluster0  0      0
cluster1  0      0
cluster2  0      0
-----
```

DISTRIBUTED LIBRARY VIEW								
LIBRARY	STATE	RUN-COPY-QUEUE		DEF-COPY-QUEUE		LSTATE	PT	FAM
		NUM	MB	NUM	MB			
cluster0	ON	0	0	0	0	A	Y	-
cluster1	ON	0	0	0	0	A	N	-
cluster2	ON	0	0	0	0	A	N	-

```
      ACTIVE-COPIES
LIBRARY  RUN    DEF
cluster0  0      0
cluster1  0      0
cluster2  0      0
-----
```

```
LIBRARY  CODE-LEVELS
cluster0 8.52.200.109
cluster1 8.52.200.109
cluster2 8.52.200.109
-----
```

LIBRARY REQUEST,composite_library,LVOL,volser,INFO(,FLASH)

The **LIBRARY REQUEST,composite_library,LVOL,volser,INFO(,FLASH)** command allows an operator on the host to obtain the information the grid stores about a particular logical volume. This command can be used when you do not want to use the TS7700 MI to display the logical volume information.

The response lines are formatted as shown in the following example:

```
LOGICAL VOLUME INFO V3 .0
LOGICAL VOLUME           : DLE001
MEDIA, FMT, MAX(MB), CWRAP : ECST, 6,    800, N
SIZE(MB) COMP, CHAN, RATIO : 0, 0,  1.45:1(FICON)
CURRENT OWNER, TVC LIB    : cluster0, cluster0
MOUNTED LIB/DV, MNT STATE : -/-, -
CACHE PREFERENCE, CATEGORY : PG1, 115F (PRIVATE)
LAST MOUNTED (UTC)       : 2022-01-07 15:16:12
LAST MODIFIED LIB/DV, UTC(UTC) : cluster0/0001, 2022-01-07 15:16:11
KNOWN KNOWN CPYS, REQ, REMOVED   : 1, 1, 0 (N)
DEL EXP, WHEN (UTC)             : N, -
HOT, FlashCopy                 : N, NOT ACTIVE
LWORM RET STATE, TIME(UTC)      : N, NA
-----
LIBRARY RQ CA P-PVOL S-PVOL CPS CPQ CPP RM CP CD CC
cluster0 N Y ----- ----- CMP - DEF N 0 0 N
cluster1 N N ----- ----- NOR - NOC N 0 0 N
cluster2 N N ----- ----- NOR - NOC N 0 0 N
```

12.1.4 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. The exit 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 (D SMS,VOLUME(volser)).

The syntax of the **LIBRARY LMPOLICY** command to assign or change volume policy names is shown in Example 12-12.

Example 12-12 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.
- ▶ 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

12.1.5 Useful DEVSERV commands

Some of the most useful **DEVSERV** 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 is defined in the input/output definition file (IODF). With the **RDC** operand, it is useful for viewing the Composite Library ID and libport ID associated with a device.

Note: The **DS QT** command with the **QHA** operand can display which systems and SYSplexes are connected to a specific device in a cluster. For more information, see 12.9.2, "Other commands built to support CUIR functions" on page 655.

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 12-1 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
  
```

Figure 12-1 Sample DEVSERV QT command output

Clarification: The distributed library number or cluster index number for a specific logical drive can be determined by using the **DS QT** command. As shown in Figure 12-1, the response shows LIBPORT-ID 01 for logical drive 1C01. LIBPORT-ID 01 is associated with Cluster 0. For more information about the association between distributed libraries and LIBPORT-IDs, see 6.1.1, “Defining devices through HCD” on page 241.

From the **DS QT** command that is shown in Figure 12-1, you can derive the LIBRARY-ID for the composite library and the LIBPORT-ID of the logical control unit (LCU) that is presenting the logical device.

The LIBID field identifies the composite library ID associated with the device. This command cannot be used to obtain information about devices within a backend TS3500/TS4500 that is connected to a TS7700T as the host has no knowledge of the backend physical library.

Tip: You can get the device type of the backend physical drives of a distributed library from the following **LI REQ** command:

```
LI REQ,<distributed library name>,PDRIVE
```

On the z/OS host, the **QHA** (Query Host Access to volume) option can be used on the **DEVSERV QTAPE** command. This option allows the command to surface which systems and sysplexes are connected to a specific device in a cluster. For more information about the syntax and output of the **DS QT** command by using the **QHA** option, see “DEVSERV QTAPE,xxxx,QHA” on page 657.

DEVSERV QLIB,CATS command

The command **DS QLIB,CATS** allows you to view and change logical VOLSER categories without needing to initial program load (IPL) the system. Example 12-13 shows how to list all of the categories that are in use on a system.

Example 12-13 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
```

You can change the categories that are in use after you view. To perform this task, you can change the first three digits of the category. However, the last digit must remain unchanged because it represents the media type. Example 12-14 shows the command that changes all categories to 111 for the first three digits.

Example 12-14 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 are in DEVSUPxx. For more information about changing categories, see 12.6, “Effects of changing volume categories” on page 652.

DEVSERV QLIB,LIST command

Example 12-15 shows how to list all of the active libraries by using the **DS QL,LIST** command.

Example 12-15 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 more information about all of the **DS QLIB** commands, see Appendix D, “DEVSERV QLIB command” on page 919.

12.1.6 Scratch volume recovery for 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 is ENABLED for the category in the TS7700 MI (click **Virtual** → **Categories** → **Modify Category** → **Expire Hold Settings**), the volume cannot be reused as a scratch volume before the EXPIRE time is reached.

The method to recover depends on the TMS that is used. In general, change the volume's status from scratch to private and change the expiration date in the TMS by adding at least one week. Doing so prevents the TMS from returning the volume to scratch during the next few days.

There are several steps to recover a logical volume that was mistakenly scratched.

Checking a VOLSER to determine whether it was reused

The first step in determining whether a volume can be recovered is to check that the VOLSERS that you want to recover were not reused to satisfy a SCRATCH mount. Issue the **D SMS, VOL(volser)** command for each volume on which you want to check (see Example 12-16).

Example 12-16 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 *SCRTCH* HYDRAG S N N NOERROR SCRMED2
```

A 'USE ATR' of 'S' indicates that the volume is still in SCRATCH status and was not yet reused. Therefore, you can recover the volume contents if a consistent copy exists in the TS7700. If the display for this volume indicates a 'USE ATR' of 'P', the volume was 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 include a consistent copy that is in the TS7700 grid. The best command to use to check field-known consistent copies is the **LI REQ,composite-library,LVOL,volser,INFO** command, as shown in Example 12-17.

Example 12-17 Check for a consistent copy

```
LI REQ,HYDRAG,LVOL,A0000P,INFO

LOGICAL VOLUME INFO V3 .0
LOGICAL VOLUME : Z10000
MEDIA, FMT, MAX(MB), CWRAP : ECST, 6, 800, N
SIZE(MB) COMP, CHAN, RATIO : 14, 95, 6.38:1(FICON)
CURRENT OWNER, TVC LIB : Arabian, Lipizzan
MOUNTED LIB/DV, MNT STATE : -/-, -
CACHE PREFERENCE, CATEGORY : PG1, 0001 (SCRATCH)
LAST MOUNTED (UTC) : 2018-08-28 03:34:31
LAST MODIFIED LIB/DV, UTC(UTC): Lipizzan/0000, 2018-08-28 03:34:25
KNOWN KNOWN CPYS, REQ, REMOVED : 1, 1, 0 (N)
```

If 'KNOWN CPYS' is '0', you cannot recover this volume because it was delete-expired.

Changing STATUS of scratched volumes to PRIVATE

The next step in recovering the volume is to change the status from SCRATCH back to PRIVATE in the TS7700 and the TCDB. The next steps vary depending on your TMS. If your TMS is DFSMS Removable Media Manager (DFSMSrmm), you can also change the status of the volume to MASTER in the RMM CDS.

To start this process, use DFSMSrmm to search on a volume string. Then, put all of the scratch volumes that match 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 12-18.

Example 12-18 Change status

```
//STEPS 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 RMMCLIST DD features the following output:

```
READY
RMM CHANGEVOLUME A0000P STATUS(MASTER) EXPDT(14355)
```

Use the job control language (JCL) that is shown in Example 12-19 to run the previously generated CLIST. This process 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 were delete-expired.

Example 12-19 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
```

Setting the volume back to MASTER status in RMM causes the CBRUXCUA (Change Use Attribute) installation exit to be run. If the change is approved, OAM updates the volume category in the TS7700 first to match the private category that is used by this host. After this process is successful, it changes the volume from SCRATCH to Private in the TCDB. These three sources (the library, TCDB, and TMS) must feature consistent information for the status of the volume for a future mount of that volume to succeed.

The **D SMS,VOL** command can now be used to verify that the VOLSER was changed from S to P, as shown in Example 12-20.

Example 12-20 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
```

After the volume has the PRIVATE use attribute again in the TCDB, you might need to use the **LIBRARY LMPOLICY** command to reassign the appropriate Data Class, Storage Class, and Management Class constructs to the volume (because the constructs were set to blank when the volume status was changed to SCRATCH). For more information, see 12.1.4, “Library LMPOLICY command” on page 637.

12.1.7 Ejecting virtual volumes

Virtual volumes are not physical entities that can be individually removed from the library. They can also exist on physical stacked volumes that contain many other logical volumes. An EJECT should be issued only for a virtual volume in the TS7700 if the intent is to delete the volume from the TS7700.

Because of the permanent nature of EJECT, the TS7700 allows you to EJECT only a virtual volume that is in the INSERT or SCRATCH category. If a virtual volume is in any other category, the EJECT fails. If you eject a scratch volume, you cannot recover the data on that virtual volume.

Tip: If a virtual volume is in the error category (which is by default 000E), it must first be moved back to a scratch category before an EJECT can be successful. To move it, use **ISMF ALTER** to move it from the SCRATCH to SCRATCH category.

Volumes that are in the INSERT category can be directly deleted through the TS7700 MI by selecting the **Virtual Volumes** → **Delete Virtual Volumes** option. For more information, see “Delete Virtual Volumes window” on page 433.

Volumes that are in a SCRATCH category must be EJECTed from the host rather than from the TS7700 MI. They must also have the SCRATCH use attribute assigned in the TCDB and if a TMS is present, be assigned as “SCRATCH” in that TMS. These volumes can be EJECTed by using one of the following methods:

- ▶ 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
- ▶ Use ISMF to EJECT each volume.
- ▶ Use the **LIBRARY EJECT** command, as described in “MVS system commands” on page 639.
- ▶ Use the CBRXLCS macro to EJECT each volume.
- ▶ Use the CBR SPLCS SAMPLIB member to EJECT each volume.

If the volume is not in the INSERT category or is not in a SCRATCH category, the EJECT fails. If the failure is caused by the TS7700 failing the EJECT, the following message should be surfaced to indicate why the failure occurred:

```
CBR3726I Function incompatible error code 6 from library <library-name> for volume <volser>
```

If this failure occurs and the TMS now indicates that the volume is outside of the library, contact the vendor of the TMS to determine how to correct the discrepancy. The CBRUXEJC (cartridge eject) exit allows for automatic notification of the TMS when an EJECT fails; therefore, this discrepancy indicates that your TMS is not providing support for the notification function.

If an EJECT failure occurs on the host, review the surfaced messages to determine the cause of the failure and take the appropriate actions to correct it. After the cause of the failure is corrected, reissue the EJECT.

Issuing many EJECTs at the same time from a host can cause storage shortages because OAM sends only 100 requests to the library at a time and queues the remaining requests. Try to limit EJECT submission on a single host to a few thousand at a time.

The TS7700 supports a maximum of 1,000 queued EJECTs from all attached hosts. When EJECTs are issued from a host, the initial status is returned from the TS7700 to the host. This process indicates that the request was scheduled in the library.

Only after successful completion of the EJECT within the TS7700 is the host notified and the request is no longer counted in the 1,000 queued EJECTs limit for the TS7700. If a 1,001th EJECT request is issued from any connected host, the TS7700 fails the request and states too many EJECTs are queued with ERA29 and Sense Byte 8 value of X'0E'.

The following commands can be used on IBM Z hosts to list the outstanding and the active requests:

```
F OAM,QUERY,WAITING
F OAM,QUERY,ACTIVE
```

It is recommended to send no more than 1,000 EJECTs across all attached hosts to a TS7700 because only 1,000 are accepted by a TS7700. When those 1,000 EJECTs are complete, the next 1,000 can be issued. Otherwise, the requests can be rejected by the TS7700 and must be reissued.

Ejecting used virtual volumes without a TCDB entry

When a virtual volume is moved from the insert category into a scratch or private category, it can no longer be deleted from the TS7700 MI. To delete these volumes, you must verify that a TCDB volume entry exists for the volume to be deleted and verify that it is in scratch status.

After these conditions are verified, you can use the EJECT function from a host that is connected to the TS7700. Because these volumes are virtual, the virtual volumes are deleted rather than being ejected by the host EJECT function.

The following steps describe how to create the volume entry for the volume (if it is not present in the TCDB), ALTER the volume to SCRATCH status, and EJECT the volume from the host by using different methods:

1. Use the following JCL to call 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(compplibname) - MEDIATYPE(mediatype) -
          LOCATION(LIBRARY)
```

2. Use **ISMF** to **ALTER** the use attribute from PRIVATE to SCRATCH. This command starts the CBRUXCUA exit to communicate with the TMS. If the TMS allows the change, the category is changed in the library to the corresponding media type category in the host's DEVSUPxx parmlib member.

3. Use one of the following methods to EJECT (delete) the volume:
 - 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
 - Use ISMF to EJECT each volume.
 - Use the LIBRARY EJECT command, as described in “MVS system commands” on page 639.
 - Use the CBRXLCS macro to EJECT each volume.
 - Use the CBR SPLCS SAMPLIB member to EJECT each volume.

These steps can also be completed after a DR test if a DR host is shut down before deleting the logical volumes created during the DR test. These volumes continue to use space on the TS7700 until they are deleted (through EJECT).

12.2 Messages from the library

This section describes TS7700 enhanced message support and relevant messages.

12.2.1 CBR3750I console message

When the host receives a message from the library that is informational or indicates an abnormal condition of some type, the host surfaces this *message* (by using OAM) within the CBR3750I message. This message uses the following format:

```
CBR3750I Message from library library-name: message.
```

When CBR3750I is issued, this indicates that a *message* was sent from library *library-name*. The operator at the library manager console entered a message that is to be broadcast to the host, or the library broadcast a message to the host to relay status information or report an error condition. For more information about the messages that can be broadcast from the library to the host, see [IBM TS7700 Series Notification Messages](#).

If more assistance is required regarding the contents of the messages that surfaced, contact TS7700 support. The only role that OAM provides is to surface on the host the library-generated *message* through the CBR3750I message.

With OAM APAR OA52376 and R4.1.2 installed on all the clusters in a grid, this message was enhanced to include the severity impact text (INFORMATION, WARNING, IMPACT, SERIOUS, or CRITICAL) that is associated with each library message.

This text prevents you from having to reference the white paper to determine the severity that is associated with a particular library message. It also results in the format for the CBR3750I message being updated as shown in the following example:

```
CBR3750I Message from library library-name: message. Severity impact: severity impact text.
```

You can modify the severity impact text for each message through the TS7700 MI if you feel a different severity impact text is more applicable in your environment. If you choose to modify the text, an “*” is added to the severity impact text as shown in the following example:

```
CBR3750I Message from library library-name: message. Severity impact: severity impact text*.
```

In addition, each library message can now be associated with a free-form customer impact text through the TS7700 MI. You can use this function to surface more information that might be relevant for a specific library message in your environment. If a customer impact text is defined to a library message, the associated CBR3750I message is surfaced as shown in the following example:

```
CBR3750I Message from library library-name: message. Severity impact: severity impact text. Customer impact: customer-provided-impact-text.
```

12.2.2 TS7700 Host Console messages

Some abnormal conditions that occur in a TS7700 are not surfaced through the CBR3750I message. Rather, they are surfaced through their own unique message number in OAM. Some of the more common messages that you might see are described in this section. For more information about the current list, see the appropriate release of *z/OS MVS System Messages, Vol 4 (CBD-DMO)*.

Incompatibility error message

In an incompatible function error, you might see the following message CBR3726I:

```
CBR3726I Function incompatible error code error-code from library library-name for volume volser.
```

In this message, an error occurred during the processing of volume *volser* in the library *library-name*. The library returned a unit check with an error code *error-code*, which indicates that an incompatible function was requested. A command was 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 at this IBM Documentation [web page](#). At this page, scroll down in the table to find the **Function Incompatible** entry.

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, the following message is generated:

```
CBR3768I VTS operations in library library-name no longer degraded
```

Warning cache use capacity (TS7700D or TS7700T)

For a TS7700D or TS7700T, the following warning and critical cache free space messages are surfaced:

```
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 surfaced:

```
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 TS7700T 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 no volumes 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 that are attached to the cluster that is 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 was enabled to borrow from the common scratch pool and empty physical volumes are available 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 specific information about the pool that is out of empty physical volumes:

```
CBR3750I Message from library library-name: OP0138 The Common Scratch Pool (Pool 00) is out of empty media volumes. Severity impact: CRITICAL.  
CBR3750I Message from library library-name: OP0139 Storage pool xx is out of scratch volumes. Severity impact: CRITICAL.
```

The OP0138 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 TS7700T is in a grid configuration, and its associated distributed library enters the out-of-empty-stacked-volume state, operations are affected in the following ways:

- ▶ All copy operations are immediately suspended in the cluster (regardless of which pool has become empty).
- ▶ If the cluster features a Copy Consistency Point of RUN, the grid enters the Immediate Mode Copy Operations Deferred state. An MVS console message is generated:

```
CBR3787E One or more immediate mode copy operations deferred in library library-name.
```
- ▶ The copy attempt fails if another cluster attempts to copy a logical volume that is not resident in the cache.
- ▶ 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 that is needed is not in the remote cluster's TVC, the recall of the data fails. If data is written to the remote cluster's TVC, the writes are allowed.

However, because empty physical volumes might not be 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, expedite the reclamation of physical stacked volumes or add 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 include 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 TS7700T or TS7700C 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 TS7700T or TS7700C 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 are to prevent excessive messages from being created.

This state surfaces the following message:

```
CBR3750I Message from library library-name:OP0160 Above threshold for uncopied data in cache, throttling possible. Severity impact: SERIOUS.
```

When the condition is no longer present, the following message is surfaced:

```
CBR3750I Message from library library-name:OP0161 Below threshold for uncopied data in cache. Severity impact: INFORMATION.
```

12.3 Expire Hold and scratch processing considerations

This section describes the interaction between SCRATCH processing and the Expire Hold setting for scratch categories on the TS7700. For more information about the Expire and Expire Hold options for a category and the conditions that must be met for a volume that is assigned to the category to be delete-expired, see "Defining scratch categories" on page 590.

When Expire Hold is set for a scratch category, it can heavily affect scratch volume availability. Consider the following cases:

- ▶ Expire Hold option is enabled and the TS7700 is low on scratch volumes
- ▶ Expire Hold option is enabled and Cache Utilization is higher than wanted in the TS7700D.

12.3.1 Expire Hold and low on scratch volumes

The Expire Hold option prevents a recently scratched volume from being reused or deleted until the specified EXPIRE time has passed. This may cause the TS7700 to run out of available SCRATCH volumes to satisfy mount requests to write new data.

If the EXPIRE time is set too long and not enough virtual volumes are released to keep a healthy level of SCRATCH volumes available, it might be necessary to reduce the EXPIRE time.

Changing the EXPIRE time affects only new volumes going into the SCRATCH pool. The existing volumes in the pool continue to be held until the original EXPIRE time passes. If Expire Hold is removed from the scratch category, these volumes can then be added to the candidate list for SCRATCH mounts. Removing the Expire Hold option immediately helps to alleviate the low on scratch condition, but no longer protects volumes that were inadvertently sent to SCRATCH. This can compromise the ability to recover user data on volumes in the SCRATCH pool,

12.3.2 Expire Hold and cache utilization in a TS7700D

When the Expire Hold option is enabled, the cache in a TS7700 is used up in part by holding data from virtual volumes that were sent to the SCRATCH pool. This creates the risk of running out of cache in a TS7700D when AUTOREMOVAL is not enabled.

If a TS7700D is consistently running high on cache usage, consider adjusting the EXPIRE time. Removing the Expire Hold setting does not immediately alleviate the high cache usage condition. This allows only volumes in the SCRATCH pool to be added to the candidate list for delete expire processing

A task runs in the library once per hour that processes some number of volumes from this list and reduces cache usage by deleting the expired volumes from cache. The maximum number of volumes that are deleted per hour by default is 1000 but customizable (up to 5000), and is controlled by using the following command:

```
LI REQ,composite-library,SETTING,DELEXP,COUNT,value
```

For more information about this command, see 12.1.3, “Host Console Request function” on page 632.

For more information, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

The EXPIRE time is the grace period that enables the recovery of the data if a procedural error occurs. Careful consideration must 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.

12.4 Scratch count mismatch

Some discrepancy often exists between the values that are reported for scratch counts from various sources (TS7700 MI, TCDB, TMS). To obtain the number of usable scratch volumes in the TS7700, issue the **D SMS, LIBRARY(libname),DETAIL** command. In addition to providing other information about the designated tape library, this command reports the number of usable scratch volumes. This value is obtained directly from the TS7700 and surfaced within the CBR1110I message.

When Expire Hold is in effect, the total number of scratch volumes in the TCDB for the composite library can differ from the total number of usable SCRATCH volumes because some volumes have not expired due to the EXPIRE time. For this reason, the most accurate source of scratch counts for a TS7700 is always the output from the **D SMS, LIBRARY(libname),DETAIL** command.

12.5 Host cartridge entry processing

To ensure that the virtual volumes are assigned to the correct library manager categories during cartridge entry, it is important to understand the host cartridge entry process and the components and interactions that are involved.

When a set of volumes is inserted (defined) to a stand-alone TS7700 or to a grid, the hardware first assigns each of these volumes to the insert category, X'FF00'. After this process, an attention is sent to all z/OS hosts attached to the TS7700s in the grid. This notifies them that there are volumes in the insert category that are ready to be processed by host cartridge entry processing.

The first host that responds to this attention has the first opportunity to run host cartridge entry processing against the volumes in the insert category. No method is available to control which host receives this attention first. For this reason, it is important to correctly configure your hosts to control which hosts can accept which volumes during host cartridge entry processing.

If you do not want a host to be eligible to respond to this attention, you can disable cartridge entry processing on that host by using the **LI DISABLE, CBRUXENT** command. When you want to re-enable cartridge entry on that host (allow the host to again receive this attention), issue the **LI RESET, CBRUXENT** command.

If the CBRUXENT cartridge entry exit is disabled on a host, that host does not respond to any attention that volumes are in the insert category in any tape library that is ready to be processed by host cartridge entry processing. If CBRUXENT is enabled on a host, the host responds to the attention and runs host cartridge entry processing.

The host requests a list of the volumes in the insert category and then passes the first VOLSER on the list to the CBRUXENT exit through OAM.

The return codes that are set by the CBRUXENT exit determine whether this host accepts the entry of this volume or moves on to the next volume in the list. By default, RMM provides its version of the CBRUXENT exit that calls RMM to approve or deny the entry of a volume. If RMM is used, see 12.5.1, "Removable Media Manager cartridge entry considerations" on page 651 before continuing with this section.

If a host uses an alternative tape management system (TMS) there should be a custom version of the CBRUXENT exit that calls the TMS on that host to approve or deny the entry of a volume, with the result being passed back to CBRUXENT.

In either condition, it is important that the TMS on each host that is connected to the TS7700 approves requests only for the volumes that are intended to belong to that host. Otherwise, an unintended library manager category can be assigned to the volume. If you are using a non-RMM TMS, contact your vendor to determine any considerations for that TMS during host cartridge entry processing.

If the CBRUXENT exit indicates that the volume entry is approved, OAM sets the library manager category of the volume to match what is specified in the DEVSUPxx parmlib member on that host for that media type (MEDIA1 or MEDIA2 for TS7700). OAM then creates the volume entry in the TCDB for the volume (or updates the entry if it is already present). After this process, host cartridge entry processing then moves on to the next volume in the list.

When one or more volumes is successfully inserted, the volumes are surfaced within a CBR3610I cartridge entry success message output to the console and SYSLOG. The output destination for success messages (CBR3610I) and volume ignore messages (CBR3621I) can be modified through the SETTLIB parameter in the CBROAMxx parmlib member.

For more information, see *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

If any volumes remain in the insert category after the host completes one iteration through the list, the next host that responded to the attention from TS7700 then can process host cartridge entry against those volumes that are still in the insert category. OAM sends a request to the library to ask for the volumes that exist in the insert category and this process repeats.

The category assigned to a volume is important during scratch mount processing because the TS7700 uses it to decide which volumes are eligible to satisfy a scratch mount request. For example, volume SAMP01 is assigned a category of x'0102' during host cartridge entry processing. Later a job, on the same host, requests a scratch mount of a MEDIA2 volume. During this mount request, the host passes the category that is assigned to MEDIA2 in the DEVSUPxx parmlib member (x'0102' in this example) to the TS7700. The TS7700 selects a volume in the x'0102' category (which might be SAMP01 if other volumes are assigned the x'0102' category as well) and passes the volser back to the host to satisfy the mount request.

12.5.1 Removable Media Manager cartridge entry considerations

Removable Media Manager (RMM) is started during the host cartridge entry process by way of the supplied CBRUXENT exit. When started, RMM must decide whether it wants to approve the cartridge entry request or not.

In many RMM environments, it is common that a REJECT ANYUSE(*) rule or PRTITION VOLUME(*) TYPE(NORMM) SMT(IGNORE) NOSMT(IGNORE) rule is in effect. If either of these rules are in effect, volumes *must* first be defined to RMM before the entry request for that volume is approved by RMM during host cartridge entry. Otherwise, RMM instructs the CBRUXENT exit to ignore the volume and the following message is surfaced:

```
CBR3621I Enter request ignored by the cartridge entry installation exit
(CBRUXENT).
```

The volume remains assigned to the insert category (x'FF00') and an attempt to process the next volume (if any is present) is made by host cartridge entry processing.

An RMM ADDVOLUME command must be used to define the volume to RMM before RMM can approve the entry of the volume during host cartridge entry processing. The following JCL provides an example:

```
//TSOBATCH EXEC PGM=IKJEFT01
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
RMM AV A00000 COUNT(100) INIT(N) STATUS(SCRATCH) -
VOLUMETYPE(LOGICAL) MEDIATYPE(ECCST) LOCATION(compositelibname)
/*
```

If the RMM environment does *not* have a REJECT ANYUSE(*) rule or PRTITION VOLUME(*) TYPE(NORMM) SMT(IGNORE) NOSMT(IGNORE) rule in effect, RMM approves the insertion of the volume and automatically creates the volume entry in RMM.

12.6 Effects of changing volume categories

The host has no knowledge of the category of a specific volume. However, this information can be surfaced on the host by using the `D SMS,VOLUME(volser)` command. The output of this command includes the LM CATEGORY field that surfaces the category that is assigned to that volume in the library manager. It should be noted that these categories are only used for scratch mounts.

The categories that are assigned by the host can be changed dynamically by using the `DEVSERV QLIB,CATS` command or by modifying the values specified in the DEVSUPxx member and IPLing.

Special consideration should be given to the effects of changing the categories in use on a host. The most common problem is that, after the change, users might forget that all of the virtual volumes in the scratch pool still belong to the initially defined categories. If the categories defined in the DEVSUPxx parmlib member on the host processing the mount do not match the categories that are assigned to the volumes you expect to be mounted to satisfy a scratch mount request, those volumes are not selected to satisfy the scratch mount request.

For example, if MEDIA2 volume SAMP01 is assigned to category x'0102' and the DEVSUPxx member on the host is changed to specify category x'0112' for MEDIA2 and the host restarted, SAMP01 is no longer eligible to be selected by the TS7700 to satisfy a scratch mount request from this host because the category that is assigned to SAMP01 (x'0102') does not match the category that is assigned to MEDIA2 in the DEVSUPxx parmlib member (x'0112').

If no volumes are assigned the new categories, requests for scratch mounts fail and OAM surfaces the CBR4105I and CBR4196D messages.

This issue can be resolved by using several methods. If the categories were changed because you want to partition the library, a scratch pool must be created for this host by adding a range of volumes that are accepted by the TMS during subsequent cartridge entry processing.

If the old scratch pool was intended to be used by this host, the category of a single volume or range of volumes can be updated by using the ISMF windows to ALTER one or the use attribute from SCRATCH to SCRATCH. This step sets the category of the volume in the library manager to match the associated newly defined category on the host (in DEVSUPxx).

If a large range of volumes must be changed, consider the use of the CBRSP LCS utility to make 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.

Note: Modifying the volume entries in the TCDB by using IDCAMS does *not* change the categories that are stored in the library manager and should not be used for this purpose.

12.7 Library messages and automation

Some messages can be useful to capture by automation. In particular, the following messages 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 restart 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 TS7700 support to determine why Call Home was being attempted. For more information about a full list of the conditions that are surfaced within the CBR3750I message, see [IBM TS7700 Series Notification Messages](#).

12.8 Mount retry

When a mount fails, specific error conditions allow for the mount to be retried. If they can be retried, the following message is issued:

```
CBR4196D Job job-name, drive device-number, volser volser, error code error-code.  
{Reply 'R' to retry or 'C' to cancel. | Reply 'R' to retry, 'W' to wait or 'C' to  
cancel.}
```

The CBR4196D message is issued along with messages describing the error condition, as shown in the following example:

```
CBR4195I LACS retry possible for job TEST01: 015  
IEE763I NAME= CBRLLACS CODE= 140169  
CBR4000I LACS MOUNT permanent error for drive OBCA.  
CBR4105I No MEDIA2 scratch volumes available in library TESTLIB.  
IEE764I END OF CBR4195I RELATED MESSAGES  
*06 CBR4196D Job TEST01, drive OBCA, volser SCRTCH, error code 140169.  
Reply 'R' to retry or 'C' to cancel.
```

OAM APAR OA52376 allows OAM to automate the mount's retry function. This function is specified by using the **SETTLIB** command in the CBROAMxx startup parmlib member. It is controlled by the following parameters:

- ▶ LACSRETRYTIMES(1-9): Specifies the number of times to automatically retry the mount
- ▶ LACSRETRYMINUTES(1-9): Specifies the number of minutes between each automatic retry
- ▶ LACSRETRYFAIL(YES|NO): Specifies whether the mount is failed or the CBR4196D surfaced

The full syntax of the command that can be specified as shown in the following example:

```
SETTLIB LACSRETRYTIMES(X) LACSRETRYMINUTES(Y) LACSRETRYFAIL(YES|NO)
```

You can use different combinations of these three parameters to control how OAM automatically responds to retry-able mount failures. The following OAM responses are possible:

- ▶ Retry the mount a specified number of times (LACSRETRYTIMES), every number of minutes LACSRETRYMINUTES. If the mount is not yet successful at the end of the retry processing, surface the CBR4196D (LACSRETRYFAIL(NO)).
- ▶ Retry the mount a specified number of times (LACSRETRYTIMES), every number of minutes LACSRETRYMINUTES. If the mount is not yet successful at the end of the retry processing, fail the mount request (LACSRETRYFAIL(YES)). No CBR4196D is issued.
- ▶ Fail the mount immediately (LACSRETRYFAIL(YES)). No CBR4196D is issued.
- ▶ Issue the CBR4196D as is done today. LACSRETRYFAIL(NO) or no parameters specified.

12.8.1 Enhanced mount retry defaults

The only default is LACSRETRYFAIL(NO), which is set regardless of whether any of these parameters are defined in the CBROAMxx parmlib member. The LACSRETRYTIMES and LACSRETRYMINUTES parameter values receive only default values if one of these parameters is specified without the other. For example, if only LACSRETRYMINUTES is specified, LACSRETRYTIMES is set to 6. Alternatively, if only LACSRETRYTIMES is specified, LACSRETRYMINUTES is set to 2. If you are using this function, specify LACSRETRYTIMES *and* LACSRETRYMINUTES to avoid confusion.

12.8.2 Enhanced mount retry example

The following message is displayed (as an example) if the values that are specified are LACSRETRYTIMES(6) LACSRETRYMINUTES(2) and either LACSRETRYFAIL(YES/NO):

```
CBR4195I LACS retry possible for job TEST01: 015
IEE763I NAME=CBRLACS CODE=140169
CBR4000I LACS MOUNT permanent error for drive OBCA
CBR4105I No MEDIA2 scratch volumes available in library TESTLIB.
IEE764I END OF CBR4195I RELATED MESSAGES
CBR4197D Job TEST01, drive OBCA, volser SCRTCH, code 140169.
Retrying every 2 minutes, 6 times. Reply 'C' to cancel.
```

As shown in this example, a new message (CBR4197D) is surfaced allowing you to cancel the automatic retry. If the mount is canceled, the LACSRETRYFAIL specification is not used and the mount is failed immediately.

If the mount is not satisfied after the six retries every two minutes (for a total of 12 minutes of retry processing), the following message is issued if LACSRETRYFAIL is set to YES:

```
IEE763I NAME=CBRLACS CODE= 140169
CBR4000I LACS MOUNT permanent error for drive OBCA.
CBR4105I No MEDIA2 scratch volumes available in library TVL10001.
IEE764I END OF IEC147I RELATED MESSAGES
```

These messages are the same set of messages that are issued if the mount is not considered to be retry-able, or if the CBR4196D message is responded to with "C". Otherwise, if LACSRETRYFAIL is set to NO, the CBR4196D message is issued.

For more information about this enhancement, see OA52376.

12.9 CUIR for tape

Before a TS7700 cluster can be placed into service, in addition to other manual tasks, a customer must vary offline the devices in that cluster to each host that is connected to that cluster. Because this process can be time-consuming, R4.1.2 introduced a feature that is called *Control Unit Initiated Reconfiguration* (CUIR) that helps with this process. The CUIR feature was enhanced with R5.2.1 to include the automatic varying offline and online of devices in unhealthy clusters.

When R5.2.1 is installed on each cluster in a grid, CUIR can detect that a cluster in the grid is entering service or is unhealthy (fenced). Each host connected to that cluster is signaled to vary offline any devices that the host has online to that cluster.

If the z/OS host has OA52376 installed (OA60929 must also be installed to include the R5.2.1 enhancement), it receives this signal and attempts to vary offline the devices in the cluster. Any long-running jobs must be canceled or swapped to another cluster to ensure that the device moves from “pending offline” to “offline”.

When a device is varied offline by this CUIR function, it is known as being “offline for CUIR reasons”. After the cluster is no longer in service mode or is no longer unhealthy, it can notify each connected host that the status changed and each host can vary those devices that were “offline for CUIR reasons”, back online.

12.9.1 LIBRARY REQUEST commands to enable or disable CUIR

Signaling the host when the cluster is entering service mode or becomes unhealthy is enabled or disabled by using the **LIBRARY REQUEST** command with the following syntax:

```
LIBRARY REQUEST,library-name,CUIR,SETTING,{SERVICE|FENCE|ALL},{ENABLE|DISABLE}
```

Signaling the host when the cluster is no longer in service mode or is no longer unhealthy is enabled or disabled by using the **LIBRARY REQUEST** command with the following syntax:

```
LIBRARY REQUEST,library-name,CUIR,AONLINE,{SERVICE|FENCE|ALL},{ENABLE|DISABLE}
```

In each command, *library-name* is the name of the composite library as defined in SMS on the z/OS host.

12.9.2 Other commands built to support CUIR functions

The following commands were added to support the CUIR functions that were introduced in 4.1.2:

```
LIBRARY REQUEST,library-name,LDRIVE  
LIBRARY REQUEST,library-name,LDRIVE,group,index  
DS QT,xxx,QHA  
LIBRARY DISPDRV,library-name  
VARY device-number,ONLINE,RESET
```

This section discusses each of these commands and provides the syntax and output for each command.

LIBRARY REQUEST,*library-name*,LDRIVE

This keyword allows logical drive information (including CUIR-specific information) associated with a cluster (distributed library) or grid (composite library) to be surfaced. It can be used to better understand where manual vary commands might need to be issued to track the status of a CUIR request from a TS7700, and to verify the functions enabled on the cluster or grid. This command has the following syntax:

```
LIBRARY REQUEST,library-name,LDRIVE
```

The output from the command for a composite library is:

```
Logical Drive Status Information V1 .0
Composite Library View
Current Time (UTC): yyyy-mm-dd hh:mm:ss
Service Vary: Enabled|Disabled Auto Online: Enabled|Disabled
Unhealthy Vary: Enabled|Disabled Auto Online: Enabled|Disabled
XXXXX (CL0)
Assigned/Grouped/Total Devices: nnnn/nnnn/nnnn
Assigned/Grouped/Total LPARS: nnnn/nnnn/nnnn SSV: nnnn SUV: nnnn
Active Service Vary: Y|N
Active Unhealthy Vary: Y|N
XXXXX (CL1)
Assigned/Grouped/Total Devices: nnnn/nnnn/nnnn
Assigned/Grouped/Total LPARS: nnnn/nnnn/nnnn SSV: nnnn SUV: nnnn
Active Service Vary: Y|N
Active Unhealthy Vary: Y|N
...
XXXXX (CL7)
Assigned/Grouped/Total Devices: nnnn/nnnn/nnnn
Assigned/Grouped/Total LPARS: nnnn/nnnn/nnnn SSV: nnnn SUV: nnnn
Active Service Vary: Y|N
Active Unhealthy Vary: Y|N
```

The output for the command against a distributed library is:

```
Logical Drive Status Information V1 .0
Distributed Library View
Current Time (UTC): yyyy-mm-dd hh:mm:ss
XXXXX (CLx)
Service Vary: Enabled|Disabled Auto Online: Enabled|Disabled
Unhealthy Vary: Enabled|Disabled Auto Online: Enabled|Disabled
Assigned/Grouped/Total Devices: nnnn/nnnn/nnnn
Assigned/Grouped/Total LPARS: nnnn/nnnn/nnnn SSV: nnnn SUV: nnnn
Active Service Vary: Y|N
Active Unhealthy Vary: Y|N <cluster index list>
Active Path Group Indexes
list of indexes 0 - 4095 separated by a single space (multiple rows possible)
```

The keywords for this command when issued against a distributed library are the same as when issued against a composite library, with the addition of the Active Path Group Indexes, which is the list of Active PGIDs.

For a full description of the surfaced output, see [IBM TS7700 Series Control Unit Initiated Reconfiguration \(CUIR\) User's Guide](#).

LIBRARY REQUEST,*library-name*,LDRIVE,GROUP,*index*

When the GROUP parameter is specified (by using an associated index from the distributed library output), more information can be obtained about an index (LPAR) that is attached to that distributed library. This version of the command has the following syntax:

```
LIBRARY REQUEST,library-name,LDRIVE,GROUP,index
```

The output from this command is:

```
Logical Drive Path Group ID Information V1. 0
Distributed Library View
Current Time (UTC): yyyy-mm-dd hh:mm:ss
XXXXX (CLx)
Path Group ID Index: nnnn
SSV: YES|NO SUV: YES|NO
Path Group ID: xxxxxxxxxxxxxxxxxxxxxxxx
CSSID: xx LPARID: xx CPU SERIAL#: xxxx CPU Type: xxxx GUEST: YES|NO
System Name: {xxxxxxxx|UNKNOWN} Sysplex Name: {xxxxxxxx|UNKNOWN}
WWNN(CEC): xxxxxxxxxxxxxxxxxxxx Channel ID: xx
FICON Ports: FxPy FxPy FxPy FxPy FxPy FxPy FxPy FxPy
Paths [nnnn] nnnn nnnn nnnn ...
Assigned/Grouped/Total Devices: nnnn/nnnn/nnnn
Assigned & Grouped Devices
list of devices in assigned and grouped state
Grouped Only Devices
list of devices still in a grouped state (online or pending offline)
```

For a full description of the surfaced output, see [IBM TS7700 Series Control Unit Initiated Reconfiguration \(CUIR\) User's Guide](#).

DEVSERV QTAPE,*xxxx*,QHA

This option allows the **DEVSERV QTAPE** command to surface which systems and SYSplexes are connected to a specific device in a cluster. The output from the command is:

```
DS QT,xxxx,QHA
```

The output for the command features the following syntax:

```
DEVSERV QTAPE,xxxx,QHA
IEE459I hh:mm:ss DEVSERV QTAPE
UNIT DTYPE DSTATUS CUTAPE DEVTYPE CU-SERIAL DEV-SERIAL ACL LIBID
xxxx aaaaa bbbbbb cccccc dddddd eeee-eeee ffff-ffff g hhhh
  QUERY HOST ACCESS TO DEVICE
PATH-GROUP-ID          FLAGS STATUS SYSTEM SYSPLEX
iiiiiiiiiiiiiiiiiiii jjjj kkk  sysname plexname
```

The FLAGS field contains information about how the host configured the device and whether the host supports CUIR.

For more information, see [IBM TS7700 Series Control Unit Initiated Reconfiguration \(CUIR\) User's Guide](#).

LIBRARY DISPDRV,library-name

The **LIBRARY DISPDRV,library-name** command is updated with a new field, **CU**, that was added to account for the new offline reason code for CUIR. The output from the command is:

CBR1220I Tape drive status:

DRIVE	DEVICE	LIBRARY	ON	OFFREASN	LM	ICL	ICL	MOUNT
NUM	TYPE	NAME		LI OP PT CU	AV	CATEGRY	LOAD	VOLUME
devnum	devtyp	libname	b	c d e f g		hhhhhhh	i	mntvol

If a device is offline for CUIR reasons, the field contains a 'Y'; otherwise, it contains an 'N'.

VARY device-number,ONLINE,RESET

Occasionally, a device can be stuck offline for CUIR reasons. This situation can occur if, for example, the cluster that varied the device offline for CUIR reasons is no longer attached to the host. If the cluster is no longer attached, it cannot send the signal to the host to vary the device online. If this situation occurs, you must use this command to bring online the device that is being kept offline for CUIR reasons.

12.10 Cloud Storage tier considerations (R4.2 enhancement)

With R4.2, the TS7700 can store logical volumes on an attached Cloud Storage Tier. Although no host support is needed to use this new Cloud Storage Tier function, applying OA55481 to each attached z/OS host is recommended.

For more information about this APAR and host considerations that involve this new function, see *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-55733.

One feature of OA55481 allows for OAM to detect if a copy of a volume exists in a cloud storage tier and if it does, surfaces this information in new text in the output from the **D SMS,VOLUME(volser)** command. This function is enhanced by OA59161, which documents the bit representing this status in the CBRTVI mapping and allows a user application to query this information with the CBRXLCS FUNC=QVR macro.

For more information, see this [IBM Support web page](#).

12.11 Return-to-scratch enhancement (OA48240)

When a volume is returned to scratch, the following I/O calls are made to the library for each volume:

- ▶ Move the volume from private to scratch
- ▶ Obtain the number of volumes in the scratch category

APAR OA48240 (z/OS V1R13+ and included in base z/OS V2R3) can eliminate the second I/O call. After the APAR is applied, this change can be done by using the **LIBRARY DISABLE,CATCOUNT** command. This change can decrease the overall duration of return-to-scratch processing. Reenabling the second I/O call can be done by using the **LIBRARY RESET,CATCOUNT** command.

When CATCOUNT is disabled, OAM continues to update the current count of scratch tapes in the library by using a monitoring task in the OAM address space. This task queries the library for the current scratch count every 10 minutes. In addition, OAM continues to immediately update the scratch count when a volume is changed from scratch to private.

12.12 OAM Object SYSZTIOT enhancement

With OA56836 (available on z/OS 2.2 and later), OAM Object added the ability to use XTIOOT support. OAM's use of XTIOOT reduces SYSZTIOT contention in the OAM address space during high levels of object store or retrieve activity to and from tape libraries. This function is enabled by specifying the NON_VSAM_XTIOOT=YES statement in the DEVSUPxx parmlib member that is used for the system.

Note: If NON_VSAM_XTIOOT is set to NO on your system, review the implications of setting the parameter to YES on your system before doing so as it can affect applications other than OAM. For more information, see the Coding the DEVSUPxx option in *z/OS DFSMS Using Datasets*, SC23-6855.

If a non-IBM tape vendor is used, verify with the vendor that their software supports XTIOOTs before enabling this function. When NON_VSAM_XTIOOT=YES is specified in the DEVSUPxx parmlib member, OAM always uses the XTIOOT function. Because this configuration might not always be wanted, OAM's support was enhanced by way of OA58048 to allow OAM's use of the XTIOOT function to be controlled by a new CBROAMxx parmlib member option, as shown in the following example:

```
SETOAM SYSZTIOTUSEXTIOOT(ENABLED|DISABLED)
```

With OA58048 applied, NON_VSAM_XTIOOT=YES must be set in the DEVSUPxx parmlib member and SETOAM SYSZTIOTUSEXTIOOT(DISABLED) must *not* be specified in the CBROAMxx parmlib member for OAM to use its XTIOOT support. By default, the SETOAM SYSZTIOTUSEXTIOOT parameter is enabled.

12.13 Enhanced SSMHONOR support

During normal device allocation for an SMS managed tape mount request, every device in the Tape Storage Group that is compatible with the media type and recording technology passed on the request is eligible for selection. In some circumstances, such as during problem diagnosis, it is useful to select one specific device, or one of a subset of devices that belong to the Tape Storage Group.

This type of device selection is achieved by using SSMHONOR support. This support is enabled by specifying the SSMHONOR keyword in the JCL UNIT parameter along with either one or more specified devices or an esoteric. It allows a user to direct an SMS-managed allocation to one of the specified devices or a device in the esoteric (if that device or esoteric is a subgroup of the eligible devices that SMS returned for the request).

This support can also be requested by an application by using the Dynamic Allocation (SVC 99) text unit (DALSMHR).

OA59161 (for z/OS 2.3 and 2.4) extends this support by allowing a device or esoteric to be associated with a Tape Storage Group. This is done by specifying the "SSMHONOR=" string within the description field of the Tape Storage Group, followed by the required device or esoteric.

Once specified and an SMS-managed allocation request uses the Tape Storage Group, Allocation selects the device (or a device within the esoteric) that is defined by the "SSMHONOR=" string to satisfy the request (if that device or esoteric is a subgroup of the eligible devices that SMS returned for the request).

Enabling this support through SMS policies enables customers to use their ACS routines (and the assignment of Tape Storage Groups) to the direct certain workloads to a limited set of devices. This ability leaves the full set of devices to other (more critical) workloads, which provides a device fencing capability.

For more information about the syntax and uses of the enhanced SMSHONOR function for scratch and specific allocations, see this [IBM Support web page](#).

12.14 DFSMSHsm RECYCLE Enhancement for TS7700C

With the introduction of the TS7700C, data that is written to a TS7700C by DFSMSHsm can be offloaded to an attached cloud store. If an operator runs the **DFSMSHsm RECYCLE** command to consolidate virtual tapes and one or more of those tapes is in the attached cloud store, the TS7700 must recall the entire tape to the TS7700 cache before it can be processed by the RECYCLE function. As external cloud providers typically charge to recall data from the cloud, recalling DFSMSHsm tapes back into the TS7700 cache for RECYCLE processing can incur cost.

To address this issue and provide an option to reduce the number of DFSMSHsm-owned tapes that are stored in an attached cloud store from being selected during a RECYCLE function, OA59919 (for z/OS 2.3 and 2.4) was built. OA59919 (for z/OS 2.3 and 2.4) provides the ability for an operator to further specify whether a RECYCLE process chooses all eligible volumes in the TS7700, only those in TS7700 cache, or only those that are not in TS7700 cache.

The APAR provides the RECYCLETAPERESIDENCE(BOTH | INCACHE | NOTINCACHE) keywords that can be added by using SETSYS.

TAPERESIDENCE(BOTH | INCACHE | NOTINCACHE) can also be specified on the RECYCLE command within the SELECT parameter.

For more information about the full syntax, see this [IBM Support web page](#).

12.15 DFSMSHsm RECYCLE Considerations when using zEDC

If available on a z/OS host, DFSMSHsm recommends the use of zEDC compression because it significantly reduces the data that DFSMSHsm processes and writes to tape volumes. This setting can be toggled for migration and backup processing through the SETSYS **ZCOMPRESS(ALL | NONE)** command and for dump processing through the **ZCOMPRESS(YES | NO)** parameter on the **DEFINE DUMPCLASS** command in the ARCCMDxx parmlib member. For example:

```
SETSYS ZCOMPRESS(ALL)
DEFINE DUMPCLASS(MYCLASS ZCOMPRESS(YES))
```

When the ZCOMPRESS function is used, DFSMSHsm recommends ensuring that the data class that is selected for the tape data set specifies **COMPACTION = N**. This specification prevents the hardware from attempting to further compress the compressed data.

After DFSMSHsm-owned data is written to tape by using zEDC compression, a **RECYCLE** command that is issued later might have the potential to select old, non-EDC compressed tape volumes and the new, zEDC compressed tape volumes.

If non-EDC compressed tape volumes can be processed during RECYCLE processing, DFSMSshm recommends ensuring that a different data class is assigned during that processing (one that specifies COMPACTION = Y) than during the original migration or backup. Assigning a data class with COMPACTION = Y ensures that old tape data sets that were compressed by using tape hardware compression during migration or backup are compressed again during RECYCLE processing. If a data class is selected with COMPACTION = N, any previously hardware-compressed data is written decompressed during RECYCLE processing.

As the data class that is selected during RECYCLE specifies COMPACTION=Y, the zEDC compressed tape data sets also attempt to be compressed by tape hardware compression. However, this “double compression” of zEDC compressed tape data sets occurs during RECYCLE processing only.

This recommendation is accomplished by ensuring that a different esoteric is defined for RECYCLE than for migration and backup processing; for example:

```
SETSYS UNITNAME(esoteric1)
SETSYS MIGUNITNAME(esoteric1)
SETSYS RECYCLEOUTPUT(BACKUP(esoteric2) MIGRATION(esoteric2))
```

The data class ACS routine must be coded to select the data class based on these esoterics, which are passed in the &UNIT parameter to the data class ACS routine.

12.16 LWORM retention changes

When the new LWORM retention function that is included in R5.2 is enabled, the TS7700 settings for LWORM retention can be set in the TS7700 data class to prevent premature return to scratch processing. This support is used when specific commands are issued to return a volume to scratch or where normal expiration criteria are reached in the defined rules of the tape management system and the TMS attempts to return the volume to scratch.

If the conditions that are specified in the TS7700 settings for LWORM retention are not met and DFSMSrmm is in use, any call to move the volume to SCRATCH fails with the following message:

```
EDG2433I ERROR REQUESTING LCS FUNCTION CUA RETURN CODE 12 REASON CODE 0312.
```

In addition, the volume is left in MASTER/PRIVATE status in the host and in the private category (xxxF) in the TS7700.

Return code 12 with reason code 312 is defined as “Requested function is incompatible with the library”. It is set by OAM when Error Recovery Action (ERA) 29 is returned by the TS7700 back to the host directly when the attempt to change the category code is rejected by the TS7700.

The LWORM retention settings in the TS7700 cannot be overridden by the host. In addition, the volume cannot change categories until the specified TS7700 data class LWORM retention settings for the volume are satisfied.

The volume remains available for read activity and can be appended to as normal for LWORM processing. No error condition is set on the host software side and the volume is not moved to the error category (xxxE) by the TS7700.

If this condition occurs, the reason the TMS attempted to move the volume to SCRATCH must be analyzed and addressed.

For more information about the TS7700 LWORM Retention function, see [TS7700 LWORM Retention Function User's Guide](#). If the TMS in use is DFSMSrmm, consider applying the PTF for OA64338 to obtain the latest changes to the LWORM function that have been made in DFSMSrmm.

12.16.1 Adjusting a volume's expiration date

In DFSMSrmm, the volume can have its expiration date adjusted to match the criteria that are specified in the TS7700 settings and, if needed, the retention method changed to EXPDT by using the following TSO command:

```
RMM CV volser RM(EXPDT) EXPDT(YYYY/DDD)
```

The volume then moves to SCRATCH during the return to scratch processing on the specified date, assuming that the EXPDT value was set to the date that is allowed by the TS7700 LWORM retention settings or later. When this command is used, you might want to set the expiration date value to one day past what is specified in the TS7700 settings. Doing so avoids the situation in which the return to scratch attempt occurs earlier in the day than the TS7700 timestamp for volume creation allows.

Normally, when a volume is created this extra day is not needed because the timestamp on the host and the TS7700 for the expiration values are the same.

12.16.2 Mirroring retention settings

One method to manage data sets on LWORM volumes for which retention settings were specified for a TS7700 data class is to mirror those settings in the TMS. Using this method, the TMS does not attempt to scratch the volumes before the date specified in the TS7700 LWORM retention settings.

For example, in DFSMSrmm, you can specify the use of the retention method EXPDT and specify a number of days for retention that matches the **FIXED DURATION** value set in the TS7700 data class that is used for allocating LWORM data sets.

With SMS ACS routines and **EDRMMxx** parm **MCATTR**(VRSELXDI or ALL), the allocations of the data sets that are written to LWORM volumes can be driven to a data class and management class pair that specifies the LWORM designation and the **EXPIRE AFTER DATE/DAYS** value that matches the required specifications for retention in the TS7700 data class.

Alternatively, the expiration date can be specified in the SMS data class or passed in the JCL rather than the management class. Tape expiration criteria can also be set up by using the RMM DEFAULT TABLE to assign the same attributes that are available in the management class.

For example, when using management class MC00FIX2 and data class D00YFIX1, if D00YFIX1 is configured within the TS7700 by the IBM SSR to have a **FIXED DURATION** of 7 days, the SMS management class definition for MC00FIX2 can be set to mirror this value by setting the **EXPIRE AFTER DATE/DAYS** value to 7 days, as shown in Figure 12-2 on page 663.

```

MANAGEMENT CLASS DEFINE                                     Page 1 of 8
SCDS Name . . . . . : SYS2.ITSCPLEX.DFSMS.SCDs
Management Class Name : MC00FIX2

To DEFINE Management Class, Specify:

Description ==> _____
                ==> _____

Expiration Attributes
Expire after Days Non-usage . . . 7 (1 to 93000 or NOLIMIT)
Expire after Date/Days . . . . . NOLIMIT (0 to 93000, yyyy/mm/dd or
                                         NOLIMIT)

Retention Limit . . . . . NOLIMIT (0 to 93000 or NOLIMIT)

```

Figure 12-2 Page 1 of the SMS management class definition showing an Expire after Date/Days value of 7 days

The **RETENTION METHOD** value of EXPDT is specified on page 8 of the management class, as shown in Figure 12-3.

```

MANAGEMENT CLASS DEFINE                                     Page 8 of 8
SCDS Name . . . . . : SYS2.ITSCPLEX.DFSMS.SCDs
Management Class Name : MC00FIX2

To DEFINE Management Class, Specify:

Tape Volume Attributes
Retention Method . . . . . EXPDT (VRSEL, EXPDT or blank)
Volume Set Management Level _____ (VOLUME, FIRSTFILE, SET or blank)

Tape Data Set Attributes
Exclude from VRSEL . . . . . _ (Y, N or blank)
Retain While Cataloged . . . _____ (ON, OFF, UNTILEXPIRED or blank)

```

Figure 12-3 Page 8 of the SMS management class definition

These values correlate with the data class, D00YFIX1, that is to be used. On the TS7700, this data class Logical WORM Yes, as shown in Figure 12-4, and was configured by the IBM SSR to have a FIXED DURATION of 7 days.





Name:	D00YFIX1
Virtual Volume Size (Device MiB):	Insert Media Class 
Compression Method:	ZSTD Compression 
3490 Counters Handling:	Surface EOT 
Logical WORM:	Yes 
Description:	LWORM for DC D00YFIX1

Figure 12-4 A data class within the TS7700 MI that is defined to use LWORM

The **FIXED DURATION** value of 7 days can be verified by issuing the **LI REQ** command with the **LWORMR,SHOW,1** keywords, and noting the output for the D00YFIX1 data class (see Example 12-21).

Example 12-21 LWORM settings LI REQ display

```
LI REQ,D0001,LWORMR,SHOW,1
CBR1020I Processing LIBRARY command: REQ,D0001,LWORMR,SHOW,1.
CBR1280I Library D0001 request. 622
Keywords: LWORMR,SHOW,1
-----
LWORMR SHOW V1 .0
INDEX:1
ID:DTCLASS ,FIXDUR ,APPDUR, FLG      ID :DTCLASS ,FIXDUR, APPDUR, FLG
1:D00YFFH ,30 ,30 ,1      2 :D00YFIX2 ,-1 , -1 ,8A
3:D00YFIX1. ,7 ,0 ,208A   4 :D00YFIX3 ,7 ,7 ,1
```

By using this method, the data sets and the volumes they are on expire in DFSMSrmm on a specific date 7 days after creation. This expiration date can be extended by using the **EXPIRE AFTER DAYS NON-USAGE** value in addition to or in place of **EXPIRE AFTER DATE/DAYS** if needed.

DFSMSrmm uses the longest of the two expiration dates that are determined by these two values.

12.16.3 Using the WHILECATALOG function

Another method to manage LWORM volumes in DFSMSrmm is to make retention decisions that are based on the catalog entry for the data sets on the LWORM volume. This function is called the **WHILECATALOG (WC)** function. The EXPDT retention method in RMM features the following choices:

- ▶ **WHILECATALOG(EXPIRED)**, which is abbreviated as **WC(UX)**

WC(UX) retains the data sets on a volume until such a time as the expiration date is reached or the data sets on the volume are no longer cataloged. With **WC(UX)**, the catalog entry that is deleted or the expiration date that is met can cause the data set and volume to expire.

WC(UX) is not a good choice for managing LWORM volumes because when the data sets are no longer cataloged, their expiration date is dynamically adjusted to be the current date plus the value of the **CATLGDAYS** parameter as specified in the EDGRMMxx parmli member (defaults to 2 if not specified).

Therefore, when the catalog entry is deleted for the last data set on a volume that is managed by using **WC(UX)**, the expiration date for that data set and the volume are moved up, possibly before the **FIXED DURATION** that is specified in the TS7700 settings for the LWORM data class.

- ▶ **WHILECATALOG(ON)**, which is abbreviated as **WC(ON)**

WC(ON) is a logical choice for managing data sets on LWORM volumes. This function retains a data set (and the volume that it is on) if the data set is cataloged and the expiration date is not yet passed.

With **WC(ON)**, after the catalog entry is deleted, DFSMSrmm compares the current expiration date and the value of **CATLGDAYS** and uses the greater of the two values to set the expiration date.

After that expiration date passes, the data set is no longer retained and the volume can be expired. The expiration date can be set by data class, management class, or JCL. If that date is equal to or greater than the **FIXED DURATION** value for retaining LWORM volumes as specified in the TS7700, **WC(ON)** can be used effectively.

Optionally, the DFSMSrmm DEFAULT TABLE can be used with RETPD and OVERRIDE to overrule any JCL passed expiration date to ensure a value greater than the TS7700 **FIXED DURATION** is used.

Note: The settings in the TS7700 for LWORM retention are a minimum value. The settings in the TMS might keep a volume for longer than what is specified in the TS7700 data class.

12.16.4 Keeping LWORM volumes permanently

With the TS7700 LWORM retention function, it is possible to configure the retention settings within the TS7700 to keep LWORM volumes permanently. In such a case, DFSMSrmm can also be configured also to keep a data set and volume permanently by using any of the following methods:

- ▶ Pass an expiration date of 99365 or 99366 in the JCL or specified in Data Class
- ▶ Use MCATTR in DFRMM and assign a Management Class with **EXPIRE AFTER DATE/DAYS** and **EXPIRE AFTER DAYS NON-USAGE** set to **NOLIMIT**
- ▶ Use a DEFAULT TABLE entry to set **RETPD(PERMANENT)** for LWORM data sets that are based on job name, data set name, or program name

The older method of retention in DFSMSrmm can also be used to manage LWORM volumes. A VRS can be assigned to retain a data set and volumes permanently, retain for a specific number of days, or retain with a combination of days and catalog entry.

The VRS must not include the **UNTIL EXPIRED** option or the **IGNORE EXPIRY DATE** option if it is to manage LWORM data with the TS7700 LWORM retention functions.



Monitoring

This chapter describes the methods of monitoring information of the IBM TS7700. It also describes the contents of the information.

This chapter includes the following topics:

- ▶ 13.1, “Overview” on page 668
- ▶ 13.2, “Base information: Types of statistical records” on page 670
- ▶ 13.3, “Web-based Monitoring method” on page 671
- ▶ 13.4, “Bulk Volume Information Retrieval” on page 687
- ▶ 13.5, “IBM Tape Tools” on page 703
- ▶ 13.6, “Host Console Request Commands for monitoring” on page 730
- ▶ 13.7, “IBM z/OS commands for monitoring” on page 735
- ▶ 13.8, “Alerts and exception and message handling” on page 738

13.1 Overview

The IBM TS7700 series is part of a line of tape virtualization products that revolutionized the way that mainframes use their tape resources. As the capability of tape virtualization grows, so too does 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 Management Interface (MI) that is based on open standards through which a web-based storage management application can request specific information that the TS7700 maintains. In addition, several methods on z/OS can provide the information to mainframe applications. Since R5.3, TS7700 supports a REST API to get TS7700 system information. Refer to Chapter 17, “RESTful API” on page 855.

You can use the following interfaces, tools, and methods to monitor the TS7700:

- ▶ Web-based Monitoring method:
 - TS7700 MI
 - IBM TS4500 Management GUI and TS3500 tape library Specialist (TS7700T only)
- ▶ Bulk Volume Information Retrieval function (BVIR)
- ▶ IBM Tape Tools: VEHSTATS, VEHAUDIT, and VEHGRXCL
- ▶ Host Console Request Commands
- ▶ REST API

The GUI/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, performance evaluation tools are now 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.

TS7700 supported Transparent Cloud Tiering since R4.2. For more information about monitoring the Cloud Attached model, see Chapter 12, “Monitoring the TS7700C”, in *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573. It includes the following contents about monitoring TS7700C:

- ▶ TS7700 object name format
- ▶ BVIR example
- ▶ VEHSTATS
- ▶ MI windows
- ▶ **LI REQ** commands
- ▶ Capacity monitoring of cloud storage

Since R5.2.2, TS7700 supports TS7700 Advanced Object Store. For more information about monitoring TS7700 Advanced Object Store, see *IBM TS7700 R5 DS8000 Object Store User's Guide*, REDP-5583.

For more information about interfaces, such as TS4500 Management GUI, TS3500 Tape Library Specialist and TS7700 MI, see 13.3, “Web-based Monitoring method” on page 671.

For more information about BVIR, see 13.4, “Bulk Volume Information Retrieval” on page 687.

For more information about IBM Tape Tools, see 13.5, “IBM Tape Tools” on page 703.

For more information about Host Console Request Command, 13.6, “Host Console Request Commands for monitoring” on page 730.

For more information about REST API, see Chapter 17, “RESTful API” on page 855. An overview of these interfaces, tools, and methods is shown in Figure 13-1.

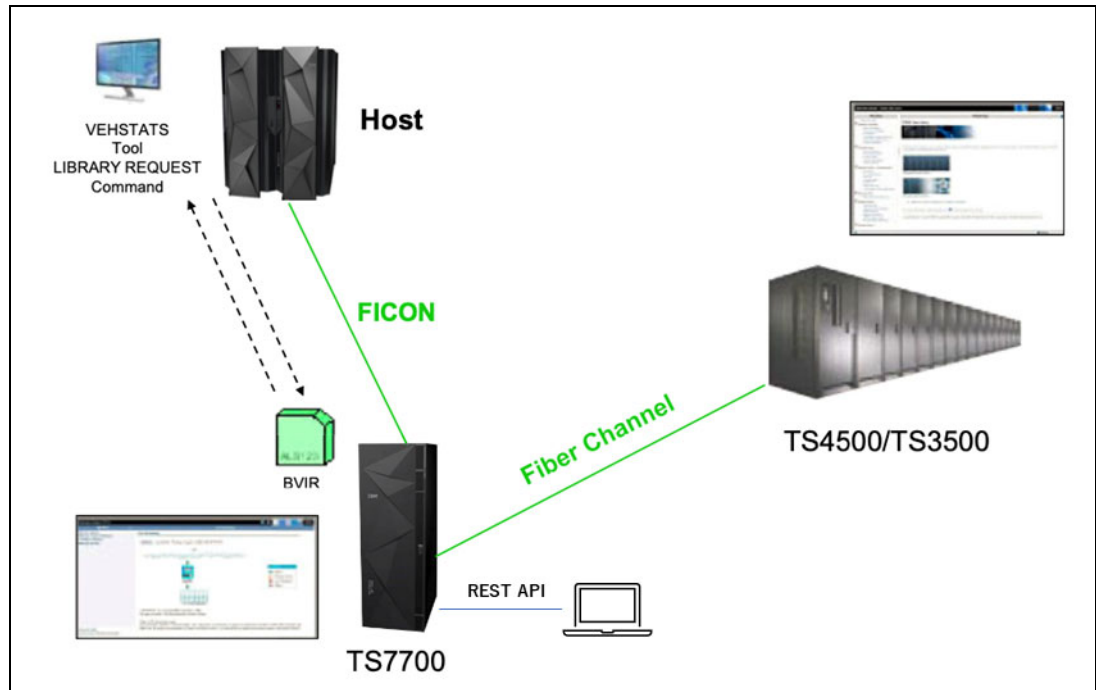


Figure 13-1 Interfaces, tools, and methods to monitor the TS7700

The information that can be acquired by each interface, tool, and method is listed in Table 13-1.

Table 13-1 Information acquired by each interface, tool, and methods

Interface, tool, and method	UI	Acquired Information
TS7700 MI	Web-based	Virtual/physical mounts, host throughput, cache throttling, cache utilization, grid network throughput, and so on.
TS4500 Management GUI	Web-based	Information about TS4500, such as physical configuration, events, logical libraries, state of drives and FC/Ethernet ports, locations of cartridges, and users or roles. Available with TS7700T only.
TS3500 Tape Library Specialist	Web-based	Information about TS3500, such as cartridges, libraries, drives, ports, and accesses. Available with TS7700T only.
BVIR	JCL (z/OS)	Information for all of the logical volumes in TS7700 at once, such as volume status, cache content, physical to logical volume mapping, physical media pools, physical volume status, and copy audit.

Interface, tool, and method	UI	Acquired Information
IBM Tape Tools (VEHSTATS, and so on)	JCL (z/OS)	These tools format binary data that is acquired with BVIR into readable reports that include activities about virtual devices, host adapters, cache partitions, and so on. Note: Not all data that is acquired with BVIR is in binary format; some are in readable format.
Host Console Request Command	z/OS Console or TS7700 MI	Provides a simple way to acquire information about TS7700 resources, such as logical/physical volume, cache, grid controls, and copies. Other methods might be suitable for further investigation sometimes.
REST API	REST API	A platform-independent programming interface that provides a way for external applications to manage TS7700. The TS7700 system information can be retrieved through a REST API by using the TS7700 Management Interface user credentials and security settings.

13.2 Base information: Types of statistical records

All of the interfaces and tools process the following types of statistics that are provided by the TS7700:

- ▶ Point-in-time statistics
- ▶ Historical statistics

The user can retrieve these 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.

13.2.1 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.

13.2.2 Historical statistics

Historical 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.

Tips: Consider the following points:

- ▶ If a cluster is not available at the time that the historical statistics are recorded (except for the headers), all the data fields for that cluster 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 data or the logical volumes that are used to obtain the statistics by BVIR HISTORICAL STATISTICS function.

Point-in-time statistics and historical statistics are recorded. The point-in-time records present data from the most recent interval, which provides speedometer-like statistics. The historical statistics provide data that can be used to observe historical trends.

For more information about the records, see the following resources:

- ▶ [IBM Support: TS7700 Statistical Data Format V53a](#)
- ▶ [IBM TS7700 Series VEHSTATS Decoder](#)

13.3 Web-based Monitoring method

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 are not yet completed 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 by using a supported web browser.

For more information about supported browsers, see [Connecting to the management interface](#).

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, the following web interfaces are available:

- ▶ TS7700 Management Interface
- ▶ IBM TS4500 Management GUI and TS3500 tape library Specialist (TS7700T only)

This section describes these web-based monitoring methods.

13.3.1 TS7700 Management Interface: Performance page

To view performance information and statistics, click **Monitor** → **Performance**, as shown in the upper left of Figure 13-2.

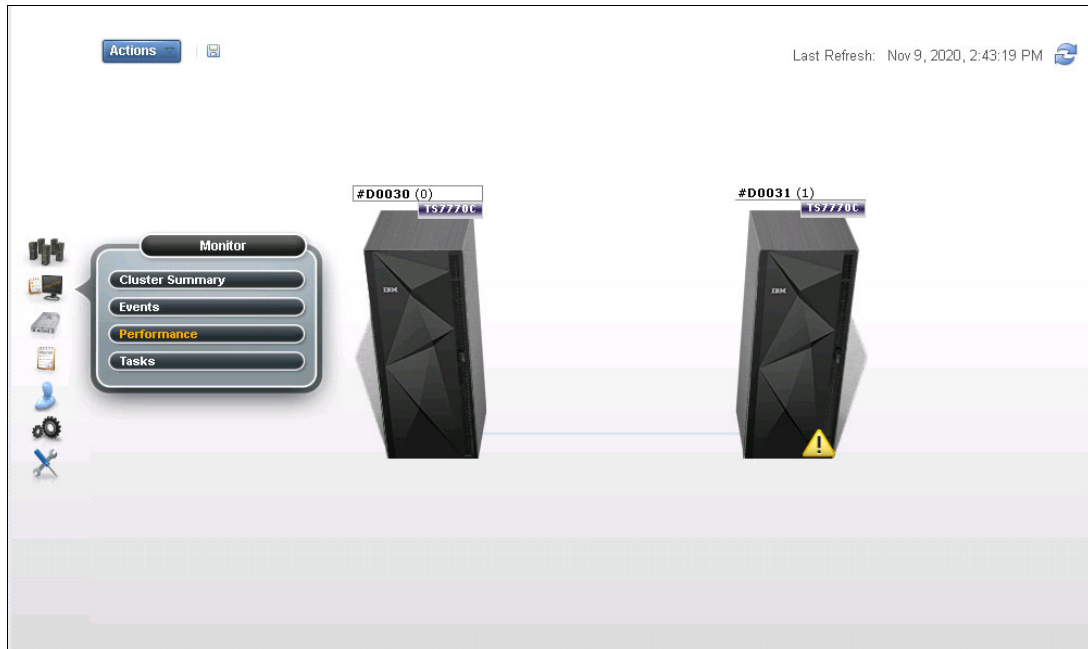


Figure 13-2 TS7700 MI Performance

The graphical views display snapshots of the processing activities from the last 15 minutes if nothing else is stated when describing the windows.

The following windows are available under the Performance page:

- ▶ Historical Summary
- ▶ Virtual Mounts
- ▶ Physical Mounts
- ▶ Host Throughput
- ▶ Cache Throttling
- ▶ Cache Utilization
- ▶ Cache Partitions
- ▶ Grid Network Throughput
- ▶ Pending Updates

These windows are described next.

Historical Summary window

This window (see Figure 13-3) 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.

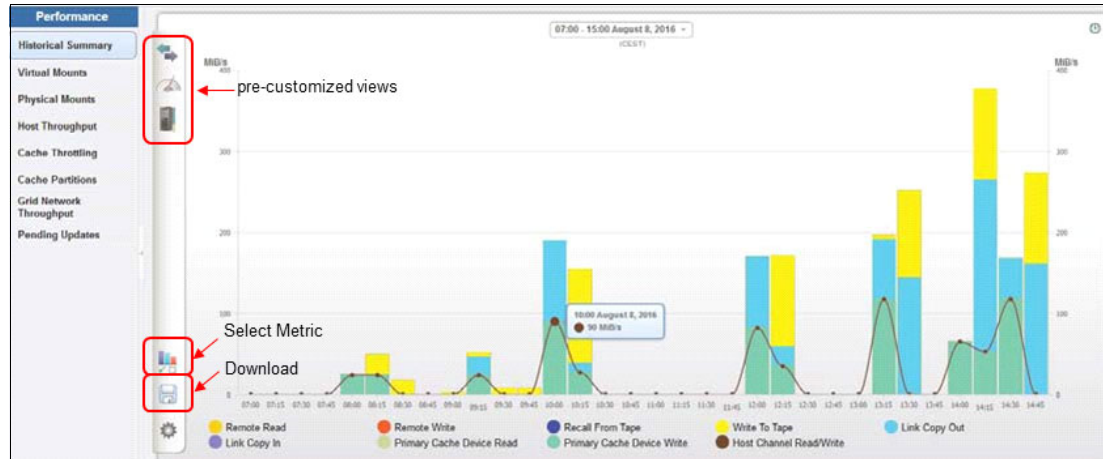


Figure 13-3 TS7700 MI Historical Summary of TS7760T

Multiple customized views can be selected, which depend on the installed cluster type, such as displaying for the maximum of an entire day:

- ▶ Throughput performance
- ▶ Throttling information
- ▶ Copy Queue

The value that is shown in the performance graph can be changed by selecting the different metrics. To do so, click **Select Metric** and choose the requested values, as shown in Figure 13-4. The maximum number of values you can view in one picture is 10. You can also save the graph by clicking the Download Button (disc).

Select Metrics		X	
Throughput		8 Selected	
I/O			
Host Channel	<input type="checkbox"/> Read/Write (MiB/s)		
Primary Cache	<input type="checkbox"/> Device Read (MiB/s)	<input type="checkbox"/> Device Write (MiB/s)	
Remote	<input type="checkbox"/> Read (MiB/s)	<input type="checkbox"/> Write (MiB/s)	
Copies			
Link Copies	<input type="checkbox"/> In (MiB/s)	<input type="checkbox"/> Out (MiB/s)	
Copy Queue Size	<input checked="" type="checkbox"/> Total (MiB)	<input checked="" type="checkbox"/> Immediate (MiB)	
	<input checked="" type="checkbox"/> Synchronous Deferred (MiB)	<input checked="" type="checkbox"/> Immediate Deferred (MiB)	
	<input checked="" type="checkbox"/> Deferred (MiB)	<input checked="" type="checkbox"/> Family Deferred (MiB)	
	<input checked="" type="checkbox"/> Copy Refresh (MiB)	<input checked="" type="checkbox"/> Time Delayed (MiB)	
Copy Life Span	<input type="checkbox"/> Immediate (seconds)	<input type="checkbox"/> Deferred (seconds)	
	<input type="checkbox"/> Time Delayed (seconds)		
Storage		0 Selected	
System		0 Selected	
<input type="button" value="Set"/> <input type="button" value="Cancel"/> <input type="button" value="Clear"/>		8 Total Selected	

Figure 13-4 Select the Host I/O metrics

The pre-customized “Throughput” graph is shown in Figure 13-3 on page 673. The graph shows the following cache bandwidth-relevant information:

- ▶ Host channel read/write
- ▶ 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 13-5. It shows which type of throttling happened during the selected interval, and the throttling effect in milliseconds.




Figure 13-5 TS7700 MI Historical Summary Throttling overview

A pre-customized “Copy Queue” is shown in Figure 13-6.



Figure 13-6 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. You can select only 10 items for one graph, as shown in Figure 13-7.

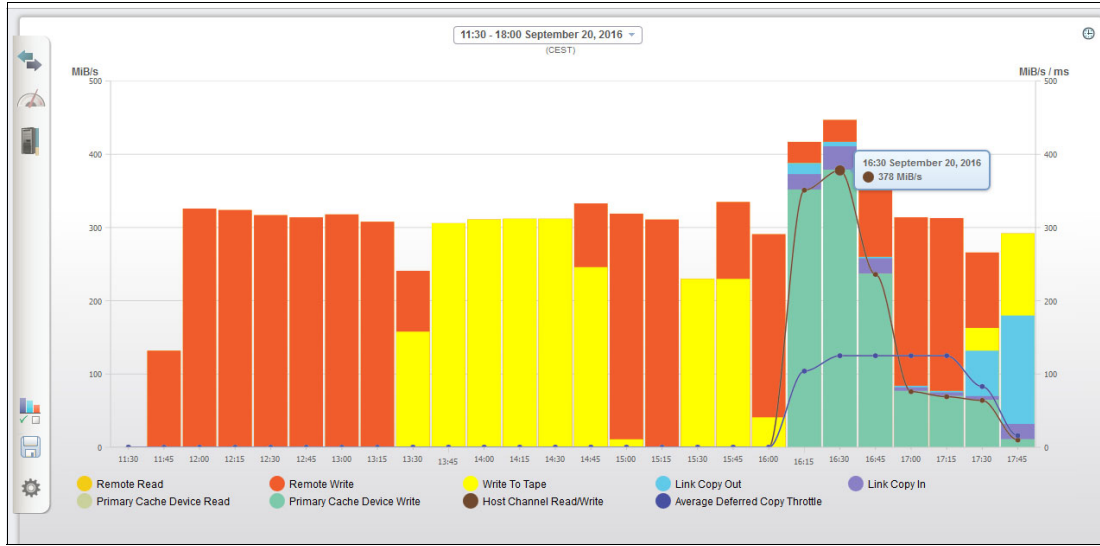


Figure 13-7 Historical summary overview

Virtual Mounts window

Use this window (see Figure 13-8) 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. The example that is shown in Figure 13-8 is from a TS7700 Cluster that is part of a two-cluster grid configuration.

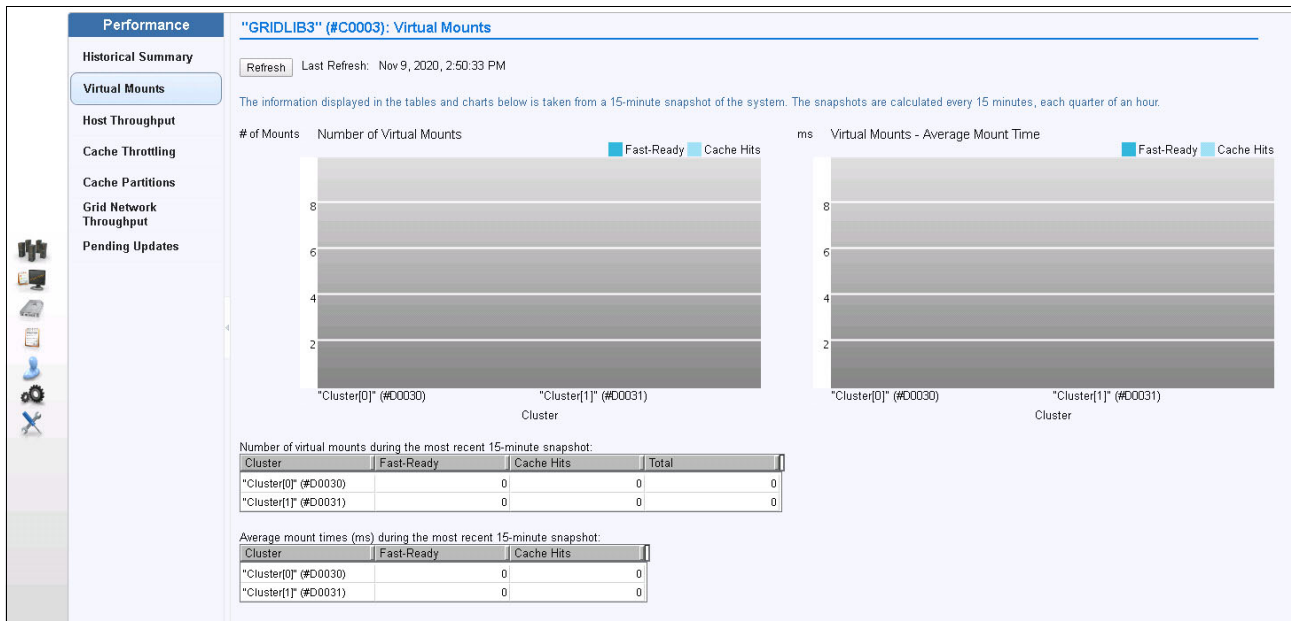


Figure 13-8 TS7700 MI Virtual Mounts window

The “Number of Virtual mounts during last 15-minute snapshots” table features 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
Total	Total number of logical mounts

The “Average mount times (ms) during last 15-minute snapshots” table includes 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

Physical Mounts window

Use this window (see Figure 13-9) to view physical mount statistics for the TS7700T. 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 that is shown in Figure 13-9 is from a TS7700T cluster that is part of a multi-cluster grid configuration (six-cluster grid).

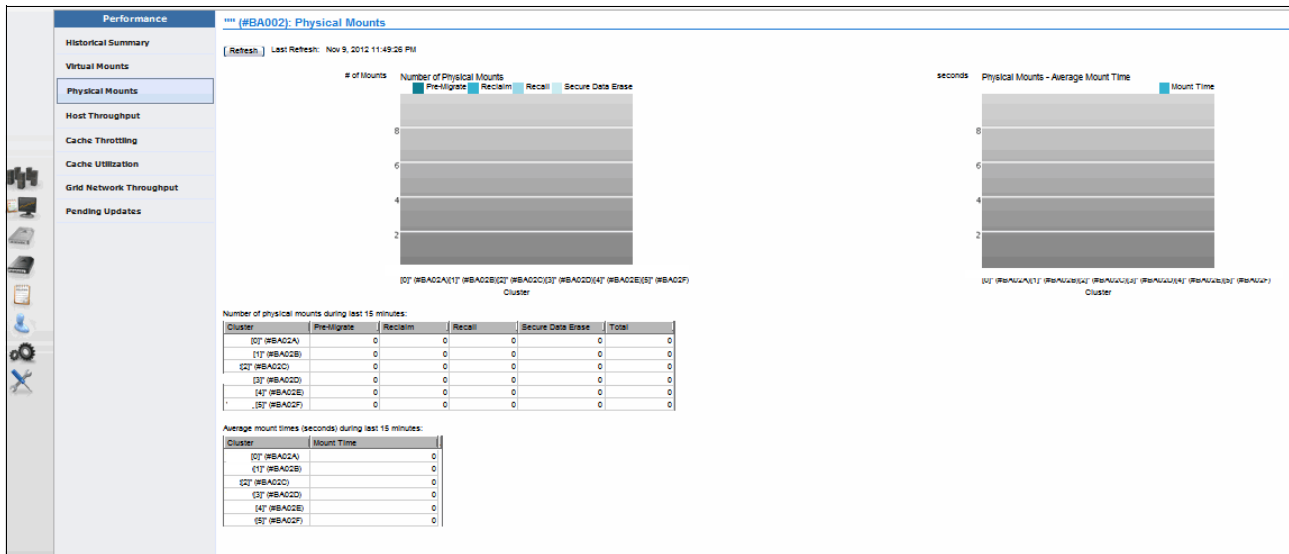


Figure 13-9 TS7700T MI Physical Mounts window

The table cells show the following items:

Cluster	The cluster name
Pre-migrate	Number of premigrated mounts
Reclaim	Number of reclaimed mounts
Recall	Number of recall mounts
Secure Data Erase	Number of Secure Data Erase mounts
Total	Total physical mounts
Mount Time	Average mount time for physical mounts

Host Throughput window

You can use this window (see Figure 13-10) 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 can 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 that is shown in Figure 13-10 is from a TS7700 cluster that is part of a multi-cluster grid configuration (two-cluster grid).

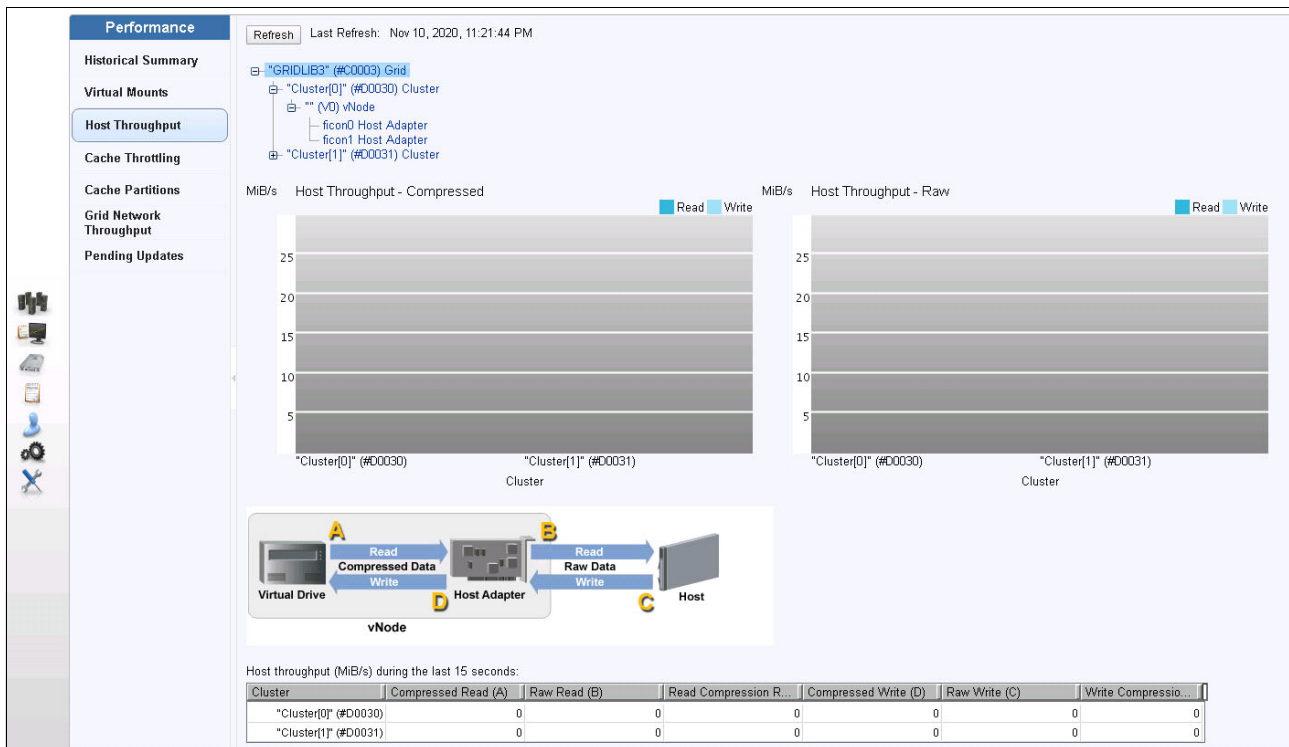


Figure 13-10 TS7700 MI Host Throughput window

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 that is 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 raw data read to compressed data read.

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 raw data written to compressed data written

Cache Throttling window

You can use this window (see Figure 13-11) to view cache throttling statistics for the TS7700. The example that is shown in Figure 13-11 is from a TS7700 cluster that is part of a multi-cluster grid configuration (two-cluster grid).

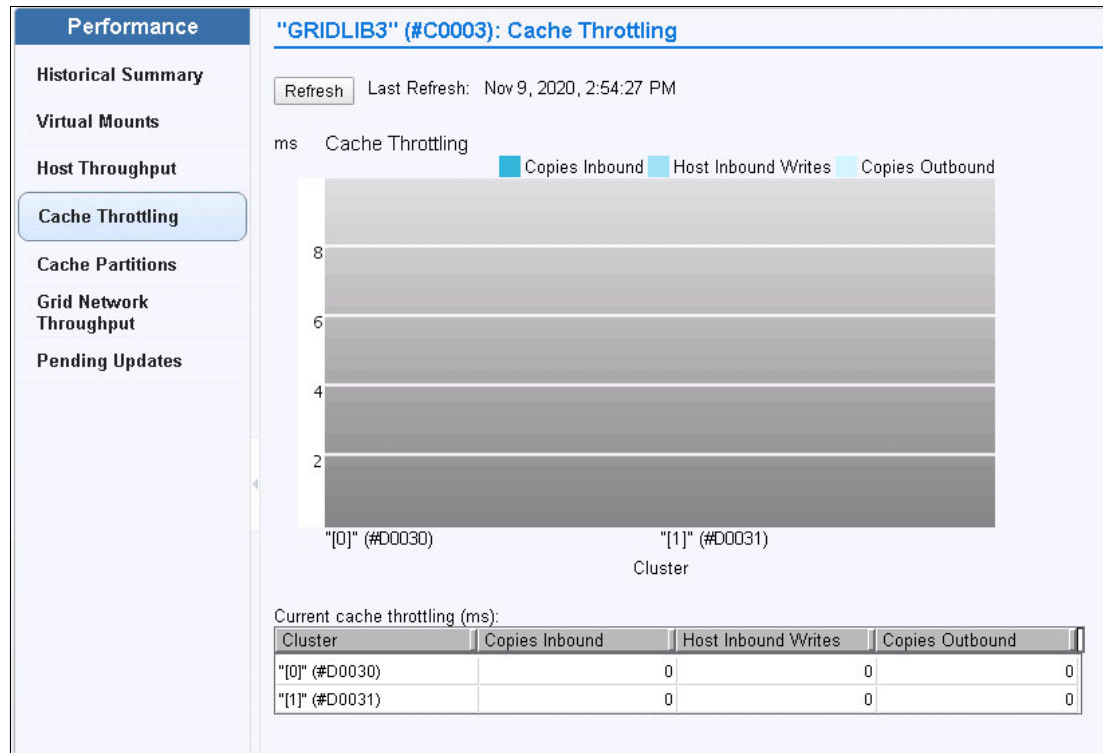


Figure 13-11 TS7700 MI Cache Throttling window

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	Cluster name
Copies Inbound	Amount of time that is inserted between internal inbound copy operations
Host Inbound Writes	Amount of time that is inserted between host write operations
Copies Outbound	Amount of time that is inserted between internal outbound copy operations

Apart from the Cache Throttling window, a throttling indicator is also included in the head of the cluster in the Grid Summary and the Cluster Summary pages. Hover over the indicator to see the throttling type and the throttling effect. This view is only a snapshot view.

Cache Utilization window

You can use this window (see Figure 13-12) to view cache utilization statistics for the TS7700D. If a TS7700T or TS7700C is installed, this selection is called “Cache Partitions” and is explained in the next section.

The example that is shown in Figure 13-12 is from a TS7700T cluster that is part of a multi-cluster grid configuration (four-cluster grid).




Figure 13-12 TS7700D 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 bar graph and table form, and can be displayed 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 that is present in the cache preference group.
- ▶ Total amount of data currently in the cache: Total amount of data that is present in volumes that are assigned to the cache preference group.
- ▶ Median duration that volumes 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). A bar graph is not used.

Clarification: Median Duration in Cache and Number of Logical Volumes Migrated statistics include a table column for each of the periods that are in parentheses.

Cache Partitions window

You can use this window (see Figure 13-13) to view Tape cache partitions and their usage. Depending on your filter list, you might see the following output that is shown in Figure 13-13.

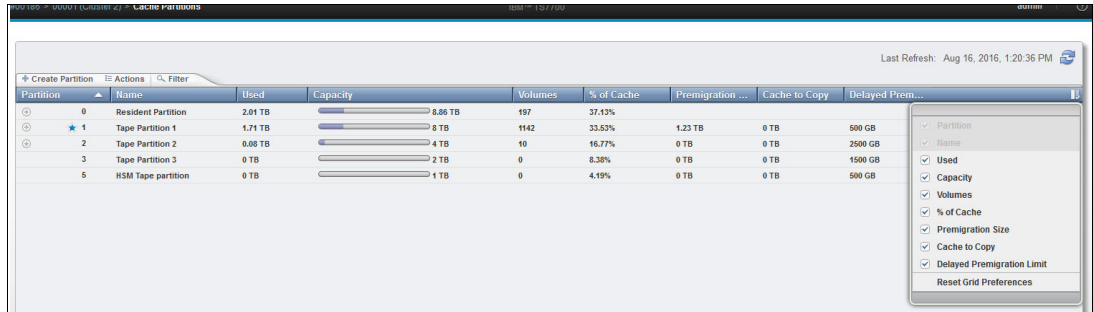


Figure 13-13 Cache Partition view in a TS7700T or TS7700C

For more information, see “Cache Partition” on page 400.

Grid Network Throughput window

Use this window (see Figure 13-14) to view network path usage (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, no cross-cluster data transfer is available through the Ethernet adapters.

The example that is shown in Figure 13-14 is from a TS7700 Cluster that is part of a multi-cluster grid configuration (six-cluster grid).

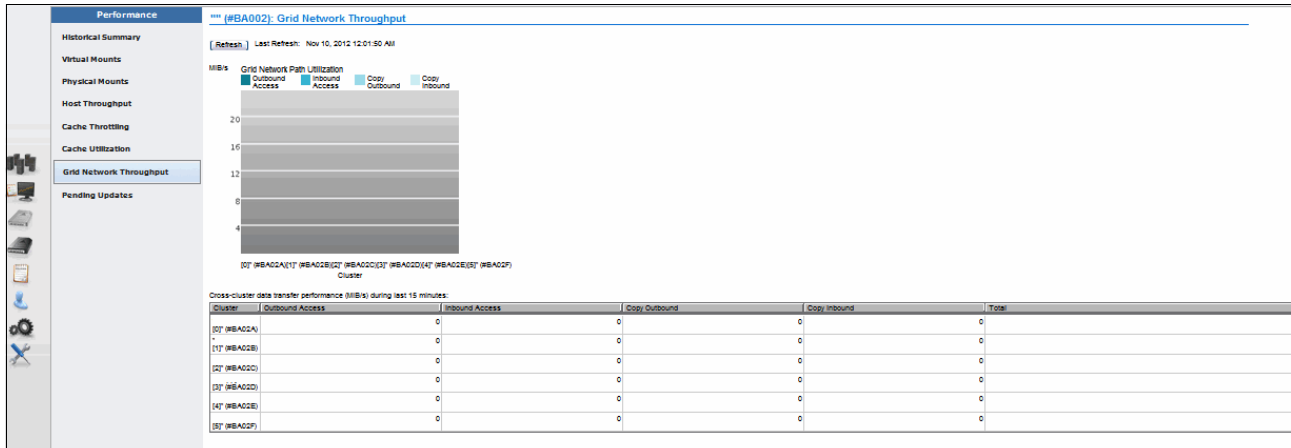


Figure 13-14 TS7700 MI grid network throughput in a six-cluster grid

The table displays data for cross-cluster data transfer performance (MBps) during the last 15 minutes. The table cells show the following items:

Cluster	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
Total	Total data transfer rate for the cluster

Pending Updates window

Use this window (see Figure 13-15) 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 pending updates can take effect, all clusters must be online.

"GRIDLIB3" (#C0003): Pending Updates

Last Refresh: Nov 10, 2020, 8:10:00 PM

Takeover ownership updates required in each cluster (.csv file)

Number of updates required against each cluster:

Cluster	Standard	Read/Write Takeover	Read-Only Takeover	Service Takeover
"[0]" (#D0030)	0	0	0	0
"[1]" (#D0031)	0	0	0	0

Number of **Immediate Deferred** copies for each cluster:

Cluster	Quantity	Size (GiBs)
"[0]" (#D0030)	0	0
"[1]" (#D0031)	0	0

Number of **Synchronous Deferred** copies for each cluster:

Cluster	Quantity	Size (GiBs)
"[0]" (#D0030)	0	0
"[1]" (#D0031)	0	0

Figure 13-15 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 and Synchronous Deferred copies, which, as with pending updates, normally result from changes that are made while a cluster is Offline or in service prep mode or service mode.

Remember: Pending Immediate Deferred and Synchronous Deferred copies must be avoided. They might be a result of grid network problems.

With Release 3.2, the download section also includes the tape with a *hot token*. Hot tokens are volumes that were changed during an unavailability of the cluster and now need a reconciliation. The reconciliation is run during the cluster online setting.

13.3.2 TS7700 Management Interface: Other windows

Other pages are available in the TS7700 MI.

Virtual Tape Drives window

To view the information about volumes that are mounted by virtual devices, select the **Virtual** → **Virtual Tape Drives** page. In the MI, the synchronous mode mounts are visible only in the Virtual Tape Device view of the Host I/O cluster. The selected second synchronous target cluster does not show any mount that an internal device is used because for the synchronous mount, which is not shown in the Virtual Tape drive window (see Figure 13-16).

Address	Mounted V...	Time On Drive	Cache Mount Cl...	Bytes Read	Bytes Written
vt000C	A78548	0 hours, 0 minutes, 2 seconds	TS7740 (1), 00001 (2)	0 KIB (0 KIB)	18.54 KIB (18.06 KIB)
vt000A	A78550	0 hours, 0 minutes, 1 seconds	TS7740 (1), 00001 (2)	0 KIB (0 KIB)	0.18 KIB (0.03 KIB)
vt000E	A78729	0 hours, 6 minutes, 41 seconds	TS7740 (1), 00001 (2)	0 KIB (0 KIB)	121.1 KIB (120.69 KIB)
vt0006	A78779	0 hours, 2 minutes, 2 seconds	TS7740 (1), 00001 (2)	0 KIB (0 KIB)	106.77 KIB (99.02 KIB)
vt0002	A78799	0 hours, 0 minutes, 2 seconds	TS7740 (1), 00001 (2)	0 KIB (0 KIB)	19.42 KIB (18.95 KIB)
vt0010	A78801		TS7740 (1), 00001 (2)	0 KIB (0 KIB)	0 KIB (0 KIB)

Address	Mounted V...	Time On Drive	Cache Mount Cl...
vt0000			
vt0001			
vt0002			
vt0003			

Figure 13-16 Virtual Tape Drive view for synchronous mode copy

Active Data Distribution window

Use this window (see Figure 13-17) to view the usage of back-end cartridges. To view the active pools, select the **Physical** → **Physical Volumes** → **Active Data Distribution** window. Figure 13-17 shows the active pools and correspondent data distribution (number of cartridges by occupancy percentage range).

Pool	Media Type	0+	10+	20+	30+	40+	50+	60+	70+	80+	90+
3	JA(ETC)	0	0	0	0	0	0	0	0	0	21
5	JA(ETC)	0	0	0	0	0	0	0	0	0	5
10	JA(ETC)	0	0	0	0	0	0	0	0	0	8

Figure 13-17 Pool Active Data Distribution

Click a pool link to display information about that pool, as shown in Figure 13-18.




Figure 13-18 Information Display of selected Pool

Review your Active Data Distribution. A low usage 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 defined multiple physical pools, you might need to check this number on a per pool basis, depending on your Borrow/Return policies. In the example that is shown in Figure 13-18, Pool 3 features the **borrow, return** parameter enabled.

13.3.3 TS4500 Management GUI

The TS4500 Management GUI includes a new design, and now features the same design as the TS7700 series. To gain an overview of the system, mouse over the TS4500 icon. Depending on your position, the health status is provided, but configuration information is also shown (see Figure 13-19).



Figure 13-19 TS4500 Overview example

In the TS4500 Logical Library view, you can find the information about the number of cartridges, drives, and maximum cartridges. Use the FILTER option to select the columns.

Figure 13-20 shows this window and the selected properties.

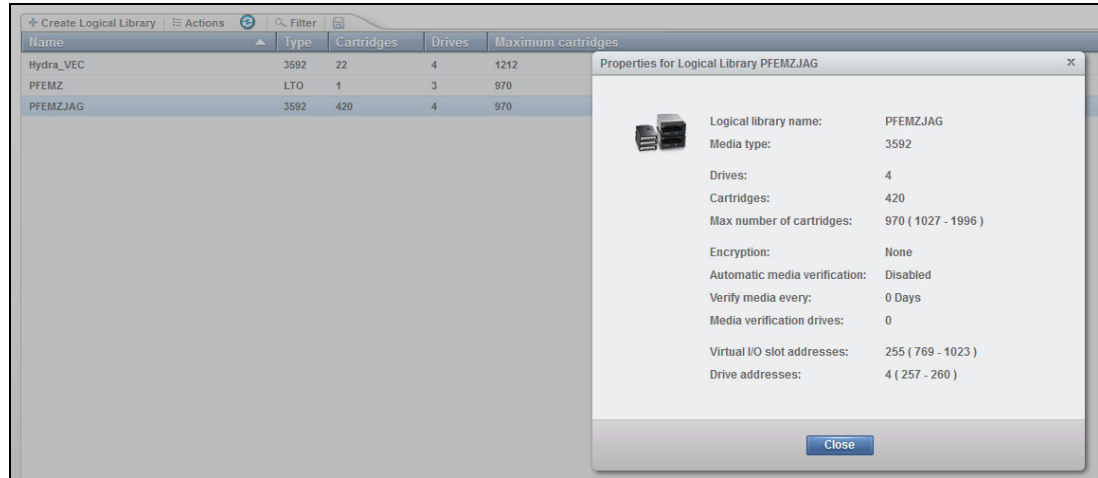


Figure 13-20 TS4500 logical library display with properties

For more information about the IBM TS4500 tape library, see *IBM TS4500 R9 Tape Library Guide*, [SG24-8235](#).

13.3.4 TS3500 Tape Library Specialist

The Tape Library Specialist (TS3500 Tape Library Specialist), which is available with the TS7700T only, 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 supported only 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.

The TS3500 Tape Library Specialist session times out after a default setting of 10 minutes. This value is different from the TS7700 MI.

Figure 13-21 shows the TS3500 Tape Library System Summary window.

IBM System Storage™ TS3500 Tape Library

Work Items

System Summary

- Cartridges
 - Data Cartridges
 - Cleaning Cartridges
 - I/O Station
 - Cartridge Assignment Policy
 - Barcode Encryption Policy
 - Key Label Mapping
 - Insert Notification
- Library
 - Logical Libraries
 - Accessor
 - Preferred Accessor Zones
 - ALMS
 - Virtual IO
 - Date and Time
- Drives
 - Drive Summary
 - Drive Assignment
 - Control Paths
 - World Wide Names
 - Cleaning Mode
- Ports
 - Fibre Channel Summary
- Access
 - Web Security
 - Operator Panel Security
 - Key Manager Addresses
 - SNMP Settings
 - SNMP Destinations
 - SNMP System Data
 - SMI-S Agent
 - Library IP Addresses
 - Secure Socket Layer
- Service

Library Status:

Accessors	OK
3592 Capacity Utilization	63%

View Configuration and Cartridge Counts:

All Frames

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 13-21 TS3500 Tape Library Specialist System Summary window

You can change the default values through the TS3500 Tape Library Specialist by selecting **Manage Access** → **Operator window Security**, which opens the window that is shown in Figure 13-22.

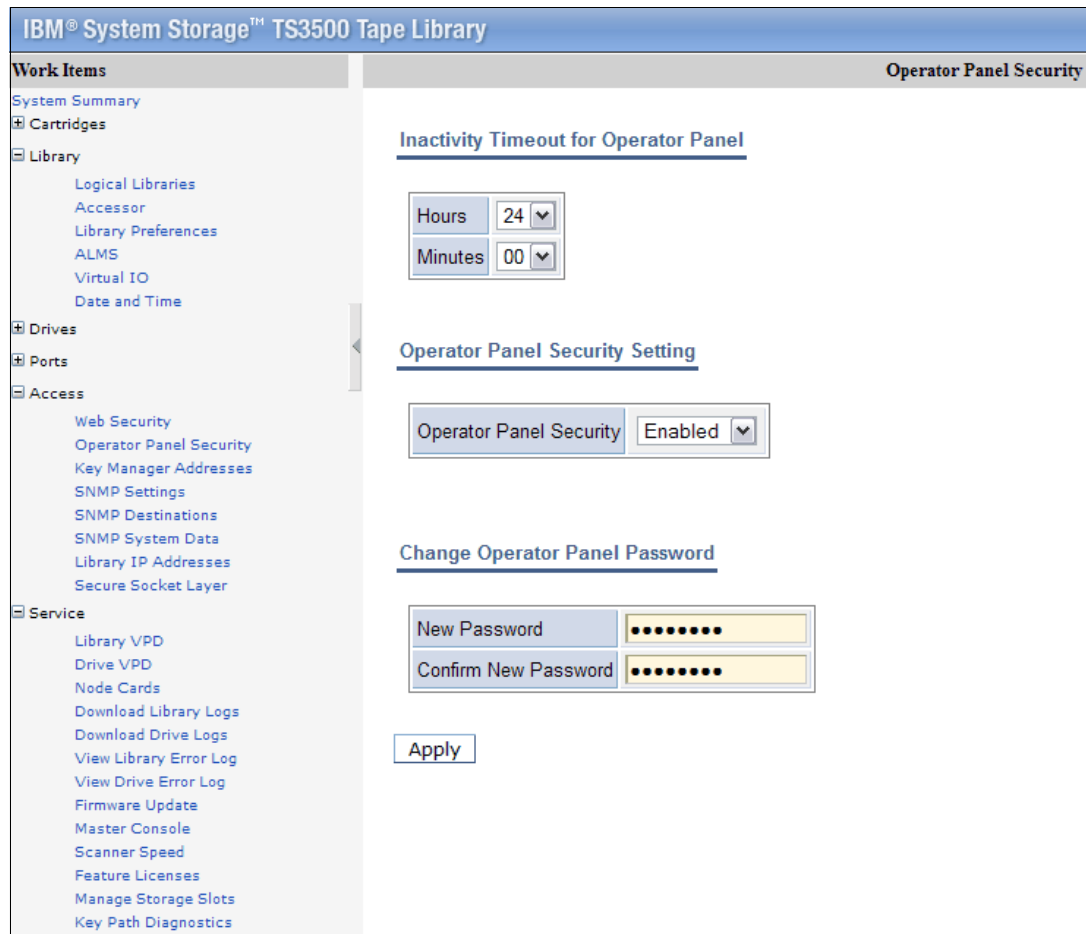


Figure 13-22 TS3500 Tape Library Specialist Operator window Security window

Some information that is provided by the TS3500 Tape Library Specialist is in a display-only format and no option is available 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.

Display only

The following information is available in a display-only format:

- ▶ Accessor Usage:
 - Activity of each Accessor and gripper
 - Travel meters of Accessors
- ▶ Drive Status and Activity

Download only

The following information can be viewed when downloaded to a workstation:

- ▶ Drive Statistics:
 - Last VOLSER on this drive
 - Write and Read megabytes (MB) per drive
 - Write and Read errors that are corrected per drive
 - Write and Read errors uncorrected per drive
- ▶ Mount History for cartridges:
 - 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

- ▶ Fiber Port statistics:

The Fiber 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:
 - 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 that are processed by the TS7700.

13.4 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. Conveniently, 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. By using BVIR, information can be obtained 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 Information
- ▶ Historical Statistics Information
- ▶ Physical Media Pools Information
- ▶ Physical Volume Status Information

- ▶ Copy Audit Information
- ▶ GGM Information
- ▶ Cloud Volume to Logical Volume Mapping Information
- ▶ Cloud Content Information
- ▶ Object Copy Queue Information
- ▶ Object Copy Audit Information
- ▶ Cloud Export Backup information

For more information, see [IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#).

Clarification: Records that are listed in 13.4, “Bulk Volume Information Retrieval” on page 687 show an initial record of “1234567890123...”. This record does not exist; rather, it is provided to improve readability.

13.4.1 BVIR overview

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 start 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 show how information is obtained 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.

A scratch or specific volume request can be used. The data set contains a minimum of two records that specify the type of data that is being requested. The records are in human-readable form and contain lines of character data.

The data set can be cataloged or uncataloged (although cataloging the data set can make it simpler 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 perspective, is part of the mount processing time.

After the TS7700 completes 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 is not unmounted and remounted between job steps in a JES3 environment.

After the response data set is 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: Because of the two-step approach, BVIR volumes cannot be written with LWORM specifications. Assign a Data Class without LWORM for BVIR volumes in R5.1 or earlier. In R5.1 PGA1 or later, LWORM volumes can be used as BVIR volumes.

Figure 13-23 shows the process flow of BVIR.

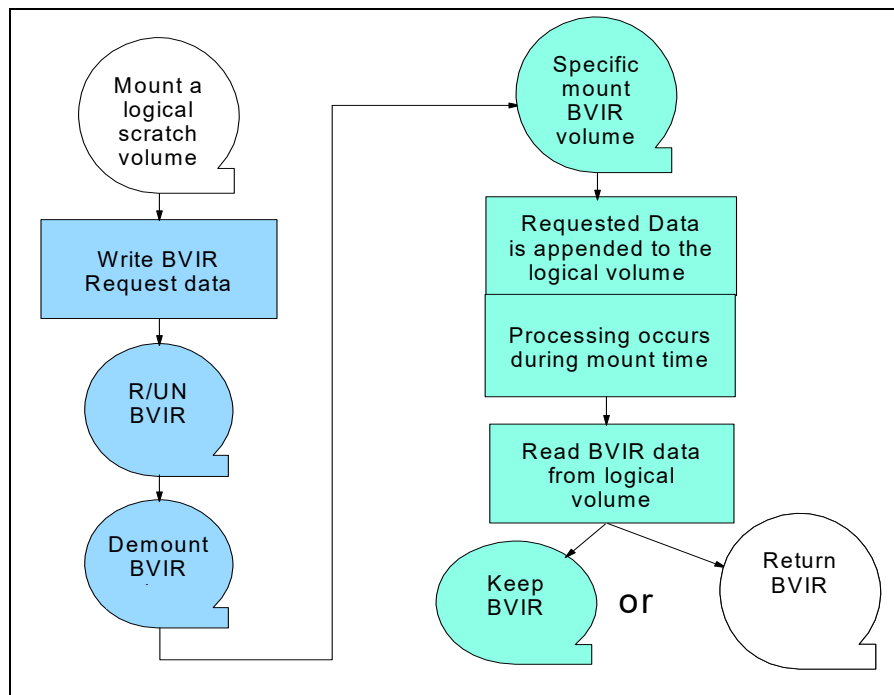


Figure 13-23 BVIR process flow

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 issue only one request at a time. Certain requests (for example, the volume map) might take several minutes to complete when many logical volumes are on physical volumes.

When a BVIR request does not finish in the allotted time, BVIR times out and returns no response data. Before R4.1.2, the timeout value is fixed at 300 minutes. At 4.1.2 or later release, the timeout value can be modified by using host console request command **LI REQ,dist1ib,SETTING2,BVIR,TIMEOUT**. For more information, see [TS7700 Library Request Command](#).

To prevent “locking” out another request during that time, the TS7700 handles 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, the data that is returned from the TS7700 can be in human-readable or binary form depending on the request. See the response sections for the specifics of the returned data.

The general format for the request/response data set is shown in Example 13-1.

Example 13-1 BVIR output format

```
123456789012345678901234567890123456789012345
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

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. Consider the following points:

- ▶ Records 1 and 2 contain the data request commands.
- ▶ Record 3 contains the date and time when the report was created, and the version of BVIR.
- ▶ Record 4 contains the hardware serial number and the distributed library ID of the TS7700.
- ▶ Record 5 contains all blanks.

Records 6–*N* and higher contain the requested data. The information is described in general terms. For more information about these records, see [TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#).

13.4.2 BVIR 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. No restrictions exist 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. The normal method is typically by deletion of the data set on the volume and a return-to-scratch policy based on data set deletion.

13.4.3 BVIR 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 with the first character position of the record and padded with blank characters on the right to complete the record.

Important: Consider the following points:

- ▶ The request fields must be as shown in Table 13-3 on page 692. Not starting with the first character position of the record or the use of extra blanks between words results in a failed request.
- ▶ The file must be written in uncompressed format so that it is 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. For more information, see Appendix E, “Sample job control language” on page 923.

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 that is assigned to the volume must specify the specific cluster that is to include the copy of the request volume.

The format for the request data set records is described in the following sections. Starting with code level R5.2.2, a change was made to allow more than one line in a BVIR request. With this new implemented method, BVIR requests are categorized as traditional or non-traditional format.

Traditional Request Format

This section describes the traditional request format.

Record 1

Record 1 must contain the command exactly as shown in Example 13-2.

Example 13-2 BVIR request record 1

```
1234567890123456789012345678
VTS BULK VOLUME DATA REQUEST
```

The format for the request’s data set records is listed in Table 13-2.

Table 13-2 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

- ▶ VOLUME MAP PRIMARY
- ▶ VOLUME MAP BACKUP
- ▶ 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
- ▶ GGM COPY RESULT
- ▶ GGM COPY STATUS
- ▶ CLOUD VOLUME MAP (ALL) xx
- ▶ CLOUD BACKUP LIST
- ▶ COPY AUDIT CLOUDA INCLUDE/EXCLUDE libids
- ▶ COPY AUDIT CM CLOUDA INCLUDE/EXCLUDE libids
- ▶ CLOUD CONTENT xxxxxxxx (yy)
- ▶ OBJECT COPY QUEUE

The format for the request's data set records is listed in Table 13-3.

Table 13-3 BVIR request record 2

Record 2: Request identifier		
Bytes	Name	Contents
1-80	Request	'VOLUME STATUS zzzzzz' or 'CACHE CONTENTS' or 'VOLUME MAP' or 'VOLUME MAP PRIMARY' or 'VOLUME MAP BACKUP' 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' or 'GGM COPY RESULT' or 'GGM COPY STATUS' or 'CLOUD VOLUME MAP (ALL) xx' or 'CLOUD BACKUP LIST' 'CLOUD AUDIT CLOUDA INCLUDE/EXCLUDE libids' 'COPY AUDIT CM CLOUDA INCLUDE/EXCLUDE libids' 'CLOUD CONTENT xxxxxxxx (yy)' 'OBJECT COPY QUEUE' 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 were 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.

Since R5.1 PGA1, the response contains fields that are related to the LWORM retention function.

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 `xxx-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 issue is indicated in the response record.

This issue 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 was 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 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 no physical volumes are assigned to the specified pool, that issue is indicated in the response record. Data can be requested for pools 0–32.

For the Cloud Backup List request, a list of cloud backups is returned, which is introduced at R5.1.

For the Cloud Volume Map request, if the total number of cloud volume map entries exceeds 4 million, the use of `xx` can provide the next 4 million entries. For example, when '0' is specified (which is equivalent to no `xx`), the first million entries that are ordered by `volser` are provided in response.

Since R5.1, the response contains other fields to show if a volume is retained version. An optional parameter 'ALL' is introduced to show all volume-cloud mapping including old versions retained in the cloud.

Since R5.1 PGA1, the response contains cloud volume version recovery fields.

For point-in-time and historical statistics requests, any other 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 on only the specific cluster from which the data is to be obtained.

For the Cloud Content requests that were introduced at R5.2.1, `xxxxxxx` specifies a cloud pool nickname to be used. The additional keyword `yy` is optional and must be a number equal to or larger than 0.

You can use the Object Copy Queue request that was introduced at R5.2.2 to view the current queue for object copies.

Non-Traditional Request Format

This section describes the non-traditional request format.

Record 1

This record is the same as Traditional Request Format.

Record 2 and after

With records 2 and 3, you can specify which data you want to obtain. The following options are available:

► Object Status

- Record 2: OBJECT STATUS <Cloudname> <hours>
- Record 3: OBJET ID <object 128 bytes>
- Record 4-N:
 - CONTAINER <container name 256 bytes> (up to 4 lines long)
 - OBJPREFIX <object prefix> (up to 2 lines long)

For object status requests that were introduced at R5.2.2, large amounts of object-related data can exist. The idea behind this command is to provide users with search capability to return large data sets of object information.

► Object Copy Audit

- Record 2: OBJECT COPY AUDIT INCLUDE|EXCLUDE <lib-id0>,...,<lib-id7>
- Record 3: DATASET 0
- Record 4: COPYMODE YES|NO

The object copy audit request that was introduced in R5.2.2 audits the databases on a set of specified TS7700 distributed libraries. This request is made to determine whether any objects exist that do not have a valid copy on at least one of them.

Use a specific MC that includes a copy policy that is defined to indicate that only the wanted cluster is to have a copy of the data. By ensuring that a sole copy of the request volume exists, any virtual device address on any of the clusters in the same grid configuration can be used to request and access the data.

It is not necessary for host connectivity to exist to the specific cluster. If an MC is used that indicates that more than one cluster is to include a valid copy of the request volume, unpredictable response data results can occur.

That MC must have only one consistency point that is defined on the code level R5.0 or later, or it fails the BVIR request.

Running code level R5.0 or later fails the BVIR requests in a Grid configuration when the Management Class has more than one consistency point that is defined. For more information, see [Info Letter: TS7700 running code R5.0 may expectedly fail BVIR requests](#).

13.4.4 BVIR Response data format

When the request data set was written to the volume and then closed and unmounted, the TS7700 validates the contents of the request volume when mounted again 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 80–1256 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 that are 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 issue is indicated. The record length is fixed; the record length of each response data is listed in Table 13-4.

Table 13-4 Record length of response data

BVIR request	Record length in bytes
VOLUME STATUS <i>vvvvvv</i>	643 (Changed at R5.3)
CACHE CONTENTS	80
VOLUME MAP (PRIMARY or BACKUP)	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>	600 (Changed at R5.4)
PHYSICAL VOLUME STATUS POOL <i>xx</i>	600 (Changed at R5.4)
COPY AUDIT COPYMODE INCLUDE/EXCLUDE <i>libids</i>	80
GGM COPY STATUS/RESULT	400
CLOUD VOLUME MAP (<i>ALL</i>) <i>xx</i>	1344 (Changed at R5.3)
CLOUD BACKUP LIST	7600
CLOUD CONTENT <i>xxxxxxxx yy</i>	240–1196
OBJECT COPY QUEUE	856
OBJECT STATUS	1256
OBJECT COPY AUDIT	796

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 TS7700 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 request records that are described in 13.4.3, “BVIR Request data format” on page 690 and the response data set contains three explanatory records (Records 3–5) and starting with Record 6, the response to the data request.

For more information about the record formats of the response record, see the following white papers:

- ▶ [IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#)
- ▶ [TS7700 Statistical Data Format](#)

The response data set features the following 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 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 that are based on the request. If the request cannot be understood or was invalid, that issue is indicated.

13.4.5 BVIR Response data

This section explains how to interpret each BVIR Response Data Set for 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 information
- ▶ Physical media pools information
- ▶ Physical volume status information
- ▶ Copy audit information

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 to which the requested volume was written. In a TS7700 grid configuration, separate requests must be sent to each cluster to obtain the volume status information for the individual clusters. By 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 [IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide](#). Fields are presented in the order that is defined in the table and are comma-separated. The overall length of each record is 643 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 13-3.

Example 13-3 BVIR volume status information

```
123456789012345678901234567890123456789012345678901234567890123
ABC123,0,2009-04-22-11.56.45.871263,0,0,32,0,N,2548,N,8719,N...
```

The following 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 was modified and it is not yet copied.

Suppose that you use the deferred Copy Consistency Point (which is typically when is significant distance exists between the TS7700s in the grid configuration). In this situation, some number of volumes are inconsistent between the TS7700s at any point.

If an event 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 event 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 are accessed and closed.

This field can be used to determine which volumes in each TS7700 cluster were not copied to build a set of jobs to access them, and force the copy. The PRESTAGE program from the IBM Tape Tools website can support you in doing that job efficiently. You can use the VESYNC job to identify volumes that need to be copied to a new cluster.

The COPYRFSH program is one of the most recent additions to the tape tools suite. You can use COPYRFSH to queue up copies to one or more new clusters or existing clusters that currently do not have copies.

The COPYRFSH tool incorporates similar intelligence to PRESTAGE allowing the source data to be efficiently recalled from physical tape as part of the replication process.

For more information about various tools that are available for monitoring your TS7700 and helping with MES scenarios, see 13.5.3, “VEHSTATS tool overview” on page 710. You can also access the [IBM Tape Tools website](#).

► 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 completed their copy operations to specific clusters.

► Volume Ownership and Volume Ownership Taken

At any 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 is not transferred until the volume is unloaded. Ownership can be transferred by using one of the following methods:

- Through communication with the current owning cluster
- Through a recovery process that is called *ownership takeover*

Normally, if the cluster that is receiving a mount command does not own 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 that receives 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 failed or only the grid network links to it failed. Operator intervention is required to indicate that the owning cluster failed and that ownership takeover by the other clusters is allowed. The following types of ownership takeover are available:

- Write ownership takeover (WOT): The cluster that is 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, including 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 was deleted on the cluster.

► **Removal State**

As part of normal processing in a TS7700 Grid configuration where a mixture of both TS7700T 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 timestamp, can be used to determine whether the data that is associated with the volume was removed.

► **Hot**

This field represents the cluster's view of which clusters include 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 Insert or Eject pending updates, or regular pending updates.

Insert/Eject updates are related to volumes that are 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 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 are typically all private volumes. However, some might be 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 inconsistencies exist in the mapping information that is provided with this function. This inconsistency 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 issue 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 that are 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 available for the TS7700T only.

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 information

The 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.

Because pool 0 (common scratch pool) contains only empty volumes, only the empty count is returned. Volumes that were 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 are marked as full, is returned. The count for empty includes physical volumes that were assigned to the pool and volumes that were borrowed from the common scratch pool but were not yet 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 include 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 information

A database is maintained on each 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, which is 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 13-4 for Volume A03599. The overall length of each record is 400 bytes with blank padding after the last field, as needed.

Example 13-4 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: Generating the response might take several minutes to complete, depending on the number of volumes that are requested and how busy the TS7700 cluster is at the time of the request.

Copy audit information

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. No validation occurs regarding whether that cluster should have a copy. To consider the copy modes, 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 includes 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 the following 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 issue might occur 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 must be on the disaster recovery TS7700 only 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 were not yet made from a source TS7700 to one or more of the specified clusters. This issue can result because the source TS7700 or the links to it are unavailable, or because a copy policy of Deferred was specified and a copy was not yet completed when the audit was run.
- ▶ Each of the specified clusters contained a valid copy at one time, but since removed it as part of the TS7700 hybrid automated removal policy function. Automatic removal can occur 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 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, which TS7700 clusters are to be audited must be specified. The clusters are specified by using their associated distributed library ID (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.
- ▶ Valid requests include the following examples (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 include a valid copy on any of the specified clusters. Volumes that were never used, had their associated data deleted, or were 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 13-5. The response records are written in 80-byte FB format.

Example 13-5 BVIR message when Copy Audit is requested

```
1234567890123456789012345678901234567890123456789012
L00001,R,D,D,N,N,N,N,N,N,R,N,N,N,N,N,R
```

Tips: The output for Copy Audit includes Copy Consistency Points and Removal States for up to eight TS7700 clusters. The Copy Audit output format 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. The duration depends on the number of logical volumes that were 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 13-6.

Example 13-6 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, then the response file indicates in record 6 that the request is invalid, as shown in Example 13-7.

Example 13-7 BVIR message when an invalid variable is specified

```
12345678901234567890123456
INVALID VARIABLE SPECIFIED
```

13.5 IBM Tape Tools

IBM provides Tape Tools, which are available to monitor your tape environment. Several of these tools are specific to the TS7700 and based on BVIR data, such as VEHSTATS. The following tools are included in IBM Tape Tools:

- ▶ VEHSTATS tool
- ▶ VEHGRXCL tool
- ▶ VEHAUDIT tool
- ▶ Other tools

13.5.1 IBM Tape Tools overview

All the TS7700 monitoring and evaluating tools that are described in this section are available at [IBM Tape Tools website](#).

The content of the Readme.txt file that provides basic information about the tape tools is listed in Example 13-8 on page 704.

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, the results will be unpredictable. To avoid this situation, you need to be informed of these enhancements so you can stay current.

To be informed of major changes to any of the tools distributed by using this FTP site, send an email message to:

Tape.Tools@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 contains a chronological history of all changes that are made to the tools.

Review that file regularly, at least monthly, perhaps weekly, so you can see if any changes apply to you.

Look in the file, OVERVIEW.pdf, for an overview of all currently available tools.

The JCL, CNTL, and LOAD libraries for all the tools are stored in the files ibmjcl.xmi, ibmcntl.xmi and ibmload.xmi.

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 comments 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.

Register the change distribution list that is referenced in the Readme.txt file by sending an email to Tape.Tools@ibm.com. You are updated about any significant changes that are made to any of the tools that are distributed by way of this website after you register.

Most of these tools are z/OS based and can be downloaded from [IBM Tape Tools website](#). At the web location, you find the overview.pdf file that contains a list of the available tools for you to download.

Tools that might be of interest are listed in Table 13-5 on page 705.

Table 13-5 Tape tools selection

Tool (job name)	Major use	Benefit	Inputs	Outputs
The program to extract BVIR data from TS7700				
BVIRHSTx	Get historical stats from TS7700	Statistics file (BVIR historical file)	TS7700	Creates U, VB, SMF format
BVIRMES	Get BVIR Volume Status information	Volume Status file	TS7700	Creates VOLFILE data set
BVIRPHY	Get BVIR Physical Volume Status information	Physical Volume Status file	TS7700	Creates PHYFILE data set
BVIRPIT	Get Point in Time stats from TS7700	Statistics file (BVIR PIT file)	TS7700	Creates BVIRPIT data file
BVIRPOOL	Identify available scratch by pool	Reports all pools at once	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 and Tape Catalog (CA1, TLMS, RMM, ZARA, CTLT)	Logical volumes by job name or dsname, logical to physical reports
BVIRVTS	Get Cache Content information, Physical Volume to Logical Volume Mapping Information, Copy Audit Information	Cache Content, Volume Map, Copy Audit	TS7700	Creates CACHFILE, VOLFILE, AUDFILE
BVPITRPT	Get Point in Time stats from TS7700 and then run VEPSTATS	Immediately available	TS7700	Point in Time stats
Programs for creating different reports				
EXPDIST	Quantify the number of volumes expiring n days from now	Determine the rate of return to scratch	CA1, TLMS, RMM, ZARA, CTLT	Volume count distribution by media
FINDLRG*	Identify multi-volume tape data sets	Creates filter list to separate 3590 workload for Batch Magic	CA1, TLMS, RMM, ZARA, CTLT	Data set length
GETVOLS	Get VOLSERS from list of dsns	Automate input to PRESTAGE	CA1, TLMS, RMM, ZARA, CTLT	VOLSERS for requested dsns
GRPDSN	Generalize dsname lists	Speeds up the process of generalizing dsname lists, which makes tape study more accurate	Dsname list	Generalized dsname list

Tool (job name)	Major use	Benefit	Inputs	Outputs
LASTLIST	Report volsers last referenced between SDATE/EDATE	Determine how many volsers were last referenced during a period	CA1, TLMS, RMM, ZARA, CTLT	List of volsers and dsnames
OFFSITE*	Identifies data sets sent offsite	Creates filter list to separate offsite workload	CA1, TLMS, RMM, ZARA, CTLT	Report plus disk file filter list of offsite dsnames
ORPHANS*	Identify orphan data sets in Tape Management Catalog	Clean up tool	CA1, TLMS, RMM, ZARA, CTLT	Listing file showing all multi-occurrence GDGs that have not been created in the last nn days
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, TGROUPE, and system reporting
TMCREUSE	Identify data sets with create date equal to last ref date	Get candidate list for VTS PG0	CA1, TLMS, RMM, ZARA, CTLT	Filter list of potential PG0 candidates
TVCBYDSN	List data sets in VTS cache	Verify that cache management policies are working as expected	BVIR data and CA1, TLMS, RMM, ZARA, CTLT	List of dsns in the cache and the order they are flushed
VEHAUDIT	Shows where each LVOL resides in a TS7700 grid	Shows where each LVOL resides in a TS7700 grid	BVIR data: CACHFILE, VOLFILE, MESFILE and CA1, TLMS, RMM, ZARA, CTLT	<ul style="list-style-type: none"> ▶ DTLRPT showing the attributes and location of lvols in TS7700 ▶ COPYDIST shows copy times for peer copies
VEHSTATS_M ODEL (xls & ppt)	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 records and interval
VEHSTATS*	TS7700 historical performance reporting	Show activity on and performance of TS7700	BVIRHSTx file	Reports showing mounts, data transfer, box usage, and additional information
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

Tool (job name)	Major use	Benefit	Inputs	Outputs
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 tape management catalog. 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
VOLREUSE*	Show volume reuse activity	Identify data sets for VTS cache management	SMF 14, 15, 21, 30	OUTRPT reuse distribution
GGM – Grid-to-Grid Migration				
BVIRGGM	Provides the response of all the GGM copy activity records for already completed GGM copies	Records provide information for each GGM copy volume	TS7700	Creates GGMFILE data set
GGMTOOL1	It is used to create the lists of “GGM COPY” commands for the volumes with STATUS(MASTER)	List of commands for execution	RMM	The list of command – inputs for the job ISSUECMD
GGMTOOL2	It is used to create the lists of “GGM COPY” commands for the volumes with STATUS(SCRATCH)	List of commands for execution	RMM	The list of command – inputs for the job ISSUECMD
GGMTOOLA	It is used to create the lists of “GGM COPY” commands for the volumes with STATUS(MASTER)	List of commands for execution	RMM	The list of command – inputs for the job ISSUECMD
GGMTOOLB	It is used to create the lists of “GGM COPY” commands for the volumes with STATUS(SCRATCH)	List of commands for execution	RMM	The list of command – inputs for the job ISSUECMD
GGMSORT	Sorts command list (GGMQUE created by GGMTOOL1/2/A/B job) by physical volumes from the “volume map” (VOLFILE)	This is needed if source TS7700 is with tapes (TS7740 or TS7720T or TS7760T) to reduce migration time	Command list (GGMQUE created by GGMTOOL1/2/A/B job)	Sorted command list by physical volumes

Tool (job name)	Major use	Benefit	Inputs	Outputs
BVIRGGM	Provides the response of all the GGM copy activity records for already completed the GGM copies	Records provide information for each GGM copy volume	TS7700	Creates GGMFILE data set – input for the job GGMINFO
Programs for copying volumes from one cluster to another				
COPYRFSH	Cause peer copies to be made in the newly added TS7700 cluster	No host mounts involved	List of lvols needing copy	LIB REQ commands
PRESTAGE*	Recall lvols to VTS	Ordered and efficient	BVIR VOLUME MAP	Jobs submitted to recall lvols
Other programs				
FORMCATS	Create a common record format	Single format for all tape catalogs	CA1, TLMS, RMM, ZARA, CTLT	Common extract file from different tape catalogs
ISSUECMD	Run “Library Request” commands	Possibility to send commands from host to TS7700	The list of commands	Command are running on TS7700
MOUNTMON*	Monitor mount pending and volume allocations	Determine accurate mount times and concurrent drive allocations	Samples tape UCBs	Detail, summary, distribution, hourly, TGROUPE, system reporting
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

13.5.2 IBM Tape Tools installation

Public access is provided to the IBM Tape Tools [website](#), which contains various tools that can help you analyze your tape environment. Figure 13-24 on page 709 shows several tools that are available from the [website](#).

Index of /storage/tapetool/		
Name	Size	Date Modified
[parent directory]		
A_License_Agreement_for_IBM_Tape_Tools.pdf	21.9 kB	1/28/11, 9:00:00
Synchronizing_cluster_copies_in_TS7700.pdf	68.7 kB	2/13/18, 9:00:00
TS7680_Statistics_Report_Install.doc	46.0 kB	1/19/11, 9:00:00
TS7700.VEHSTATS.Decoder.V22a.pdf	1.5 MB	4/15/19, 9:00:00
TS7700.VEHSTATS.Decoder.V23.pdf	2.1 MB	12/20/19, 9:00:00
TS7700_Starting_with_VEHSTATS.pdf	2.2 MB	5/22/17, 9:00:00
TapeWise.PDF	325 kB	8/17/05, 9:00:00
VEHAUDIT_How_to_read_the_report_CPYDST.doc	45.0 kB	6/27/16, 9:00:00
VEHAUDIT_How_to_read_the_report_DTLRPT-V2018-04-03.pdf	297 kB	4/3/18, 9:00:00
VEHGRXCL.txt	3.0 kB	11/8/14, 9:00:00
VEHSTATS_MODEL.ppt	1.6 MB	12/3/14, 9:00:00
VEHSTATS_MODEL.xls	1.9 MB	12/3/14, 9:00:00
VEHSTATS_MODEL_Day.ppt	909 kB	12/3/14, 9:00:00
VEHSTATS_MODEL_v21.ppt	2.3 MB	4/19/17, 9:00:00
VEHSTATS_MODEL_v21.xls	3.7 MB	4/19/17, 9:00:00
VEHSTATS_user_manual.pdf	1.0 MB	5/2/18, 9:00:00
VTSANAL9.XLS	1.5 MB	1/31/15, 9:00:00
VTSGRPH9.ppt	736 kB	1/31/15, 9:00:00
VTSGRXCL.TXT	1.5 kB	2/2/15, 9:00:00
badblksz.txt	1.8 kB	8/5/04, 9:00:00
eximcalc.123	10.6 kB	3/2/07, 9:00:00
export.txt	539 B	8/5/04, 9:00:00
findlr.txt	1.0 kB	5/7/02, 9:00:00
fsrtmm.txt	972 B	5/7/02, 9:00:00
ftpcust.txt	1.0 kB	3/21/02, 9:00:00
ggmintro.txt	2.5 kB	11/3/15, 9:00:00
grpdsn.txt	1.7 kB	8/5/04, 9:00:00
ibmcntl.xmi	2.9 MB	7/23/20, 11:41:00
ibmjcl.xmi	1.8 MB	10/16/20, 10:41:00
ibmload.xmi	3.5 MB	10/16/20, 10:41:00
ibmpat.xmi	317 kB	9/6/14, 9:00:00
ibmtools.txt	7.0 kB	4/16/18, 9:00:00
ifasmfdp.txt	2.0 kB	8/5/04, 9:00:00
libmangr.pdf	150 kB	1/21/05, 9:00:00
mountmon.pdf	277 kB	8/16/05, 9:00:00
mountmon.txt	1.7 kB	5/7/02, 9:00:00
mountrec.txt	7.8 kB	5/14/07, 9:00:00
mountsmf.txt	19.9 kB	5/14/07, 9:00:00
offsite.txt	589 B	8/5/04, 9:00:00
orphans.txt	572 B	8/5/04, 9:00:00
overview.pdf	172 kB	11/1/17, 9:00:00
overview_v2012.pdf	18.1 kB	10/31/17, 9:00:00
prestige.txt	2.8 kB	8/29/06, 9:00:00
pt_hourly_report_install.bsx	4.3 MB	3/3/11, 9:00:00
readme.txt	1.5 kB	4/16/18, 9:00:00

Figure 13-24 Tape tools catalog

IBM Tape Tools are accessible at [IBM Tape Tools website](#).

- ▶ Text files are available for most tools. In addition, each job to run a tool contains a detailed description of the function of the tool and the parameters that must be specified. To obtain the tape tools, download the XMI files to your computer through FTP from Time Sharing Option (TSO) on your z/OS system:
- ▶ 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 needed only if you run the QSAMDRVR utility.

The `ibmtools.txt` file contains detailed information about how to download and install the tools libraries.

After you create 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($$CPYLIB)` to create a JCL library that has a unique second node (&SITE symbolic). This step creates a private JCL library for you from which you can submit jobs while leaving the original as is. CNTL and LOAD can then be shared by multiple users who are running jobs from the same system.
2. Edit and submit `userid.SITENAME.IBMT00LS.JCL($$TAILOR)` to tailor the JCL according to your system requirements.

The `updates.txt` file contains all fixes and enhancements that are made to the tools. Review this file regularly to determine whether any of the programs that you use were modified.

To ensure that you are not working with outdated tools, the tools are controlled through an EXPIRE member. Several times per year (like every quarter), a new EXPIRE value is issued that is valid for the next 12 months. When you download the current tools package anytime during the year, you have at least nine months remaining on the EXPIRE value.

If your IBM tools jobs stop running because the expiration date passed, download from IBM Software Tape Tools again to get the current `IBMT00LS.JCL(EXPIRE)` member.

13.5.3 VEHSTATS tool overview

VEHSTATS is a tool that is included in IBM Tape Tools.

The TS7700 activity is recorded in the subsystem. The following types of statistics are available:

- ▶ Point-in-time: A snapshot of activity in the last 15 seconds
- ▶ Historical: 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 923).

Because both types of statistical data are delivered in binary format from the BVIR functions, you must convert the content into a readable format.

This conversion can be done manually by using the information that is provided in the following documents:

- ▶ [TS7700 Statistical Data Format](#)
- ▶ [IBM Virtualization Engine TS7700 Series VEHSTATS Decoder](#)

Alternatively, you can use an existing automation tool. IBM provides a historical statistics tool called VEHSTATS. For more information about interpreting the reports, see 13.5.5, “VEHSTATS reports” on page 712.

You can use VEHSTATS to monitor TS7700 virtual and physical tape drives, and TVC activity to produce 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 13-32 on page 726 highlights the following 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 files `VEHSTATS_Model.ppt` and `VEHSTATS_Model.xls` can be used 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. The VEHGRXCL.txt file contains the description for the graphical package.

A tool that is called VEPSTATS is also available. It is used for working with Pint-in-time statistics.

13.5.4 Running the VEHSTATS jobs

The [IBM TS7700 Series Tape Tools VEHSTATS user manual](#) contains detailed information how to run VEHSTATS. Several output options are available for VEHSTATS, and you must submit separate jobs, depending on your requirements. The IBMT00LS.JCL member VEHSTATS (see Example 13-9) provides guidance about which job to choose.

Example 13-9 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 IBMT00LS 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 use the disk flat file method. The flat files can be recorded as RECFM=U or RECFB=VB.

The following specific jobs in IBMT00LS.JCL are designed to fit your particular needs:

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 include 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 13-10 must be specified.

Example 13-10 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
```

The fields that are shown in the various reports depend on which ORDER member in IBMT00LS.JCL is being used. Complete the following steps to ensure that the reports and the flat file contain all of the 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 were deactivated as indicated by an asterisk in the first column. Example 13-11 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 the cache to be displayed in MB or in GB.

Example 13-11 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 [IBM Tape Tools](#), specify the ORDERV12 member because it contains all the fields that are used when creating the graphics. Also, verify that all statements are activated for all clusters in your environment.

Note: Current VEHSTATS does not accept SMF records as input.

13.5.5 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 occurring. Also, comparative analysis can be used to determine whether an upgrade, such as adding physical tape drives, might improve the overall performance of the TS7700T. VEHSTATS is not a projection tool, but it provides the basis for an overall health check of the TS7700.

VEHSTATS provides much information. The following most important reports are available for the TS7700, and the results and analysis that can help you understand the reports better:

- ▶ H20VIRT: vnode Virtual Device Historical Records
- ▶ H21ADP0x: vnode Adapter Historical Activity
- ▶ H21ADPXX: vnode Adapter Historical Activity combined (by adapter)
- ▶ H21ADPSU: vnode Adapter Historical Activity combined (total)
- ▶ H30TVCx: hnode HSM Historical Cache Partition:
 - For a TS7700D, this report represents the cache in H30TVC1.
 - For a TS7700T or TS7700C, multiple TVCs (TVC2 and TVC3) are presented. TVC1 contains the data from CP0, and TVC2 contains the data of CP1, and so on.
- ▶ H31IMEX: hnode Export/Import Historical Activity
- ▶ H32TDU12: hnode Library Historical Drive Activity
- ▶ H32GUPXX: hnode Library Historical General Use Pools 01/02–31/32
- ▶ H32CSP: hnode Library Historical Scratch Pool Activity
- ▶ H33GRID: hnode Grid Historical Peer-to-Peer (PTP) Activity
- ▶ H35CLOCL: hnode Cloud Historical Activity by Clusters (new at R5.1)
- ▶ H35CLOID: hnode Cloud Historical Activity by Pool IDs (new at R5.1)
- ▶ H36OBJSG: hnode Object Store General (new at R5.2.2)
- ▶ H37CLOSN: hnode Object Store Activity by Clusters (new at R5.2.2)

- ▶ H37OSNCL: hnode Object Store Activity by Store Names (new at R5.2.2)
- ▶ H38OSNPT: hnode Object Store by Name and Partition (new at R5.2.2)
- ▶ AVGRDST: Hrs Interval Average Recall Mount Pending Distribution
- ▶ DAYMRY: Daily Summary
- ▶ MONMRY: Monthly Summary
- ▶ COMPARE: Interval Cluster Comparison
- ▶ HOURFLOW: 15-minute interval or 1-hour interval overall throughput report
- ▶ DAYHSMRY: Daily flat file

The following sections describe the contents of some of these reports.

Tip: Be sure that you have a copy of the white paper [TS7700 Statistical Data Format](#) and the [IBM TS7700 Series VEHSTATS Decoder](#) available as you become familiar with the VEHSTATS reports.

H20VIRT: vnode Virtual Device Historical Records

Example 13-12 on page 714 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:

- ▶ MIN: The minimum number of mounted virtual drives.
- ▶ AVG: The average number of mounted virtual drives.
- ▶ MAX: The maximum number of mounted virtual drives.
- ▶ MAX THRPUT: The enabled throughput limit with FC 5268.
- ▶ ATTMPT THRPUT: The amount of MBps the host wanted to deliver to TS7700.
- ▶ Delay MAX: The maximum total throughput delay over the 15-minute interval. The total throughput delay is accumulated for each 15-second period within the 15-minute interval.
- ▶ Delay AVG: The average total throughput delay over the 15-minute interval. The total throughput delay is accumulated for each 15-second period within the 15-minute interval.
- ▶ PCT OF 15 sec INTVLS: How the delay occurred each second during the 15-second sample interval.
- ▶ The number of channel blocks that are 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 13-12 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 13-13 shows the second half of the report.

Example 13-13 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 (called *behind*). Normally, the numbers in both columns are shown.

Use the ratio of both numbers to understand the performance of your TS7700. In the 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 block sizes, but this rule is not fixed. For example, DFSMSshm writes a 16,384 block size, and Db2 writes block sizes up to 256,000. The report contains more differences for block sizes, but these are not shown in the example.

H21ADP: vnode Adapter Historical Activity

The next example report provides details about the vnode Host Adapter Activity. Although a large amount of information is 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 13-14 shows a sample report for Adapter 3 of Cluster 0.

Example 13-14 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
19JUL10M0 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
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
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
04:00:00  4 0  0 0  69 0  0 0  24 2.87  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
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
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
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
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

```

The following fields are the most important fields in this report:

- ▶ GBS_RTE: This field shows the 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, as shown in Example 13-15.

Example 13-15 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. It also shows how many one-hour intervals showed which throughput.

For example, consider the following points regarding the second line of the report data:

- ▶ The throughput was 51–100 MBps in 191 intervals.
- ▶ A total of 191 intervals are 25.8% of the entire measurement period.
- ▶ In 90.2% of the measurement intervals, the throughput was less than 100 MBps.

H30TVCx: hnode HSM Historical Cache Partition

This report provides details of Cache Partitions Activity in the TS7700. The total TVC used is reported in TOTAL TVC_GB USED. For a TS7700D, only one H30TVC1 is provided. For a TS7700T or TS7700C, multiple H30TVCx can be shown. The H30TVC1 for a TS7700T or TS7700C contains the information of the resident partition. If you 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 number of logical volumes that were removed by using autoremoval
- ▶ The value and reason of the throttle

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 more cache capacity is needed, or cache management policies must be adjusted. For a TS7700T or TS7700C, you might also reconsider the Tape or Cloud Partition sizes.
- ▶ FAST_RDY_AVG_SECS and CACHE_HIT_AVG_SECS must 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 13-16.

Example 13-16 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

Synchronous mode monitoring

The information for synchronous mode is shown in H30TVCx, independently for each partition. As in the MI, all numbers that are related to the synchronous mode operation are shown only in the Host I/O cluster for these mounts. The “secondary” synchronous mode does not reflect synchronous mounts.

In the upper part of Figure 13-25 on page 717, you see that there were 44 scratch mounts (FAST NUM MNTS) at 07:15:00. In addition, you see the same number in the SYNC NUM MNTS field, which means that they were all scratch mounts and the same number of mounts to a remote cluster was run.

The receiver of the synchronous copy reports nothing, as shown in the lower part of Figure 13-25 on page 717.

1(C) IBM REPORT=H30TVC1 (16040)															HNODE HSM HISTORICAL CACHE PARTITION														
GRID#=-00186 DIST_LIB_ID= 2 VNODE_ID= 0 NODE_SERIAL=CL2H6395 VE_CODE_LEVEL=008.033.000.00															TVC_SIZE= 23858GB														
PARTITION SIZE= 9858GB															TVC_SIZE= 23858GB														
24APR16SU															24APR16SU														
RECORD	AVG	MAX	AVG	MAX	PART	TOTAL		FAST_RDY		CACHE_HIT		CACHE_MIS		SYNC_MODE		P-MIG													
END_TIME	CPU_UTIL	DISK_UTIL	HIT%			NUM	AVG	NUM	AVG	NUM	AVG	NUM	AVG	NUM	AVG	NUM	AVG												
						MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS												
	R2.0	>R3.0	.0062												R2.1	R1.5													
07:00:00	4	11	0	1		0		0	.00	0	.00	0	.00	0	.00	0	.00												
07:15:00	13	38	31	89	100	88	.71	44	.71	0	.00	0	.00	44	.71	1000													
07:30:00	23	33	67	95	100	194	.64	97	.64	0	.00	0	.00	97	.64	1000													
07:45:00	30	39	87	100	100	150	.85	75	.85	0	.00	0	.00	75	.85	1000													
08:00:00	23	29	64	91	100	128	.58	64	.58	0	.00	0	.00	64	.58	1000													
08:15:00	14	29	33	87		0		0	.00	0	.00	0	.00	0	.00	1000													
08:30:00	11	34	23	99	100	64	.77	32	.77	0	.00	0	.00	32	.77	1000													
08:45:00	29	41	85	100	100	192	.81	96	.81	0	.00	0	.00	96	.81	1000													
09:00:00	25	33	68	100	100	76	.81	38	.81	0	.00	0	.00	38	.81	1000													
09:15:00	13	29	35	76	100	54	.62	27	.62	0	.00	0	.00	27	.62	1000													
09:30:00	14	27	42	86	100	94	.58	47	.58	0	.00	0	.00	47	.58	1000													
09:45:00	4	9	2	17		0		0	.00	0	.00	0	.00	0	.00	1000													

1(C) IBM REPORT=H30TVC1 (16040)															HNODE HSM HISTORICAL CACHE PARTITION														
GRID#=-00186 DIST_LIB_ID= 1 VNODE_ID= 0 NODE_SERIAL=CL1H3833 VE_CODE_LEVEL=008.032.003.00															TVC_SIZE= 5999GB														
PARTITION SIZE= 5999GB															TVC_SIZE= 5999GB														
24APR16SU															24APR16SU														
RECORD	AVG	MAX	AVG	MAX	PART	TOTAL		FAST_RDY		CACHE_HIT		CACHE_MIS		SYNC_MODE		P-MIG													
END_TIME	CPU_UTIL	DISK_UTIL	HIT%			NUM	AVG	NUM	AVG	NUM	AVG	NUM	AVG	NUM	AVG	NUM	AVG												
						MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS	MNTS	SECS												
	R2.0	>R3.0	.0062												R2.1	R1.5													
07:00:00	4	10	0	1		0		0	.00	0	.00	0	.00	0	.00	0	.00												
07:15:00	13	31	23	69		0		0	.00	0	.00	0	.00	0	.00	0	.00												
07:30:00	24	38	47	82		0		0	.00	0	.00	0	.00	0	.00	0	.00												
07:45:00	28	35	60	85		0		0	.00	0	.00	0	.00	0	.00	0	.00												
08:00:00	23	35	44	74		0		0	.00	0	.00	0	.00	0	.00	0	.00												
08:15:00	13	24	22	55		0		0	.00	0	.00	0	.00	0	.00	0	.00												
08:30:00	10	28	14	49		0		0	.00	0	.00	0	.00	0	.00	0	.00												
08:45:00	25	34	46	64		0		0	.00	0	.00	0	.00	0	.00	0	.00												
09:00:00	24	39	49	87		0		0	.00	0	.00	0	.00	0	.00	0	.00												
09:15:00	16	37	30	78		0		0	.00	0	.00	0	.00	0	.00	0	.00												
09:30:00	18	24	27	40		0		0	.00	0	.00	0	.00	0	.00	0	.00												
09:45:00	5	9	3	18		0		0	.00	0	.00	0	.00	0	.00	0	.00												

Figure 13-25 Synchronous mode copies shown in VEHSTATS

No other information (such as the number of concurrent sync copies and inconsistency at interval) for the synchronous mode is available because none of NUM MNTS is applicable on the receiver cluster.

Only if the synchronous mode copy cannot be processed and sync-deferred is activated are reports written. However, these copies are reported with “DEFERRED” and a drill-down cannot be done.

Throttle reason monitoring

In the H30TVCx report for each interval the throttling is reported for each of the different throttling types. To interpret these values, review the throttling reasons in the following figures. Figure 13-26 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 13-26 Host Write Throttling reasons

Figure 13-27 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 13-27 Copy Throttling reasons

Each value is treated as a bit so if a value of “x03” is shown in the H30TVC, which 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 occurred, and the relative throttling impact value (%RLTV IMPAC VALUE). Even if throttling occurred, this issue might occur 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 run time.

Deferred Copy Throttle monitoring

As shown in Figure 13-28, the H30TVC report contains the following information about deferred copy 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 13-28 Deferred copy throttling

- ▶ NUM 15 MIN INTVL: Shows the number of intervals in the hour that 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 occurred.
- ▶ NUM 30 SEC SMPLS: Shows in how many 30-second samples in the reported intervals the throttling occurred. Therefore, in an hour report, a maximum of 120 samples (60 minutes * 2 samples of 30 seconds each) are available. If 1 is reported, a deferred throttling occurred only in 30 seconds of the whole hour.
- ▶ AVG SEC INTVL: Shows the amount of penalty in seconds that is given for each deferred copy action in this interval.

Looking at Figure 13-28 on page 718, you find one interval where deferred copy throttling occurred in all 120 samples. The maximum value of 0.125 s penalty for each copy operation result. However, in the report that is shown, no interval existed in which no deferred copy throttling occurred, but limited throttling was measured in some intervals.

Be aware that depending on your network, a throttling higher than 20 ms normally results in little or no deferred copy action. For more information how to influence the deferred copy throttling, see 14.7.6, “Tuning possibilities for copies: Deferred copy throttling” on page 773

H32TDU12: hnode Library Historical Drive Activity

The report H32TDU12 gives you an overview about the usage of the backend drives. Example 13-17 shows the following information:

- ▶ 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
 - TOT: Total number of STG, MIG, RCM, and SDE

Example 13-17 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, it 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 result might indicate a need to evaluate the cache size or cache management policies. However, this rule is not a fixed rule and more analysis is required. For example, if HSM migration is into a TS7700T, you might see high recall activity during

the morning, which can be driven by temporary development or user activity. This result is normal and not a problem.

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 issue might be performance-related. To confirm whether that issue is a bottleneck, the overall situation must be carefully reviewed.

To do so, first review which mounts are run. If a reclaim is still processed, no performance issue occurred (except if a panic reclaim occurred).

If no reclaim was run, review how long in average a physical cartridge was mounted. To calculate this value, divide the total mounts by the number of installed drives and the number of intervals sample. That result shows how long a physical tape cartridge can be mounted on a physical tape drive. If this value is less than 10 minutes, more investigation should be done. If this value is less than 4 minutes, a performance issue is likely.

H32GUPXX: hnode Library Historical General Use Pools

The H32GUPXX (General Pool Use) report is shown in Example 13-18. 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:

- ▶ Number of active logical volumes
- ▶ Amount of active data in GB
- ▶ Amount of data written in MB
- ▶ Amount of data read in MB
- ▶ Current reclamation threshold and target pool

Example 13-18 VEHSTATS report for General Pool Use

(C) IBM REPORT=H32GUP01(10210)										HNODE LIBRARY HIST GUP/POOLING ACTIVITY										RUN ON 18JUL2010 @ 16:57:51				PAGE 03	
GRID#=CC001 DIST_LIB_ID= 0 VNODE_ID= 0										NODE_SERIAL=CLOABCDE VE_CODE_LEVEL=008.006.000.0110										-				UTC NOTCHG	
18JUL10 POOL 01 3592-E05 3592JA										READ UN- POOL 02 3592-E05										READ UN-					
RECORD	ACTIVE	ACTIVE	MB	MB	VOL_COUNT	RECLAIM-	ONLY	AVAI		ACTIVE	ACTIVE	MB	MB	VOL_COUNT	RECLAIM-	ONLY	AVAI								
TIME	LVOLS	GB	WRITTN	READ	SCR	PRIV	PCT	POOL	VOLS	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	0	0	0	0	0	0	25	02	00	00						
4:30:00	65079	18052	37888	0	2	56	25	01	00	0	0	0	0	0	0	25	02	00	00						
4:45:00	65079	18052	83895	0	2	56	25	01	00	0	0	0	0	0	0	25	02	00	00						
5:00:00	65630	18206	94721	0	2	57	25	01	00	0	0	0	0	0	0	25	02	00	00						
5:15:00	65630	18206	98630	0	2	57	25	01	00	0	0	0	0	0	0	25	02	00	00						
5:30:00	65630	18206	124490	0	2	57	25	01	00	0	0	0	0	0	0	25	02	00	00						
5:45:00	65630	18206	119979	0	2	57	25	01	00	0	0	0	0	0	0	25	02	00	00						
6:00:00	67069	18610	108854	0	2	57	25	01	00	0	0	0	0	0	0	25	02	00	00						
6:15:00	67069	18610	108854	0	2	57	25	01	00	0	0	0	0	0	0	25	02	00	00						
6:30:00	67069	18610	97126	0	2	57	25	01	00	0	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.

H32CSP: hnode Library Historical Scratch Pool Activity

You can use the H32GUPXX report to view the cartridges on a per pool base. In addition, use the H32CSP report to see the trend of empty cartridges. This report provides an overview of the empty cartridges in the common scratch pool on a cartridge type basis, as shown in Example 13-19.

Example 13-19 VEHSTATS report for Common Scratch Pool

18JUL10 -----SCRATCH_STACKED_VOLUMES_AVAILABLE_BY_TYPE-----										
RECORD	3590J	3590K	3592JA	3592JJ	NONE	NONE	3592JB	NONE		
TIME	MEDIA0	MEDIA1	MEDIA2	MEDIA3	MEDIA4	MEDIA5	MEDIA6	MEDIA7		
4:15:00	0		0		42		0		0	0
4:30:00	0		0		42		0		0	0
4:45:00	0		0		42		0		0	0
5:00:00	0		0		41		0		0	0

It is *not* sufficient to check only the scratches in the common scratch pool. In addition, you must check that all pools can borrow from the CSP and return the empty cartridges to the CSP. If a pool is set to **no borrow**, ensure that enough empty cartridges are always in inside this pool. This number is reflected in the H32GUPXX reports.

In a heterogeneous environment, back-level cartridges (JA/JB) can be used only for read purposes and *not* write purposes.

H33GRID: hnode Historical Peer-to-Peer (PTP) Activity

The Peer-to-Peer Activity report that is shown in Example 13-20 provides various performance metrics of grid activity. This report can be useful for installations that are 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 13-20 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	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO	TO									
RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE	RECEIVE									
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---									
MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---	MINUTES---									
01:00:00	1	13	1	0	139077	38.6	43	1	346	61355	17.0	746	0.2	156	0.0											
02:00:00	6	1518	7	0	150440	41.7	84	462	11410	64536	17.9	4448	1.2	1175	0.3											
03:00:00	2	3239	3	0	88799	24.6	38	8	44	57164	15.8	1114	0.3	166	0.0											
04:00:00	2	574	4	0	241205	67.0	4	82	29	109850	30.5	1409	0.3	401	0.1											
05:00:00	3	1055	2	0	70637	19.6	9	390	136	51464	14.2	2488	0.6	0												
06:00:00	16	9432	2	0	187776	52.1	33	1519	491	100580	27.9	2526	0.7	463	0.1											
07:00:00	0	0	0	0	86624	24.0	19	63	12649	50139	13.9	6036	1.6	1988	0.5											
08:00:00	1	484	0	0	46314	12.8	26	30	12292	23216	6.4	9563	2.6	1971	0.5											

For the time of the report, you can identify the following items in 15-minute increments:

- ▶ Number of logical volumes to be copied (valid only for a multi-cluster grid configuration)
- ▶ Amount of data to be copied (in MB)
- ▶ Average age of copy jobs on the deferred and immediate copy queue
- ▶ Amount of data (in MB) to and from the TVC that is driven by copy activity
- ▶ 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 13-20, you see three active clusters with write operations from a host. Although this configuration might not be common, 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 that 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.

RUN copies monitoring

The information about the RUN copies is contained in the H33GRID report in the receiving cluster, as shown in Figure 13-29. To see the entire grid performance, review each cluster individually. To understand how much RUN copies (amount of lvols and MB) were processed in an interval, look to LVOLS TO_TV_C_BY_RUN_COPY.

To understand if RUN copies were queued for processing in that interval, review AV_RUN_QUEAGE -- MINUTES--. Having numbers here means that RUN cannot be processed. Having RUN lvols waiting also means that the job cannot process further. If the report shows multiple indications for this behavior, closely review 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 13-29 shows the H33GRID report.

08NOV15SU	LVOLS TO RECEIVE	MiB TO RECEIVE	AV_DEF QUEAGE ---MINUTES---	AV_RUN QUEAGE	#_LVOLS TIM_DLY CPY_QUE	LVOLS TO_TV_C_BY RUN_COPY	MB_ TO_TV_C_BY DEF_COPY	LVOLS TO_TV_C_BY DEF_COPY	MB_ TO_TV_C_BY DEF_COPY
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 13-29 H33Grid report to see copy behavior for RUN and Deferred copies

Deferred copies monitoring

For deferred copies, you have only the AV_DEF QUEAGE -- MINUTES --- field in the H33GRID report. 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, review 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. For more information about deferred copy throttle in the H30TVCx report, see “Deferred Copy Throttle monitoring” on page 718

HOURLFLOW: 15-minute or 1-hour interval overall throughput report

In the VEHSTATS reports, you find a report that is called HOURLFLOW. This report gives you per cluster the overall compressed throughput on a 15-minute or 1-hour interval, and which task is using the cache bandwidth. Figure 13-30 shows such a report.

REPORT=HOURLFLOW (16090)		DATA FLOW IN MiB/sec BY CLUSTER										RUN ON 09MAY2016 @ 14:12:40				PAGE 1						
F DIST_LIB_ID=00		NODE_SERIAL=CL0H7044 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	MiB/s Fr_TVC Dev_Wr	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 DCThrt	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	19	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	0	7	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	25.7	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 13-30 Hourflow Report of VEHSTATS

The MiB/s Total Xfer is an average value. Notice that some peaks maybe higher.

In this example, the maximum total data transfer value is 497 MBps (MiB/s Total Xfer), but that value is mainly driven by the premigration read task (MiB/s Fr_TVC PreMig). If premigration causes a performance issue, review the PMPRIOR and PMTHLVL settings for tuning.

Summary reports

In addition to daily and monthly summary reports per cluster, VEHSTATS 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.

13.5.6 Performance evaluation tool for VEHSTATS reports

After you collect the statistics data from your clusters, you can use available tools to format and plot the binary data.

For more information, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

When evaluating performance, a graph that reveals a significant amount of information succinctly is the *cache throughput* for a cluster graph.

Performance tools are available on the IBM Techdocs website 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 material does not contain any information about the specifics of a Tape Attach model. However, the other information is still valuable, and did not change, especially how to create the Excel spreadsheet and the charts.

For more information about performance tools and how to use them, see the following Techdocs web pages:

- ▶ [Tools](#)
- ▶ [Class replay](#)

The 24-hour, 15-minute data spreadsheets include the cache throughput chart. Figure 13-31 shows the cache throughput that was plotted from VEHSTATS.




Figure 13-31 Cache throughput plotted from VEHSTATS

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 MiBps):

- ▶ Compressed host write: The MiBps of the data that is written to the cache. This bar is hunter green.
- ▶ Compressed host read: The MiBps of the data read from cache. This bar is lime green.
- ▶ Data that is copied out from this cluster to other clusters: 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.

Consider the following points:

- For a two-cluster grid, this value is a single value.
- For a three-cluster grid, two values are used: one for copies to each of the other clusters.
- For a four-cluster grid, three values are used: one for copies to each of the other clusters.

The descriptions of the following 5 items are for plotting the VEHSTATS data source Cluster'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: 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: The rate at which data is read from cache and written to physical tape. This bar is yellow.

- ▶ Compressed data recalled from tape to cache: The rate at which data is read from tape into a cache for a mount that requires a recall. This bar is dark blue.
- ▶ Compressed remote reads from this cluster: The rate that other clusters use this TVC as an I/O cache for read. This bar is orange.
- ▶ Compressed remote writes to this cluster: 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. An accompanying data collection guide is available for each. The first worksheet in each spreadsheet includes instructions for populating the data into the spreadsheet. A guide to help with the interpretation of the 90-day trends is also included.

Separate one-week spreadsheets are available for two-, three-, four-, five-, and six-cluster grids. The spreadsheets use the one-hour interval data to produce charts for the one-week period. A data collection guide is also available.

Separate 24-hour spreadsheets are available for two-, three-, four-, five-, and six-cluster grids. The spreadsheets use the 15-minute interval data to produce charts for the 24-hour period. A data collection guide is also available.

These spreadsheets are intended for experienced TS7700 users. A detailed knowledge of the TS7700 is expected, and familiarity with the use of spreadsheets.

13.5.7 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.txt file contains the description of VEHSTATS_Model.ppt and VEHSTATS_Model.xls. You can use these files to create graphs on cluster activity that is 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 overall process that is used 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 is to 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 show 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 13-32 gives an overview of all of the sections that are 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](#)
 - [Block sizes](#)
 - [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 13-32 Sample VEHGRXCL: Agenda

Figure 13-33 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 13-33 Sample VEHGRXCL: Overview

Figure 13-34 is an example throughput, which is expressed in MBps.

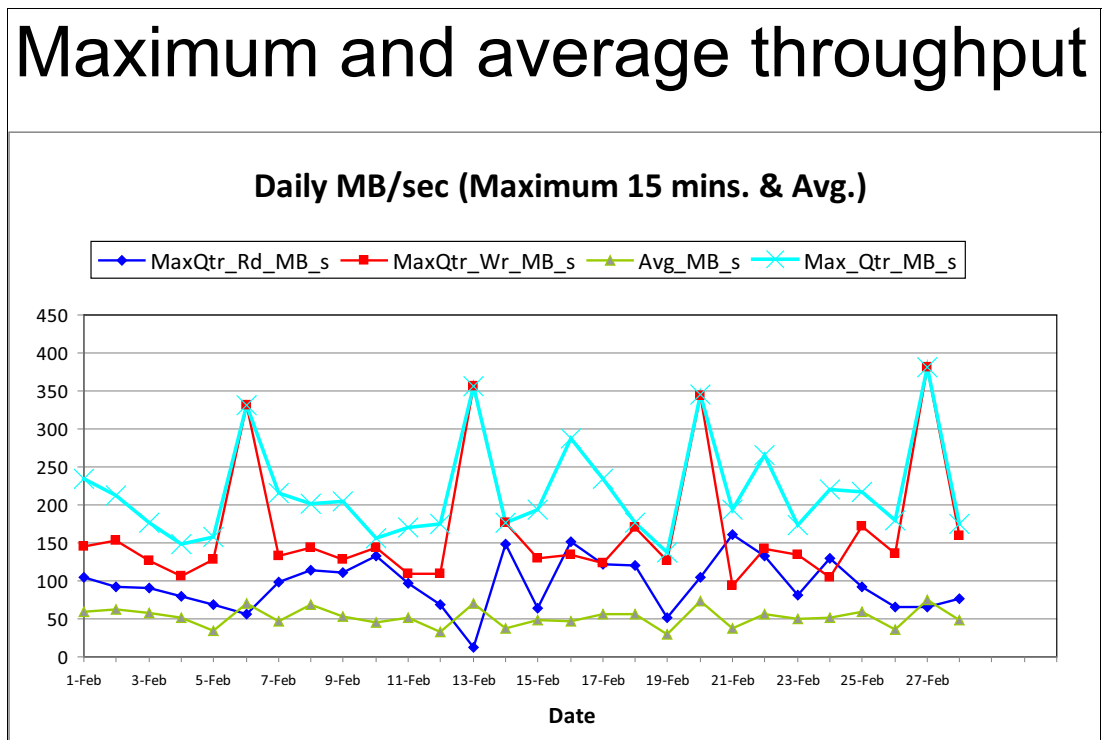


Figure 13-34 Sample VEHGRXCL: Maximum and average throughput

Figure 13-35 is an example of physical mounts.

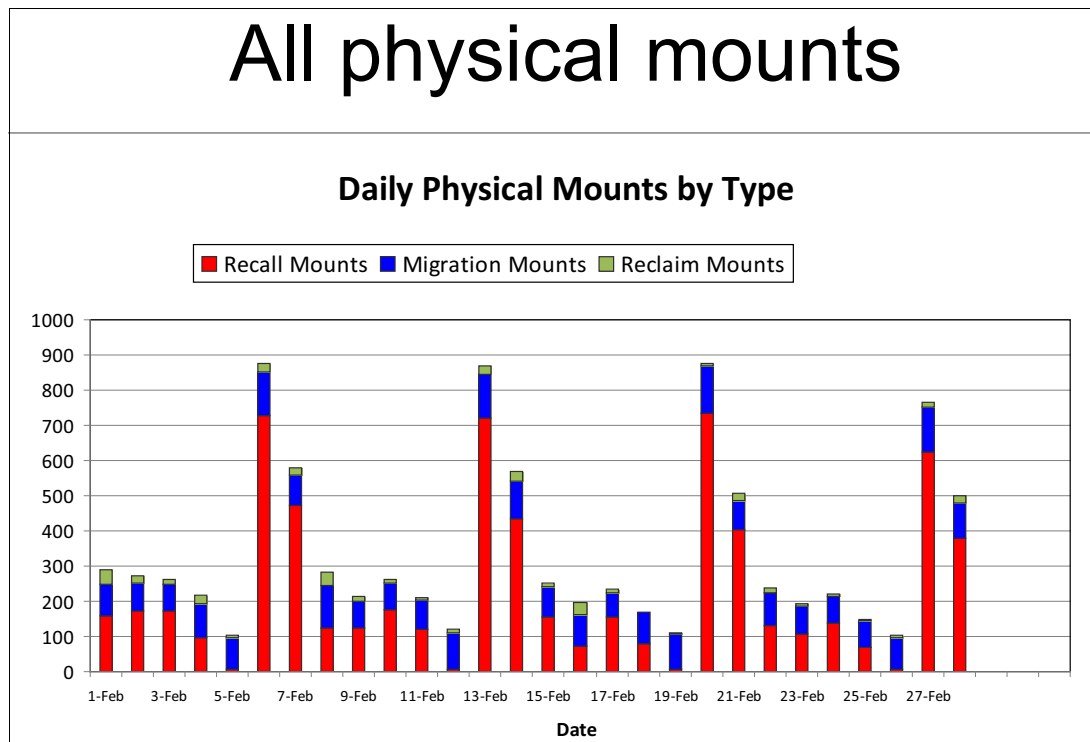


Figure 13-35 Sample VEHGRXCL: All physical mounts

13.5.8 VEHAUDIT tool overview

VEHAUDIT is also based on the BVIR information and a part of the tape tool suite. It includes several independent reports, which are described next.

Copy Distribution report (CPYDST)

This report shows how long copies taken from one cluster to the other clusters took in the grid. Family relationships can also be reflected. This report can be used to identify the RPO, and how long it takes to create copies.

For more information, see [this document](#).

Detail Report (DTLRPT)

The DTLRPT provides an audit function and information about which logical volumes include a deviation from the requested copy consistency policy on the specified cluster. As opposed to the BVIRAUD function, it does *not* compare several clusters to determine whether 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.

For more information, see [this document](#).

The report contains missing copies and 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 expiration date)
- ▶ Category
- ▶ Size
- ▶ Data set name
- ▶ Physical volume placement (only TS7700T)
- ▶ 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 can eliminate the scratches. The list that is produced in this report can then be used as an input for copy refresh processing.

13.5.9 The Other IBM Tape Tools

Several tape tools that can be used to help you better understand your tape processing regarding TS7700 operation and migration are described in this section.

IOSTATS

IOSTATS tool can be downloaded from [IBM Tape Tools website](#).

You can use the 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 specific period before the hardware is installed. 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.

TAPEWISE

As with IOSTATS, the TAPEWISE tool is available from the [IBM Tape Tools website](#). Based on input parameters, TAPEWISE can generate the following reports:

- ▶ Tape activity analysis, including reads and Disp=mod analysis
- ▶ Mounts and MBps 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 website](#). MOUNTMON runs as a started task or batch job and monitors all tape activity on the system. The program must be an 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

Descriptions of the other tools can be found in Table 13-5 on page 705 and the [overview.pdf](#) in the [IBM Tape Tools website](#).

13.6 Host Console Request Commands for monitoring

This section describes some of the host console request commands that are useful for monitoring. For more information about the host console request commands, see [TS7700 Library Request Command](#).

13.6.1 LI REQ,distlib,CACHE2

R5.0 deprecates the CACHE request and introduces the CACHE2 request. The **LI REQ** command with the subcommand **CACHE2** shows you the cache use of the cluster that is defined in the command. Example 13-21 shows an example output.

Note: When CACHE instead of CACHE2 is still issued to the 8.50.x.x cluster, the following text is shown at the end of the output to encourage using CACHE2:

“CACHE REQUEST DEPRECATED; SUPERSEDED BY CACHE2.”

Example 13-21 Example output of LI REQ,distlib,CACHE2 command

```
TAPE VOLUME CACHE STATE V5 .0
TOTAL SPACE INSTALLED/ENABLED: 157TB/ 140TB
TOTAL ADJUSTED CACHE SPACE USED: 335GB
CACHE ENCRYPTION STATUS: ENABLED-INTERNAL
OVERCOMMITTED CACHE PARTITIONS: NONE
CACHE RESIDENT ONLY PARTITION
PRIVATE CACHE SPACE USED: 329GB
SCRATCH CACHE SPACE USED: 0.0GB
CP ALLOC USED PIN PKP PRM COPY CPYT
0 133TB 329GB 0.0GB 299GB 29.0GB 0.0GB 0
```

FLASH COPY INFORMATION		
INDEX	ENABLED	SIZE
1	NO	0.0GB
2	NO	0.0GB
3	NO	0.0GB
4	NO	0.0GB
5	NO	0.0GB
6	NO	0.0GB
7	NO	0.0GB
8	NO	0.0GB

13.6.2 LI REQ,complib,STATUS,GRID

The **LI REQ,complib,STATUS,GRID** command examines the current state of the Grid connections and operation. Example 13-22 shows an output example.

Example 13-22 Example output of LI REQ,complib,STATUS,GRID command

```

GRID STATUS V3 .0
COMPOSITE LIBRARY VIEW
      IMMED-DEFERRED  OWNERSHIP-T/O  RECONCILE HCOPY
LIBRARY  STATE  NUM  MB  MODE  NUM  NUM  ENB
Lipizzan  ON    0   0   -    0    0   Y
Arabian   ON    0   0   -    0    0   Y
  Oak     ON    0   0   -    0    0   Y
Palomino  ON    0   0   -    0    0   Y
-----
COMPOSITE LIBRARY VIEW
      SYNC-DEFERRED
LIBRARY  NUM  MB
Lipizzan  0   0
Arabian   0   0
  Oak     0   0
Palomino  0   0
-----
DISTRIBUTED LIBRARY VIEW
      RUN-COPY-QUEUE  DEF-COPY-QUEUE  LSTATE PT  FAM
LIBRARY  STATE  NUM  MB  NUM  MB
Lipizzan  ON    0   0   0    0   A  N  -
Arabian   ON    0   0   0    0   A  Y  1
  Oak     ON    0   0   0    0   A  N  1
Palomino  ON    0   0   0    0   A  Y  -
-----
      ACTIVE-COPIES  CODE-LEVELS
LIBRARY  RUN  DEF
Lipizzan  0   0   8.50.0.134
Arabian   0   0   8.50.0.134
  Oak     0   0   8.50.0.134
Palomino  0   0   8.50.0.134

```

13.6.3 LI REQ,distlib,STATUS,GRIDLINK

The **LI REQ,distlib,STATUS,GRIDLINK** command shows the grid link performance that is available with the Host Console Request function. Example 13-23 shows an example output. The configurations that are related to grid link are also shown in the output.

Example 13-23 Example of LI REQ,distlib,STATUS,GRIDLINK output

```
GRIDLINK STATUS V2 .2
CAPTURE TIMESTAMP: 2019-12-02 10:35:20
LINK VIEW
LINK NUM      CFG  NEG   READ  WRITE  TOTAL  ERR  LINK STATE
MB/S         MB/S  MB/S          01234567
0   1000 1000   0.0   0.0   0.0   0   -   A
1   1000 1000   0.0   0.0   0.0   0   -   A
2     0    0    0.0   0.0   0.0   0   -
3     0    0    0.0   0.0   0.0   0   -
-----
LINK PATH LATENCY VIEW
LIBRARY          LINK 0          LINK 1          LINK 2          LINK 3
LATENCY IN MSEC
cluster0          0              0              0              0
cluster5          0              0              0              0
-----
CLUSTER VIEW
DATA PACKETS SENT:          38516
DATA PACKETS RETRANSMITTED: 0
PERCENT RETRANSMITTED:     0.0000
-----
LOCAL LINK IP ADDRESS
LINK 0 IP ADDR:          9.11.219.204
LINK 1 IP ADDR:          9.11.219.205
LINK 2 IP ADDR:
LINK 3 IP ADDR:
-----
SECURE DATA TRANSFER
LIBRARY    PROTOCOL  CIPHER KEY
cluster0   TLS1.2    DISABLED
cluster5   TLS1.2    AES-256
```

13.6.4 LI REQ,distlib,STATUS,GRLNKACT

The **LI REQ,distlib,STATUS,GRLNKACT** command shows the point-in-time Grid link activity to all of the available clusters in the Grid configuration and cloud if the cluster is attached to cloud (even if the cluster is a stand-alone configuration, the request is accepted and provides the Grid link activity to cloud). Figure 13-36 on page 733 shows an output example.


```

GRLNRACT STATUS V1 .0
CLUSTER INDEX: 0 LINK COUNT: 4 Time: Wed Nov 11 13:40:38 CUT 2020
GRID LINK ESTABLISHED SOCKET CONNECTIONS-----
LN INTF IP          CO C1 C2 C3 C4 C5 C6 C7 MQ RFA CLD
L0 en0 10.30.1.20    0 132 0 0 0 0 0 0 0 132 0 0
L1 en8 10.31.1.20    0 0 0 0 0 0 0 0 0 0 0 0
L2 en1 10.32.1.20    0 0 0 0 0 0 0 0 0 0 0 0
L3 en9 10.33.1.20    0 0 0 0 0 0 0 0 0 0 0 0
LN INTF IP          GGM OBJ
L0 en0 10.30.1.20    0 0
L1 en8 10.31.1.20    0 0
L2 en1 10.32.1.20    0 0
L3 en9 10.33.1.20    0 0
NET ACTIVITY -----TCP REC/SND ADAPTER BUFFER ACTIVITY BYTES-----
LN TxMBs RxMBs MQ_REC MQ_SND GFA_REC GFA_SND CLD_REC CLD_SND
L0 0 0 0 0 0 0 0 0
L1 0 0 0 0 0 0 0 0
L2 0 0 0 0 0 0 0 0
L3 0 0 0 0 0 0 0 0
TOT 1 1 - - - - -
LN GGM_REC GGM_SND OBJ_REC OBJ_SND
L0 0 0 0 0
L1 0 0 0 0
L2 0 0 0 0
L3 0 0 0 0
GRID LINK THROUGHPUT ESTIMATES-MB/s-----
DIR C1 C2 C3 C4 C5 C6 C7 GRD_TOT CLD GGM OBJ TOT
Tx 0 0 0 0 0 0 0 0 0 0 0 0
Rx 0 0 0 0 0 0 0 0 0 0 0 0
GRID CLOUD TIER EXPORT AND IMPORT ACTIVITY-----
ACTIVE EXPORT VOLUME COUNT: 0
ACTIVE IMPORT VOLUME COUNT: 0

```

Figure 13-36 EXAMPLE of LI REQ,distlib,STATUS,GRLNRACT output

13.6.5 LI REQ,distlib,PDRIVE

The **LI REQ,distlib,PDRIVE** command gives you a snapshot about the physical tape drive environment. It shows the installed drive, model, and status. In addition, you can see what action is performed, for which pool the drive is working, and the mounted physical volume. If the status is not idle, the logical volume that is in use is also provided. This command is supported only with the TS7700T.

Example 13-24 shows a sample output of this command.

Example 13-24 Example output of LI REQ,distlib,PDRIVE 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

13.6.6 LI REQ,{complib | distlib},DIAGDATA

The **LI REQ,complib,DIAGDATA** command was introduced in R3.3 PGA2. This data is used by the Grid Resiliency function, which was supported since R4.1.2. For more information about Grid Resiliency function, see 2.4.36, “Grid resiliency functions” on page 101.

This command is supported by a grid configuration only and not with a stand-alone cluster.

This command helps you monitor the health of clusters in the grid communication, which provides diagnostic information for handshake between clusters in the grid.

LI REQ,complib,DIAGDATA,SHOW,{AVG|MAX|MIN|TMO|ERR}

The **LI REQ,complib,DIAGDATA,SHOW,{AVG|MAX|MIN|TMO|ERR}** command is used to display the average, maximum, or minimum elapsed time, or the timeout or error count for handshaking of all clusters in the grid.

Example 13-25 shows a sample output of this command.

Example 13-25 Example of LI REQ,complib,DIAGDATA,SHOW

```

DIAGTIME STATUS V2 .0
COMPOSITE LIBRARY VIEW
CURRENT TIME (UTC):                2017-10-26 08:13:45
#SCRATCH MOUNTS:                    80 #PRIVATE MOUNTS:                758
AVG DATA (UNIT IS SEC):
D0020 (CL0)   CLO  CL1  CL2  CL3  CL4  CL5  CL6  CL7
  SCR-AVG     NA   0    0   NC   NC   NC   NC   NC
  PRI-AVG     NA   0    0   NC   NC   NC   NC   NC
  TOK-AVG     NA   1    1   NC   NC   NC   NC   NC
DIAGNOSTIC TIME LAST RESET (UTC): 2017-10-26 07:46:30
D0021 (CL1)   CLO  CL1  CL2  CL3  CL4  CL5  CL6  CL7
  SCR-AVG     1   NA   1   NC   NC   NC   NC   NC
  PRI-AVG     1   NA  84   NC   NC   NC   NC   NC
  TOK-AVG     1   NA   1   NC   NC   NC   NC   NC
DIAGNOSTIC TIME LAST RESET (UTC): 2017-10-26 07:46:37
D0022 (CL2)   CLO  CL1  CL2  CL3  CL4  CL5  CL6  CL7
  SCR-AVG     0   0   NA   NC   NC   NC   NC   NC
  PRI-AVG     0   0   NA   NC   NC   NC   NC   NC
  TOK-AVG     1   1   NA   NC   NC   NC   NC   NC
DIAGNOSTIC TIME LAST RESET (UTC): 2017-10-26 07:46:42

```

LI REQ,distlib,DIAGDATA,SHOW

The **LI REQ,distlib,DIAGDATA,SHOW** command shows detailed information about the specific cluster.

LI REQ,distlib,DIAGDATA,RESET

The **LI REQ,distlib,DIAGDATA,RESET** command resets statistics and starts new interval. Run this command periodically to correct interval statistics data.

13.7 IBM z/OS commands for monitoring

The previously introduced methods and options for monitoring the TS7700, which are described from 13.3, “Web-based Monitoring method” on page 671 to 13.6, “Host Console Request Commands for monitoring” on page 730. Other subsystem monitoring options are described in this section.

13.7.1 DISPLAY SMS command

Several **DISPLAY SMS** commands are available 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 13-26, separated by a dashed line.

Example 13-26 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 that are defined.

Category count scratch transition ENABLED.

CBRUXCUA processing ENABLED.

CBRUXEJC processing ENABLED.

CBRUXENT processing OPERATOR-DISABLED.

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		
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		
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 COMPACTON: 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 12, “IBM z/OS host console operations” on page 625 and *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

13.7.2 LIBRARY command

The **LIBRARY** command can be used to check for missing virtual drives or for the status of the grid links. Example 13-27 shows the output of the **LIBRARY DD** command that you can use to verify whether all virtual drives are available.

Example 13-27 Sample response for LI DD,libname command

```

LI DD,GRIDLIB3
CBR1220I Tape drive status: 152
DRIVE DEVICE LIBRARY ON OFFREASON LM ICL ICL MOUNT
NUM TYPE NAME LI OP PT CU AV CATEGRY LOAD VOLUME
4600 3490 GRIDLIB3 Y N N N N A NONE N
4601 3490 GRIDLIB3 Y N N N N A NONE N
4602 3490 GRIDLIB3 Y N N N N A NONE N
4603 3490 GRIDLIB3 Y N N N N A NONE N
4604 3490 GRIDLIB3 N N Y N N - --N/A-- -

```

For more information about the **LIBRARY** command, see Chapter 12, “IBM z/OS host console operations” on page 625, and *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC23-6867.

13.7.3 DEVSERV command

The **DEVSERV QLIB,CATS** command can be used to confirm the current category code settings by the z/OS DEVSUPxx parameter. Example 13-28 shows the output of this command.

Example 13-28 Sample response for DEVSERV QLIB,CATS command

```

DS QL,CATS
IEE459I 19.15.32 DEVSERV QLIB 000
0221 0222 0223 0224 0225 0226 0227 0228 0009 000A 022E 022F

```

For more information about the **LIBRARY** command, see Chapter 12, “IBM z/OS host console operations” on page 625 and *z/OS MVS System Commands*, SA38-0666.

13.8 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.

13.8.1 Alerting of specific events

The TS7700 offers various threshold and timeout alerts. Most of them can be set by using the host console request command `LI REQ,distlib,SETTING,ALERT` and `LI REQ,distlib,SETTING2,ALERT`. This section includes a brief introduction.

For more information, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide](#).

Most threshold alerts allow two thresholds per setting. The first setting issues the “alarm” message that the threshold is now crossed. A second, lower limit setting informs you when the condition is resolved, and that the amount of data fell beyond the threshold.

Before you adjust the values of the thresholds alerts, evaluate the 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

The following alerts send messages to the host system log if the pending inbound copy backlog indicates that the amount of uncopied data exceeded the low or the critical warning limit specified (in GB):

- ▶ Inbound copy backlog (PCPYLOW and PCPYCRIT)
- ▶ Uncopied data in cache (COPYLOW and COPYHIGH)

The following alerts send messages to the host system log if the age of pending inbound copy backlog exceeded the low or the critical warning limit specified (in hours):

- ▶ All copy backlog age (CAGALOW and CAGAHIGH)
- ▶ Non time-delayed copy backlog age (CAGLOW and CAGHIGH)

The following alerts send messages to the host system log if the total amount of data on the cache exceeded the low or the critical warning limit specified (in GB):

- ▶ CPO data in cache (RESLOW and RESDHIGH)
- ▶ CPx data in cache (RSDTLOW and RSDSTDHIGH)
- ▶ Object data in cache (RSDOLOW and RSDOHIGH)

The following alerts send messages to the host system log if number of available physical tape drives or scratch tape cartridges became under the low or the critical warning limit specified (in number):

- ▶ Physical drive availability (PDRVLOW and PDRVCRIT)
- ▶ Number of physical scratch tapes (PSCRLOW and PSCRCRIT)
- ▶ Sunset physical drive availability (PDRVSLOW and PDRVSCRT)

The following settings control whether messages are sent or when messages are sent when free space in the disk cache is running out:

- ▶ Automatic Removal message (REMOVMSG)
- ▶ Limited Cache threshold (LMTDTHR)

All parameters can be set independently, but the values include some dependencies that must be acknowledged.

The default value for each parameter is 0, which indicates that no warning limit is set and messages are not generated.

Alerts for inbound copy backlog (PCPYLOW and PCPYCRIT)

These values specify the threshold in GB of volumes that are 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 by using the following commands:

```
LI REQ,distributed library,SETTING,ALERT,PCPYLOW,value
LI REQ,distributed library,SETTING,ALERT,PCPYCRIT,value
```

Alerts for uncopied data in cache (COPYLOW and COPYHIGH)

These values specify the threshold in GB of volumes that are waiting to be copied to another cluster in the TS7700 grid configuration. The COPYLOW value defines the first level of warning. The COPYHIGH 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 by using the following commands:

```
LI REQ,distributed library,SETTING,ALERT,COPYLOW,value
LI REQ,distributed library,SETTING,ALERT,COPYHIGH,value
```

Alerts for all copy backlog age (CAGALOW and CAGAHIGH)

These values specify the threshold (in hours) that the longest inbound copy backlog age on all copy jobs. The CAGALOW value defines the first level of warning. The CAGAHIGH represents a critical state of inbound copy backlog age. These alerts are applicable for all TS7700s, and can be set for each distributed library independently.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING2,ALERT,CAGALOW,value
LI REQ,distributed library,SETTING2,ALERT,CAGAHIGH,value
```

Alerts for non time-delayed copy backlog age (CAGLOW and CAGHIGH)

These values specify the threshold (in hours) that the longest inbound copy backlog age on non-time delayed copy jobs. The CAGLOW value defines the first level of warning. The CAGHIGH represents a critical state of inbound copy backlog age. These alerts are applicable for all TS7700s, and can be set for each distributed library independently.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING2,ALERT,CAGLOW,value
LI REQ,distributed library,SETTING2,ALERT,CAGHIGH,value
```

Alerts for CP0 data in cache (RESLOW and RESHIGH)

These values specify the amount of data in the cache. The RESLOW value defines the first level of warning. The RESHIGH represents a critical state of data in the cache.

The following values are measured:

- ▶ TS7700D: All data in cache
- ▶ TS7700T and TS7700C: Resident data in CP0

Change the value by using 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 a TS7700D and TS7700T/C CP0, if RESLOW/RESHIGH was exceeded, check the amount of Pinned data and the amount of data that is subject to auto removal. You might run out of cache if the data is filled with Pinned data and the data is not expired by the host soon.

Alerts for CPx data in cache (RSDTLOW and RSDTHIGH)

These alerts are similar to other alerts, but apply to the TS7700T tape partitions and TS7700C cloud partitions. The RSDTLOW value defines the first level of warning. The RSDTHIGH represents a critical state of data in a cache of non-premigrated data from all tape partitions.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING,ALERT,RSDTLOW,value
LI REQ,distributed library,SETTING,ALERT,RSDTHIGH,value
```

Alerts for object data in cache (RSDOLOW and RSDOHIGH)

These alerts are similar to other alerts, but apply to the object cache partitions for TS7700 Advanced object store. The RSDOLOW value defines the first level of warning. The RSDOHIGH represents a critical state of data in object partitions. They were introduced in R5.0 under the keyword SETTING2 and moved to the new keyword OBJSET1 in R5.2.2.

Change the value by using the following commands:

```
LI REQ,distributed library,OBJSET1,ALERT,RSDOLOW,value
LI REQ,distributed library,OBJSET1,ALERT,RSDOHIGH,value
```

Alerts for physical drive availability (PDRVLOW and PDRVCRT)

These values specify the number of available backend drives in a TS7700T. The PDRVLOW value defines the first level of warning. The PDRVCRT represents a critical state. They are applicable for TS7700T only and can be set for each distributed library independently.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING,ALERT,PDRVLOW,value
LI REQ,distributed library,SETTING,ALERT,PDRVCRT,value
```


Alerts for number of physical scratch tapes (PSCRLOW and PSCRCRIT)

These values specify the number of available physical scratch cartridges in a TS7700T. The PSCRLOW value defines the first level of warning. The PSCRCRIT represents a critical state.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING,ALERT,PSCRLOW,value  
LI REQ,distributed library,SETTING,ALERT,PSCRCRIT,value
```

Alerts for sunset physical drive availability (PDRVSLOW and PDRVSCRT)

These values specify the number of available sunset drives in a TS7700T. The PDRVSLOW value defines the first level of warning. The PDRVSCRIT represents a critical state. They are applicable for TS7700T only and can be set for each distributed library independently.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING2,ALERT,PDRVLOW,value  
LI REQ,distributed library,SETTING2,ALERT,PDRVCRIT,value
```

Setting for Automatic Removal message (REMOVMSG)

ENABLE or DISABLE can be specified. When ENABLE is specified, a message is sent to the host system log and a MI event is created when automatic removal of logical volumes starts or stops. When DISABLE is specified, no host message is sent and no MI event is created. This value is applicable for TS7700D and CP0 of TS7700T/C and can be set for each distributed library independently.

Change the setting by using the following commands:

```
LI REQ,distributed library,SETTING,ALERT,REMOVMSG,[ENABLE|DISABLE]
```

Setting for Limited Cache threshold (LMTDTHR)

This value specifies the amount of free space in the disk cache when TS7700 enters the limited cache warning state (in GB). It is applicable for TS7700D and CP0 of TS7700T/C and can be set for each distributed library independently.

Change the value by using the following commands:

```
LI REQ,distributed library,SETTING2,ALERT,LMTDTHR,value
```

13.8.2 Handling Replication Exceptions

In specific conditions, the selected replication mode cannot be run, which is true for Synchronous mode copies and Immediate (RUN) copies.

To ensure that the production is not affected, 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 changes to a “SYNC-Deferred” state, or if the job fails.

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 the 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 that is greater than standard deferred copies, but lower than non-deferred immediate copies.

A volume might enter the immediate-deferred state for many reasons. 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 entered the immediate-deferred state.

The reasons why a volume moves to the immediate-deferred state are included 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 (see Example 13-29).

Example 13-29 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
```

The following new failure content is introduced into the CCW(RUN) ERA35 sense data:

- ▶ Byte 14 FSM Error. If set to 0x1C (Immediate Copy Failure), the extra 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: A total of 40 minutes have elapsed and one or more copies have timed out.
 - 0x04: Reverts to immediate-deferred because of health/state of RUN target clusters.
 - 0x05: The 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 extra 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 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 message can be reported in the system log. Although that message allows a message automation on the host to detect possible problems, the thousands of messages that are generated might overflow the host log. This issue is especially true in maintenance situations. You can use the Host Console Request (HCR) command to handle all OS000I triggered by “IMMED-Deferred”:

```
LI REQ,distributed library,SETTING,COPY,IMMSNS,[ALL|NONE|UNEXP]
```

The following options are available:

- ▶ 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 option is the default.
- ▶ UNEXP: Only unexpected failures are presented to the host. All messages during a maintenance situation or during “Gridlink Disable” are not 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 a write failure occurs, 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, Cluster 0 and Cluster 1 must be available for write operations. If either cluster becomes unavailable, write operations fail. With the option enabled, if Cluster 0 or Cluster 1 becomes unavailable, write operations continue. The second “S” copy becomes a synchronous-deferred copy.

In this example, if the host is attached to Cluster 2 only and the option is enabled, the write operations continue, even if 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 [IBM Virtualization Engine TS7700 Series Best Practices - Synchronous Mode Copy](#).

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 affect your disaster recovery capability, the production jobs might not be affected at all. Degradation of composite library is reported but the production jobs might not be affected also in the case where the cluster is not available because of maintenance purposes.

Therefore, you can use a Host Console Request Command is provided 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.

The warning message uses the following format:

```
CBR3750I, G0030, Library xx Primary, Alternative, Primary2, and Alternative2 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, which is described in 13.6.3, “LI REQ,distlib,STATUS,GRIDLINK” on page 732, and on the TS7700 MI.

The grid link degraded threshold also includes the following 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 must be changed. If you receive intermittent warnings, allow the SSR to change the threshold and iteration to the suggested values from IBM Support.

For example, suppose that clusters in a two-cluster grid are 3218 km (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 25% and the iterations are set to 12 and 15.

A Host Console Request Command is provided so that you can define whether these conditions report the composite library as grid link degraded or not:

```
LI REQ,composite library,SETTING,ALERT,LINKDEG, [ENABLE|DISABLE]
```




Performance considerations

This chapter describes the factors that determine and influence the performance of the IBM TS7700. It also describes what actions to take to improve TS7700 performance when necessary. This chapter also covers the possible settings, alerts, and messages that can be considered for exception handling and automation.

Note: For more information, see the following IBM white papers that are related to TS7700 performance:

- ▶ [IBM TS7700 R5.4 Performance](#)
- ▶ [IBM TS7700 R5.3 Performance](#)
- ▶ [IBM TS7700 R5.1 Performance](#)
- ▶ [IBM TS7700 R5.0.1 Performance](#)
- ▶ [IBM TS7700 R4 \(TS7760\) Performance](#)
- ▶ [IBM Virtualization Engine TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#)
- ▶ [IBM TS7700 - R4.1.2 Performance](#) (with the software-based compression and the 16 Gb FICON)

This chapter includes the following topics:

- ▶ 14.1, “Overview” on page 748
- ▶ 14.2, “TS7700 performance history” on page 750
- ▶ 14.3, “Basic performance overview” on page 752
- ▶ 14.4, “TS7700 throughput: Host I/O increments” on page 762
- ▶ 14.5, “Considerations for Virtual Device Allocation” on page 764
- ▶ 14.6, “Cache throughput and cache bandwidth” on page 765
- ▶ 14.7, “Grid link and replication performance” on page 769
- ▶ 14.8, “Considerations for the backend TS7700T” on page 776
- ▶ 14.9, “Cloud Tiering” on page 780
- ▶ 14.10, “TS7700 Advanced Object Store for DS8000” on page 781

14.1 Overview

R5.4 introduced support for 65000 MiB logical volume size.

R5.3 PGA1 increased the maximum number of 3.84 TB SAS SSDs drawers from four to ten which is a full base frame (one controller drawer and nine expansion drawers).

R5.3 increased the maximum number of 3.84 TB SAS SSDs drawers from two to four (one controller drawer and three expansion drawers), which is an enhancement to solid-state drive (SSD)-based TVC configuration introduced in R5.2.1.

R5.2.2 introduced TS7700 Advanced Object Store for DS8000 that is an enhancement to the previous DS8000 Object Store that was introduced in R5.0. TS7700 Advanced Object Store supports grid-wide replication of DS8000 objects that are stored in TS7700 TVC.

R5.2.1 PGA1 introduced the support of IBM TS1160 tape drives that are installed on TS4500 and TS3500 tape libraries as back-end tape drives for TS7700.

R5.2.1 introduced solid-state drive (SSD)-based TVC by using 3.84 TB SAS SSDs with 60 TB usable capacity per drawer, up to two drawers (one controller drawer and one expansion drawer).

R5.1 introduced DS8000 Object Store enhancement with encryption/compression, and cloud support enhancement such as grid awareness, multiple cloud pools, volume version retention, cloud export and recovery, and others.

R5.0 introduced the TS7770 that is built on the new Power9 platform with new cache disks. The TS7770 models (TS7770D and TS7770T and TS7770C) replace all previous models.

R4.2 introduced the TS7700C, which supports the ability to store virtual tape volumes in a Cloud Object Storage.

R4.1.2 introduced the 16 Gbps Fibre Channel connection (FICON) adapters, the 8 TB hard disk drives (HDD) in the Tape Volume Cache that replaced the 4 TB HDD, and software-based compression.

Figure 14-1 on page 749 shows the activities between the following components:

- ▶ TS7700
- ▶ IBM Z host regarding read/write I/O through the FICON channels
- ▶ Attached tape library regarding migration/recall
- ▶ Connected Cloud Object Storage regarding migration/recall through the grid network
- ▶ Grid configuration regarding replication copy through grid network
- ▶ DS8000 regarding object offload through the grid network

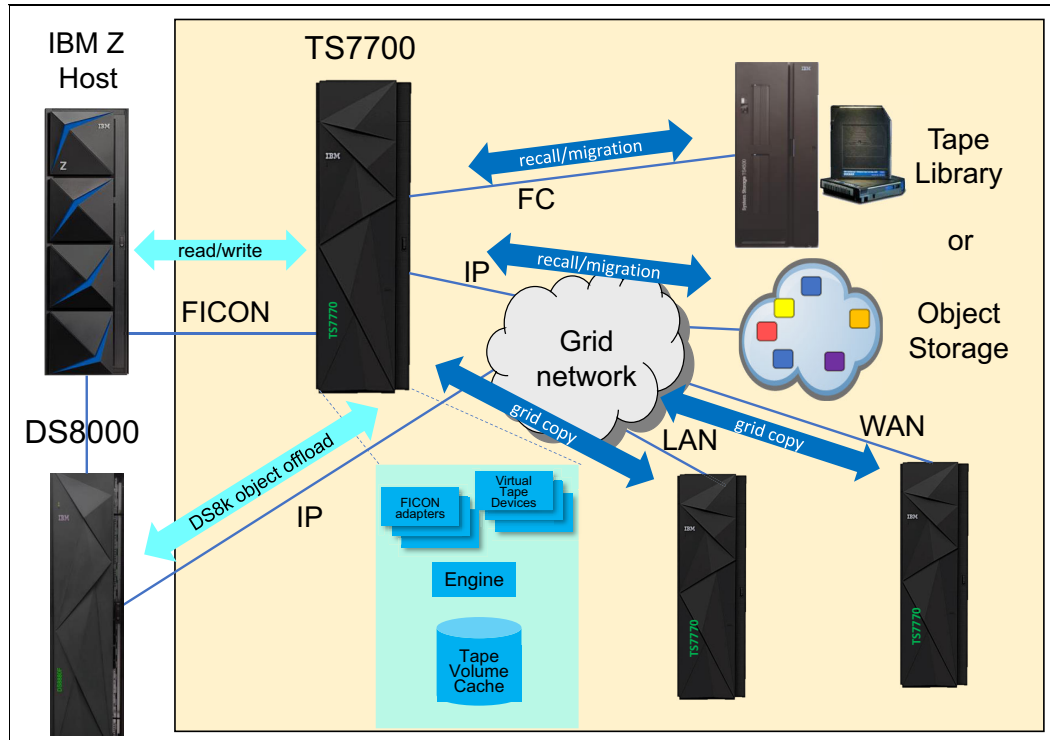


Figure 14-1 Activities between the TS7700 and connected components that are related to performance

Figure 14-1 also shows the related components in the TS7700.

This chapter includes the following newest overall performance information for the TS7700:

- ▶ For understanding the overall performance characteristics for the TS7700, see 14.2, “TS7700 performance history” on page 750
- ▶ For understanding the aspects that influence the performance, see 14.3, “Basic performance overview” on page 752
- ▶ For information about Host I/O throughput handled in the TS7700, see 14.4, “TS7700 throughput: Host I/O increments” on page 762
- ▶ For Virtual Device Allocation performance considerations, see 14.5, “Considerations for Virtual Device Allocation” on page 764
- ▶ For understanding the tape volume cache (TVC) performance with TS7700D, TS7700T, and TS7700C, see 14.6, “Cache throughput and cache bandwidth” on page 765
- ▶ For the grid replication copy and the remote mount through the grid network, see 14.7, “Grid link and replication performance” on page 769
- ▶ For understanding the specialty for the TS7700T regarding the migration effects on the attached physical tape library, see 14.8, “Considerations for the backend TS7700T” on page 776
- ▶ For understanding the specialty for the TS7700C regarding the migration effects on the Cloud Object Storage, see 14.9, “Cloud Tiering” on page 780
- ▶ For understanding the DS8000 object offload activity effects on the grid network, see 14.10, “TS7700 Advanced Object Store for DS8000” on page 781

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.

Chapter 13, “Monitoring” on page 667 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.

14.2 TS7700 performance history

This section describes the performance history of previous TS7700 models, such as TS7740, TS7720, TS7760, and some of the 3494 virtual tape server models. It does *not* include information about TS7770.

The TS7700 can provide significant benefits to a tape processing environment. In general, performance depends on several factors, such as total system configuration, Tape Volume Cache (TVC) capacity, the number of physical tape drives that is 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 in your environment from the figures that are presented in this chapter. 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 14-2 on page 751 shows the evolution in the write performance with the TS7700 family. For more information, see [IBM TS7700 R4 \(TS7760\) Performance](#).

The charts that are shown in Figure 14-2 on page 751 and Figure 14-3 on page 752 are for illustrative purposes only. Always use the most recently published performance white papers, which are available at [IBM Support](#).

Figure 14-2 on page 751 shows the evolution of performance in the TS7700 IBM family that is 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 that use 32 kibibyte (KiB) blocks and queued sequential access method (QSAM) BUFNO = 20.

Before R 3.2, volume size was 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 14-2 on page 751 shows write performance history.




Figure 14-2 VTS and TS7700 maximum host write throughput

Figure 14-3 on page 752 shows the read hit performance numbers.

The numbers that are shown in Figure 14-3 on page 752 were obtained with 128 concurrent jobs in all runs, each job uses 32 KiB blocks, and QSAM BUFNO = 20. Before R 3.2, the volume size was 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 the following important characteristics:

- ▶ With the selection of IBM Db2 as the central repository in the TS7700, 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.

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.




Figure 14-3 VTS and TS7700 maximum host read hit throughput

14.3 Basic performance overview

The performance of TS7700 includes several characteristics and is influenced by many aspects. The following section gives a brief overview of those TS7700 performance aspects, and describes some of the dependencies.

For more information about TS7700 performance, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

The following major aspects 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
- ▶ Terminology of throughput

- ▶ Throttling in the TS7700
- ▶ Compression methods
- ▶ Internal SSD/HDDs for TS7700 engine
- ▶ SSD/HDDs for TVC

14.3.1 TS7700 components and task distribution

While writing scratch volumes or premigrating and recalling virtual volumes to/from physical stacked volumes, hardware components are shared by tasks that are 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.

An overview of all of the tasks is shown in Figure 14-4. 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 next.

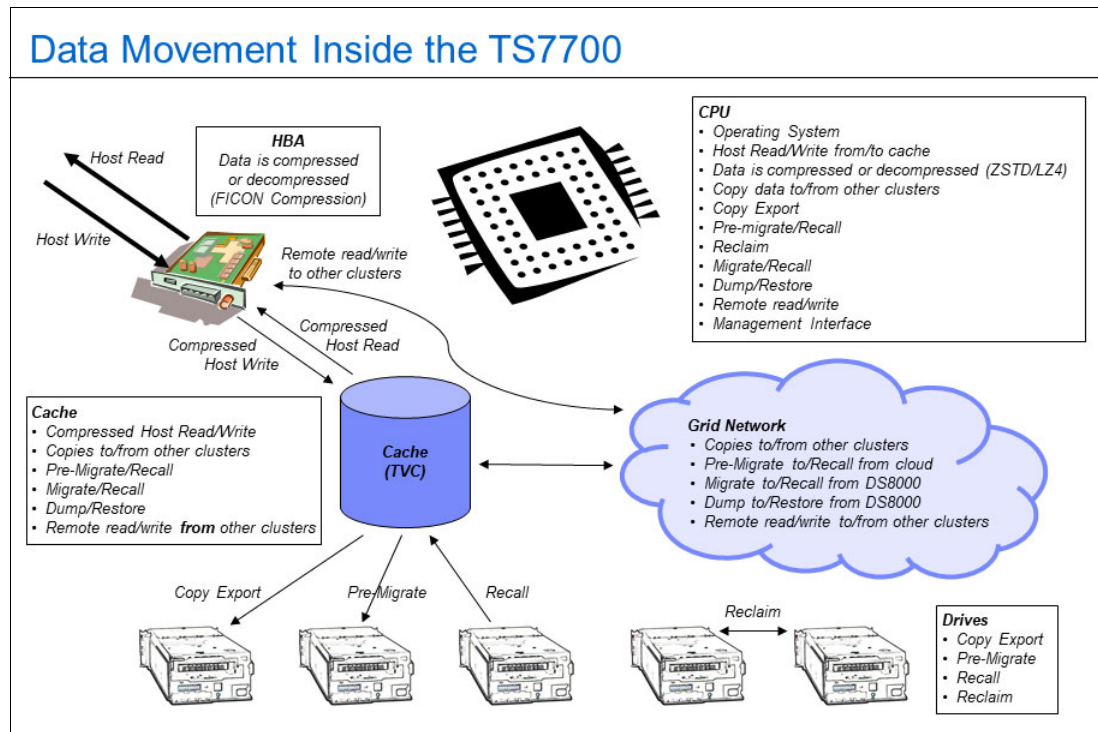


Figure 14-4 Tasks that are performed by the TS7700

In general, the tasks are the same for all models of the TS7700. For the TS7700D, the backend tasks are not applicable.

All of these tasks share resources, especially the TS7700 Server processor, TVC, and physical tape drives that are attached to a TS7700T. 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 effect 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 TS7700T detects that the number of empty physical volumes that dropped below the minimum value, and reclaims that must be done immediately to increase the count.
- ▶ The cache fills with copy data: To protect from uncopied volumes being removed from cache, the TS7700T/TS7700C throttles data that is coming into the cache. For the TS7700T/TS7700C, this type of throttling occurs to only Host I/O that is related to the CPx partitions. Data that is written to CP0 is *not* throttled in this situation.

For more information about task processing, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

Note: The TS7700T/TS7700C is not yet reflected in this version of the white paper.

14.3.2 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. Each of these copy modes features specific characteristics and influences your RPO for specific application workload.

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 nonpeak time (all types of deferred copies), other copy modes are enforced to run before job end (Synchronous, RUN).

These characteristics of the copy modes imply that resources from the TS7700 (cache bandwidth, CPU, and grid link bandwidth) must be available now in Synchronous and RUN copy modes.

The following replication modes are available:

- ▶ *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. Because of this behavior, synchronous mode copy features some performance benefits over RUN copy modes. Typically, when the Rewind Unload (RUN) is sent, the Synchronous copy is completed.
- ▶ With *RUN Copy mode (Immediate Copy mode)*, the second copy is not started until the **Rewind Unload** 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) RUN Copy mode is not supported with 25 GB or 65 GB logical volumes.
- ▶ In *Deferred Copy mode*, data is queued for copying when the **Rewind Unload** command is received, 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 14.7.6, “Tuning possibilities for copies: Deferred copy throttling” on page 773.

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. For more information, see 14.7.4, “Tuning possibilities for copies: COPYCOUNT Control” on page 772.

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. Delaying deferred copies might affect your RPO.

- ▶ The *Time Delayed Replication mode* is useful to produce copies to multiple TS7700 clusters only after a predefined timeline expires. This copy mode might reduce the number of needed copies as no copy is produced at all when the data expires before the predefined timeline expires; however, the copies can be produced in a different timeline (if you specify hours after create/access). You cannot specify a specific start.

In specific situations, requested Synchronous mode copy or RUN copies cannot be processed. In this case (depending on your customization), the jobs do 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 13.8.2, “Handling Replication Exceptions” on page 741.

The deferred copies are processed in the following order:

1. Synchronous-Deferred PG0
2. Synchronous-Deferred PG1
3. Immediate-Deferred PG0
4. Immediate-Deferred PG1
5. Deferred PG0
6. Deferred PG1

In the grid, the following 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 transfer
 - 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

Consider a composite library with two or more clusters at a local site and two or more clusters 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, which optimizes the use of the data link and shortens the copy time. Also, the family members cooperate among themselves, each pulling a copy of a 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 process prevents family copies from stagnating in the copy queue.

For more information about cluster families, see [IBM Virtualization Engine TS7700 Series Best Practices - TS7700 Hybrid Grid Usage](#).

14.3.3 Workload profile from your hosts

In general, the following types of workloads are available:

- ▶ **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:** This workload is the user-driven workload; for example, hierarchical storage management (HSM) or object access method (OAM) processing requests. This workload also consists of the event-driven workload; for example, database log archiving or System Management Facilities (SMF) processing exists.

Unplanned read workload might have peaks that can affect the response times of these actions (read/recall times). However, these actions can also influence the deferred copy times and in a TS7700T, 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.

14.3.4 Lifecycle Management of your data

This specific aspect is important for a TS7700T. Depending on your amount of data (and logical volumes) with a short expiration date, the TS7700T must run more reclamation. This process affects your back-end drives and TS7700T processor cycles.

In a hybrid grid, such data can be placed in the TS7700D, and can be replicated with the Time Delayed copy mode, which can lead to the reduction of the backend activities in the TS7700T. The use of the delay premigration queue on a TS7700T can also reduce the back-end activities for migration and reclaim.

14.3.5 Parameters and customization of the TS7700

The TS7700 offers various tuning possibilities, especially for cache management and replication and backend activities.

TS7700 tuning activities include the following examples:

- ▶ Preference group of the data (data is preferably in cache or not) in TS7700T/TS7700C
- ▶ Number of the tape/cloud cache partitions in TS7700T/TS7700C
- ▶ Use of the premigration delay feature in TS7700T

- ▶ Premigration threshold and control of premigration tasks for TS7700T
- ▶ Deferred Copy Throttling (to prioritize the host workload)
- ▶ Number of concurrent copy tasks
- ▶ Schedule for reclamation in TS7700T
- ▶ Number of physical volume pools in TS7700T

Consider that some of these activities include dependencies.

If you change the preconfigured values, review your adjustment with the performance monitoring tools.

For more information, see [IBM TS7700 Series Best Practices - Understanding, Monitoring, and Tuning the TS7700 Performance](#).

14.3.6 Throughput terminology

A TS7700 disk-cache-only cluster has a fairly consistent workload throughout.

Because the TS7700T/TS7700C contains physical tapes/cloud to which the cache data is periodically written, recalls from tape/cloud to cache occur, and Copy Export and reclaim activities occur for TS7700T.

The TS7700T/TS7700C exhibits the following basic throughput rates:

- ▶ Peak write
- ▶ Sustained write
- ▶ Read-hit
- ▶ Recall

These rates are described next.

Peak and sustained write throughput

For the TS7700T, a measurement is not begun until all previously written data is copied or premigrated from the disk cache to physical tape. Starting with this initial condition, data from the host is first written into the TS7700T disk cache with little (if any) premigration activity occurring. This approach allows for a higher initial data rate, and is known as the *peak* data rate.

After a pre-established threshold of nonpremigrated 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 another threshold of nonpremigrated 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 that is coming in from the host and the amount of data that is copied to physical tape. The resulting data rate for this mode of behavior is called the *sustained* data rate. Theoretically, it can continue 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. The priority and throttling thresholds can be increased by using a host CLI request, which is described in 14.6.2, “Premigration and premigration throttling values” on page 766.

For TS7700C, the premigration *priority* threshold and premigration *throttling* threshold are not subject to customization by LI REQ command; the premigration throttling threshold is fixed at a value of the size of the premigration queue.

Read-hit and recall throughput

Similar to write activity, the following types of TS7700T/TS7700C read performance are available:

- ▶ *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 requested data is no longer in the disk cache and must be first read in from physical tape/cloud.

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, including the following examples:

- ▶ Number of physical tape drives
- ▶ Spread of requested logical volumes over physical volumes
- ▶ Location of the logical volumes on the physical volumes
- ▶ Length of the physical media for TS7700T
- ▶ Network bandwidth to cloud for TS7700C

14.3.7 Throttling in the TS7700

Throttling is the mechanism that was 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 usage, where the components are used almost to their maximum capacity and bottlenecks start to show. The criteria balance the user needs with the vital resources that are needed for the TS7700 to function.

This control is accomplished by delaying the start 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 process alleviates stress on an overloaded component, leaves extra central processor unit (CPU) cycles to another needed function, or waits for a slower operation to finish.

The subsystem includes 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/TS7700C specific throttling behaviors

From a throttling perspective, the TS7700T/TS7700C is different than a TS7700 disk-only model. Resident partition and tape/cloud partitions feature two independent throttling measurements and are treated differently. Reaching a throttling limit on the tape partitions (for example, PMTHLVL) does not affect the workload that is directed to the resident partition and vice versa.

The following rules apply:

- ▶ Throttling that is started by reaching the maximum Host Throughput applies to all partitions (resident and tape/cloud partitions).
- ▶ Throttling that is started by reaching any premigration limit (PMTHLVL or maximum size of the premigration feature queue) affects only tape/cloud partitions, but not the workload that is directed to CP0.
- ▶ Copy Throttling and Deferred copy throttling have a common measurement regardless of whether the workload is created in CP0 or CP1–CP7.

Important: Resident partition (CP0) and Tape Partitions (CP1–CP7) are monitored and handled separately in a TS7700T/TS7700C.

However, even if the PMTHLVL throttling does not apply to the CP0 of a TS7700T/TS7700C, an indirect influence still exists because of the shared cache modules.

Consider the following points when you configure or monitor a TS7700T/TS7700C resource usage:

- ▶ Workloads that create data (host I/O, remote writes, or copy processes from other clusters) in the CP0 use 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 no throttling occurs to the jobs that are creating data in CP0. Therefore, resources for premigration might be limited even after PMTHLVL is reached, especially in small configurations (for example, up to four drawers) where the jobs running to CP0 are exhausting the cache bandwidth resources.
- ▶ If the unpremigrated amount of data still increases, throttling the workload into the CPx also increases. This issue might cause several delays to the jobs that are creating data in CPx.

Note: PMTTLVL setting is applicable to TS7700T only. The equivalent setting for TS7700C is fixed at a value of the size of the premigration queue.

TS7700T tape drives usage considerations

The physical tape drives are managed by the TS7700T internal management software, and cannot be accessed from any other attached host. These drives are used exclusively by the TS7700T for the mounts that are required for copying virtual volumes to stacked volumes, which recall virtual volumes into the cache and reclaims stacked volume space.

The availability of TS7700T physical tape drives for certain functions can significantly affect 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 TS7700T delays these housekeeping tasks to preference host I/O while no thresholds are reached (PMTHLVL).

The 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 free space low: The TS7700T increases the number of drives that are available to the premigration function. It also reduces the number of drives available for recalls.
- ▶ Premigration threshold crossed: The 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 that is 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 used by other processes, it can lead to limiting the number of virtual volumes in the cache to be premigrated. This limited premigration might lead to premigration throttling (host I/O is throttled), and later it can lead to available 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 result.

Recall performance is highly dependent on 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 configuration 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. The use of 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 process can take seconds to minutes because the process of moving a logical volume must be finished, and then the cartridge must 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 physical volumes must 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 TS7700T.

14.3.8 Compression methods

The compression methods affect performance in terms of host I/O throughput. With compression, the amount of data flow to and from TVC and grid links decreases. In general, a higher compression ratio leads to higher host I/O throughput, and at the same time, higher data transfer rate in a grid, in measurement of uncompressed data size.

TS7700 supports different compression methods. ALDC compression in the FICON adapter is an older algorithm, which was supported since its first release and produces lower average compression than others. With R4.1.2, the TS7700 introduced two new software-based compression methods (LZ4 and ZSTD) that achieve higher average compression with cost of TS7700 CPU.

Of these two software-based methods, LZ4 is faster (in terms of compression) and lower CPU usage with improved compression ratio, whereas ZSTD is higher CPU usage, but fast enough with highest compression ratio. To use the newer compression methods, all clusters in a grid must have R4.1.2 or later microcode level.

Because the methods use TS7700 CPU, consider that CPU power in TS7700 is not exhausted, especially with ZSTD compression that requires higher CPU. Clients with older TS7700 models, such as VEB/V07, are recommended to conduct performance testing *before* moving into production usage. VEHSTATS performance reports might also be useful.

Performance improvement varies depending on many aspects, including data patterns that affect compression ratio; performance metrics, such as read/write or peak/sustained; TS7700 models that affect CPU power; configuration, such as stand-alone or grid; job concurrency, and others.

Consider the following points regarding the results of performance benchmarks with specific conditions:

- ▶ For workloads that do not compress, performance between LZ4 and ZSTD is equal. Performance is lower for FICON compression because of data expansion.
- ▶ For workloads that do compress, FICON compression is the slowest in all cases. When cache is not the bottleneck, LZ4 and ZSTD performed fairly equally. When cache is the bottleneck, the method that compresses the most efficiently provides the highest performance. However, with the older TS7700 model where the CPU overhead is a significant factor, the software-based methods (especially ZSTD) might result in poor performance.
- ▶ ZSTD results in lower performance with less job concurrency because of the overhead that the ZSTD algorithm incurs. A single stream LZ4/FICON can run up to two times faster than a single stream ZSTD workload.
- ▶ For batch periods where few active devices run in parallel, the use of LZ4 is ideal if performance is most important. If compression is more important, ZSTD is a better choice. It is also feasible to mix workloads so that some use LZ4 while others use ZSTD.
- ▶ When many workloads are expected to run in parallel, ZSTD is often the best choice for cumulative performance and compression. If some of the workloads require higher performance among others, they can choose LZ4 so that they can run independently faster, but with less compression.

Note: All of these comparisons are relative to the maximum achievable throughput in the configuration. Small drawer configurations can also produce results where lower job counts run faster with ZSTD because the improved compression puts less strain on the limited disk cache.

For more information about results of performance benchmarks by different compression methods, see the following IBM Support web pages:

- ▶ [IBM TS7700 – R5.0.1 Performance](#)
- ▶ [IBM TS7700 R4.1.2 Performance](#)

14.3.9 Internal SSD/HDDs for TS7700 engine

For TS7770 Server model (3957-VED), SSDs now are an option at the initial configuration for the server, and SAS HDDs that were an option for previous models. SSDs lead to improved code load and response times, and Db2 in TS7700 maintenance times.

14.3.10 SSD/HDDs for TVC

For TS7770 Server models (3957-VED) with microcode level of R5.2.1, SAS SSDs now are an option at the initial configuration for TVC use, and SAS HDDs that were an option for previous releases. The maximum number of SSD drawers has been increased from two to four in R5.3, and four to ten in R5.3 PGA1.

SSD-based TVC delivers a higher level of performance with lower latency. A 2 SSD drawer configuration offers performance that is comparable to a 10 HDD drawer configuration. It provides a high-performance option to customers with the bottleneck in I/O drawers as described in the following examples:

- ▶ TS7700T/TS7700C configurations without high TVC capacity requirements that require much data offload to physical tape/cloud
- ▶ TS7700 configurations without high TVC capacity requirements that process numerous low compression ratio data

Note: When I/O drawers are not the bottleneck, SSD drawer configuration does not necessarily improve the performance.

For more information about SSD-based TVC performance benchmark results, see [IBM TS7770, TS7770T, TS7770C, Release 5.4 Performance White Paper](#).

14.4 TS7700 throughput: Host I/O increments

In a TS7700, the read/write throughput from the Host is limited by the total number of Host I/O increment features that are installed.

The following Host I/O increment features are available:

- ▶ FC 5268 100 MBps increment
- ▶ FC 9268 Plant installation of 100 MBps throughput

At least one FC 9268 (default feature) is installed. This installation can be from one default increment (100 MBps uncompressed data) up to 40 increments (unlimited) by adding FC 5268. A TS7760 with an 8 Gbps FICON adapter that is installed is limited to 25 increments.

To understand how many MBps Host I/O a TS7700 can absorb as a maximum, the following aspects must be considered:

- ▶ Configuration:
 - Amount and type of drawers that are installed
 - FICON attachment
- ▶ Compression ratio
- ▶ Replication mode
- ▶ Blocksize of the I/O
- ▶ Read/write ratio

Especially in a TS7700T, the cache bandwidth can also be a limit for the Host I/O.

For more information about monitoring the Host I/O, see Chapter 13, “Monitoring” on page 667.

14.4.1 Host Throughput Feature Codes

If the TS7700 is equipped with 16-Gb FICON cards, consider that more throughput increments might need to be considered to unleash the full data transfer capabilities.

The previous 8-Gb FICON system featured a maximum of 25 throughput increments. Any data rates greater than 25 GBps were given for free. With the new 16-Gb cards (or after an upgrade occurs), the new throughput increment limit is 40 GBps.

If a cluster can achieve speeds faster than 25 GBps with 8-Gb FICON cards and 25 throughput increments in the past, that speed no longer can be true because the TS7700 limits them to exactly 25 GBps, supposing that 16-Gb FICON cards were installed and the same 25 throughput increments were licensed. Therefore, consider purchasing enough throughput increments (up to 40) to allow the TS7700 cluster to run at unthrottled speeds.

For more information, see Chapter 7, “Hardware configurations and upgrade considerations” on page 261.

Figure 14-5 shows the feature code license entry picture.



Figure 14-5 Feature Code license entry picture

Figure 14-6 shows the installed increments (FC 5268). In this example, four FC 5268 increment features are installed. The throughput is limited to 500 MBps because of the total number of the installed FC 5268 and default one FC 9268.

The screenshot displays the 'Yacko[0] (#BA02A): Feature Licenses' interface. At the top, there is a 'Refresh' button and a timestamp 'Last Refresh: Nov 9, 2012 11:02:51 PM'. Below this, the 'Current Available Resources' section shows 'Cluster common resources' with 'Cluster-Wide Disk Cache Enabled' at 3 TB and 'Cross-Cluster Communication (Grid)' at Enabled. A 'Peak data throughput' table shows 'vNo.' and 'Peak Data Throughput' with a value of 500 MB/s. The 'Currently activated feature licenses' table lists several licenses, including four '100-MB/sec Increment' licenses (FC 5268) and one 'Grid Enablement' license (FC 4015).

Select	Feature Code	Feature Description	License Key	Node	Node Serial Number	Activated	Expires
<input type="radio"/>	5267	1-TB Cache Enablement	dbb460c739e5ed31e0402a0d31...	(h0) hNode	78-00C9A	Oct 19, 2012 11:38:21 PM	Never
<input type="radio"/>	5267	1-TB Cache Enablement	dbb460c739e5ed31f533e47510...	(h0) hNode	78-00C9A	Oct 19, 2012 11:38:23 PM	Never
<input type="radio"/>	5267	1-TB Cache Enablement	dbb460c739e5ed31fe22b1e2e...	(h0) hNode	78-00C9A	Oct 19, 2012 11:38:25 PM	Never
<input type="radio"/>	4015	Grid Enablement	dbb460c739e5ed31f0a13af...	(h0) hNode	78-00C9A	Oct 19, 2012 11:38:39 PM	Never
<input type="radio"/>	9900	Encryption Configuration	dbb460c739e5ed31f1d1485a6c...	(h0) hNode	78-00C9A	Oct 19, 2012 11:38:50 PM	Never
<input type="radio"/>	4015	Grid Enablement	dbb460c739e5ed31f0a13af...	(v0) vNode	78-00C9A	Oct 19, 2012 11:38:39 PM	Never
<input type="radio"/>	5268	100-MB/sec Increment	dbb460c739e5ed317ba14e24d5...	(v0) vNode	78-00C9A	Oct 19, 2012 11:38:13 PM	Never
<input type="radio"/>	5268	100-MB/sec Increment	dbb460c739e5ed31f99ca13139...	(v0) vNode	78-00C9A	Oct 19, 2012 11:38:16 PM	Never
<input type="radio"/>	5268	100-MB/sec Increment	dbb460c739e5ed31f898d6abc2...	(v0) vNode	78-00C9A	Oct 19, 2012 11:38:18 PM	Never
<input type="radio"/>	5268	100-MB/sec Increment	dbb460c739e5ed31e0415e03a...	(v0) vNode	78-00C9A	Oct 19, 2012 11:38:20 PM	Never

Figure 14-6 Feature Code licenses example

14.4.2 Tuning for Host I/O

Tuning for the Host I/O is limited. If the increments are exhausted, the issue cannot be solved in the TS7700. 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 at the limit, see 14.6.1, “Tuning Cache bandwidth: Premigration” on page 766.

14.5 Considerations for Virtual Device Allocation

In the TS7700 GRID and z/OS environment, device allocation assistance (DAA) can be used in private (specific) mount to select a preferred cluster virtual device in grid, which maximizes TVC cache hit and optimizes performance and workload by minimizing remote mounts and remote cache accesses.

When DAA is enabled (default setting), in the grid configuration with TS7700T or TS7700C the virtual device is selected preferentially at the cluster where the volume already cached in the TVC. Conversely, when DAA is disabled, z/OS randomly selects a virtual device. It can cause a cache miss on TVC and need longer time to access data by recalling from tape or cloud or by remote-mounted if other cluster have the data in the TVC. If you have TS7700T or TS7700C in the grid configuration, the DAA function prevents cache miss and leads to maximizing performance.

For more information about DAA, see 2.4.15, “Device Allocation and Allocation Assistance” on page 78.

14.6 Cache throughput and cache bandwidth

The TS7700 cache features a finite bandwidth, depending on your actual environment (number of installed cache drawers and models). Other influences are the compression ratio and block sizes that are used. For example, the difference in cache bandwidth (depending on the number of cache drawers that are installed) is shown in Figure 14-7.




Figure 14-7 Performance versus Drawer Counts

For more information, see [IBM TS7700 R5.1 \(TS7770\) Performance](#).

Always use the most recently published performance white papers that are available at the [IBM Support website](#).

This TVC bandwidth is shared between the host I/O (compressed), copy activity (from and to other clusters), premigration activity, and 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.

For more information about monitoring the Cache throughput, see Chapter 13, "Monitoring" on page 667.

14.6.1 Tuning Cache bandwidth: Premigration

Several methods are available to tune the use of cache bandwidth. This tuning is applicable only to the TS7700T:

- ▶ When premigration is run (PMPRIOR and PMTHLVL)
- ▶ Number of drives for premigration
- ▶ Use delayed premigration to never premigrate data or delay the premigration to a more suitable time slot

For more information about tuning premigration in the TS7700C, see the **LIBRARY REQUEST,distributed_library,CLDSET** section in *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-55733.

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 MiBps, the CPU is more than 99% busy, and the total I/O rate (read and write) against the cache is greater than 200 MBps (*not* MiBps).

The circle on the graph that is shown in Figure 14-8 shows 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 premigrated tasks are allowed to start again.




Figure 14-8 Fast Host Write premigration algorithm

14.6.2 Premigration and premigration throttling values

The premigration and premigration throttling value parameters are used to define actions, which are triggered by the amount of nonpremigrated data in the cache. Both values are applicable to the TS7700T only.

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 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 must be premigrated:
 - Per pool
 - In total

Premigration Throttling threshold (PMTHLVL)

When this threshold is crossed, the host write throttle and copy throttle are started. The purpose is to slow incoming data to allow the amount of nonpremigrated data to be reduced and not rise above this threshold. The default value is 2000 GB.

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, several dependencies and consequences exist. Review these parameters to ensure that the parameters are adjusted, especially after hardware is replaced.

No guideline is available regarding the values of PMPRIOR and PMTHLVL. The following aspects must be considered:

- ▶ Installed a number of:
 - FC for Cache enablement
 - FC 5274 and FC 5279 for premigration queue size
- ▶ Workload profile
- ▶ Requirement regarding how long limited premigration data should stay in cache unpemigrated

Determine the balance between the value of PMPRIOR and PMTHLVL. If throttling occurs, it can be monitored with the MI or the VEHSTATS reports. Review the values periodically.

To adjust the parameters, use the **Host Console Request** command. When you attempt to define a value that cannot be used, the TS7700 automatically uses a suitable value. For more information about these settings, see this [IBM Support web page](#).

Use the following keywords:

- ▶ SETTING, CACHE, PMPRIOR
- ▶ SETTING, CACHE, PMTHLVL

PMPRIOR setting

If the value of PMPRIOR is crossed, the TS7700T starts the premigration. The premigration might decrease the resources that are available for other tasks in the TS7700T and shorten the peak throughput period of a workload window.

Raising this threshold increases the timeline where the TS7700T can run in Peak mode. However, the exposure is more for nonpremigrated data in the 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. If a value is specified that is higher than the premigration throttling threshold value (PMTHLVL), it is set to the premigration throttling threshold value.

If TS7700T has only one FC 5274 (that is, the maximum amount of queued premigration content is 1 TB), an internal PMPRIOR of 600 is used, even though the use of the **SETTING** command continues to show 1600 (the default value).

PMTHLVL setting

After the PMTHLVL is crossed, the Host I/O, remote writes, and copies into a TS7700T are throttled. In a TS7700T, only the 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 continues to increase.

Raising the threshold avoids the application of the throttles, and keeps host and copy throughput higher. However, the exposure is more for nonpremigrated data in the cache.

The default value is 2000 GB. The maximum is the total number of installed premigration queue sizes (FC 5274 and FC 5279).

A value that is greater than the total amount of FC 5274 and FC 5279 cannot be set. If a value is specified that is not less than the total amount of FC 5274 and FC 5279, the threshold is set to the total amount of FC 5274 and FC 5279. If TS7700T has only one FC 5274 (that is, the maximum amount of queued premigration content is 1 TB), an internal PMTHLVL of 1000 is used, even though the **SETTING** command continues to show 2000 (default value).

In a TS7700T, the recommendation is to set the PMTHLVL equal to the total amount of FC 5274 and FC 5279. If this recommendation is used, this value is the amount of FC that is installed.

14.6.3 Performance consideration for a cache DDM's rebuild

When a DDM fails and must be replaced, a REBUILD process of the cache DDM arrays occurs after the failed DDM is replaced. Access to the cache can result in some performance degradation, especially under many host accesses, migrations, recalls, or grid replications that lead to high usage of cache disks.

Several **LI REQ SETTING RBPRIOR/RBTHLVL** commands are available that help control the amount of resources that is devoted to the DDM REBUILD operation. These resources mitigate performance effects that might affect sensitive workloads from this degradation.

To mitigate this effect, the recommendation is to use one of the following **LI REQ** commands:

- ▶ Use the **LI REQ SETTING CACHE RBPRIOR 1** command to keep Cache Controller to always rebuild the cache DDM array at the slowest pace.

- ▶ Use the `LI REQ SETTING CACHE RBTHLVL 5000` command (only if `RBPRIOR=0`) to consider 5000 MBps as the required upper average disk I/O threshold in which the TS7700 allows the rebuild to run at high priority level. That is, when the average disk I/O rate is 100–5000 MBps, the cache rebuild or copyback is done at the lowest priority level 1 to allow more resources that are dedicated to host, copy, and premigration I/O.

When the average disk I/O rate is over 5000 MBps, the cache rebuild or copyback is done at the high priority level 4 for the CSA cache controller. When the average disk I/O rate is over 5000 MBps, the cache rebuild or copyback is done at the highest priority level 5 for the CSB cache controller.

You can monitor the cache disk I/O throughput by the disk utilization percent (in the historical statistics or the performance page in MI) and the primary cache device read/write (in the performance page in MI).

14.7 Grid link and replication performance

This section includes a description of the aspects and definitions that influence the grid link performance and provides information about 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 number of concurrent copy tasks
- ▶ Number of remote read/write operations
- ▶ Cloud-attached configuration
- ▶ DS8000 Object Store configuration

Cache bandwidth is always an influencing factor and was described in 14.6, “Cache throughput and cache bandwidth” on page 765.

14.7.1 Mixing different grid link adapters and traffic from Cloud attach or DS8000 object store considerations

Different grid link adapter types cannot be included in a single cluster. However, a situation can exist in a grid in which some clusters are connected to the grid link by using 10 Gb adapters, and other clusters are connected by using 1 Gbps adapters. This situation is especially true for migration or upgrade scenarios.

In the TS7700 grid, a 1:1 relationship exists between the primary and primary adapters, and the secondary and secondary adapters. Because of this relationship, in a mixed environment of two 10 Gbps and four 1 Gbps adapters, the clusters with the four 1 Gbps links cannot use the full speed of the installed grid link adapters.

In the current TS7700C configuration, grid link adapters are shared by communication traffic between clusters *and* between a cluster and cloud. When cloud traffic is high (especially in on-premises Object Storage), grid link adapters are highly used. This issue might affect the grid link data replication throughput between clusters in the grid.

R5.1 introduced enhanced cloud storage tier with grid awareness. With this support, any clusters that have no valid copy can access the cloud copy when the cluster that has a valid copy is not available. Therefore, if you expect the logical volumes to exist only in the cloud, you do not have to replicate them to other clusters. In this configuration, you can reduce the required network bandwidth.

Also, if you configure DS8000 Object Store function that connects to DS8000, data traffic between DS8000 and TS7700 might affect the grid link data replication throughput. Therefore, carefully plan to fit the performance requirements.

You can control DFSMSHsm data migration from DS8000 to TS7700 by setting the migration schedule in the z/OS side to avoid the grid link adapter competition between usual tape I/O and object store traffic. DS8000 object compression that was introduced in R5.1 can also be used to reduce the amount of data that goes through the grid link.

14.7.2 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 must 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
- ▶ Network traffic quality of service (QoS) capabilities

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 higher than the network infrastructure can 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 can match the capabilities of the underlying network, which results 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, which results 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 support flow control between the TS7700s and the switches, which allows 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 can pace with each other's rate.

- ▶ Ensure that the performance of the switch can handle the data rates that are expected from all of the network traffic.
- ▶ Ensure that the network traffic QoS capabilities are configured, such as IBM SAN42B-R with IP-extension or other appliances for QoS control.

Latency between the sites is the primary factor. However, because of bit error rates or the network cannot support the maximum capacity of the links, packet loss causes TCP to resend data, which multiplies the effect of the latency.

14.7.3 Selected replication mode

The number of concurrent running copy tasks and the replication mode influence the overall performance of a TS7700 grid. The following shared resources are used:

- ▶ CPU cycles of the TS7700
- ▶ Grid bandwidth
- ▶ Cache bandwidth

Synchronous mode copy

Synchronous mode copy can have a positive effect on the cache bandwidth. As opposed to all other replication modes in which the data must 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 not adhere to the same rules as the other copy modes, as listed in Table 14-1.

Table 14-1 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 are not throttled.	These copies can be throttled.
The number of concurrent copies can be controlled by a setting.	No	Yes
Copies can be halted by a <code>LI REQ GRIDCTL COPY DISABLE</code> command.	No	Yes

Synchronous mode copy considerations

Synchronous mode copy waits for the completion of remote write operation when explicit or implicit synchronous CCW is issued by the mainframe or TS7700 internal buffer becomes full.

For a single application perspective, the elapsed time might become longer than other copy modes when network latency is relatively long. For example, the required time for DFSMSHsm secondary space management might be affected in the longer network latency.

14.7.4 Tuning possibilities for copies: COPYCOUNT Control

Tuning the counts of the number of concurrent copy jobs over the grid links can be done for several reasons,

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, CPYCNT, RUN
- ▶ SETTING, CPYCNT, DEF

Increasing the copy count

Increasing the copy count can be beneficial if the following conditions exist:

- ▶ The grid links are not saturated.
- ▶ The number of 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.

Often, one grid link 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.

Consider that if too many copies are running concurrently, the grid link is overflowed. This issue can result in package loss and retries, and lead overall to a lower performance of the grid link environment.

If you want to increase the number, do so in smaller steps (5 or 10) to become familiar 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.

Decreasing 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 MiBps).
- ▶ If the bandwidth is limited and running too many copies in parallel prolongs the single copy time. This issue can result in timeouts. When you cross this threshold, the system switches from RUN to IMMED-Deferred. For a deferred copy, the copy is deleted from the copy tasks and is scheduled back to the copy queue.
- ▶ If packet loss is reported and a 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 timeout, which is reported as packet loss.

14.7.5 Tuning to avoid the throttling

Several tuning possibilities can be considered for the different types and reasons for throttling.

Some of these reasons are the result of parameter changes (PMTHLVL adjustments or ICOPYT); other issues can be solved only by providing higher bandwidth or more resources (cache or drives).

Although adjusting a parameter might be beneficial for a specific situation, it might affect other behaviors in the grid. Therefore, discuss such tuning measurements with your IBM representative.

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 20 or more minutes old. The TS7700 changes an immediate copy to immediate-deferred if the immediate copy is 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. For more information, see 13.8.2, “Handling Replication Exceptions” on page 741.

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 was observed when the host write throttled because immediate copies are turned off. You are trading off host I/O for the length of time that is required to complete an immediate copy.

For more information about enabling and disabling host write throttle because of immediate copies, see [TS7700 Library Request Command V5.4](#). Search for the following keywords:

- ▶ SETTING
- ▶ THROTTLE
- ▶ ICOPYT

14.7.6 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 users, host throughput is more important than the deferred copies, but for others, deferred copies are as important. You can tune the performance of the deferred copies by adjusting the DCT value and threshold.

DCT 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 greatly influences 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 must be adjusted from the default value, use 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 equally treated as host I/O. Depending on your RPO requirements, that setting is also feasible.

The DCT value can be set by using the **Host Console Request** command. For more information about setting this throttle, see [TS7700 Library Request Command V5.4](#) and search for the following keywords:

- ▶ SETTING
- ▶ THROTTLE
- ▶ DCOPYT

DCT value threshold

The DCT 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 because of the increased CPU power on the IBM Power7 processor. The CPU usage is often less than 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 DCTAVGTD:

- ▶ SETTING, THROTTLE, DCOPYT
- ▶ SETTING, THROTTLE, DCTAVGDT

Note: The recommendation is to *not* change DCTAVGDT and instead use for the tuning the DCOPYT only. Changing DCTAVGDT might not reduce the throttling as expected.

Application of DCT

The use of DCT is shown in this section. In Figure 14-9, 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 greater than the DCT threshold, which is set to the default of 100 MBps.

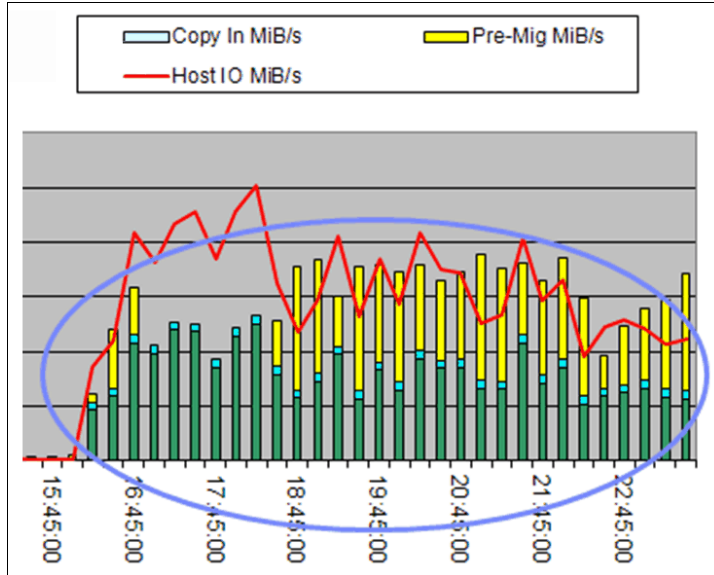


Figure 14-9 DCT being applied

Figure 14-9 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 14-10 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 14-10 DCT turned off

In Figure 14-10, 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).

14.8 Considerations for the backend TS7700T

In the backend, monitor the two different resources, backend drives and backend cartridges. Several possibilities are available to determine whether sufficient backend resources are available, how they are used, and how the use increases over time. The most difficult question is if sufficient backend drives are available.

14.8.1 Number of back-end drives

It is important to ensure that enough back-end drives are available. If not enough back-end drives are available, you might encounter the following issues:

- ▶ Recalls are too slow because no free back-end drive was available.
- ▶ Premigration cannot be processed sufficiently; therefore, throttling occurred because the limit of premigration queue sizes was reached.
- ▶ More backend cartridges are needed because no drives were available to run the reclaim.

The performance of the TS7700T diminishes if not enough back-end drives are available.

In an older TS7740, a direct dependency existed between Host I/O Increments throughput and the number of backend drives. This dependency 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 in a TS7700T. Part of the Host I/O can be written in the CP0, and this data is never premigrated. Also, a part of the data might be expired in the cache, even if they were written to CPx by using the Delay Premigration parameter.

Therefore, such a strict relationship no longer exists.

As a guideline for the TS7700T, use the ranges of back-end drives that are listed in Table 14-2 based on the host throughput that is configured for the TS7700T. The lower number of drives in the range is for scenarios that have few recalls. The upper number is for scenarios that have numerous recalls. These ranges are guidelines, *not* rules.

Table 14-2 Performance increments versus back-end drives

Throughput (MiBps)	Back-end drives
Up to 400	4–6
400–800	6–8
800–1200	8–12
1200–1600	10–16
1600 or higher	16

Therefore, if the Host increments cannot be used for guidance, the question is how to determine the number of needed backend drives.

Although no overall rule is available to make this determination, consider the following points:

- ▶ The more physical backend pools that 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 high. Therefore, a reclaim must copy a high amount of data from one physical tape to another tape. This copy process uses two drives, and the more data that must be transferred, the longer these drives are occupied. Reclaim is often a low-priority task, but the number of necessary tapes increases if not enough reclaims can be run.
- ▶ 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.

For more information about setting the Available Physical Drive Low and High warning levels, see [TS7700 Library Request Command V5.4](#). Use the following keywords:

- ▶ SETTING, ALERT, PDRVLOW
- ▶ SETTING, ALERT, PDRVCRT

14.8.2 Tune back-end drive usage

If you must tune the back-end drive usage, review the following information:

- ▶ Usage of delay premigration for data with short lifecycle
- ▶ Number of physical pools
- ▶ Reclaim operations
- ▶ Amount of premigration drives
- ▶ Copy Export operation

Delay premigration usage

To eliminate the necessity to premigrate and reclaim data, you might consider the use of delay premigration for some of your data. To determine whether this solution is viable, analyze the information from your tape management system. This information shows if any data has such a short lifetime cycle, and how much cache space is needed to set premigration delay time in a storage class.

Contact your IBM representative for assistance.

Number of physical pools

Each physical pool is treated independently regarding the backend drive usage, which 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 do not use any cache bandwidth resources, because the data is copied from physical tape to physical tape directly. If needed, the TS7740/TS7700T can allocate pairs of idle drives for reclaim operations, which ensure leaving 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.

The reclaim threshold and the amount of data that must be moved (depending on the stacked tape capacity and the reclaim percentage) are listed in Table 14-3. When the threshold is reduced from 40% to 20%, only half of the data must be reclaimed. This change reduces by half the time and resources that are needed for reclaim. However, it raises the needed number of backend cartridges and slots in the library.

Table 14-3 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	3000 GB	4000 GB

Inhibit Reclaim schedule

Use the Inhibit Reclaim schedule to inhibit reclaims during your busy periods, which leaves 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 be 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. 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.

By using the **Host Library Request** command, you can limit the number of reclaim tasks in the TS7700T. The second keyword RECLAIM can be used along with the third keyword of RCLMMAX. This expansion applies only to the TS7700T. Also, the Inhibit Reclaim schedule is still acknowledged.

The maximum number of reclaim tasks is limited by the TS7700T based on the number of available back-end drives, as listed in Table 14-4. These values changed during the evolution of the product, and might be different in previous releases.

Table 14-4 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 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 or later, premigration can benefit from having 8 - 10 drives available for premigration, the default value is 10. No benefit is gained by using more than 10 running pre-migrations.

The limit setting is in the 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.

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.

14.8.3 Number of back-end cartridges

The number of needed back-end cartridges is determined not only by the amount of data that is stored on the cartridges. The following parameters can also influence the number of cartridges you need:

- ▶ Reclaim value
- ▶ Number of physical pools and pool properties
- ▶ Delete Expire setting

To monitor your situation and the trend of the cartridge usage, several possibilities are available, which are described next.

14.8.4 Tuning of the usage of Back-end cartridges with VEHSTATS

Several possibilities are available to influence the number of used backend cartridges.

Reclaim value

As explained in the “Defining reclamation settings in a TS7700T” on page 581 back-end 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. At times, it is sufficient to change to a slightly higher reclaim value.

Amount of physical pool and pool properties

For each physical pool, two empty cartridges are often kept inside the pool. Especially for smaller capacity configurations with large capacity cartridges such as JD media, a larger number of cartridges are needed even if you do not want extra capacity.

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 is still be treated as valid data. For more information about delete expire, see 3.3.9, “Logical Volume Delete Expire Processing versus previous implementations” on page 132.

Note: A short delete expire value might reduce your cartridge usage, but a short value does not enable you to rescue any unintentionally deleted data. We suggest not to use a value below 5 days. A best practice is to use 7 days.

14.9 Cloud Tiering

In the cloud-attached configuration, monitor the two different resources, network bandwidth and premigration queue size. Several possibilities are available to determine whether you have sufficient network bandwidth and premigration queue, how they are used, and how the use increases over time. The most difficult question is if you have sufficient network bandwidth.

14.9.1 Network bandwidth and premigration queue size

Network bandwidth to public cloud is often limited compared to private cloud. It is also much slower than 3592 physical tape drives; therefore, consider the premigration queue size.

If the premigration backlog that is not yet premigrated to cloud reaches the configured premigration queue size based on installed FC 5274 (1 TB Active Premigration Queue) and FC 5279 (5 TB Active Premigration Queue) numbers, you see premigration throttling. After premigration throttling is triggered, host I/O throughput might be negatively affected.

If not enough network bandwidth and premigration queue size are available, select the suitably sized logical volumes to be migrated to the cloud.

Grid awareness of volumes in the cloud that was introduced in R5.1 might reduce the grid link usage between clusters. For more information, see 14.7.1, “Mixing different grid link adapters and traffic from Cloud attach or DS8000 object store considerations” on page 769.

14.9.2 Logical volume size

If network bandwidth is limited, premigration to cloud and recall from cloud for a logical volume requires more time compared to 3592 tape drives, especially when larger logical volume size such as 25 GB or 65 GB is used. For example, a 25 GB logical volume requires almost 40 minutes if only 100 Mbps bandwidth is available.

If recall cannot be completed within the specified mount wait time with MIH MOUNTMSG=YES,(MNTS=mm:ss in IECIOSxx parmlib), you see IOS070E MOUNT PENDING messages. This issue might trigger some error recovery procedures, depending on your message monitoring environment. Therefore, larger logical volume sizes might need to be accounted for in your message monitoring.

14.9.3 Premigrate and Recall time out

You can customize premigrate and recall time out value by using **LIBRARY REQUEST CLDSET** commands. You can also set the maximum concurrent tasks to premigrate and recall by using **LIBRARY REQUEST CLDSET** commands. If network bandwidth is narrow, you might need to set a longer timeout value or a smaller number of concurrent tasks.

14.10 TS7700 Advanced Object Store for DS8000

In this section, general considerations about the TS7700 Advanced Object Store that were introduced in R5.2.2 are provided. For more information about the Object Store function in earlier code level, see *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#).

In the TS7700 Advanced Object Store configuration, monitor the network bandwidth, especially when the object store function, usual logical volume replication, or cloud attach functions coexist within a TS7700 cluster.

Several possibilities are available to determine whether you have sufficient network bandwidth, how they are used, and how the use increases over time. The most difficult question is whether you have sufficient network bandwidth.

For more information, see [Transparent Cloud Tiering Performance R5.4](#).

14.10.1 Network bandwidth

TS7700 Advanced Object Store function uses the grid link adapters as communication paths between TS7700 and DS8000. Therefore, if you configure Object Store functions on the same TS7700 cluster that are used for tape I/Os or cloud attach functions (especially on-premises attach), the traffic might affect the object store throughput because the network bandwidth and adapters are shared by those functions.

Consider the following points:

- ▶ In R5.0, you can display the point-in-time grid link activity by using the **GRLNKACT** command, including Object Store and GGM functions.
- ▶ In R5.1, you can use DS8000 object compression with z/OS setting. The amount of traffic between the DS8000 and the TS7700 (Object Store) can be reduced by the compression.
- ▶ In R5.2.2, TS7700 Advanced Object Store supports object replication to any clusters in a Grid that supports FC 5283 by Synchronous Mode Copy or Deferred Mode Copy.

14.10.2 Network latency

Each DS8000 targets one or two TS7700s in a Grid for load-balancing. If TS7700 clusters are in the local site with DS8000, performance is not affected because the network latency is almost zero.

However, if TS7700 clusters that receive objects are remotely located from DS8000 and Synchronous Mode Copy is used for replication, data migration from DS8000 to TS7700 requires more time to complete. This extra time is required because synchronous mode write waits for completion from remote TS7700 during the network latency and this synchronization occurs in relatively small buffer size. Therefore, single task performance might be limited when the network latency is long, even if the sufficient network bandwidth is provided.

If you need the higher total throughput even in longer latency, consider the use of Deferred Mode Copy for replication. In addition, raising the multi-task level can be effective. For example, you can specify the maximum migration tasks for DFSMSHsm automatic migration by using **SEYSYS MAXMIGRATIONTASKS**. This consideration is similar to TS7700 remote grid replication in synchronous mode for DFSMSHsm ML2 migration.



Copy Export

This chapter describes the Copy Export and Copy Export Recovery functions and how to use them.

This chapter includes the following topics:

- ▶ 15.1, “Copy Export overview and considerations” on page 784
- ▶ 15.2, “Implementing and running Copy Export” on page 800
- ▶ 15.3, “Using Copy Export Recovery” on page 810
- ▶ 15.4, “Using Copy Exported tape for damaged volume recovery” on page 818

15.1 Copy Export overview and considerations

Copy Export enables a copy of selected logical volumes that is written to the IBM TS7700 to be removed and taken offsite for disaster recovery (DR) 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. TS7700 supported DS8K Object Store since R5.0, but Copy Export is applicable to logical volumes only, not to objects.

15.1.1 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 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 by using the TS7700 Management Interface (MI).

An option on the MI window enables a secondary pool to be designated 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.

15.1.2 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 DR purposes, performing it soon after the data is created minimizes the time for the recovery point objective (RPO).

When the time comes to start a Copy Export, a Copy Export job is run from the production host. Logical volumes that are assigned to the Copy Export pool and written before the Copy Export job is issued are candidates to be exported.

After the Copy Export job is issued, TS7700 continues pre-migrating such candidate logical volumes. It waits for completion of their pre-migrations, then proceeds to export them.

The TS7700 pre-migrates any logical volumes in the Copy Export pool that were not pre-migrated. Any new logical volumes that are written after the Copy Export operation is started are not included in the Copy Export set of physical volumes. These volumes are 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 a limited number of physical volumes. For recovery, 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 mounted during a Copy Export operation are excluded from the export set, as are any volumes that are not in the Tape Volume Cache (TVC) of the export cluster.

The host that starts 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.

At the R5.0 or earlier code level, the export list volume always features version number 03. Since R5.1, a new format export list volume that features version number 04 is supported.

Starting with R5.1, the new format export list volume supported version number 04. For more information about export list volume format, see [IBM TS7700 Series Copy Export Function User's Guide](#).

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 must 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 DR 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.

Copy Export Merge

If the copy-exported tapes are used for recovery, the recovery TS7700 cannot contain any previous data and the recovery process cannot merge data from more than one source TS7700 together. However, for the customers running R2.1 or later code on a TS7700, a service offering called Copy Export Merge is available. With Copy Export Merge, a customer can merge data from a copy export backup tape into a TS7700 (stand-alone or grid).

15.1.3 General considerations for Copy Export

Consider the following points when you are planning to use the Copy Export function for DR:

- ▶ TS7700T supports 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 features equal or later 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 a TS7700T. A TS7740 Copy Export set can also be restored into TS7700T.
- ▶ 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 also must be set up to handle encrypted volumes and access 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 generation 3592 model drives.
- ▶ When the Copy Export acceleration (LMTDBPVL) option is used, the database backup is appended to only the first two and last two in ascending order of name of 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 a failure occurs 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 by using of the following JCL:
 - **OPTIONS1,COPY,LMTDBPVL** (volumes are marked as export hold)
 - **OPTIONS1,COPY,EJECT,LMTDBPVL** or **OPTIONS1,COPY,LMTDBPVL,EJECT** (volumes are 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. For more information about exportable physical volumes, see Table 15-1 on page 788. 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 to be exported next time are the volumes for that pool that were newly written or modified since the last export began.
- ▶ Previously exported volumes that were changed are no longer active copy-exported volumes. For recovery, all exported physical volumes that still contain active data from a source TS7700 must 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 logical volumes that are assigned to the secondary pool that is specified in the export list volume that is on a physical volume of the pool or in the cache of the TS7700 performing the export operation are considered for export.
- ▶ Logical volumes to be exported that are only in the cache and not mounted when the Copy Export operation is started are copied to stacked volumes in the secondary pool as part of the Copy Export operation.
- ▶ Logical volumes that are compressed by the new Compression Method (LZ4 or ZSTD) can be exported in the same manner as they were previously.
- ▶ 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 starts is not part of the export regardless of whether the lvol is written from host or copied from another cluster. Such lvol is written to a physical volume in the pool but is not exported.

- ▶ Logical volumes that are mounted when the Copy Export operation is started cannot be 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 be defined previously 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 process occurs regardless of whether that secondary pool is involved in the current Copy Export operation.

The library that is associated with the TS7700 that is 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 ended by the TS7700 when fewer than four physical tape drives are available.

In heterogeneous drive configurations, the operation is ended when fewer than four drives are available for each installed drive type. If the “not enough tape drives” condition occurs, the processing for the physical stacked volume in progress is completed. Also, the export status file records reflect what was completed before the operation was ended.

- ▶ 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 include 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 issue 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 the secondary pool that is used for Copy Export. You can change it by 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.
- ▶ An available service offering that is called Copy Export Merge allows a customer to merge data from more than one source TS7700 Copy Export backup. However, the recovery TS7700 cannot contain any previous data, and a client-started recovery process cannot merge data from more than one source TS7700. As a part of the Copy Export Recovery, an option is provided to erase any previous data on the TS7700. This process enables a

TS7700 that is used for DR testing to be reused for testing of a different source TS7700's data.

- ▶ If Copy Export is issued while the maintenance mode of the physical tape library is enabled by the **LI REQ PHYSLIB MAINT ENABLE** command, it is suspended until the maintenance mode is disabled. Copy Export must be canceled to exit from the suspended state. For more information, see 15.2.4, "Canceling a Copy Export operation" on page 808.
- ▶ 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 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.

The exportable physical volumes that are based on tape drives, export format, and the LMTDBPVL option are listed in Table 15-1.

Important: Support situation of TS1160 is a little complicated:

- ▶ R5.2.1 PGA1 supports the TS1160 tape drive but cannot copy export tapes in the TS1155 recording format (55F format).
- ▶ R5.2.1 PGA2 and R5.3 has full support for the TS1160 tape drive including copy exporting JD/JL media in the TS1155 recording format (55F format)

Table 15-1 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
60F only	Default	Yes or no	▶ JC/JK/JD/JL in E08 format
			▶ JD/JL in 55F format
	60F		▶ JE/JM in 60F format
	55F		▶ JD/JL in 55F format
	E08	▶ JC/JK/JD/JL in E08 format	
60F and E08	Default	Yes or no	▶ JC/JK in E07 format
			▶ JC/JK/JD/JL in E08 format
			▶ JD/JL in 55F format
			▶ JE/JM in 60F format
	60F	▶ JE/JM in 60F format	
	55F	▶ JD/JL in 55F format	
	E08	▶ JC/JK/JD/JL in E08 format	

Installed tape drives	Export format of copy export pool	LMTDBPVL enabled?	Exportable physical volumes
60F and E07	Default	Yes	<ul style="list-style-type: none"> ▶ JB in E05 format ▶ JB in E06 format ▶ JB/JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format
	60F		▶ JE/JM in 60F format
	55F		▶ JD/JL in 55F format
	E08		▶ JC/JK/JD/JL in E08 format
	Default	No	<ul style="list-style-type: none"> ▶ JB in E06 format ▶ JB/JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format
	60F		▶ JE/JM in 60F format
	55F		▶ JD/JL in 55F format
	E08		JC/JK/JD/JL in E08 format
60F and E06	Default	Yes	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JB in E05 format ▶ JB in E06 format ▶ JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format
	60F		▶ JE/JM in 60F format
	55F		▶ JD/JL in 55F format
	E08		▶ JC/JK/JD/JL in E08 format
	Default	No	<ul style="list-style-type: none"> ▶ JB in E05 format ▶ JB in E06 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format
	60F		▶ JE/JM in 60F format
	55F		▶ JD/JL in 55F format
	E08		▶ JC/JK/JD/JL in E08 format

Installed tape drives	Export format of copy export pool	LMTDBPVL enabled?	Exportable physical volumes	
60F and E05	Default	Yes	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JB in E05 format ▶ JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format 	
	60F		▶ JE/JM in 60F format	
	55F		▶ JD/JL in 55F format	
	E08		▶ JC/JK/JD/JL in E08 format	
	Default	No	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JB in E05 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format 	
	60F		▶ JE/JM in 60F format	
	55F		▶ JD/JL in 55F format	
	E08		▶ JC/JK/JD/JL in E08 format	
	60F and J1A	Default	Yes	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format
		60F		▶ JE/JM in 60F format
		55F		▶ JD/JL in 55F format
		E08		▶ JC/JK/JD/JL in E08 format
Default		No	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JC/JK/JD/JL in E08 format ▶ JD/JL in 55F format ▶ JE/JM in 60F format 	
60F			▶ JE/JM in 60F format	
55F			▶ JD/JL in 55F format	
E08			▶ JC/JK/JD/JL in E08 format	
E08 only		Default	Yes or no	<ul style="list-style-type: none"> ▶ JC/JK in E07 format ▶ 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	<ul style="list-style-type: none"> ▶ JB in E05 format ▶ JB in E06 format ▶ JB/JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format
	E08		▶ JC/JK/JD/JL in E08 format
	E07		▶ JB/JC/JK in E07 format
	Default	No	<ul style="list-style-type: none"> ▶ JB in E06 format ▶ JB/JC/JK in E07 format ▶ 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	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JA/JJ/JB in E05/E06 format ▶ JC/JK in E07 format ▶ JC/JK/JD/JL in E08 format
	E08		▶ JC/JK/JD/JL in E08 format
	E07		▶ JC/JK in E07 format
	Default	No	<ul style="list-style-type: none"> ▶ JA/JJ/JB in E05/E06 format ▶ JC/JK in E07 format ▶ 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	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JA/JJ/JB in E05 format ▶ JC/JK in E07 format ▶ 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	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JC/JK in E07 format ▶ 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	<ul style="list-style-type: none"> ▶ JB in E06 format ▶ 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

Installed tape drives	Export format of copy export pool	LMTDBPVL enabled?	Exportable physical volumes
E05 only	Default	Yes or no	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format ▶ JA/JJ/JB in E05 format
J1A only	Default	Yes or no	<ul style="list-style-type: none"> ▶ JA/JJ in J1A format

15.1.4 Copy Export grid considerations

Copy Export is supported in grid and stand-alone environments. In this section, we review 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 that is 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. A service offering that is called Copy Export Merge is available, but you cannot merge copy-exported volumes from more than one source TS7700 in the recovery TS7700.
- ▶ Define an MC name for each specific set of data to export. On the TS7700 that is used to export that data, define a secondary physical volume pool for that MC name and ensure that you indicate that it is an export pool. Although you must 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 include a secondary copy of the data in case the primary copies on other TS7700s are inaccessible.

- ▶ While you are defining the MC name for the data, ensure that the TS7700 to perform the Copy Export operation includes a copy policy that specifies that it is to have a copy.
- ▶ If a logical volume is to be copied to the TS7700 that is performing the Copy Export operation and that copy did not yet complete when the export is started, it is not included in the current export operation.
- ▶ When the Copy Export operation is run, the export list volume must be valid on only the TS7700 that is performing the operation. 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 on only 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 started is on more than one TS7700, the Copy Export operation fails.

Tip: If the MC specified for export list volume is defined to have copies on 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 the following Copy Export example in a three-way grid environment:

- a. A Copy Export with the export list volume EXP000 is started 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], which indicates that the only copy of EXP000 is on C2.
- c. If Copy Policy Override is activated on the C0 and the Copy Export is started from the host that is attached to C0, a copy of EXP000 is created on the C0 *and* C2.
- 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 following example, assume that the TS7700 that is to perform the Copy Export operation is Cluster 1 in a two-way grid environment. The pool on that cluster to export is pool 8. 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 started by using any virtual tape drive in the TS7700 grid configuration. It does not need to be started on a virtual drive address in the TS7700 that will 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 statuses are broadcast to all hosts that are 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 started will be exported by the next Copy Export. You can check the copy status of the logical volumes on the TS7700 that is to perform the Copy Export operation before starting 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 a service offering is available called Copy Export Merge where a customer can merge data from more than one source TS7700 Copy Export backup, Copy Export Recovery is always performed to a stand-alone empty TS7700. As part of a client-started recovery process, the recovery TS7700 processes all grid-related information in the database, converting it to resemble 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 the following 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,volsers,DELETE
```

15.1.5 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 that is based on your provided criteria for its pool.

Figure 15-1 on page 795 shows how the Reclaim Threshold Percentage is set in Physical Volume Pool Properties. If the ratio between active data size and total bytes that are 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 the Inhibit Reclaim Schedule. If more than one volume is 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.

Another Offsite Reclaim type is 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,volsers,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 the Inhibit Offsite Reclaim Schedule.

If the primary logical volume copies are on the physical volume that is in the Read Only Recovery (ROR) state, Offsite Reclaim skips re-creating secondary copies of them. Copy Exported physical volumes that include their secondary copies must wait for the ROR process completion to be empty. Also, while the ROR process for such physical volume is running, the next Copy Export can be kept waiting until the ROR process completes.

Figure 15-1 shows the Reclaim Threshold Percentage for a normal Offsite Reclaim.

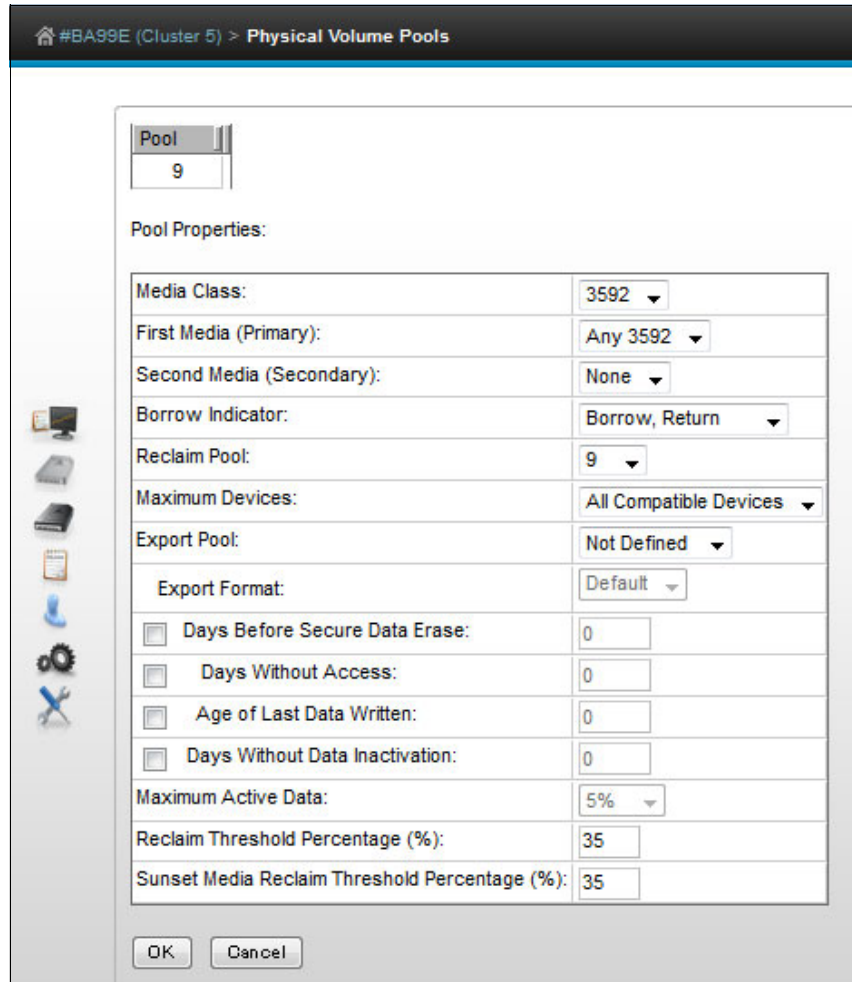


Figure 15-1 Reclaim Threshold Percentage is set in Physical Volume Pool Properties

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 the 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.

15.1.6 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 listed in Table 15-2.

Note: All messages are prefaced with CBR3750I.

Table 15-2 SYSLOG messages from the library

Message description	Action needed
<p>E0000 EXPORT OPERATION STARTED FOR EXPORT LIST VOLUME XXXXXX</p> <p>This message is generated when the TS7700 begins the Copy Export operation.</p>	None
<p>E0002 OPENING EXPORT LIST VOLUME XXXXXX FAILED</p> <p>This message is generated when opening the export list volume failed during the Copy Export operation.</p>	Check whether the export list volume or cache file system is in a bad state.
<p>E1007 OPENING EXPORT LIST VOLUME XXXXXX FAILED. TAPE EXPORT CONTAINS EXPP00L</p> <p>This message is for version 04 Export List File only. It is generated when opening Export List File Volume fails during copy export operation. Opening Export List File Volume fails because it contains Cloud Pool in Export Pool record even though OPTYPE is TAPE.</p>	Check whether the EXPPOOLS record of Export List Volume contains any Cloud Pool (contains record header "EXPPOOLS" only) in the case of OPTYPE TAPE.
<p>E1008 OPENING EXPORT LIST VOLUME XXXXXX FAILED. TAPE EXPORT CONTAINS DBP00L</p> <p>This message is for version 04 Export List File only. It is generated when opening Export List File Volume fails during copy export operation. Opening Export List File Volume fails because it contains Cloud Pool in DB Pool record even though OPTYPE is TAPE.</p>	Check whether the DBPOOLS record of Export List Volume contains any Cloud Pool (contains record header "DBPOOLS" only) in the case of OPTYPE TAPE.
<p>E0005 ALL EXPORT PROCESSING COMPLETED FOR EXPORT LIST VOLUME XXXXXX</p> <p>This message is generated when the TS7700 completes an export operation.</p>	None
<p>E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL IN EXPORT-HOLD</p> <p>This message is generated during Copy Export operations when an exported stacked volume 'YYYYYY' is assigned to the export-hold category. The 'LLLLLLLL' field is replaced with the distributed library name of the TS7700 performing the export operation.</p>	None
<p>E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL IN EJECT</p> <p>This message is generated during Copy Export operations when an exported stacked volume 'YYYYYY' is assigned to the eject category at R3.1 or earlier code level. 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.</p>	Remove ejected volumes from the convenience I/O station.

Message description	Action needed
<p>E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL QUEUED FOR EJECT</p> <p>This message is generated during Copy Export operations when an exported stacked volume 'YYYYYY' is assigned to the eject category at code level from R3.2 to R4.1.1. 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.</p>	<p>Remove ejected volumes from the convenience I/O station.</p>
<p>E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL IN EJECT-QUEUE</p> <p>This message is generated during Copy Export operations when an exported stacked volume 'YYYYYY' is assigned to the eject category at R4.1.2 or later code level. 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.</p>	<p>Remove ejected volumes from the convenience I/O station.</p>
<p>E0013 EXPORT PROCESSING SUSPENDED, WAITING FOR SCRATCH VOLUME</p> <p>This message is generated every 5 minutes when the TS7700 needs a scratch stacked volume to continue export processing and there are none available.</p>	<p>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.</p>
<p>E0014 EXPORT PROCESSING RESUMED, SCRATCH VOLUME MADE AVAILABLE</p> <p>This message is generated 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.</p>	<p>None</p>
<p>E0015 EXPORT PROCESSING TERMINATED, WAITING FOR SCRATCH VOLUME</p> <p>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.</p>	<p>The operator must make more TS7700 stacked volumes available, analyze the export status file on the export list volume, and reissue the export operation.</p>
<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>	<p>None</p>
<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>	<p>None</p>
<p>E0018 EXPORT TERMINATED, EXCESSIVE TIME FOR COPY TO STACKED VOLUMES</p> <p>The export process 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>	<p>Call for IBM support.</p>
<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>	<p>None</p>

Message description	Action needed
<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 was 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 is completed.</p>	Analyze 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 that is performing the export operation determines 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 that is resident on the TS7700. The 'LLLLLLLL' field is replaced by the distributed library name of the TS7700 that is performing the export operation. The export operation continues with the valid copies.</p>	When the export operation completes, analyze the export status file on the export list volume to determine the logical volumes that were not exported. Ensure that they completed their copy operations and then perform another export operation.
<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>	The logical volume and the physical volume are 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>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>	None

Message description	Action needed
<p>E0030 STACKED VOLUME XXXXXX RETURNED TO THE LIBRARY LLLLLLLL</p> <p>This message is generated when Copy Exported physical volume 'XXXXXX' is inserted back to the physical tape library. 'LLLLLLLL' field is replaced by the distributed library name of the TS7700 to which the physical tape library is attached.</p>	None
<p>R0000 RECLAIM SUCCESSFUL FOR EXPORTED STACKED VOLUME XXXXXX - YOU CAN RETURN IT TO THE SPECIFIED LIBRARY LLLLLLLL</p> <p>This message is generated when the TS7700 successfully completes reclaim processing for an exported stacked volume 'XXXXXX' 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>	The exported physical volume no longer contains active data and can be returned to a physical tape library that is attached to distributed library 'LLLLLLLL' from its offsite location for reuse. If it is placed in export-hold, it should be returned when the next copy export is completed.

15.1.7 Copy Export and DFSMSrmm

For users of DFSMSrmm, DFSMSrmm automatically handles and tracks the stacked volumes that are created by Copy Export when stacked volume support is enabled. However, which logical volume copies are on the stacked volume cannot be tracked. Retain the updated export list file that you created and the library that was updated so that a record exists of the logical volumes that were exported, and in which exported stacked volume that they are stored.

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 15-2 on page 796). Removable Media Management (RMM) intercepts this message and performs one of the following actions:

- ▶ If the stacked volume is predefined to RMM, RMM marks the volume as ejected or in-transit, and sets AUTOMOVE.
- ▶ 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), marks the volume as in-transit, and sets AUTOMOVE.

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 15-1.

Example 15-1 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 movement procedures can be used to move and track them.

A copy exported stacked 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 (**LI REQ,libname,COPYEXP,volser,RECLAIM**).

When it becomes eligible for reclaim, the exported stacked volume no longer contains active data and can be returned from its offsite location for reuse.

When the TS7700 Virtualization Engine successfully completed reclaim processing for a previously exported stacked volume, the following message is issued by the library and appears as part of the message text for CBR3750I:

```
CBR3750I Message from library library-name: R0000 RECLAIM SUCCESSFUL FOR EXPORTED
STACKED VOLUME volser.
```

Note: This message is generated for only reclaims that are started by LIBRARY REQUEST. This message is *not* issued for reclaims of copy-exported volumes that were made eligible through normal pool policy.

RMM intercepts this R0000 status message and, if a move is possible and required, marks the volume to be moved back to the library. The results of volume processing for a reclaimed volume are as though the following command was issued:

```
RMM CV volser HOME(library-name)
```

Then, if the volume is not library resident and is not moving, the following command is issued:

```
RMM CV volser LOCATION(HOME)
```

Finally, the volume's required location is set to the library-name to ensure that if any further movement is needed, the correct destination can be set by inventory management DSTORE processing and the volume is placed under manual move control. The use of the MANUALMOVE setting ensures that the volume is not moved again by VRS policies until reused for a later copy export. Reclaim support is provided regardless of whether stacked volume support is enabled.

For more information and error messages that are related to the Copy Export function in RMM, see *z/OS DFSMSrmm Implementation and Customization Guide*, [SC23-6874](#).

15.2 Implementing and running Copy Export

Implementing and running Copy Export are described in this section. For more information and error messages that relate to the Copy Export function, see [IBM TS7700 Series Copy Export Function User's Guide](#).

15.2.1 Setting up data management definitions

To set up the data management definitions, complete the following steps:

1. Decide the MC construct name (or names).

As part of the plan for the use of the Copy Export function, you must decide on at least one MC construct name. A preferred practice is to make the name meaningful and related 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 DR 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 more information 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.

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 more information about how to add an MC, see 10.2.2, “Management Classes window” on page 478.

4. Define the VOLSER ranges for the 3592 media.

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 help simplify the task of knowing the volumes to use in recovery. Also, it can help simplify the task of returning a volume that no longer has active data on it to the TS7700 that manages it.

For more information about how to define the VOLSER ranges, see “Defining VOLSER ranges for physical volumes” on page 578.

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 specified 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 579.

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.

15.2.2 Validating before activating the Copy Export function

Before the logical volumes are exported, you must perform several general validations. Before you start 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.

Complete the following 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 command 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.

The output of the **PDRIVE** request is shown in Example 15-2. This command is valid only when run against a distributed library.

Example 15-2 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 0.1
  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 15-2, 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 S00006 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 TS7700 MI with a maximum of 10,000 volumes. You can use the Library Request host console command that specifies the **POOLCNT** request. The response to the **LI REQ, <distributed library-ID>, POOLCNT** command is shown in Example 15-3 on page 803.

Example 15-3 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 0.0
  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 15-3 shows the command **POOLCNT**. The response is listed per pool:

- Media type used for each pool
- Number of:
 - Empty physical volumes that are available for Scratch processing
 - Physical volumes in the filling state
 - Full volumes
 - Physical volumes that were reclaimed, but need to be erased
 - Physical volumes in read-only recovery (ROR) state
 - Volumes unavailable or in a destroyed state (1 in Pool 1)
 - 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 (see Figure 15-2).

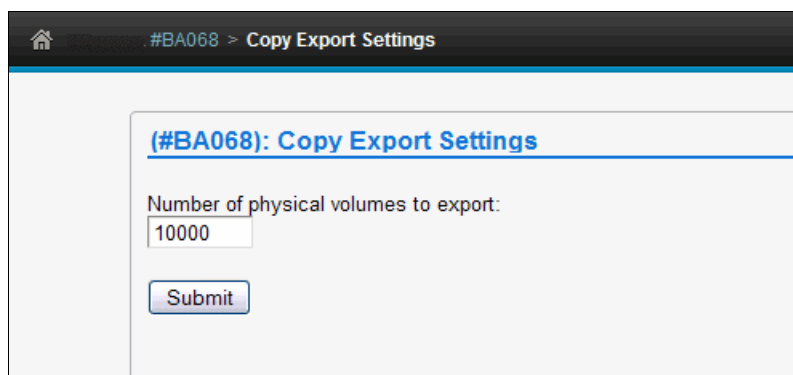


Figure 15-2 Maximum allowable number of volumes in a copy-exported state

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 12.1.3, "Host Console Request function" on page 632.

15.2.3 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 (see Example 15-4).

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 is run. If copies from other clusters exist that are not 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 the following steps to run the Copy Export operation:

1. Create the export list volume JCL (see Example 15-4).

Example 15-4 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 the following files:
 - An export list file, as created in step 1 in Example 15-4 on page 804. 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 15-4 on page 804. This file is reserved for future use.
 - An export status file, as created in step 3 in Example 15-4 on page 804. In this file, the information is stored from the Copy Export operation. Keep this file because it contains information that is 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 have the ACS routines assign this MC, or you can specify it in the JCL. These files must have the same expiration dates as the longest of the logical volumes you export because they must be kept for reference.

Figure 15-3 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 (N) on the other cluster.




Figure 15-3 Management Class settings for the export list volume

- The Copy Export operation is started 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 15-5.

Example 15-5 Library export command

```
LIBRARY EXPORT,logical VOLSER
```

- The host sends a command to the composite library. From there, it is routed to the TS7700 where the export list volume is stored.
- The running TS7700 validates the request, checking for required resources, and if all is acceptable, the Copy Export continues.
- Logical volumes that are related to the exported pool that are still only in the cache can delay the process. They are copied to physical volumes in the pool as part of the Copy Export run.
- Messages about the progress are sent to the system console. All messages are in the format that is shown in Example 15-6. For more information about Library Message Text, see Table 15-2 on page 796.

Example 15-6 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 15-4 shows the physical volume details.

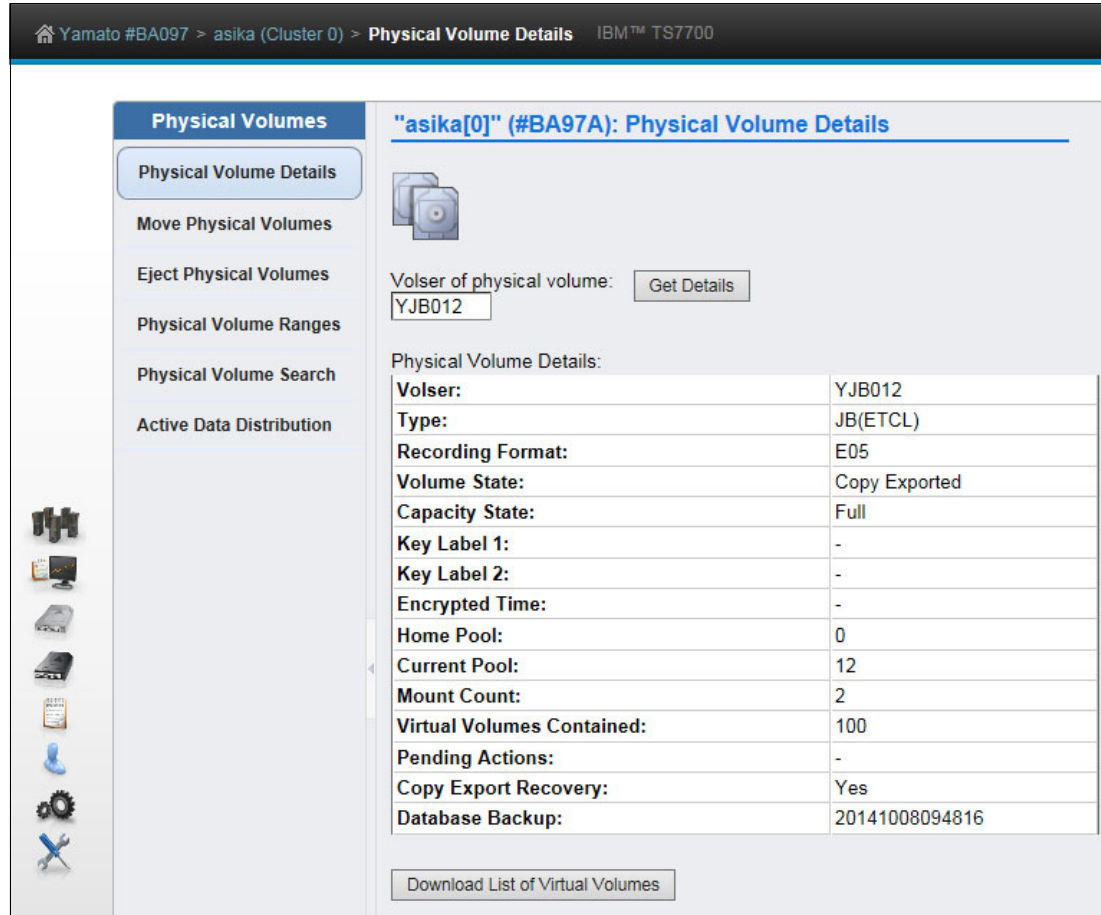


Figure 15-4 Physical volume details selection for list of exported volumes

Note: The copy export can also be started through JCL by using the CBRXLCS FUNC=EXPORT programming interface. For more information, see SAMPLE members CBRSPX03 and CBRSPX03, which is provided in SYS1.SAMPLIB and documented in [DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries](#).

Sample member CBRSPX03 writes the three required files on the export list volume by using a private volume and export list format 03 and has a fourth step (STEP4) that starts CBRSPX03 to start the copy export. CBRSPX03 is an example program that starts the CBRXLCS programming interface and must be modified to suit your business needs. When modified, it must be assembled and link-edited on your system for it to be usable through JCL.

15.2.4 Canceling a Copy Export operation

Examine the export status file records to see what was processed before the cancellation request. Any physical volumes that completed the export process must be processed as though the export operation was completed.

A Copy Export operation can be canceled for many reasons, including the following examples:

- ▶ After starting a Copy Export operation, you might realize that the pool that is being processed for export is incorrect.
- ▶ After starting a Copy Export operation, you might realize that the maintenance mode of the physical tape library is enabled.
- ▶ 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, no physical scratch volumes are 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 started 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 service CBRXLCS.

If an export operation must be canceled and no host 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 reached the point where it was 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 that indicates that the export was canceled by a host request. The message contains information about how much export processing was completed before the execution of the cancellation request.

15.2.5 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. Consider the following points:

- ▶ If no errors or exceptions were encountered during the operation, message CBR3855I is generated. The message features the format that is shown in Example 15-7.

Example 15-7 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 features the format that is shown in Example 15-8.

Example 15-8 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 provide statistics about what was processed during the operation. The following statistics are reported:

- ▶ Requested-number: 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: 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 exported are not considered to be resident in the TS7700.
- ▶ Exported-number: The number of logical volumes that were successfully exported.
- ▶ Stacked-number: The number of physical volumes that were successfully exported.
- ▶ MBytes Exported: The number of megabytes (MB) that are 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 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 effect is primarily because of the optimization of the operation by using multiple available drives that are configured by Maximum Devices in Pool Properties in the MI for the Copy Export pool.

15.3 Using Copy Export Recovery

The recovery process can be done in a test mode for DR testing purposes. This process enables a test restore without compromising the contents of the Copy Export sets. This section provides an example of how to use a Copy Export Recovery process.

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, an IBM service offering is available 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.

15.3.1 Planning and considerations for testing Copy Export Recovery

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 in this section.

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.

Restoring a TS7700 in an existing grid by using this way allows data that might exist 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.
- ▶ A limited heterogeneous configuration might be required if TS1150 or TS1160 drives are required. Drive format is also a prerequisite for copy export recovery. The recovery TS7700T must use one or more drive models that can support the formats that are used to source all or any part of the copy export set.
- ▶ If the source copy-exported volumes are encrypted, the recovery TS7700 must access a key manager that has the Encryption Keys for the data.
- ▶ Enough library storage slots must exist 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.
- ▶ Copy export recovery only occurs only within a TS7700T that has a code level equal to or higher than the code level of the TS7700 that exported the last subset of physical volumes. Since copy export was introduced in R1.3, copy exported sets from R1.3 or later can be supported, but with limitations. Because various Db2 versions existed throughout the TS7700's release cycle, each of these versions must be included with a TS7700 installation image for any potential export restore that might occur.

For example, the Db2 version that is used during R1.3 must be temporarily installed to restore a DB from a R1.3 exported tape. To reduce the code image size, R3.3 began limiting the number of Db2 levels that are included. By limiting which Db2 versions are present, it limits how far back a copy export set can be supported.

As of R5.2, a copy export set can be restored if the last export operation of the entire set occurred on a R4.1.2 or later TS7700. Tapes that were created at levels *before* R4.1.2 might exist in the set from an earlier export, but the tape that is used for restoring the DB must be from a R4.1.2 or newer export request. The supported level correspondences are listed in Table 15-3.

Table 15-3 Code level at recovery site, Db2 level, and export site code level supported

Code Level At Recovery Site	Db2 Level	Copy Export Recovery Level Supported
R5.3, R5.2.2, R5.2.1	11.5	R4.1.2 and higher
R5.1, R5.0, R4.2, R4.1.2	11.1	R3.3 and higher
R4.1.1, R3.3	10.5	R2.1 and higher

- ▶ Recovery sites must be able to support the media of the copy exported physical volume so Lifecycle Management also needs to be considered. See 15.3.2, "Lifecycle Management" on page 812
- ▶ Because the code levels on the recovery TS7700 must be at the same or later code level as the source TS7700, you do not need to care about whether logical volumes on copy-exported physical volumes are compressed by new Compression Method (LZ4 or ZSTD) or not. The recovery TS7700 must read compressed logical volumes.
- ▶ 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.

- ▶ Setting “TVCWDEG=EQUAL or LOWER” and “PRETHDEG=DISABLE” by **LI REQ** “SETTING2, PHYSLIB” to copy export recovery cluster is recommended to prevent workloads from stopping because of a physical library degraded state, such as out of scratch state.

As for the TVCDEG setting, avoid setting a value of DISABLE. Although the cluster can enter the out of scratch state, reads from tape and writes to CP0 continue to work without changing the PHYSLIB settings. However, if the host attempts to write data to CP1–CP7, the host jobs are suspended because of the out of scratch condition.

For more information about **LI REQ** “SETTING2, PHYSLIB, TVCWDEG” and **LI REQ** “SETTING2, PHYSLIB, PRETHDEG”, see [TS7700 Library Request Command white paper](#).

- ▶ 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 was 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.

15.3.2 Lifecycle Management

Lifecycle Management must also be considered when managing exported physical volumes. The physical tape drives that are used at recovery sites must support the media of the copy exported physical volume.

Because hardware is no longer available or has reached end of service, the hardware at the recovery sites might not be able to read an exported volume. For example, a copy exported tape on JB media is not readable at a recovery site with only TS1150 or TS1160 drives. Tapes that are JB media can be read only by TS1140 tape drives that reached end of service in February 2022. Table 15-4 lists TS7700 supported tape drives and media.

Table 15-4 TS7700 Supported Tape Drives/Media

	TS1140 (E07) (end of service February 2022)	TS1150 (E08)	TS1160 (60F)
Write	<ul style="list-style-type: none"> ▶ 4 TB (JC-E07) ▶ 1.6 TB (JB-E07) ▶ 700 GB (JK-E07) 	<ul style="list-style-type: none"> ▶ 10 TB (JD-E08) ▶ 7 TB (JC-E08) ▶ 2 TB (JL-E08) ▶ 900 GB (JK-E08) 	<ul style="list-style-type: none"> ▶ 20 TB (JE-60F) ▶ 5 TB (JM-60F) ▶ 15 TB (JD-55F) ▶ 3 TB (JL-55F)
Append	1 TB (JB-E06)	<ul style="list-style-type: none"> ▶ 4 TB (JC-E07) ▶ 700 GB (JK-E07) 	<ul style="list-style-type: none"> ▶ 10 TB (JD-E08) ▶ 7 TB (JC-E08) ▶ 2 TB (JL-E08) ▶ 900 GB (JK-E08)

	TS1140 (E07) (end of service February 2022)	TS1150 (E08)	TS1160 (60F)
Read	<ul style="list-style-type: none"> ▶ JC ▶ JB ▶ JK 	<ul style="list-style-type: none"> ▶ JD ▶ JC ▶ JL ▶ JK 	<ul style="list-style-type: none"> ▶ JE ▶ JD ▶ JC ▶ JM ▶ JL ▶ JK

Physical tape has a limited life span and must be moved to a newer tape approximately every seven years (at a minimum) to prevent encountering unreadable tapes.

For more information about managing exported tapes, see “Checking the Status of Copy Exported Physical Volumes” in [IBM TS7700 Series Copy Export Function User’s Guide](#).

If a tape was exported over seven years before or is on older media, see “Reclamation of Copy Exported Physical Volumes” in [IBM TS7700 Series Copy Export Function User’s Guide](#) for instructions about how to reclaim the data on older media and create a secondary copy.

If the only copy of data is on secondary exported physical volumes, when these volumes age or require a media refresh, follow the process that is described in “Creating a New Primary/Secondary Copy after Disaster Recovery” in [IBM TS7700 Series Copy Export Function User’s Guide](#) to create secondary copies on newer media.

15.3.3 Performing Copy Export Recovery

Complete 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 were granted 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 15-5 shows the window that opens to inform you that the TS7700 contains data that must be erased.

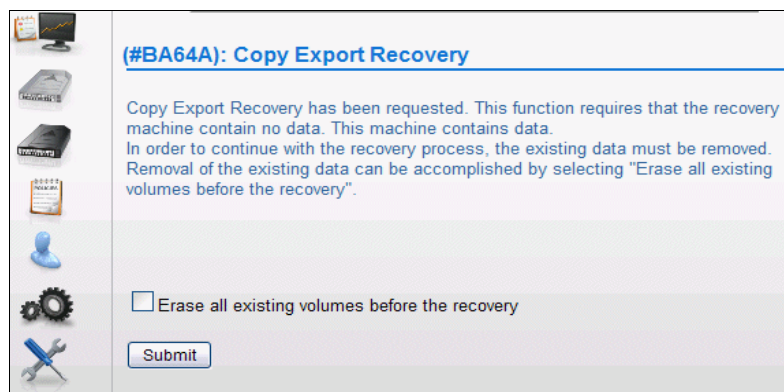


Figure 15-5 Copy Export Recovery window with erase volume option

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 approximately 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 must 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 that are 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 DR plans and have kept **Disaster Recovery Test Mode** 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 likely were 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: Consider the following points:

- ▶ Before continuing the recovery process, be sure that all of the copy-exported physical volumes were added. Any volumes that are not known to the TS7700 when the recovery process continues are not 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 completes 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 simplest 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 DR test, keep the **Disaster Recovery Test Mode** option 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 that are associated with the volume. If the volume was written to while it was mounted, the actions that are defined by the storage management constructs are taken.

If the volume was not modified, actions are taken only if a change occurred 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 last had the volume written to pool 4. Furthermore, the SG was not explicitly defined on the recovery TS7700 or it specified a different pool.

In this case, on the unloading of the volume, a new copy of it is written to the pool that is 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 option 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 requires scratch physical volumes and enables the copy-exported volumes to be reused after they are reclaimed. By keeping the **Disaster Recovery Test Mode** option selected, the reclaim operation is not performed.

With the **Disaster Recovery Test Mode** option 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 Copy Export Recovery is used because of a real disaster, clear the option.

Enter the volume serial number, select the option, 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 approximately 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 started, 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:

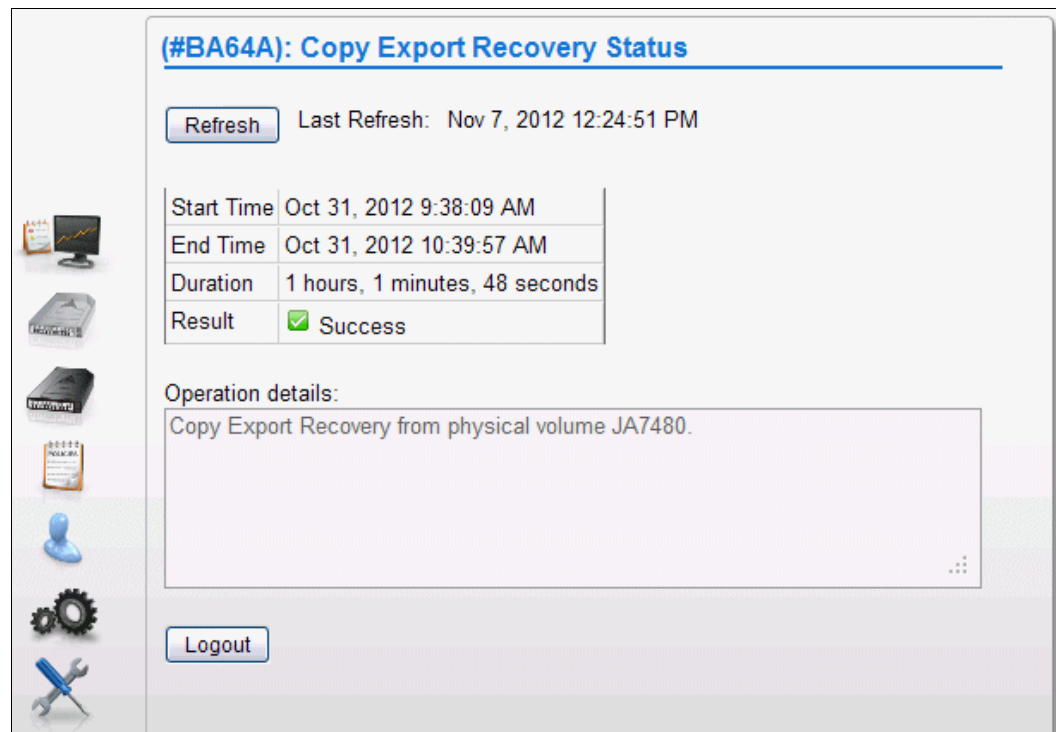
- a. The TS7700 is taken offline.
- b. The requested recovery tape XXXXXX is being mounted on device YYY.
- c. The database backup is being retrieved from the specified recovery tape XXXXXX.
- d. The requested recovery tape is being unmounted following the retrieval of the database backup.
- e. The database backup that is retrieved from tape is 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. 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 578.
12. If you ran Copy Export Recovery because of a real disaster (you cleared the **Disaster Recovery Test Mode** option), 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 (see Figure 15-6).



(#BA64A): Copy Export Recovery Status

Last Refresh: Nov 7, 2012 12:24:51 PM

Start Time	Oct 31, 2012 9:38:09 AM
End Time	Oct 31, 2012 10:39:57 AM
Duration	1 hours, 1 minutes, 48 seconds
Result	<input checked="" type="checkbox"/> Success

Operation details:
Copy Export Recovery from physical volume JA7480.

Figure 15-6 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 TS7700 Series Copy Export Function User's Guide](#) white paper.

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 213.

15.3.4 Restoring the host and library environments

Before you can use the recovered logical volumes, you must also 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 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.

15.4 Using Copy Exported tape for damaged volume recovery

In the case where secondary copies are copy-exported, returning a Copy Exported physical volume to the physical tape library can be used to recover logical volumes. Unlike Copy Export Recovery, this method does not need a stand-alone TS7700 with an empty cache.

Typically, this method is used when the primary physical volume is damaged and the primary copies on it cannot be recovered, even by Read Only Recovery (ROR) process in the stand-alone TS7700 environment. In such a case, the logical volume copies *cannot* be re-created in the TS7700 unless the corresponding Copy Exported physical volume is returned to the physical tape library to allow the TS7700 to re-create the logical volumes by using the secondary logical volume copies that are on the Copy Exported physical volume.

Note: More than one Copy Exported physical volume might be required to recover all the logical volumes that were stacked to the damaged primary physical volume. After all of the logical volumes are recovered to the cache, the logical volumes are then premigrated to new physical volumes.

It is suggested that a Copy Export is then run so that new Copy Exported physical volumes can be taken offsite for future recovery actions.



Disaster recovery testing in a grid configuration

This chapter describes disaster recovery (DR) testing in a TS7700 grid configuration and includes the following topics:

- ▶ 16.1, “DR testing overview” on page 820
- ▶ 16.2, “DR testing methods” on page 820
- ▶ 16.3, “DR testing general considerations” on page 824
- ▶ 16.4, “DR for FlashCopy concepts and command examples” on page 831
- ▶ 16.5, “DR testing methods examples” on page 842
- ▶ 16.6, “Expected failures during a DR test” on page 852

16.1 DR testing overview

In a perfect world, DR testing is not needed. However, in reality, many factors exist that can lead to a disaster that prevents the use 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.

Often, recovering from a disaster is simple and requires fewer steps than having to simulate a disaster and then clean up your disaster environment as though the simulation never occurred. Although general DR concepts are described in Chapter 5, “Disaster recovery” on page 213, this chapter focuses on concepts that are related to DR testing specifically. Examples are provided where needed and include step-by-step walkthroughs for the following methods that clients can use to accomplish DR testing in a TS7700 grid environment:

- ▶ By using FlashCopy
- ▶ By using Write Protect Mode on DR clusters
- ▶ Without the use of Write Protect Mode on DR clusters
- ▶ By breaking the grid links to DR clusters

All of these methods have advantages and disadvantages and you need to evaluate them against your environment and resources. You can then choose which method best fits your DR testing needs and ability.

The description of each method assumes that you are familiar with the DR concepts that are presented in Chapter 5, “Disaster recovery” on page 213. The end of this chapter includes instructions for how to perform a DR test by using each method. Although it might be tempting to jump right to these instructions, it is recommended that you review this chapter in its entirety to ensure that you are familiar with the concepts and options that are available in a TS7700 grid environment before you start the testing process.

16.2 DR testing methods

This section describes four methods that can be used to test DR in a TS7700 grid environment. Although the Copy Export and, new in R5.1, Cloud Export functions also can be used to test DR for a TS7700T or TS7700C (by using the physical copy-exported tapes or cloud exported logical tapes to rebuild a production environment on an empty TS7700T), this chapter focuses on DR testing that uses a grid environment that consists of production clusters and DR clusters.

For more information about Copy Export, see Chapter 15, “Copy Export” on page 783.

For more information about Cloud Export, see Chapter 17, “Cloud Storage Tier export, recovery, and testing”, in *IBM TS7700 R5.3 Cloud Storage Tier Guide*, REDP-5573.

16.2.1 Method 1: DR Testing by using FlashCopy

This method of DR testing uses the FlashCopy function that was introduced in R3.1. This function enables a DR host to perform testing against a point-in-time consistency snapshot while production operations and replication continue. Production data continues to replicate during the entire DR test and the same virtual volume can be mounted at the same time by a DR host and a production host.

With FlashCopy and the implicit 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 virtual volumes. Any DR host continues to have read access to any production volumes that were returned to scratch while the FlashCopy is active.

During a DR test, production volumes might need to be mounted from the DR and production hosts. Without FlashCopy for DR testing, volume mounts are serialized so that only one host can access a volume at any one time and the other hosts receive an IN USE exception. This exception can be a problem when the production host is attempting the mount and the mount fails.

FlashCopy enables virtual 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 features the following advantages and disadvantages:

Note: Although a TS7720/TS7740 can exist as part of a grid that uses FlashCopy functions and retain the same functions as they did in pre-R5.0 releases, references to them were removed from this section. These references are not included here because both were removed from service on April 30, 2020. If your grid contains a TS7720 or TS7740, see *IBM TS7700 Release 4.2 Guide*, SG24-83666.

► Advantages:

- After the FlashCopy is enabled, all read activity against volumes that are included in the FlashCopy (the ones in write-protected categories on one or more DR clusters) are from that point-in-time copy. This scenario closely simulates a real disaster scenario in which one or more production clusters are no longer accessible and the disaster clusters can access the production data from a point-in-time. Volumes that belong to categories that are excluded from write-protection on the DR clusters can continue to include data that is written to them during the DR test.
- Data that is written from a production host to the production clusters can continue to be copied to the disaster clusters without the risk of a disaster host accessing the live data. While the FlashCopy is active, the disaster host can access only the point-in-time copy of the production data that is present on the DR clusters at the time of the FlashCopy. With R4.1.1, the live copy of production volumes can be accessed by the DR host selectively by using the new **LIVEACC** option on the **LI REQ** command. For more information, see 16.4.5, “LIVEACC option” on page 835.
- Write Protect Mode can be enabled permanently for the clusters in the DR Family (not just during the duration of the DR test), which can provide continuous protection from the overwriting/return-to-scratch processing on the DR host of volumes that are assigned to categories that are not in the Excluded from Write Protect list.

- ▶ Disadvantages:
 - The FlashCopy on the DR clusters helps ensure that if a virtual 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 issue leads to a higher cache use on the DR clusters.
 - The Write Protect Mode and Write Protect Exclude categories must be configured correctly for any data to be written to the DR clusters during a DR test. If they are configured incorrectly (by defining production categories as being excluded from write-protect), production data might be overwritten during the DR test.
 - When a FlashCopy for DR is active anywhere within a grid, eject processing is not allowed (failed with an ERA29 mod x'10' or function incompatible). This situation includes ejects issued to clusters inside the DR family and outside of the DR family and is independent of write-protect mode exclusion categories. When FlashCopy for DR is disabled, ejects are allowed if no other eject denial state, such as write protect against the targeted volume's category, is active.

16.2.2 Method 2: DR Testing by using Write Protect Mode on DR clusters

This method uses the Write Protect Mode function in TS7700 clusters 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 are 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), are write-protected on the DR clusters.

This method features the following advantages and disadvantages:

- ▶ Advantages:
 - By enabling Write Protect Mode on the DR clusters, the cluster prevents the write at the hardware level, even if a job on the DR host tries to mount a production volume for write on a DR cluster.
 - Production data can still be written to the production clusters and those clusters can still copy data to the DR clusters so that if a real disaster occurs, the data on the DR clusters is more up to date than if the copying did not occur.
 - Write Protect Mode can be enabled permanently on the DR cluster (not just during the duration of the DR test), which can provide continuous protection from the overwriting/return-to-scratch processing on the DR host of volumes that are assigned to categories that are not in the Excluded from Write Protect list.
- ▶ Disadvantages:
 - The Write Protect Mode and Write Protect Exclude categories must be configured correctly for any data 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.
 - No point-in-time simulation occurs. The data on the volumes that are 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 DR clusters. Jobs that are 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, the use of this method is the recommended alternative.

16.2.3 Method 3: DR testing without the use of Write Protect Mode on DR clusters

This method is similar to the previous method. However, instead of the use of the Write Protect Mode function in the DR clusters to prevent any writes that are issued from the DR host to a production volume on the DR clusters, this method relies on the ability (and correct configuration) of the tape management system (TMS) on a DR host to prevent volumes in the production volume range from being written to by the DR host.

This method features the following advantages and disadvantages:

- ▶ Advantage: Production data can still be written to the production clusters. Those clusters can still copy data to the DR clusters so that if a real disaster occurs, the data on the disaster clusters is more up to date than if the copying did not occur.
- ▶ Disadvantages:
 - No hardware-enabled write protection that prevents a DR host from writing to a production volume on a DR cluster. The TMS on the DR host *must* be configured to prevent any writes that are directed toward production volumes.
 - No point-in-time simulation occurs. The data on the volumes that are 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 DR clusters. Jobs that are 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. For more information, see “Return-to-scratch Considerations” on page 827.

If your choice is between using Write Protect Mode and not that uses Write Protect Mode, it is suggested to use Write Protect Mode (Method 2). Write Protect Mode provides another level of write-protection if the TMS on the DR host is not configured correctly to prevent writes and return-to-scratch processing from the DR host of production volumes on a DR cluster.

16.2.4 Method 4: Breaking the grid links

The fourth method is to simulate a real disaster by breaking the grid links between the production clusters and DR clusters in a TS7700 grid.

This method features the following advantages and disadvantages:

- ▶ Advantages:
 - After the grid links are broken, you are assured that any production data that is accessed from a DR cluster by the DR host is data that was copied to the DR cluster before the grid links were broken.
 - Return-to-scratch processing that is started by a production host against production volumes on production clusters does not affect the copy of the production 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 DR clusters must be allowed to takeover read-only volume ownership from the production clusters. DO NOT enable write-ownership takeover. The takeover write ownership function is only used if a real disaster occurs.
- After the grid links are broken, cartridge entry cannot occur in the grid. In addition, any scratch volumes that are owned by the DR cluster cannot be used by the production host by way of the production cluster while the grid links are broken. Only when the grid links are reestablished can these functions process normally.
- Breaking the grid links must be done by your CE (SSR). Do not disable a grid link by using only the **Library Request** command to run this method. Disabling the grid link by using 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 most significant drawback to the use of this DR method. Because of this concern, it is advised that one of the other methods that were described in this section (FlashCopy or Write Protect Mode) are used if it is possible.

Note: Do not use virtual drives in the DR clusters from the production host during the DR test. Doing so allows the production host to allocate drives on the DR cluster to read production volumes (with ROT enabled on the DR cluster). Allowing the production host to allocate drives on the DR cluster to read production volumes can lead to data integrity issues because the production host might write new data to a volume on the production cluster that is not seen by the DR cluster.

If you decide to break links during your DR test, you must carefully review 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. For more information about throttling, see 14.3.7, “Throttling in the TS7700” on page 758.

After the test ends, many virtual volumes might exist 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 restored the grid links, you can enable the copy again during a low activity workload to avoid or minimize performance degradation. For more information, see 12.1.3, “Host Console Request function” on page 632.

16.3 DR testing general considerations

Before running any of the DR test methods that described later in this chapter, it is important to understand the concepts that are involved in DR testing that are applicable to all of the methodologies. This section describes those concepts in detail, and the risks that are involved during DR testing.

16.3.1 Setup and restore of the DR Host tape environment

This section describes the setup and restore of the DR Host tape environment.

The following control data sets must be obtained from a point-in-time copy from a production host and restored and activated on the DR host:

- ▶ Tape management system (TMS) CDS
- ▶ DFSMS data catalogs, including the tape configuration database (TCDB)
- ▶ Input/output definition file (IODF)
- ▶ SMS source control data set (SCDS)
- ▶ Security Facilities data set (RACF ACF2)

Over the duration of the DR test, the production host continues to function as normal and can change these control data sets. Those changes are not reflected in the point-in-time snapshot in use on the DR hosts.

Although it is possible for a production host and DR host to share live versions of these control data sets, this chapter focuses on DR testing from the perspective of a point-in-time copy of these data sets.

Defining Scratch Categories for the DR Host

Before any volumes are inserted into the DR cluster to be accepted by the DR host during host cartridge entry processing, you must first decide which scratch, error, and private categories are used for these DR volumes.

The categories to be used on a host are defined in the DEVSUPxx parmlib member. For DR test usage with the TS7700, be sure to assign at minimum a unique category for MEDIA1, MEDIA2, ERROR, and PRIVATE. These categories must be unique from *any* other host's categories that use the grid to ensure that a unique volume pool is used by the DR host.

16.3.2 Protecting Production Data

When a DR test is run, be sure that safeguards are in place so that production volumes cannot be written to or scratched by the DR host. This safeguard can be put in place by using one of the following methods:

- ▶ Hardware Protection (Write Protect Mode on a TS7700)
- ▶ Host Software Protection (provided by the TMS)

This section describes the safeguards that are available on the hardware and software sides to accomplish this task, the DR testing methods that use them, and return-to-scratch considerations that apply to the DR host and production host.

Production Data creation during a DR test

While a DR test is underway in each method described *except* Method 4 (Breaking the grid links), the production host can continue writing data to production volumes on the production cluster. The data can then be copied to the DR cluster by using the copy policy settings on the production cluster.

The production host can also run return-to-scratch processing of production volumes. In either case, the control data sets on the DR host have no knowledge of these changes. For this reason, your DR test must account for either of these possibilities when Method 2 (Using Write Protect Mode on DR clusters) or Method 3 (DR testing without Write Protect Mode) are used.

Write Protect Mode

The purpose of Write Protect Mode is to prevent (at the cluster level) volumes being overwritten or returned to scratch. Only those volumes having categories that are explicitly excluded from write protection, that is, the ones in the Exclude from Write Protect list in the TS7700 MI (**Settings** → **Cluster Settings** → **Write Protect Mode**) can be written to or returned to scratch through a device in that cluster.

With R5.0, up to 128 categories can be included in the Excluded from Write Protect list. If any host attempts to write to a volume that belongs to a category *not* in this list, the TS7700 fails the write attempt and an error message is surfaced at the host.

For the methods that are described in this chapter, only Method 1 (DR Testing by using FlashCopy) and Method 2 (Using Write Protect Mode on DR clusters) use the Write Protect Mode function. However, only Method 2 enables and disables Write Protect Mode by way of the TS7700 MI. When Write Protect Mode is enabled or disabled by using Method 1, the Write Protect Mode *must* be enabled/disabled by using the **LI REQ DRSETUP** interface.

When Write Protect Mode is started on the DR cluster in Method 2, the production host is *not* prevented from continuing to write or return-to-scratch the production volumes on the production cluster.

The main effect of this is that if a production host returns to scratch a production volume during a DR test when using Method 2 (Using Write Protect Mode on DR clusters) or Method 3 (DR Testing without Write Protect Mode), the DR host has no knowledge that the volume is scratch. If the DR host attempts to read the volume that is now in a scratch category (because the DR host believes that it is still PRIVATE and includes data), the read fails unless all of the following conditions are true:

- ▶ Write Protect Mode is enabled on the DR cluster.
- ▶ Production categories are *not* in the Excluded from write protect list.
- ▶ The “Ignore fast ready characteristics of write protected categories” setting is enabled.
- ▶ TS7700 did not yet delete the data from the production volume.

If these conditions are true, the DR host can still read the data from the production volume.

TMS Write Protection

When hardware write protection is unavailable (by using Method 3: DR Testing without Write Protect Mode and Method 4: Breaking the grid link connections), the TMS can be used on a DR host to prevent the DR host from writing to volumes that do not belong to the DR**** range.

In RMM, the following statements can be specified in the EDGRMMxx parmlib member on the DR host to provide this protection:

- ▶ **OPENRULE VOLUME(*) TYPE(RMM) OUTPUT(REJECT) INPUT(ACCEPT)**
This rule allows for any volumes that are defined to RMM to be read, but not written to.
- ▶ **OPENRULE VOLUME(DR*) TYPE(RMM) OUTPUT(ACCEPT) INPUT(ACCEPT)**
This rule allows for volumes with DR* mask that is defined to RMM to be read and written to.

If your DR host does not use RMM, consult the vendor of your TMS to determine how that TMS can prevent DR hosts from writing to non-DR volumes.

Return-to-scratch Considerations

When Method 2 (Using Write Protect Mode) or Method 3 (DR Testing without Write Protect Mode) are used, it is recommended that enough scratch volumes are made available on the production host before the DR test so that return-to-scratch processing on the production host is not needed. Return-to-scratch processing on the DR host is run at the end of the DR test by processing *only* DR volumes and *not* production volumes.

If return-to-scratch processing must be run on the production host or DR host during the DR test, one of the following scenarios can occur:

- ▶ The DR host returns to scratch production volumes
- ▶ The production host returns to scratch production volumes that were needed by the DR host during the DR test

To prevent the first scenario, regardless of the DR test method that is chosen, including Method 2 (Using Write Protect Mode on DR clusters), it is recommended to define to your TMS statements that prevent the TMS on the DR host from being able to return-to-scratch production volumes. For RMM, the statements are defined:

```
PRITITION VOLUME(DR*) TYPE(NORMM) SMT(ACCEPT) NOSMT(IGNORE)
PRITITION VOLUME(DR*) TYPE(RMM) SMT(ACCEPT) NOSMT(IGNORE)
PRITITION VOLUME(*) TYPE(NORMM) SMT(IGNORE) NOSMT(IGNORE)
PRITITION VOLUME(*) TYPE(RMM) SMT(IGNORE) NOSMT(IGNORE)
```

These statements are also used by RMM to control which volume ranges are approved or rejected during host cartridge entry processing. For more information about this process, see 16.3.3, “Cartridge entry considerations” on page 827.

If your DR host does not use RMM, consult the vendor of your TMS to determine how that TMS can prevent DR hosts from being able to return-to-scratch production volumes.

If you suspect the second scenario might occur, set an expiration time for the production scratch categories (up to 2000 years with R5.0) that is at least equal to the duration of the DR test and enable Expire Hold to ensure those production volumes returned-to-scratch are not deleted by the production TS7700 before the end of the DR test. Be sure to restore the expiration settings for the production categories after the DR test is complete.

16.3.3 Cartridge entry considerations

After a DR host is IPL'd, one of the first steps that is taken during a DR test (regardless of the DR method chosen) is to determine the DR volume range that you want to use during the DR test and insert those volumes by using the TS7700 MI.

However, if the hosts (DR *and* production) that are connected to your TS7700 grid do not have their host cartridge entry environments configured correctly, it is possible that the DR volumes that are inserted are never processed by the DR host they are intended to be processed by, or they can be processed by the DR host but ignored or rejected because of an incorrectly setup host cartridge entry environment.

For more information about the activity that occurs during host cartridge entry processing, see 12.5, “Host cartridge entry processing” on page 650.

Whether a volume is accepted by a host during host cartridge entry processing is controlled by the CBRUXENT user exit. This exit is often supplied by the vendor (or IBM) that supplies the TMS that exists on a host, but it can also be user-written. If the exit is supplied by a TMS vendor, that exit usually calls the TMS to approve or deny the entry request of a volume.

If the DR host uses RMM as the TMS, the PRTITION statements that are specified in the EDGRMMxx parmlib member control the approval of cartridge entry and limits return-to-scratch processing for those specified volume ranges. On the DR host, specify the following statements to approve cartridge entry and allow return-to-scratch processing of DR**** volumes, as well deny cartridge entry and prevent return-to-scratch processing of non-DR**** volumes:

```
PRTITION VOLUME(DR*) TYPE(ALL) SMT(ACCEPT) NOSMT(IGNORE)
PRTITION VOLUME(*) TYPE(ALL) SMT(IGNORE) NOSMT(IGNORE)
```

If production hosts that are attached to the grid use RMM, the following statement should be specified on those hosts to deny cartridge entry and prevent return-to-scratch processing of DR**** volumes:

```
PRTITION VOLUME(DR*) TYPE(ALL) SMT(IGNORE) NOSMT(IGNORE)
```

If your TMS is not RMM, contact your vendor for the instructions about how to set up your TMS to approve/reject a range of volumes from a host during cartridge entry processing.

As an alternative to modifying the TMS definitions to approve or reject specific ranges of volumes during host cartridge entry processing, complete the following steps:

1. Disable host cartridge entry processing on all hosts that are connected to the grid *except* the DR host by issuing the **LI DISABLE, CBRUXENT** command on those hosts.
2. Insert the DR volumes by using the TS7700 MI and wait for the CBR3610I messages to surface on the DR host
3. Re-enable host cartridge entry processing on those hosts that it was disabled on by issuing the **LI RESET, CBRUXENT** command on those hosts.

16.3.4 Ownership takeover

If you perform the DR test with links that are broken between sites, you must enable Read Ownership Takeover (ROT) so that the DR site can have read access to the data on the production volumes that are owned by the production site. Volume ownership is obtained during the insert or mount process in a healthy grid (links not broken). Mounting production volumes on a production cluster assigns ownership to that cluster.

Important: While ROT is needed during a DR test with links that are broken between sites, never enable Write Ownership Takeover (WOT) mode for any DR test. WOT mode must be enabled only during a loss or failure of the production TS7700.

After the links are broken, if you attempt to mount one of those production volumes from the DR host without ROT ownership takeover enabled, the mount fails because the DR cluster cannot request ownership transfer from the production cluster. By enabling ROT, the DR host can mount the production virtual volumes and read their contents.

When ROT is enabled on the DR cluster, the DR host cannot 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 TS7700 MI that is associated with the DR cluster, that DR cluster owns those volumes. The DR host has complete read and write control of these DR volumes.

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.

16.3.5 DR Volume Copy policies

If the DR volumes in the DR cluster use the same management classes as are used on the production cluster for each method except Method 4 (Breaking the grid links), the data that was created on the DR volumes during the DR test will be copied to the production cluster. Allowing the DR volumes created during the DR test to be copied to the production cluster is a waste of space on the production cluster and on inter-site bandwidth.

It is recommended to use a copy mode of No Copy for the data that is created on the DR volumes during the DR test. If a management class is changed on the DR cluster for the DR test, be sure to change it back to the original settings when the DR test is complete.

16.3.6 Clean up phase of a DR test

When a DR test is complete, clean up the DR environment so that it is in the same condition as before you started the DR test. During this process, 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 by using ISMF) after the DR test, this unneeded data continues to occupy cache, physical tape space (TS7700T), or cloud space (TS7700C). Because the volumes that contain this data remain in a PRIVATE category, they never expire and continue to occupy space indefinitely.

For this reason, be sure to return-to-scratch those DR volumes that are PRIVATE and at the least (if you do not want to delete the volumes), ensure that the scratch category that they are assigned to includes an expiration time that is specified in the TS7700 MI. Otherwise, space on the TS7700 continues to be wasted because these virtual volumes are not overwritten.

Remember, run RMM EXPROC by using SYSIN of VOLUME or VOLUMES parameter to limit the range of volumes that is returned to scratch to only the ones that are written during the DR test. If your TMS vendor is *not* IBM, contact your vendor to determine how that TMS must be configured to limit the range of volumes that are processed during return-to-scratch processing.

Returning ownership of production volumes to the production cluster

During a DR test, production volumes can be read by the DR host through a device on the DR cluster. However, before this read can occur, the DR cluster must first obtain ownership of the volume from the production cluster. Likewise, before the production host can then read/write to or from the volume from the production cluster, the production cluster must first obtain ownership of the volume back from the DR cluster.

When the DR test (excluding Method 1: DR with FlashCopy) is complete, the DR cluster can be left in a state where it remains the owner of production volumes. This state is normally not an issue because reads from the production host on the production cluster transfer ownership of the volumes back to the production cluster over time.

However, if those volumes are expired by TMS on the production host before that transfer occurs, the production volumes end up as scratch volumes that are owned by the DR cluster. For this reason, it is important to transfer ownership of these volumes back to the production cluster during the cleanup phase of a DR test. The following command is used to transfer the ownership of up to all volumes from a source cluster to the cluster to which the command is issued:

```
LI REQ,distlibname,OTCNTL,START,id
```

In R5.1, the **OTCNTL** command was updated to provide an option to move ownership of a specific volume to another cluster. For more information about this command, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide Version 5.4](#).

16.3.7 Tier to Cloud considerations

With the introduction of the cloud storage tier, your production cluster and DR cluster can be TS7700Cs where both clusters are connected to the same cloud object store. However, for clusters that are earlier than R5.0, even if both are attached to the same cloud object store and the data for a production volume is in the cloud object store, the DR host (by way of the DR cluster) can only access that data if the copy policy on the production cluster sends a copy to the DR cluster and that copy-completed replication to the DR cluster.

However, when all the clusters in the grid are at R5.0 and later, the new **CLDGHOST** keyword can be specified on the **COPYRFSH** command. This specification allows the DR cluster that is attached to the same cloud object store as the production cluster to access the specified production virtual volume on that cloud object store without copying the data through the grid links.

The DR cluster that is attached to the same cloud object store as the production cluster does not allow it to access the data that is stored in the cloud object store by the production cluster unless the **CLDGHOST** keyword is used with the **COPYRFSH** command on a per volume basis. The best way to think about a cloud object store (as it relates to functions from a DR perspective with R5.0 or prior clusters) is that it is synonymous with a TS7700T.

R5.1 introduces the concept of Grid Awareness. With Grid Awareness, after one cluster puts a volume into the cloud object store, all clusters can see that virtual volume in the cloud object store (if they are configured to access that same cloud object store). This function allows users to think about TS7700s use of a cloud object store in its most ideal form: a location where virtual volumes can be accessed by any TS7700C that is configured to access that cloud object store.

Note: Grid Awareness is active only if all clusters in the Grid are at R5.1 or later.

16.4 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 contents of a virtual volume from time zero, while 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 links to ensure that only data from time zero is available to the DR host.

For more information, see [TS7700 Series Best Practices - FlashCopy for DR Testing](#).

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 live instance of a volume that is in a cluster 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 Exclude from Write Protect list.
- ▶ **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 be down level to the grid, or incomplete at time zero. An active FlashCopy indicates that Write Protect Mode is active.
- ▶ **DR Family:** A set of TS7700 clusters (most likely those clusters at the DR site) that serve the purpose of DR. A total of 1–7 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 TS7770.
- ▶ **Write Protect Mode (existing function):** 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 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 occurs. Customers can establish the time zero by using a host console request command.

16.4.1 Basic requirements and concepts

The FlashCopy for DR testing function is supported on TS7700 Grid configurations where at least one TS7760 or TS7770 cluster exists within the DR location.

Only volumes with categories that are *not* included in the Exclude from Write Protect list are included in the FlashCopy.

During an enabled Flash, the autoremoval process is disabled for the members of the DR Family. A TS7700 within a DR Family 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 extra TS7700 disk cache capacity. The extra capacity requirement must be considered when planning the size of the TS7700 disk cache.

If you are using Time Delay Replication Policy, also check the cache usage of the remaining production cluster TS7700. Volumes can be removed from the TS7700 only when the T copies are processed (in the complete grid or family).

16.4.2 FlashCopy and Write Protect enablement/disablement enhancement R4.1.2 and R4.2

Before R4.1.2 and R4.2, the **LI REQ, DRSETUP** commands that were used to enable or disable FlashCopy and Write Protect (**WP/FLASH/DOALL, ENABLE/DISABLE**) ran synchronously. However, these commands at times can time out and report that the operation failed. This issue typically occurs when the snapshots of many virtual volumes must be deleted while disabling the FlashCopy. To avoid the issue, these commands were modified to be run asynchronously. The asynchronous operation is supported after all of the clusters in the grid are at R4.1.2 (8.41.200.113)/R4.2 or later. With this support, more status messages are appended to the output of a few **DRSETUP** commands, as described in the following sections.

16.4.3 DR Family

In R5.3, 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 TS7770 must be part of the DR Family. The Write Protect excluded media categories need 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 can be added to a previously created DR Family, by using the following command (see Example 16-1):

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,ADD,<CLUSTER ID>
```

Example 16-1 Creating a DR Family and adding a cluster

```
-LI REQ,HYDRAG,DRSETUP,DRFAM01,ADD,1
CBR1020I Processing LIBRARY command: REQ,HYDRAG,DRSETUP,DRFAM01,ADD,1.
CBR1280I Library HYDRAG request. 939
Keywords: DRSETUP,DRFAM01,ADD,1
-----
DRSETUP V3 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 any **DRSETUP** command is used for a particular DR Family, it is good to check the 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 (see Example 16-2 on page 833):

```
LI REQ,<COMPOSITE>,DRSETUP,SHOW,<FAMILYNAME>
```

Example 16-2 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 V3 .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  SELFLIVE
CLUSTER1  DISABLED      3           TRUE       N/A         N
CLUSTER2  DISABLED      3           TRUE       N/A         N
-----
CATEGORIES EXCLUDED FROM WRITE PROTECTION WITHIN DR FAMILY  DRFAM01
CLUSTER  ACTIVE EXCLUDED CATEGORIES
CLUSTER1 0092 009F 3002
CLUSTER1 0092 009F 3002
```

If this command is received by a cluster in the grid that is running R4.1.2 or later, more status information is appended to the output that displays the active asynchronous **DRSETUP** command status under the following header:

DRSETUP CURRENT ACTIVE OPERATION STATUS

The following status lines can be displayed:

- ▶ **NO ACTIVE DRSETUP FLC/WP OP IS RUNNING**
This status surfaces when no active asynchronous operation is running.
- ▶ **ASYNC DRSETUP OP IS NOT SUPPORTED**
This status surfaces when the grid is in a mixed code configuration with R4.1.2/R4.2 and 8.4.1 or earlier code levels.
- ▶ **[Asynchronous operation] STARTED AT [Started time] IS RUNNING**
This status surfaces when the operation is started and the start time, where the asynchronous operation field is set to one of the following asynchronous operations:
 - WP ENABLE
 - WP DISABLE
 - FLASH ENABLE
 - FLASH DISABLE
 - DOALL ENABLE
 - DOALL DISABLE
- ▶ **ACTIVE DRSETUP OP CANNOT BE DETERMINED**
This status results if an unexpected error occurs internally and the current active asynchronous operation status cannot be checked. If this error occurs, contact TS7700 support.

If this command is received by a cluster in the grid that is running R5.0 or later and too many write protect exclusion categories exist to be shown within 50 lines of whole output, the following message is shown at the end of the write protect exclusion categories list:

MORE WRITE PROTECT EXCLUDED CATEGORIES TO DISPLAY EXIST

16.4.4 LIVECOPY enablement in a DR Family

A DR Family must contain at least one TS7700. If a TS7700T or TS7700C is present within a DR Family, an option is available allowing access to the “live” copy of a volume on the TS7700T or TS7700C cluster. This option is available if the TS7700 removed its copy or if the TS7700T or TS7700C was the only target of the volume. This option is called **LIVECOPY**.

LIVECOPY allows 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 TS7700 but do exist on a physical backend tape or object cloud store that is attached to a TS7700. 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.

The option is disabled by default. If you choose to enable this function, you must specifically enable the option by using the **Library Request** command with **LIVECOPY** keyword: (see Example 16-3):

```
LI REQ,<clib_name>,DRSETUP,<family_name>,LIVECOPY,FAMILY
```

Example 16-3 Enable the LIVECOPY option

```
LI REQ,HYDRAG,DRSETUP,DRFAM01,LIVECOPY,FAMILY
CBR1020IProcessingLIBRARYcommand:REQ,HYDRAG,DRSETUP,DRFAM01,LIVECOPY,FAMILY.
CBR1280I Library HYDRAG request. 154
Keywords: DRSETUP,DRFAM01,LIVECOPY,FAMILY
```

```
-----
DRSETUP V3 0.0
LIVE COPY USAGE HAS BEEN UPDATED TO FAMILY SUCCESSFULLY
```

To disable the **LIVECOPY** option, you must run the following command (see Example 16-4):

```
LI REQ,<clib_name>,DRSETUP,<family_name>,LIVECOPY,NONE
```

Example 16-4 Disable the LIVECOPY option

```
LI REQ,HYDRAG,DRSETUP,DRFAM01,LIVECOPY,NONE
CBR1020IProcessingLIBRARYcommand:REQ,HYDRAG,DRSETUP,DRFAM01,LIVECOPY,NONE.
CBR1280I Library HYDRAG request. 154
Keywords: DRSETUP,DRFAM01,LIVECOPY,NONE
```

```
-----
DRSETUP V3 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 by using the **DRSETUP, SHOW** command. The output of this command contains a column titled **LCOPY**. If the value under **LCOPY** is **FAMILY**, this value indicates that **LIVECOPY** is active for the DR Family. If the value under **LCOPY** is **NONE**, **LIVECOPY** is not enabled for the DR Family. **DRSETUP, SHOW** output that shows a DR Family where **LIVECOPY** is enabled is shown in Example 16-2 on page 833.

16.4.5 LIVEACC option

By default, after a FlashCopy is enabled for a DR Family, if the DR host attempts to read from a volume belonging to one of the categories that are *not* in the Excluded from write protect list, the read accesses the data that existed for that volume at the time the FlashCopy was created (time zero).

If the data for this volume was altered by the production host after time zero, you might want to allow the DR host to access the Live Copy data of the volume. The LIVEACC option on the **LI REQ** command provides this function by allowing the operator to change the category for a write-protected volume that is included in the FlashCopy to a category in the Excluded from Write Protect Mode list. The command features the following syntax:

```
LI REQ,<dlib_name>,DRSETUP,LIVEACC,<volser>,<category>
```

The command output is shown in Example 16-5.

Example 16-5 LIVEACC example

```
LI REQ,HYDRAG1,DRSETUP,LIVEACC,Z00000,001F
CBR1020I Processing LIBRARY command: REQ,HYDRAG1,DRSETUP,LIVEACC,Z00000,001F
CBR1280I Library HYDRAG request.
Keywords: DRSETUP,DRFAM001,Z00000,001F
-----
DRSETUP V1.1
Z00000 WAS SUCCESSFULLY ASSIGNED TO CATEGORY 001F
```

Eventually, when return-to-scratch processing occurs within the production host tape management system, the volume is moved back into the correct scratch category assigned to the owning system plex. There is no need to return the volume to the original private category when the test completes.

16.4.6 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 and Write Protect Mode individually, you must first disable the FlashCopy and later the Write Protect Mode. Alternatively, both can be disabled by using a single command.

Note: A FlashCopy cannot be enabled if Write Protect Mode was enabled by using the MI. If Write Protect Mode was enabled by using the MI, it must first be disabled by using the MI before a FlashCopy can be enabled by using **LI REQ**.

Do not enable FlashCopy if production hosts with tape processing include device allocations on the clusters where the Flash will be enabled. Failures might occur because the read-only mode does not enable subsequent mounts.

16.4.7 Starting FlashCopy and Write Protect Mode for a DR Family

After a DR Family is created, start Write Protect Mode and the FlashCopy simultaneously by using the command that is shown in Example 16-6.

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,DOALL,ENABLE
```

Example 16-6 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 V3 0.0
WRITE PROTECT STATUS HAS BEEN ENABLED SUCCESSFULLY
FlashCopy HAS BEEN CREATED SUCCESSFULLY
```

If the request is received by a R4.1.2/4.2 cluster and the entire Grid is at R4.1.2/R4.2 or later, it is handled asynchronously and the following output is surfaced instead:

```
DRSETUP V3 0.0
DRSETUP DRFAM DOALL ENABLE HAS STARTED SUCCESSFULLY
```

When the command completes successfully or unsuccessfully, a message is sent to the host and surfaced in a CBR3750I message. For more information about these messages, see [TS7700 Series Best Practices - FlashCopy for DR Testing](#).

16.4.8 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 command that is shown in Example 16-7.

Example 16-7 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 V3 0.0
WRITE PROTECT STATUS HAS BEEN DISABLED SUCCESSFULLY
FlashCopy HAS BEEN DELETED SUCCESSFULLY
```

If the request is received by a R4.1.2/R4.2 cluster and the entire Grid is at R4.1.2/4.2 or later, it is handled asynchronously and the following output is surfaced instead:

```
DRSETUP V3 0.0
DRSETUP DRFAM DOALL DISABLE HAS STARTED SUCCESSFULLY
```


When the command completes successfully or unsuccessfully, a message is sent to the host and surfaced in a CBR3750I message. For more information about messages, see [TS7700 Series Best Practices - FlashCopy for DR Testing](#).

16.4.9 Commands to check volume status during a DR test

During a DR test, you might want to check the status of the following 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 commands that are described in the following examples to identify if a FlashCopy exists for a specific volume, and the status from the Live copy and the FlashCopy.

To do so, use the **LI REQ,composite_library,LVOL,volser,INFO** and the **LI REQ,composite_library,LVOL,volser,INFO,FLASH** commands. If the live copy volume is identical to the FlashCopy volume, the status is ACTIVE. Only if the virtual volume was updated from production, and a second instance exists, does the status change to CREATED (see Example 16-8).

Example 16-8 Display of a logical volume after modification from production - Live copy

```

LI REQ,HYDRAG,LVOL,Z10000,INFO
CBR1020I Processing LIBRARY command: REQ,HYDRAG,LVOL,Z10000,INFO.
CBR1280I Library HYDRAG request. 883
Keywords: LVOL,Z10000,INFO
-----
LOGICAL VOLUME INFO V3 0.0
LOGICAL VOLUME           : Z10000
MEDIA, FMT, MAX(MB), CWRAP : ECST, 6, 800, N
SIZE(MB) COMP, CHAN, RATIO : 14, 95, 6.38:1(FICON)
CURRENT OWNER, TVC LIB    : Arabian, Lipizzan
MOUNTED LIB/DV, MNT STATE : -/-, -
CACHE PREFERENCE, CATEGORY : PG1, 0001 (SCRATCH)
LAST MOUNTED (UTC)        : 2018-08-28 03:34:31
LAST MODIFIED LIB/DV, UTC(UTC): Lipizzan/0000, 2018-08-28 03:34:25
KNOWN CPYS, REQ, REMOVED  : 1, 1, 0 (N)
DEL EXP, WHEN (UTC)       : N, -
HOT, FlashCopy            : N, ACTIVE
LWORM RET STATE, TIME(UTC) : D, 2020-11-03 00:00:00
-----
LIBRARY RQ CA P-PVOL S-PVOL CPS CPQ CPP RM CP CD
Lipizzan N Y ----- ----- CMP - DEF N 2 0
Arabian N N ----- ----- NOR - NOC N 0 0
cluster2 N N ----- ----- NOR - NOC N 0 0
Palomino N N ----- ----- NOR - NOC N 1 0

```

Example 16-9 shows the flash instance of the same logical volume.

Example 16-9 Display of a logical volume after modification from production - Flash volume

```
LI REQ,HYDRAG,LVOL,Z10000,INFO,FLASH
CBR1020I Processing LIBRARY command: REQ,HYDRAG,LVOL,Z10000,INFO,FLASH
CBR1280I Library HYDRAG request. 886
Keywords: LVOL,Z10000,INFO,FLASH
```

```
-----
LOGICAL VOLUME INFO V3 0.0
FlashCopy VOLUME           : Z10000
MEDIA, FMT, MAX(MB), CWRAP : ECST, 6, 800, N
SIZE(MB) COMP, CHAN, RATIO : 14, 95, 6.38:1(FICON)
CURRENT OWNER, TVC LIB     : Lipizzan, Lipizzan
MOUNTED LIB/DV, MNT STATE  : -/-, -
CACHE PREFERENCE, CATEGORY : ---, 0001 (PRIVATE)
LAST MOUNTED (UTC)         : 1970-01-01 00:00:00
LAST MODIFIED LIB/DV, UTC(UTC) : -/-, 2018-08-28 03:34:25
KNOWN CPYS, REQ, REMOVED   : -, -, -
DEL EXP, WHEN (UTC)        : N, -
HOT, FlashCopy             : N, -
LWORM RET STATE, TIME(UTC) : D, 2020-11-03 00:00:00
```

```
-----
LIBRARY RQ CA P-PVOL S-PVOL CPS CPQ CPP RM CP CD
Lipizzan N Y ----- ----- CMP - DEF N 2 0
```

Only the clusters from the DR Family are shown. This information is also available on the MI.

Figure 16-1 on page 839 shows 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 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 action opens a second view, as shown in Figure 16-2 on page 839.

Virtual Volume Details

Cluster [0] (#BA92A) Cluster [1] (#BA92B) Cluster [2] (#BA92C) Cluster [3] (#BA92D)

LEGEND

Virtual volume details:

Volser	ZKP003
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	4.4 MiB
Maximum Volume Capacity (Device)	800 MiB
Current Owner	[0] (#BA92A)
Currently Mounted	No
vNode	-
Virtual Drive	-
Cached Copy Used for Mount	[0] (#BA92A)
Mount State	-
Cache Management Preference Group	1
Last Accessed by a Host	Nov 29, 2013, 6:29:49 PM
Last Modified	Nov 29, 2013, 3:03:47 PM
Category	0001
Storage Group	-----
Management Class	RDDN
Storage Class	-----
Data Class	-----
Volume Data State	Scratched
Flash Copy	Created
Earliest Deletion On	-
Logical WORM	No

Cluster-specific Virtual Volume Properties:

	Cluster[0] (#BA92A)	Cluster[1] (#BA92B)	Cluster[2] (#BA92C)	Cluster[3] (#BA92D)
In Cache	Yes	No	Yes	-
Device Bytes Stored	4.4 MiB (Device)	4.4 MiB (Device)	4.4 MiB (Device)	0 MiB (Device)
Primary Physical Volume	-	E06343	-	-
Secondary Physical Volume	-	None	-	-
Copy Activity	Complete	Complete	Complete	Not Required
Queue Type	-	-	-	-
Copy Mode	Rewind unload (RUN)	Deferred	Deferred	No copy
Deleted	-	-	-	-
Removal Residency	No Removal Attempted	-	No Removal Attempted	-
Removal Time	-	-	-	-
Volume Copy Retention Group	⚠ Prefer Keep	-	⚠ Prefer Keep	⚠ Prefer Keep

"Created" is a link to the Flash Copy Detail page

Figure 16-1 Display of a logical volume with an active FlashCopy

Virtual Volume Details

Cluster [0] (#BA92A) Cluster [1] (#BA92B) Cluster [2] (#BA92C) Cluster [3] (#BA92D)

Flash Copy Details

DR Family Name: DRFAM001
Flash Copy Time: Nov 29, 2013, 3:54:50 PM

Flash Copy details:

Volser	ZKP003
Media Type	Enhanced Capacity Cartridge System Tape
Current Volume Size (Device)	4.4 MiB
Maximum Volume Capacity (Device)	800 MiB
Current Owner	[1] (#BA92B)
Currently Mounted	No
vNode	-
Virtual Drive	-
Cached Copy Used for Mount	[0] (#BA92A)
Mount State	-
Cache Management Preference Group	0
Last Accessed by a Host	Nov 29, 2013, 3:54:50 PM
Last Modified	Nov 29, 2013, 3:15:10 PM
Category	000F
Storage Group	-----
Management Class	RDDN
Storage Class	-----
Data Class	-----
Volume Data State	Active
Earliest Deletion On	-

Cluster-specific Flash Copy Properties:

	Cluster[1] (#BA92B)	Cluster[2] (#BA92C)
In Cache	No	Yes
Device Bytes Stored	4.4 MiB (Device)	4.4 MiB (Device)
Primary Physical Volume	E06343	-
Secondary Physical Volume	None	-
Copy Activity	Complete	Complete
Queue Type	-	-
Copy Mode	Deferred	Deferred
Deleted	-	-
Removal Residency	-	-
Removal Time	-	-
Volume Copy Retention Group	-	-

New Page - Contains same information as the Virtual Volume Detail but this view contains the properties of the flash copy not the live copy.

Figure 16-2 Display of the FlashCopy information of a logical volume

When you run a DR test, you can monitor the cache usage of your TS7700 clusters by using the **LI REQ** command interface or the TS7700 MI.

The following **CACHE2** command provides you with information about the space that is used by the FlashCopy in the middle of the output (see Example 16-10):

```
LI REQ,<distributed library name>,CACHE2
```

Example 16-10 Cache Consumption FlashCopy

```
TAPE VOLUME CACHE STATE V5.0
TOTAL SPACE INSTALLED/ENABLED: 162TB/ 162TB
TOTAL ADJUSTED CACHE SPACE USED: 0.0GB
CACHE ENCRYPTION STATUS; NOT CAPABLE
OVERCOMMITTED CACHE PARTITIONS: NONE
CACHE RESIDENT ONLY PARTITION
PRIVATE CACHE SPACE USED: 0.0GB
SCRATCH CACHE SPACE USED: 0.0GB
CP ALLOC USED PIN PKP PRM COPY CPYT
0 151TB 0.0GB 0.0GB 0.0GB 0.0GB 0.0GB 0
```

FlashCopy INFORMATION

INDEX	ENABLED	SIZE
1	NO	0.0GB
2	NO	0.0GB
3	NO	0.0GB
4	NO	0.0GB
5	NO	0.0GB
6	NO	0.0GB
7	NO	0.0GB
8	NO	0.0GB

TIERED CACHE PARTITIONS

CP	TCO	ALLOC	USED	PG0	PG1	PMIGR	D	PMIGR	COPY	PMT	CPYT
1	YNN	3000GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0
2	NNY	4000GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0
3	NNN	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0
4	YNN	4000GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0
5	NNN	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0
6	NNN	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0
7	NNN	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0.0GB	0	0

You can also find the same information from within the MI. You can select the following display windows:

- ▶ Monitor
- ▶ Performance
- ▶ Cache Usage

Figure 16-3 is an example of Cache Utilization output.

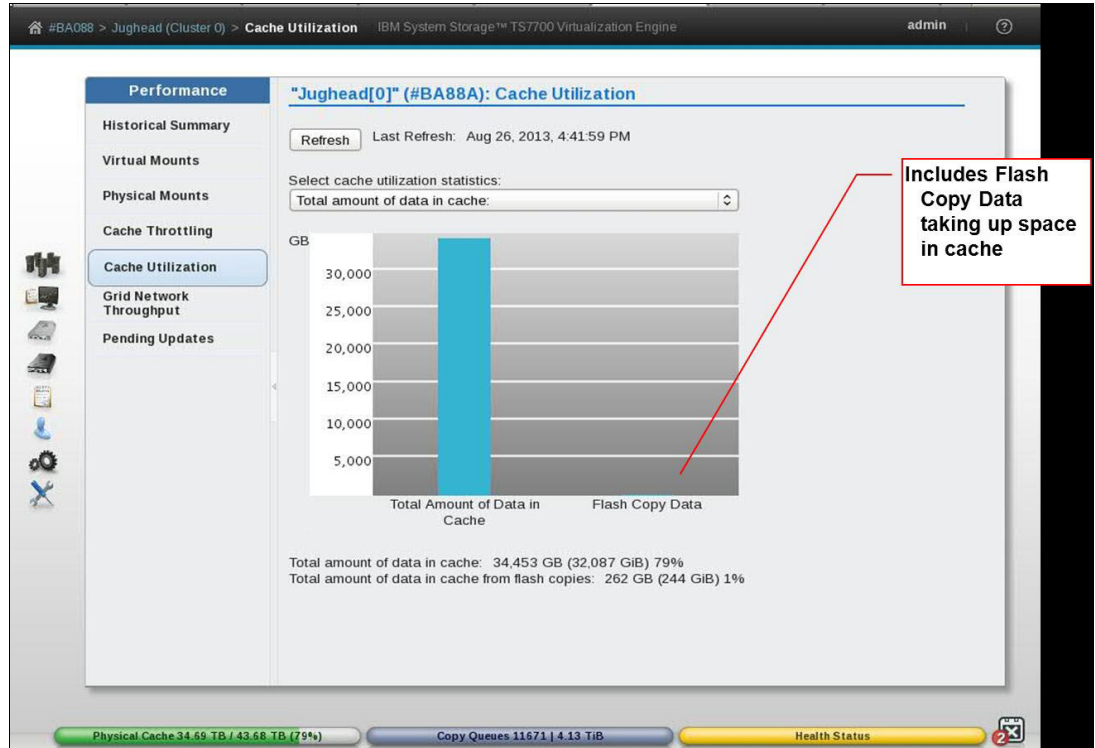


Figure 16-3 Cache usage of FlashCopy data

Also, you can control the usage of your virtual drives. You can select the following displays on the MI:

- ▶ Virtual
- ▶ Virtual Tape Drives

Figure 16-4 is an example of virtual tape drive output.

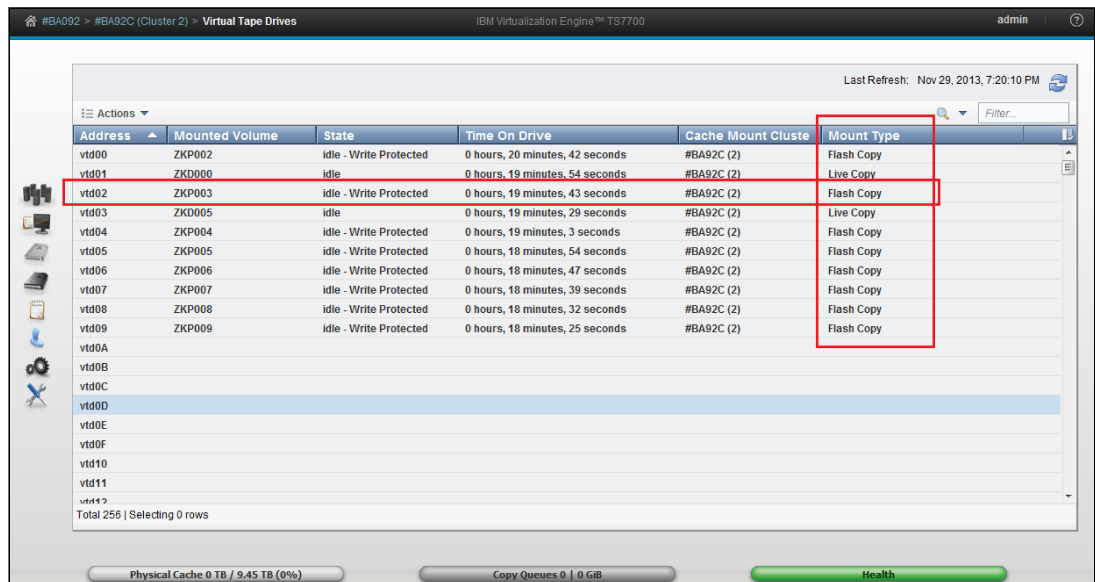


Figure 16-4 Virtual Tape Drive window during a FlashCopy for DR test

16.5 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. Although it might be tempting to skip to these lists, we suggest 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.

It is recommended to review 16.2, “DR testing methods” on page 820 before to choosing a DR testing method to use.

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 any devices online to the disaster clusters to read/write production data on those clusters.

16.5.1 Method 1: DR Testing by using FlashCopy

The next section describes the steps that can be used to run DR testing by using the FlashCopy function. For more information about these commands, see [TS7700 Series Best Practices - FlashCopy for DR Testing](#).

These steps are written in a checklist format to provide a reference of the steps that are needed to complete this method. It is advised that you review all these steps before the DR exercise and during the DR exercise.

Because the steps were written to apply to more than one TS7700 grid configuration, ensure that you understand each step and how it applies to your environment before running each step.

Method 1: DR Testing by using FlashCopy: Steps

To perform DR by using FlashCopy, complete the following steps:

1. Determine which MEDIA1, MEDIA2, ERROR, and PRIVATE categories are to 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 (DR) volumes.
2. Using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that are chosen to the TS7700 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 is in the DR Family 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 must be added only to allow a volume to be moved out of an ERROR state. MEDIA1 is needed only if MEDIA1 is used or any Automatic Class Selection (ACS) routine can result in the use of a default Data Class in which MEDIA1 is included (even if not used).

4. Issue the following command to determine the status of the DR configuration within the TS7700 grid:

```
LI REQ,<COMPOSITE>,DRSETUP,SHOW
```

If a DR Family already is defined, you must choose to use it or delete it and start with a new configuration. To delete it, you must remove each cluster and when the last cluster is

removed, the DR Family is 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 is removed, the command response confirms that the DR Family was deleted because no members exist.

Note: The steps that follow assume that you do *not* have a defined DR Family.

5. Create the DR Family that is used for the DR test and add a cluster to the DR Family by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,ADD,<CLUSTER ID>
```

Wait for the command response that confirms that the DR Family was created and the cluster was added before continuing.

After the DR Family is created, this command can be used repeatedly to add clusters to the DR Family.

6. 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 that confirms that write protect was enabled successfully.

7. 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>
```

Note: Steps 1–7 can often be completed before 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 must be disabled in the DR Family clusters if a true DR event occurs as part of the DR sequence.

8. Choose the volume serial ranges that will be created and used for input/output processing on the DR host. These ranges are known as *DR volumes*. Update the TMS on the production hosts to ignore the ranges during host cartridge entry processing (if RMM is used, see 16.3.3, “Cartridge entry considerations” on page 827) or use the **LI DISABLE, CBRUXENT** command to prevent host cartridge entry processing entirely.
9. On each DR cluster, ensure that the Management Classes that are 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 specifically for DR testing. On each DR cluster, set the Copy Mode for the Management Class that is used to No Copy for each non-DR cluster by clicking **Constructs** → **Management Classes**.
10. Start the DR host and restore the DR environment from the production environment.
11. By 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 are used by volumes that are created for use by the DR host during the DR test.

After the categories are defined, start 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 continue to be used if an IPL occurs.

12. If Live copy usage is wanted, enable Live copy by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,LIVECOPY,FAMILY
```

Wait for the command response that confirms that the live copy usage is set to 'Family'.

13. Verify that the DR Family environment is as expected by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,SHOW,<FAMILYNAME>
```

14. Update the TMS on the DR host to enable the new volume serial ranges to be accepted during host cartridge entry processing. For more information, see 16.3.3, "Cartridge entry considerations" on page 827.

15. If a new Management Class was defined for each DR cluster in Step 9, modify the ACS routines on the DR host to direct new tape allocations to this Management Class. Then, activate the new SMS configuration on the DR host.

16. On the DR host, vary online the devices to the DR Family clusters.

17. By using the TS7700 MI, insert the new volume serial ranges by selecting **Virtual** → **Virtual Volumes** → **Insert a new virtual volume range**. Verify that the DR Host successfully processed the volumes during host cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes. If the **LI DISABLE, CBRUXENT** command was used in Step 9, use the **LI RESET, CBRUXENT** command on those production hosts to re-enable the ability to run host cartridge entry processing.

18. Change the Autoremoval Temporary Threshold on the TS7700 used for DR testing to ensure that enough cache space is available for DR data and production data. This situation is only applicable for CP0 and only if more than 10 TB is available in CP0. Wait until the removal process completes.

19. When you are ready to start the DR test, enable the FlashCopy by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,FLASH,ENABLE
```

Wait for the command response that confirms that the FlashCopy was enabled successfully.

20. Verify that the DR Family environment is as expected by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,SHOW,<FAMILYNAME>
```

21. Run the DR test.

22. When the DR test is complete, **SCRATCH** the volumes that were written to by the DR host. Be careful not to scratch production volumes. For more information, see "Return-to-scratch Considerations" on page 827. This process can be done by using one of the following methods:

- 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).
- Use ISMF to **ALTER** those volumes written to by the DR host to **SCRATCH**.
- Use the CBRSPCLS SAMPLIB member to change the use attribute of each volume to **SCRATCH**

23. Keep Write Protect Mode enabled and disable FlashCopy by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,FLASH,DISABLE
```


Wait for the command response that confirms that the FlashCopy was deleted successfully.

Note: FlashCopy MUST be disabled for the volume EJECTs in the next step to process successfully.

24. After all the volumes that were 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 process can be done by using one of the following methods:

- Run a TMS job to issue the EJECTs for these volumes.
- Use ISMF to EJECT each volume.
- Use the CBRSPCLS SAMPLIB member to eject each volume.

25. Shut down the DR host.

If the Management Class that is used on the DR cluster from Step 9 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.

26. Disable Write Protect Mode by using the following command:

```
LI REQ,<COMPOSITE>,DRSETUP,<FAMILYNAME>,WP,DISABLE
```

Wait for the command response that confirms that write protect was disabled successfully.

27. Delete the DR Family by removing each cluster. When the last cluster is removed, the DR Family is 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. When the last cluster is removed, the command response confirms that the DR Family was deleted because no members exist.

16.5.2 Method 2: Using Write Protect Mode on DR clusters

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 these steps before the DR exercise, and during the DR exercise.

Because the steps were written to apply to more than one TS7700 grid configuration, ensure that you understand each step and how it applies to your environment before completing each step.

To use Write Protect Mode on DR clusters, complete the following steps:

1. Determine which MEDIA1, MEDIA2, ERROR, and PRIVATE categories are 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 (DR) volumes.
2. By using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that were chosen to the TS7700 by selecting **Virtual** → **Categories** → **Add Scratch Category**.
3. By 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 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 must be added only to allow a volume to be moved out of an ERROR state. MEDIA1 is needed only if MEDIA1 is used or any ACS routine can result in the use of a default Data Class in which MEDIA1 is included (even if not used).

4. By using the TS7700 MI, enable Write Protect Mode on each cluster that is used as a DR cluster by selecting **SETTINGS** → **Cluster Settings** → **Write Protect Mode** → **Enable Write Protect Mode** → **Submit Changes**.

Note: Steps 1–4 can often be completed before a DR test. In fact, they can be set up once and left enabled indefinitely to protect production volumes from ever being overwritten by workloads that are run on a DR host. If done far in advance, the only item to consider is that the Write Protect Mode must be disabled in the DR clusters if a true DR event occurs 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 (see 16.3.3, “Cartridge entry considerations” on page 827) or use the **LI DISABLE, CBRUXENT** command to prevent host cartridge entry processing entirely.
6. On each DR cluster, ensure that the Management Classes that are 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 specifically for DR testing. On each DR cluster, set the Copy Mode for the Management Class that is used to No Copy for each non-DR cluster by selecting **Constructs** → **Management Classes**.
7. Start the DR host and restore the DR environment from the production environment.
8. By 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 that are created for use by the DR host during the DR test.

After the categories are defined, start 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 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. For more information, see 16.3.3, “Cartridge entry considerations” on page 827.
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. By using the TS7700 MI, insert the new volume serial ranges by selecting **Virtual** → **Virtual Volumes** → **Insert a new virtual volume range**. Verify that the DR Host successfully processed the volumes during cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes.

If the **LI DISABLE, CBRUXENT** command was used in Step 5, use the **LI RESET, CBRUXENT** command on those production hosts to re-enable the ability to run host cartridge entry processing.

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. For more information, see “Return-to-scratch Considerations” on page 827. This process can be done by using one of the following methods:
 - 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).
 - Use ISMF to ALTER those volumes that were written to by the DR host to SCRATCH.
 - Use the CBRSPCLS SAMPLIB member to change the use attribute of each volume to SCRATCH.
15. When 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 process can be done by using one of the following methods:
 - Run a TMS job to issue the EJECTs for these volumes.
 - Use ISMF to EJECT each volume.
 - Use the CBRSPCLS SAMPLIB member to eject each volume.
16. Shut down the DR host.
17. If the Management Class that is used on the DR cluster from Step 6 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 want to prevent accidentally returning production volumes to SCRATCH by the DR host, keep the Write Protect Mode enabled on the DR clusters.
19. Alternatively, you can disable Write Protect Mode in each DR cluster by using the TS7700 MI by selecting **SETTINGS** → **Cluster Settings** → **Write Protect Mode** → **Disable Write Protect Mode** → **Submit Changes**.
20. If any production volumes were read by the DR host from the DR cluster during the DR test, those volumes are now owned by the DR cluster. To transfer the ownership back to the production cluster, use the **LI REQ,OTCNTL** command with the **START** keyword. For more information, see “Returning ownership of production volumes to the production cluster” on page 829.

16.5.3 Method 3: DR Testing without Write Protect Mode

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 another level of write-protection in case the TMS on the disaster DR host is not configured correctly to prevent writes to the production volumes.

As with the previous method, the steps that are described in this section are written in a checklist format to provide a reference of the steps that are needed to complete this method. It is advised that you review all these steps before the DR exercise and during the DR exercise. Because the steps were written to apply to more than one TS7700 grid configuration, ensure that you understand each step and how it applies to your environment before running each step.

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 that is

used by a production system to ensure separation of production volumes from disaster volumes.

2. By using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that were chosen to the TS7700 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 (see 16.3.3, “Cartridge entry considerations” on page 827) or use the **LI DISABLE, CBRUXENT** command to prevent host cartridge entry processing entirely.
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 specifically for DR testing. On each DR cluster, set the Copy Mode for the Management Class used to No Copy for each non-DR cluster by selecting **Constructs** → **Management Classes**.
5. Start the DR host and restore the DR environment from the production environment.
6. By 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 are used by volumes that were created for use by the DR host during the DR test.

After the categories are defined, start 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 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. For more information, see 16.3.3, “Cartridge entry considerations” on page 827.
8. Update the TMS on the DR host to reject any output that is directed toward volumes in the production volume serial range. This process is done as a safeguard to protect against production tapes being accidentally written to by the DR host. For more information, see “TMS Write Protection” on page 826.
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. By using the TS7700 MI, insert the new volume serial ranges by selecting **Virtual** → **Virtual Volumes** → **Insert a new virtual volume range**. Verify that the DR Host processed the volumes during cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes.

If the **LI DISABLE, CBRUXENT** command was used in Step 3, use the **LI RESET, CBRUXENT** command on those production hosts to re-enable the ability to run host cartridge entry processing.

12. If DFSMSHsm is active on the DR host, mark all ML2 volumes full by issuing the following command for each ML2 volume:

```
F xxxxx,DELVOL MIGRATION(MARKFULL)
```

In addition, issue the following command to prevent DFSMSHsm RECYCLE processing from occurring:

```
- F xxxxx,HOLD RECYCLE
```

13. For maximum protection, ensure that the following procedures are not run:
 - TMS housekeeping activity at the DR site
 - Short-on-scratch TMS 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. For more information, see “Return-to-scratch Considerations” on page 827. This process can be done by using one of the following methods:
 - 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).
 - Use ISMF to ALTER those volumes that are written to by the DR host to SCRATCH.
 - Use the CBRSPCLS 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 process can be done by using one of the following methods:
 - Run a TMS job to issue the EJECTs for these volumes.
 - Use ISMF to EJECT each volume.
 - Use the CBRSPCLS SAMPLIB member to eject each volume.
17. Shut down the DR host.
18. If the Management Class that was used on the DR cluster from Step 4 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.
19. If any production volumes were read by the DR host from the DR cluster during the DR test, those volumes are owned by the DR cluster. To transfer the ownership back to the production cluster, use the **LI REQ,OTCNTL** command with the **START** keyword. For more information, see “Returning ownership of production volumes to the production cluster” on page 829.

16.5.4 Method 4: Breaking the grid links

A real disaster can be simulated by breaking the grid links between the production clusters and DR clusters in a TS7700 grid.

The concern about losing data in a real disaster during a DR test is the major drawback to the use of this DR method. Because of this drawback, it is suggested that one of the other methods that was described (FlashCopy or Write Protect Mode) be used, if possible.

Important: Do not use virtual drives in the DR cluster from the production host.

As with the previous methods, these steps were written in a checklist format to provide a reference of the steps that are needed to complete this method. It is suggested that you review all these steps before the DR exercise and during the DR exercise. As the steps were written to apply to more than one TS7700 grid configuration, make sure that you understand each step and how it applies to your environment before running each step.

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. By using the TS7700 MI, add the MEDIA1 and MEDIA2 categories that are chosen to the TS7700 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 (for more information, see 16.3.3, “Cartridge entry considerations” on page 827) or use the **LI DISABLE, CBRUXENT** command to prevent host cartridge entry processing entirely.
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 specifically for DR testing. On each DR cluster, set the Copy Mode for the Management Class that is used to No Copy for each non-DR cluster by selecting **Constructs** → **Management Classes**.
5. Start the DR host and restore the DR environment from the production environment.
6. By 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 are used by volumes that are created for use by the DR host during the DR test. After the categories are defined, start the DR host to ensure that the categories are used.

Alternatively, use the **DS QL, CATS** command 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 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. For more information, see 16.3.3, “Cartridge entry considerations” on page 827.
8. Update the TMS on the DR host to reject any output that is directed toward volumes in the production volume serial range. This process is done as a safeguard to protect against production tapes being accidentally written to by the DR host. For more information, see “TMS Write Protection” on page 826.
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. By using the TS7700 MI, insert the new volume serial ranges by selecting **Virtual** → **Virtual Volumes** → **Insert a new virtual volume range**. Verify that the DR Host successfully processed the volumes during cartridge entry by reviewing the SYSLOG for the CBR3610I messages surfacing those volumes. If the **LI DISABLE, CBRUXENT** command was used in Step 3, use the **LI RESET, CBRUXENT** command on those production hosts to re-enable the ability to run host cartridge entry processing.

Note: Insert the DR volume serial range *before* completing the next step (breaking the grid links). After the grid links are broken, you cannot insert volumes to the DR cluster unless write ownership takeover mode is enabled within the DR cluster, which is *not* recommended because it can allow production volumes to be written to by the DR host.

12. Contact your CE to break the grid link connections between the production clusters and the disaster clusters. Do *not* disable a grid link by using the **Library Request** command. Disabling the grid link by using this 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 DR cluster to allow read-only access to volumes owned by each production cluster by selecting **Service** → **Ownership Takeover Mode**.
14. If DFSMSshm is active on the DR host, mark all ML2 volumes full by issuing the following command for each ML2 volume:


```
F xxxxx,DELVOL MIGRATION(MARKFULL)
```

 In addition, issue the following command to prevent DFSMSshm RECYCLE processing from occurring:
 - F xxxxx,HOLD RECYCLE
15. For maximum protection, ensure that the following procedures are not run:
 - TMS housekeeping activity at the DR site
 - Short-on-scratch TMS 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. For more information, see “Return-to-scratch Considerations” on page 827. Use one of the following methods:
 - 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).
 - Use ISMF to ALTER those volumes that were written to by the DR host to SCRATCH.
 - Use the CBRSPCLS 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 process can be done by using one of the following methods:
 - Run a TMS job to issue the EJECTs for these volumes.
 - Use ISMF to EJECT each volume.
 - Use the CBRSPCLS SAMPLIB member to eject each volume.
19. Shut down the DR host.
20. If the Management Class that was used on the DR cluster from Step 4 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 DR cluster by clicking **Service** → **Ownership Takeover Mode**. This step returns each DR cluster to its normal state regarding takeover processing.
22. Engage your CE to reestablish the link connection between the production clusters and the DR clusters.
23. If any production volumes were read by the DR host from the DR cluster during the DR test, those volumes are now owned by the DR cluster. To transfer the ownership back to the production cluster, use the **LI REQ OTCNTL** command with the **START** keyword. For more information, see “Returning ownership of production volumes to the production cluster” on page 829.

16.6 Expected failures during a DR test

Several expected failures that can occur during a DR test are described in this section.

The messages that are shown in Example 16-11 appear if you attempt to read a virtual volume that was not present at time zero in the DR Family.

Example 16-11 Expected failures during the DR 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 that is shown in Example 16-12 appear if you want to modify a volume that is in a write-protected media category.

Example 16-12 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 that is shown in Example 16-13 might appear if a job was running on the cluster while the FlashCopy was enabled.

Example 16-13 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
```

Considerations when using LWORM in TS7700 DR testing

When you use LWORM in a TS7700 environment where a copy of the tape management system's database is being used for DR testing, it is possible to run into validation errors. Some tape management systems validate the access to an LWORM volume by checking the recorded write mount count in the TMS database against the actual write mount count as recorded in the TS7700 database. If these two values do not match, access to the volume can be denied. For example, in DFRMM when the write mount count of the volume does not match, an EDG4056I is issued (Example 16-14), and access to the volume is rejected.

Example 16-14 EDG4056I message

```
EDG4056I VOLUME volser REJECTED. WRITE MOUNT COUNT count_value_read FOR VOLUME
DOES NOT MATCH RECORDED VALUE count_value_rec
```

In such cases, an access to the LWORM volume might have occurred at a point that is later than the time when the copy of the TMS database was taken. To restore access to the volume

in such a case, alter the WMC in the TMS database to match that which is known by the TS7700. In DFRMM, this can be done with following TSO subcommand:

```
RMM CV <volser> WMC(<xx>) FORCE
```

where xx is the count_value_read that was provided in the message text for the EDG4056I.



RESTful API

This chapter describes the TS7700 RESTful (REST) API.

The TS7700 RESTful API can be used as an alternative to the TS7700 Host Command Line Interface (Host Console Request function, CLI) out of the z/OS operating system or the TS7700 Management Interface (Web GUI) as a method of obtaining information from and sending commands to the TS7700 virtual tape system. It is an out-band method based on standard IP communication. The first release of this function with a limited command set was first introduced with TS7700 R5.4 release. The supported RESTful API endpoints continue to be updated.

Note: CLIs are meant to be human-readable. Application programming interfaces (APIs) are programmed into applications. Typical users rely on an interactive GUI on their devices or a command-line interface (CLI) as another route to the same end.

This chapter includes the following topics:

- ▶ 17.1, “Overview” on page 856
- ▶ 17.2, “Access Token” on page 860
- ▶ 17.3, “Query, Filter, and Sort” on page 862
- ▶ 17.3.3, “Sort function” on page 867
- ▶ 17.7, “Error Handling” on page 870

17.1 Overview

This section gives you a general overview about what a RESTful API is and its purpose with some use cases and examples. Furthermore, you learn about the URL structures being used, the data types and where to find more help and information about the available TS7700 API commands.

The TS7700 RESTful model application programming interface (API) consists of commands that are used to read system resources. Hypertext Transfer Protocol Secure (HTTPS) is used to communicate with the RESTful apiserver. The RESTful apiserver is part of the TS7700 virtual tape system internal operating system. It is part of the basic user interface to interact with the TS7700 virtual tape system, like the integrated Management Interface (Web GUI) and is active and enabled by default to receive API calls.

REST APIs provide a flexible, lightweight way to integrate applications, and have emerged as the most common method for connecting components in microservices architectures. It does provide a new user interface to the TS7700, which is standardized and platform independent. The goal is to provide all user operations that are supported by any of the current interfaces to be available in the new REST API.

Figure 17-1 shows the various communication paths to interact with the TS7700 system for administration, configuration, data replication, client queries, and support. As you can see, the RESTful API layer is also the foundation of the calls that are internally triggered by user interactions through the Management Interface (MI, Web GUI).

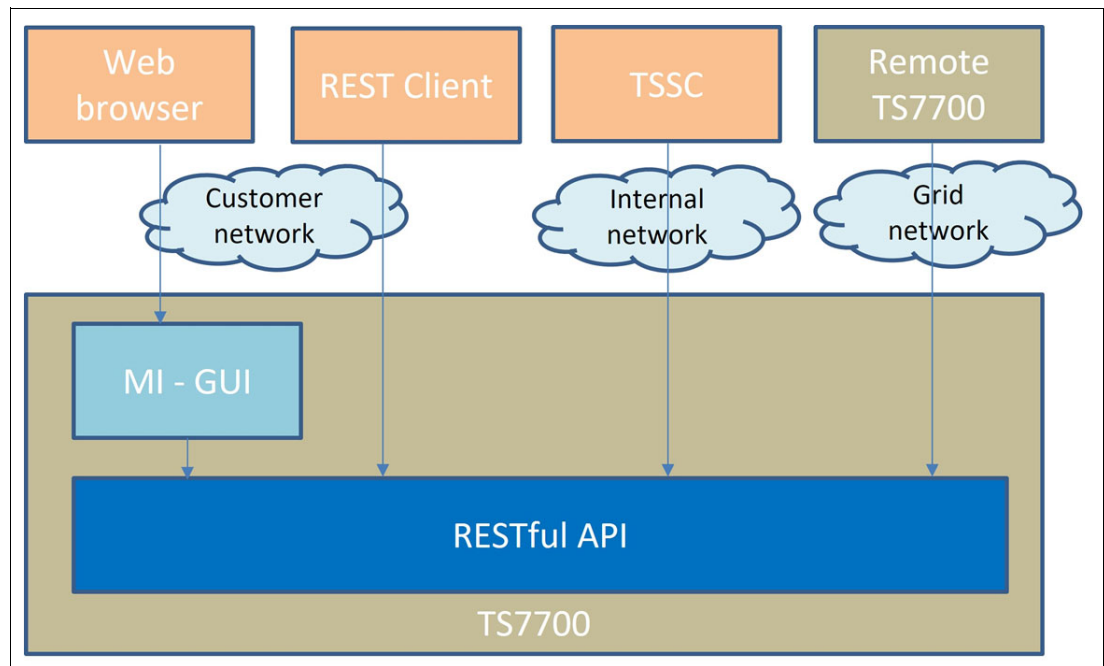


Figure 17-1 Communication paths of interacting with the TS7700 subsystem

RESTful API uses the operational method HTTP GET to query read operations.

Note: TS7700 RESTful API currently supports only the GET method for all the resources except for authentication where the POST method is used.

The URL structure for an operation is `https://<host_name>/<api>/v<x>/<endpoint>`, where:

- ▶ `host_name` is the IP address of the TS7700 storage system (same as for GUI).
- ▶ `api` is the API application name
- ▶ `x` is the version (1 or 2) of the API.
- ▶ `endpoint` is the resource that is targeted for the operation.

17.1.1 URL structure

The URL should be encoded as documented in Request for Comments (RFC) 3986. A Request for Comments is a formal document that is drafted by the Internet Engineering Task Force (IETF) that describes the specifications for a particular technology. When an RFC is ratified, it becomes a formal standards document. RFC 3986 defines the Uniform Resource Identifier (URI) encoding. A Uniform Resource Identifier (URI) is a compact sequence of characters that identifies an abstract or physical resource. This specification defines the generic URI syntax and a process for resolving URI references that might be in relative form, along with guidelines and security considerations for the use of URIs on the internet. The URI syntax defines a grammar that is a superset of all valid URIs, allowing an implementation to parse the common components of a URI reference without knowing the scheme-specific requirements of every possible identifier. This specification does not define a generative grammar for URIs; that task is performed by the individual specifications of each URI scheme.

If a version supersedes an existing version of an API, then the new version is not compatible with the previous versions. The older version will still be available for use. The default TCP/IP port number is 443, and port number 80 is disabled. The request and responses are in JavaScript Object Notation (JSON) format or Command-Separated Values (CSV) format. All the examples in this document are formatted in JSON.

Information: JSON is a text-based way of representing JavaScript object literals, arrays, and scalar data. JSON is relatively simple to read/write, and can also be easier for software to parse and generate. It is often used for serializing structured data and exchanging it over a network, typically between a server and web applications. For more information, see [What is JSON?](#).

17.1.2 Data types

The TS7700 virtual tape system RESTful API uses the following data types:

- ▶ **Boolean:** Boolean attributes are provided as true or false
- ▶ **Integer:** An integer is a whole number that can be positive or negative
- ▶ **Number:** A fractional number is formatted as a double or a float
- ▶ **String:** Alphanumeric sequence of letters or numbers
- ▶ **Array:** A list of values

17.1.3 Endpoints

The REST API for the TS7700 supports TS7700-specific endpoints. Supported endpoints, syntax, and usage for each microcode level that supports the REST API can be found through the embedded help “?” function drop-down menu within the TS7700 Web GUI. You find this option in the upper right. See Figure 17-2 on page 858.

Note: You need access to a running and online TS7700 system with R5.3 microcode or later to see that “RESTful API” option.

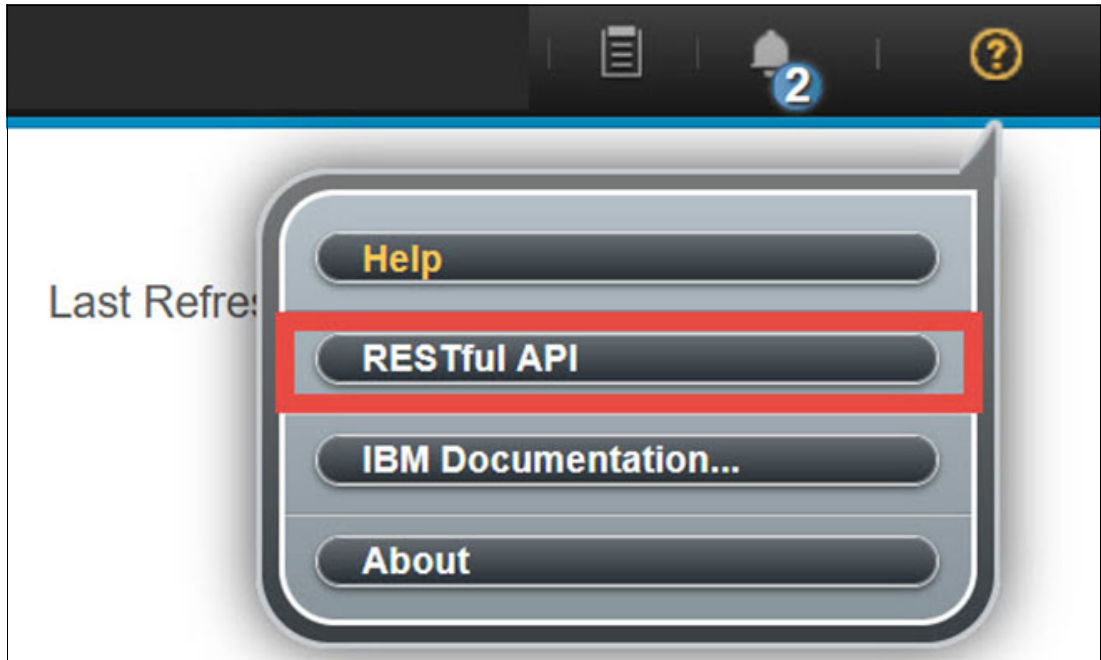


Figure 17-2 RESTful API help function and documentation within the TS7700 MI

When you click the question mark and select the “RESTful API” option in the drop-down menu, a new browser window opens as shown in Figure 17-3. This window opens to IBM Open Liberty®. Its content shows you the TS7700 REST API OpenAPI definitions. These definitions can be used by documentation generation tools to display the API, code generation tools to generate servers and clients in various programming languages, testing tools and more.

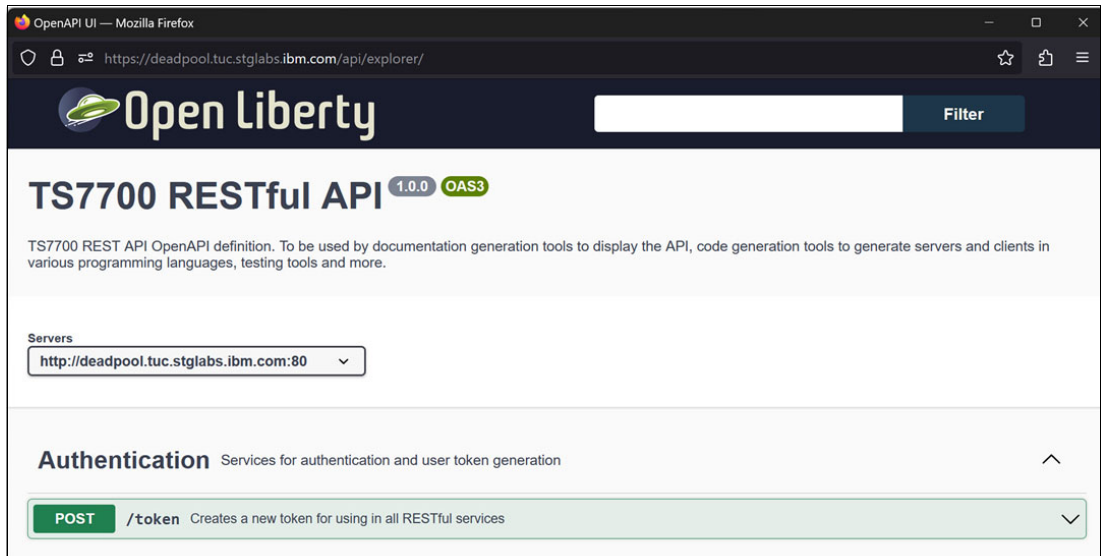


Figure 17-3 The “Open Liberty®” RESTful API help window within the TS7700 MI

By selecting a RESTful API endpoint within the Open Liberty window, you find all the needed command syntax and parameters of each available API endpoint. Figure 17-4 on page 859 shows you the **GET /clusters** endpoint details as an example.

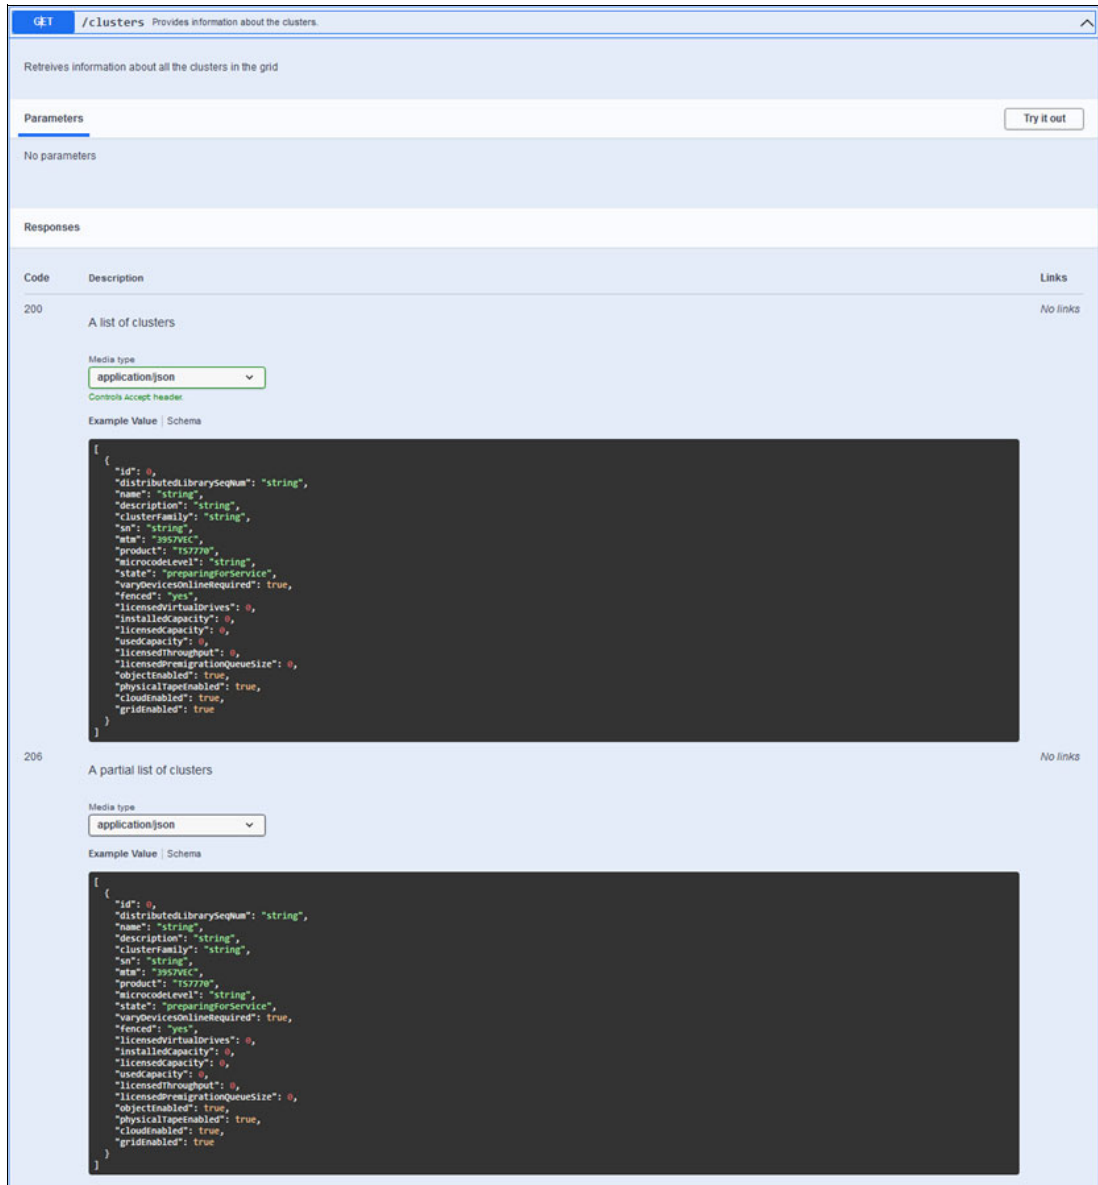


Figure 17-4 The RESTful API GET /clusters endpoint syntax and parameter help window

The output of the `GET /clusters` endpoint command shown as an example in Figure 17-4 provides information about all the clusters in the grid. The Cluster ID can be specified as a parameter, although it is not required. The command response is a list of one or more clusters of the Grid, depending on the number and members of such a grid. You can easily identify in this response important cluster characteristics like the distributed library ID, cluster serial number, machine type, microcode level, and status. Also, this cluster query can help you to see its various installed and activated license features.

Information: For more information about the TS7700 RESTful API endpoint updates of each microcode level, see [What's new](#).

The following sections within this chapter look at how to generate the initial access token, the query and response format, and coding, and includes instructions to filter a response.


```
n2gUKUhDdpBVDpU9hW15krcSemMnykV2hGBZSwRI4jj-xe1yyIaRKjLxNLWLe9ts7mJ9u1w0u0--rMwcuF
-4QnnG1bZ8_JDjnZ3vFAdizoZ8p-LGRHJNzGjZmoQt2EsDe28X4ijA"
}
]
}
```

In Example 17-4 on page 861, the individual parts of the API response, such as the “expireTime” of the requested token can be identified and read better.

Note: A TS7700 generated RESTful API session token is valid for a limited time. How soon the generated token is expired just depends on the Session Timeout value that is defined within the MI/GUI.

The token itself from this example is the sequence of numbers and characters that are shown in Example 17-5.

Example 17-5 TS7700 RESTful API token

```
eyJraWQiOiJwM0pEa2JCUMtdDNLETFvcC1idG1jTExMekZUeV9nbURrLUFTYkRlChPvIiwidHlwIjoISt1
dUIwiYXNlIjoIUmYNTYifQ.eyJ0b2t1b190eXB1IjoIUmVhcmVhIiwic3ViIjoIYWRTaW4iLCJvcmlnI
joiQkE2OEMiLCJncmlkIjoIQkEwNjgiLCJncm91cHMiOi1siQWRtaW5pc3RyYXRvcjJdLCJjbHVzdGVycyI
6IjEsMiwzIiwicG9saWN5IjoITG9jYWwiLCJpc3MiOiJjbHVzdGVyMCIsImV4cCI6MTY4NDE4MjY4NDMwM
CwiaWF0IjojNjg0MTcwMDg0fQ.TIHquqofmvymOXkwjzt-Yk4nky80rQdSXfhJCDROeMxIv_KJMCn6ENrN
KfSaCrsSiX3sRYOC01aD5LoUBXxTDnsa0dgH-zstuxKBmkucALiWyMw00Eh1v0ZmSuEQ1NDZchnEJo95V
DpZdnVJ-6LUB8-1Y1j2NiQsSVWfkTrEGNQM3aEkcLrcyMX4Qr3swaf2edqvKAMgo7avfxgrBqqI5jJCIVV
vKko-460uiI2CynyEkxBf0eE40Gcsr0A5hCHffxGVfb0swurDyPZLhQ2SYF3aK6H1I2GSqLNhERTL1UYmc
8x13emS-ZrAG8marYbtfWF8tcF1tRQv4pm-Swu1ITQH1q6QwAi40o-hTKcG1_oN-MX3d7tDc2kRvS3Bo5F
vNWvTMQ2HB52AT10HPJgwJaOfTQ988itaBC17Rd-DRcWzJ5KNXuIXQnr3Rk70sc1F7VQZP0trQHf05c-Cp
5UJkv6yb5mL_Iq3u-NGXebqfjklxeLGL0fuR0jW0U508FBgUrA0b7XhB_-7IFk9P0eTnXMheZeW4tROM3n
2gUKUhDdpBVDpU9hW15krcSemMnykV2hGBZSwRI4jj-xe1yyIaRKjLxNLWLe9ts7mJ9u1w0u0--rMwcuF-
4QnnG1bZ8_JDjnZ3vFAdizoZ8p-LGRHJNzGjZmoQt2EsDe28X4ijA
```

For the ease of working with RESTful API GET commands, it might be helpful, to store the returned TS7700 RESTful API token into a shell variable as shown in Example 17-6.

Example 17-6 RESTful API response to obtain a session token that is saved to a shell variable

```
$ hydratok=$(curl -d '{"username":"${hydrauser}"',
"password":"${hydrapassword}" }" -L -k -s -X POST -H "Content-Type:
application/json" ${hydraurl}/api/v1/token | jq -r '.data[0] | .token')
```

17.3 Query, Filter, and Sort

Topics in this section include how to query, filter, and sort the TS7700 RESTful API response through the GET method.

17.3.1 Query

You can run a query by submitting the HTTP GET method on a URL that also uses curl. Example 17-7 on page 863 shows again the query request for the properties of one or all TS7700 clusters in a grid.

Example 17-7 RESTful API GET request to obtain all cluster information

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json" ${hydraturl}/api/v1/clusters
{"metadata":{"responseSent":"2023-05-15T17:37:24Z","request":"GET https://host_name/api/v1/clusters","resources":3,"requestReceived":"2023-05-15T17:37:13Z"},"data":[{"product":"TS7760","fenced":"no","clusterFamily":"TAPE","licensedThroughput":300,"usedCapacity":39.22,"objectEnabled":false,"cloudEnabled":false,"mtm":"3957VEC","description":"","licensedCapacity":236.82,"physicalTapeEnabled":true,"installedCapacity":156.82,"licensedPremigrationQueueSize":15.0,"name":"Squint","varyDevicesOnlineRequired":false,"id":1,"sn":"78-EFC5V","state":"online","microcodeLevel":"8.51.2.12","licensedVirtualDrives":496,"gridEnabled":true,"distributedLibrarySeqNum":"BA68B"},{"product":"TS7770","fenced":"no","clusterFamily":"CLOUD","licensedThroughput":300,"usedCapacity":35.88,"objectEnabled":true,"cloudEnabled":true,"mtm":"3948VED","description":"deadpool cluster","licensedCapacity":3320.0,"physicalTapeEnabled":false,"installedCapacity":133.94,"licensedPremigrationQueueSize":60.0,"name":"Deadpool","varyDevicesOnlineRequired":false,"id":2,"sn":"78-49FD0","state":"online","microcodeLevel":"8.53.1.21","licensedVirtualDrives":496,"gridEnabled":true,"distributedLibrarySeqNum":"BA68C"},{"product":"TS7760","fenced":"no","clusterFamily":"CLOUD","licensedThroughput":300,"usedCapacity":14.88,"objectEnabled":false,"cloudEnabled":true,"mtm":"3957VEC","description":"","licensedCapacity":151.11,"physicalTapeEnabled":false,"installedCapacity":31.11,"licensedPremigrationQueueSize":3.0,"name":"Tom","varyDevicesOnlineRequired":false,"id":3,"sn":"78-EFBCV","state":"online","microcodeLevel":"8.52.0.111","licensedVirtualDrives":272,"gridEnabled":true,"distributedLibrarySeqNum":"BA68D"}]}
```

Again, this response is the native API response to the **GET /clusters** command. You can use the **jq** tool (Example 17-8) to improve readability and to format the output as shown in Figure 17-3 on page 858, where this example command was already shown from the embedded TS7700 RESTful API help (Open Liberty) window.

Example 17-8 RESTful API GET request to obtain cluster information formatted with the jq tool

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json"
${hydraturl}/api/v1/clusters |jq
{
  "metadata": {
    "responseSent": "2023-05-15T17:46:29Z",
    "request": "GET https://host_name/api/v1/clusters",
    "resources": 3,
    "requestReceived": "2023-05-15T17:46:20Z"
  },
  "data": [
    {
      "product": "TS7770",
      "fenced": "no",
      "clusterFamily": "CLOUD",
      "licensedThroughput": 300,
      "usedCapacity": 35.9,
      "objectEnabled": true,
      "cloudEnabled": true,
      "mtm": "3948VED",
      "description": "deadpool cluster",
      "licensedCapacity": 3320,
      "physicalTapeEnabled": false,
      "installedCapacity": 133.94,
      "licensedPremigrationQueueSize": 60,
```

```

    "name": "Deadpool",
    "varyDevicesOnlineRequired": false,
    "id": 2,
    "sn": "78-49FD0",
    "state": "online",
    "microcodeLevel": "8.53.1.21",
    "licensedVirtualDrives": 496,
    "gridEnabled": true,
    "distributedLibrarySeqNum": "BA68C"
  },
  {
    "product": "TS7760",
    "fenced": "no",
    "clusterFamily": "CLOUD",
    "licensedThroughput": 300,
    "usedCapacity": 14.89,
    "objectEnabled": false,
    "cloudEnabled": true,
    "mtm": "3957VEC",
    "description": "",
    "licensedCapacity": 151.11,
    "physicalTapeEnabled": false,
    "installedCapacity": 31.11,
    "licensedPremigrationQueueSize": 3,
    "name": "Tom",
    "varyDevicesOnlineRequired": false,
    "id": 3,
    "sn": "78-EFBCV",
    "state": "online",
    "microcodeLevel": "8.52.200.111",
    "licensedVirtualDrives": 272,
    "gridEnabled": true,
    "distributedLibrarySeqNum": "BA68D"
  },
  {
    "product": "TS7760",
    "fenced": "no",
    "clusterFamily": "TAPE",
    "licensedThroughput": 300,
    "usedCapacity": 39.23,
    "objectEnabled": false,
    "cloudEnabled": false,
    "mtm": "3957VEC",
    "description": "",
    "licensedCapacity": 236.82,
    "physicalTapeEnabled": true,
    "installedCapacity": 156.82,
    "licensedPremigrationQueueSize": 15,
    "name": "Squint",
    "varyDevicesOnlineRequired": false,
    "id": 1,
    "sn": "78-EFC5V",
    "state": "online",
    "microcodeLevel": "8.51.2.12",
    "licensedVirtualDrives": 496,
    "gridEnabled": true,
    "distributedLibrarySeqNum": "BA68B"
  }
]

```

```
}
```

Information: “Authorization: Bearer” for the token means the TS7700 RESTful API uses Bearer authentication (also called token authentication) which is an HTTP authentication scheme that involves security tokens that are called bearer tokens.

17.3.2 Filter

You can also get a filtered response for an API query by specifying the values of the attributes in the URL. The “?” is used behind the specific TS7700 RESTful API GET command to specify the filter attribute. The “&” can be used to filter for multiple characteristics of the response and to make your query and response more specific. Also, the filter on the nested column values has been supported since R5.4.1,

Note: At the time of writing, only the “&” operator is supported as a filter. To filter on the nested column value, combine a parent key and a child key with “.”.

For example, to filter the response for the categories RESTful API GET command attribute with 'type=scratch&category=5001', the URL is constructed as shown in Example 17-9.

Example 17-9 RESTful API call with filtered GET request for categories

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json" ${hydraurl}'/api/v1/categories?type=scratch&category=5001'|jq
```

The response is filtered for the TS7700 virtual volume and category definitions on all clusters of such a grid to only show the current state for “scratch” virtual volumes sitting in category 5001. Example 17-10 shows you the corresponding jq formatted output.

Example 17-10 RESTful API call with filtered GET response for scratch categories jq formatted

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json" ${hydraurl}'/api/v1/categories?type=scratch&category=5001'|jq
{
  "metadata": {
    "responseSent": "2023-05-15T18:32:36Z",
    "request": "GET https://host_name/api/v1/categories?type=scratch&category=5001",
    "resources": 1,
    "requestReceived": "2023-05-15T18:32:33Z"
  },
  "data": [
    {
      "virtualVolumes": 1035343,
      "expireTime": 1,
      "expireHold": true,
      "owningClusters": [
        {
          "cluster": 1,
          "virtualVolumes": 132652,
          "virtualVolumesExpired": 132560
        },
        {
          "cluster": 2,
          "virtualVolumes": 902691,
          "virtualVolumesExpired": 902625
        }
      ]
    }
  ]
}
```

```

        {
          "cluster": 3,
          "virtualVolumes": 0,
          "virtualVolumesExpired": 0
        }
      ],
      "category": "5001",
      "type": "scratch",
      "virtualVolumesExpired": 1035185
    }
  ]
}

```

Example 17-11, the final example, shows another query for `expireTime=24` and `expireHold=true`. You can see in the output that only the category C001 matches the used filtered criteria and no virtual volumes in none of the three clusters currently is in such category.

Example 17-11 RESTful API call with filtered GET request for expireTime and expireHold categories

```

$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json"
${hydraurl}'/api/v1/categories?expireTime=24&expireHold=true'|jq
{
  "metadata": {
    "responseSent": "2023-05-15T18:36:40Z",
    "request": "GET https://host_name/api/v1/categories?expireTime=24&expireHold=true",
    "resources": 1,
    "requestReceived": "2023-05-15T18:36:38Z"
  },
  "data": [
    {
      "virtualVolumes": 0,
      "expireTime": 24,
      "expireHold": true,
      "owningClusters": [
        {
          "cluster": 1,
          "virtualVolumes": 0,
          "virtualVolumesExpired": 0
        },
        {
          "cluster": 2,
          "virtualVolumes": 0,
          "virtualVolumesExpired": 0
        },
        {
          "cluster": 3,
          "virtualVolumes": 0,
          "virtualVolumesExpired": 0
        }
      ],
      "category": "C001",
      "type": "scratch",
      "virtualVolumesExpired": 0
    }
  ]
}

```

17.3.3 Sort function

You can also sort the output of your query by each field.

Example 17-12 shows how to get the output sorted by the field. The “sort=endpoint” is used behind the query command. Also, you can reverse the order using “sort=-endpoint”.

Example 17-12 RESTful API call with GET request for dataClasses sorted by virtualVolumeSize

```
$ curl -k -s -v "https://host_name/api/v2/dataClasses?sort=virtualVolumeSize" -H  
"Authorization: Bearer xxxxx"
```

Information: For more information, see [Sort function](#).

17.4 RESTful API Pagination

You can select and limit the data that is provided by the RESTful API. The syntax is **limit=x&offset=y** behind the query command. The **limit** specifies the number of results in the collection, and the **offset** specifies the starting point of the collection. The offset starts from 0. So, **limit=10&offset=9** specifies 10 results in places 10–19.” Both limit and offset always need to be used together. See Example 17-13.

Example 17-13 RESTful API call with GET request for dataClasses with limited results

```
$ curl -k -s -v "https://host_name/api/v2/dataClasses?limit=1&offset=2" -H  
"Authorization: Bearer xxxxx"
```

You can also use limit and offset with sort function as shown in Figure 17-4 on page 859.

Example 17-14 RESTful API GET request for dataClasses with limited results sorted by virtualVolumeSize

```
$ curl -k -s -v  
"https://host_name/api/v2/dataClasses?limit=2&offset=2&sort=virtualVolumeSize" -H  
"Authorization: Bearer xxxxx"
```

Information: For more information about pagination, see <https://www.ibm.com/docs/en/ts7700-virtual-tape/5.4.1?topic=api-restful-pagination>

17.5 Content type

You can set the “Content-Type” to specify the format of response data. The RESTful API for the TS7700 supports the following formats:

- ▶ application/json (default)
- ▶ text/csv

To get the response in CSV format, the GET request needs to specify **Content-Type: text/csv**.

Information: For more information about CSV response format, see [Overview](#).

17.6 API Description Document

This section describes how to make use of the TS7700 RESTful API description document.

An API description document is a machine-readable specification of a RESTful API, which in this context is the API of the TS7700. You can save or download this document directly from the TS7700, and it can be used as a reference to create a program and to parse RESTful API responses. They are language-neutral, so you can benefit from them no matter what language you choose.

17.6.1 Query for the Description Document

The TS7700 RESTful API description document is retrieved in YAML format by using the GET command. YAML is a human-readable data serialization language that is often used for writing configuration files. Depending on whom you ask, YAML stands for “yet another markup language” or “YAML ain’t markup language” (a recursive acronym), which emphasizes that YAML is for data, not documents.

Example 17-15 shows how to use curl again to retrieve this document of the TS7700 RESTful API.

Example 17-15 RESTful API GET request to receive the API description document

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json" ${hydraurl}/api/docs
```

Note: This request does not contain the version (“v1”) in the request URL: “/api/docs” versus “/api/v1/<request>”.

The system response and output of this GET command can be saved and piped into a file as shown in Example 17-16.

Example 17-16 RESTful API GET request to save the API description document into a YAML file

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json" ${hydraurl}/api/docs > TS7700RESTAPI.yml
```

The saved output (in this case the file `TS7700RESTAPI.yml`) can be easily viewed with a standard text editor like *vi*. Figure 17-5 shows the structure of the self-created TS7700 API description document.

```
openapi: 3.0.0
info:
  title: TS7700 RESTful API
  description: "TS7700 REST API OpenAPI definition. To be used by documentation generation\
  \ tools to display the API, code generation tools to generate servers and clients\
  \ in various programming languages, testing tools and more."
  version: 1.0.0
externalDocs:
  description: This document provides information about Representational State Transfer
  (RESTful) API used for the TS7700.
  url: https://www.ibm.com/docs/en/ts7700-virtual-tape/5.3?topic=reference-restful-api
servers:
- url: http://deadpool.tuc.stglabs.ibm.com:80
- url: https://deadpool.tuc.stglabs.ibm.com:443
tags:
- name: Authentication
  description: Services for authentication and user token generation
- name: Categories
  description: Services for querying and managing the categories.
- name: Constructs
  description: "Services for querying and managing Storage Classes, Management Classes,\
  \ Data Classes and Storage Groups."
- name: Data Classes
  description: Services for querying and managing the Data Classes.
- name: Grid
  description: Services for querying and managing the grid.
- name: Management Classes
  description: Services for querying and managing the management classes.
- name: Storage Classes
  description: Services for querying and managing the storage classes.
- name: Clusters
  description: Services for querying and managing the clusters.
- name: Storage Groups
  description: Services for querying and managing the storage groups.
paths:
  /token:
    post:
      tags:
      - Authentication
      summary: Creates a new token for using in all RESTful services
      description: Used to authenticate user and password valid for the currently
      assigned security policy and obtain back a limited time token that then can
      be used to access all services
      operationId: getToken
      requestBody:
        description: User name and password
        content:
          application/json:
            schema:
              $ref: '#/components/schemas/Credentials'
      responses:
        "200":
          description: Authentication succeeded.
          content:
            application/json:
              schema:
                type: array
                items:
                  $ref: '#/components/schemas/Token'
        "400":
          description: Invalid credentials were provided.
          content:
            application/json:
              schema:
                type: array
                items:
                  $ref: '#/components/schemas/Error'
      servers:
      - url: "https://{hostname}/api/{basePath}"
        description: the production API server
TS7700RESTAPI.yml (unix) (20:55 15/03/2023)
```

Figure 17-5 The TS7700 RESTful API description document as a YAML file

17.7 Error Handling

This section, includes an overview about the status codes and possible return codes of the TS7700 RESTful API calls.

The TS7700 RESTful API responds to each request with a response status code.

Information: For more information about Response status codes, see <https://www.ibm.com/docs/en/ts7700-virtual-tape/5.4.1?topic=api-error-handling>

Example 17-17 shows the response for a GET command after the access/session token has expired.

Example 17-17 RESTful API request by using an expired session token

```
$ curl -s -k -H "Authorization: Bearer ${hydratok}" -H "Content-Type: application/json" ${hydraurl}/api/v1/clusters {"message":"The token has expired, is invalid or missing.,"statusCode":401}
```

The returned status code is 401. This status code tells you that the token has expired, is invalid or missing. A new session token must be obtained by using the process that is explained in 17.2, “Access Token” on page 860.



IBM TS7700 support for zTape Air-GAP

This chapter provides information about TS7700 support for zTape Air-Gap (FC 5995).

This chapter includes the following topics:

- ▶ 18.1, “Overview” on page 872
- ▶ 18.2, “Description” on page 872
- ▶ 18.3, “TS7700 Attachment to Tape drives and Library” on page 872
- ▶ 18.4, “Methods Supported” on page 873
- ▶ 18.5, “User Interface” on page 873
- ▶ 18.6, “Reference” on page 873

18.1 Overview

The TS7700 zTape Air-Gap FC 5995 offers z Mainframe customers a cost-effective solution by using LTO8¹ tape drives in a TS4300 Tape Library for copying logical volumes from the TS7700 to physical tape for backup and copying logical volumes from physical volumes into the TS7700 for restore. This feature might be useful for clients that are mainly using physical tape media today with the IBM 3592 Model C07 tape control units and still need to export and exchange physical tape media with a non-VTS, native format.

TS7700 zTape Air-Gap is also referred to as, Simple Tape Attach (STA).

18.2 Description

The STA feature can be added to only a TS7700 disk-only solution and provides new methods for writing and reading to physical (LTO) tape media in a TS4300 library. This approach is an alternative solution to the TS7700 Tape Attach feature and copy export, and fits within a single rack. FC 5995 must be ordered and installed on this TS7700. In addition to providing a license key, this Feature Code delivers further more an attachment kit and setup instructions for the IBM SSR. The TS4300 Tape Library with its LTO drives and the physical tape media must be configured and ordered separately.

This feature is designed for limited and specific use cases and uses a set of new LIBRARY REQUEST commands (STAxxxx) to initiate and to track status. Similar to copy export, the logical volumes that are backed up using the new STA method are also copied to physical tape media enabling the logical volumes that were backed up to continue to be accessible in the TS7700.

The physical tape media that are created by the STA backup capability are only useful by the STA restore capability. The physical tape media that are used for restoring must be generated by using the STA backup capability.

Media encryption can be used, but would be configured by the TS4300 library and the LTO drives, external to the TS7700 (LME, Library Managed Encryption).

There is no corresponding host software that is needed for this solution; however, this solution can take full advantage of existing tape management system software (for example, z/OS DFSMSrmm and its stacked volume support). Check with your tape management system vendor for equivalent-type support.

18.3 TS7700 Attachment to Tape drives and Library

The LTO tape drives in the TS4300 are attached to the TS7700 by using one FCP HBA port on the TS7700 per drive. The drives that are installed in the library are configured as a control path for communication to the library. There can be one or two tape drives installed. The second tape drive is used for redundancy only.

¹ At the time of writing, only LTO8 tape drives are supported

18.4 Methods Supported

The use cases that are supported by this feature are backup, restore, and import of new data into the TS7700 disk cache. All existing functions in the TS7700 remain as they are today.

Here is a list of library request (LI REQ) commands available:

Backup (STABACK) Write logical volumes in the TS7700 disk cache to LTO physical media, stacking the logical volumes.

Restore (STAREST) Restore all logical volumes to the TS7700 disk cache from the physical volume specified.

Import (no dedicated STA command)

The same as the preceding STAREST method. The only difference in restore and import use cases is whether restoring data that was created locally (restore) or remotely (import).

Status (STASTAT) Return the status of the current or prior backup/restore, or the status of the library and drives.

Cancel (STACANC) Request the current backup/restore method to be cancelled.

18.5 User Interface

The user initiates the STA methods from the z/OS host by using the LIBRARY REQUEST command or from the TS7700 Management Interface using the same LIBRARY REQUEST command. For IBM Z (VM, VSE and TPF), the STA methods must be initiated through the TS7700 Management Interface.

18.6 Reference

For more information about STA, see [TS7700 zTape Air-Gap FC5995 Users Guide](#). The white paper covers more details on STA methods, host messages, installation, DFSMSrmm Stacked Volume Support, System-Managed Tape, and scenarios/use cases.



Part 4

Appendixes

This part offers management and operational information for your IBM TS7700.

This part contains the following appendixes:

- ▶ Appendix A, “Feature codes and requests for price quotations” on page 877
- ▶ Appendix B, “IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments” on page 883
- ▶ Appendix C, “JES3 examples and information” on page 899
- ▶ Appendix D, “DEVSERV QLIB command” on page 919
- ▶ Appendix E, “Sample job control language” on page 923
- ▶ Appendix F, “Library Manager volume categories” on page 951
- ▶ Appendix G, “IBM TS7700 parameter examples” on page 959
- ▶ Appendix H, “Extra IODF examples” on page 979
- ▶ Appendix I, “Case study for logical partitioning of a two-cluster grid” on page 991
- ▶ Appendix J, “Configuring externally managed encryption” on page 1013



A

Feature codes and requests for price quotations

This appendix lists all feature codes (FC) and requests for price quotations (RPQs) that are related to the installation of the IBM TS7700.

Exception: This appendix provides a general description of the FCs and where those FCs apply. For more information about feature code details, see [MT 3948 feature details](#).

This appendix includes the following topics:

- ▶ “Feature codes” on page 878
- ▶ “Requests for price quotations” on page 881

Feature codes

This section lists feature code descriptions for the TS7700 according to component, machine type, and model.

Clarification: The symbol “†” indicates that the specific feature was withdrawn.

3952/3948 F07 features

The 3952/3948 F07 includes the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 1776, RoHS (US only)
- ▶ FC 1903, PDU (Power Distribution Unit) †
- ▶ FC 1904, Optional AC, switching PDU †
- ▶ FC 1912, Single Phase PDU †
- ▶ FC 1913, Single Phase PDU (also to be chosen for 3 Phase Power, RPQ 8B3723 3 Phase Power Cable must be added)
- ▶ FC 2725, TS3000 System Console Rackmount † (all TS3000 related FC's are now an integrated part of the 3948-VED ship group. Since 3Q 2023 they don't show up anymore in the eConfig file.)
- ▶ FC 2748, Optical drive †
- ▶ FC 2749, F07 Door Lock and Key
- ▶ FC 2750, Top Exit Cabling
- ▶ FC 5512, TS3000 System Console KVM (Keyboard, Video, and Mouse) †
- ▶ FC 5631, Plant Install 3957 VED
- ▶ FC 5663, Plant Install 3956 CSB
- ▶ FC 5664, Plant Install 3956 XSB
- ▶ FC 5665, Field Install 3956 XSB
- ▶ FC 5666, Field Install 3956 CSB
- ▶ FC 5667, Plant Install 3948 CFC
- ▶ FC 5668, Plant Install 3948 XFC
- ▶ FC 5669, Field Install 3948 XFC
- ▶ FC 7337, TS7700 Encryption Capable Base Frame
- ▶ FC 7338, TS7700 Encryption Capable Expansion Frame
- ▶ FC 7339, TS7700 Encryption Capable Base Frame F07
- ▶ FC 9323, Expansion Frame Attach
- ▶ FC 9336, 1st Expansion Frame Attachment
- ▶ FC 9337, 2nd Expansion Frame Attachment
- ▶ FC 9906, New order install
- ▶ 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 AG00, Shipping and Handling - No charge
- ▶ FC AGGA, Shipping and Handling - F07
- ▶ FC AGGU, Shipping and Handling - No charge (3956/3948-CSB)

Server features for 3957/3948-VED

This section lists the server features for 3957/3948-VED.

Clarification: The symbol “†” indicates that the specific feature was withdrawn.

3957/3948-VED Server features

The 3957/3948-VED includes the following 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 1038, 10 Gb Longwave Ethernet Adapter †
- ▶ FC 1039, 1 Gb Quad-port Copper Ethernet Adapter
- ▶ FC 1041, 10 Gb Dual Port Optical Longwave Ethernet Adapter - replaces FC 1038
- ▶ FC 1776, US RoHS Indicator
- ▶ FC 3401, Enable 2nd FICON port
- ▶ FC 3402, 16 Gb FICON (shortwave)
- ▶ FC 3403, 16 Gb FICON (longwave)
- ▶ FC 3479, 64 GB DDR 4 memory upgrade for VED
- ▶ FC 4015, Grid Enablement
- ▶ FC 4016, Remove Cluster from Grid
- ▶ FC 4017, Cluster Cleanup
- ▶ FC 4643, Rack Mount kit VED with CSB
- ▶ FC 4647, Rack Mount Kit CSB
- ▶ FC 4648, Rack Mount Kit XSB
- ▶ FC 4649, Rack Mount Kit VED with CFC
- ▶ FC 4650, Rack Mount kit CFC
- ▶ FC 4651, Rack Mount Kit XFC
- ▶ FC 4694, Initial Order XSB Customer Rack FC 5243, Quad-port 16 Gb FC HBA
- ▶ FC 5262, 20 TB Cache Enablement
- ▶ FC 5263, 100 TB Cache Enablement
- ▶ FC 5268, 100 MBps Increment
- ▶ FC 5270, Increased logical volumes
- ▶ FC 5271, Selective Device Access Control
- ▶ FC 5272, Enable disk encryption - Local Key Management
- ▶ FC 5273, Enable Tape Attach
- ▶ FC 5274, Enable 1 TB Pending Tape Capacity
- ▶ FC 5275, Additional Virtual Drives
- ▶ FC 5276, Disk Encrypt-External Key Mgr FC 5278, Enable Cloud Storage Tier
- ▶ FC 5279, 5 TB Pending Tape Capacity
- ▶ FC 5281, Secure Data Transfer
- ▶ FC 5282, DS8000 Object Store †
- ▶ FC 5283, TS7700 Advanced Object Store for DS8000 (replaces FC 5282)
- ▶ FC 5904, Remote Code Load Exception (SSR On-site Code Load) † (since R5.3 available exclusively as an Expert Care option)
- ▶ FC 5995, zTape Air-Gap
- ▶ FC 5999, High Performance Tape/cloud Controller
- ▶ FC 8080, 600 GB HDDs
- ▶ FC 8081, 775 GB SSDs
- ▶ FC 8083, 3.84 TB SSDs in VED
- ▶ FC 9000, Mainframe Attachment
- ▶ FC 9219, TS3500/TS4500 Attach
- ▶ FC 9268, 100 MBps Throughput - Plant

- ▶ FC 9277, External Disk Encryption Certificate Plant - PKCS 12 File-Plant
- ▶ FC 9339, Replaces Existing TS7700
- ▶ FC 9350, Plant Install
- ▶ FC 9700, No Factory Cables
- ▶ FC 9900, Tape Encryption Enablement
- ▶ FC 9904, Remote Code Load †
- ▶ FC AGKV Ship with R5.2 Machine Code †
- ▶ FC AGKW, Ship with R5.3 Machine Code
- ▶ FC AGKX, Ship with R5.4 Machine Code

Note: FC 3479 64 GB DDR 4 memory upgrade for model VED requires systems to be at machine code level R5.0 or higher. The minimum quantity is 0, and the maximum quantity that is allowed is 1.

For more information on the TS7700 VED server memory amount and recommendations, check the corresponding 3957-VED sales manual publication or contact your IBM technical sales representative.

Cache Controller Drawer features for 3956/3948-CSB

Clarification: The symbol “†” indicates that the specific feature was withdrawn.

3956/3948-CSB Encryption Capable Cache Controller Drawer

The 3956-CSB includes the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 7119, 120 TB (78.75 TB Usable) SAS
- ▶ FC 7405, Encryption CSB (USB flash drives (Four Pack))
- ▶ FC 9352, Plant install
- ▶ FC 9353, Field Merge CSB
- ▶ FC AG00, No Shipping and Handling
- ▶ FC AGGU, Shipping and Handling

Cache Expansion Drawer features 3956/3948-XSB

Feature Codes that are available for the Cache Drawer are listed by model in this section.

Clarification: The symbol “†” indicates that the specific feature was withdrawn.

3956/3948-XSB Encryption Capable Cache Expansion Drawer features

The 3956-CXB includes the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 7119, 120 TB (78.75 TB Usable) SAS
- ▶ FC 9354, Plant install
- ▶ FC 9355, Field Merge XSB
- ▶ FC AG00, No Shipping Handling
- ▶ FC AGGV, Shipping and Handling

Cache Controller Drawer features for 3956/3948-CFC and 3956/3948-XFC

This section covers the following features:

- ▶ 3956/3948-CFC SSD Cache Controller Drawer features
- ▶ 3956/3948-XFC SSD Cache Expansion Drawer features

3956/3948-CFC SSD Cache Controller Drawer features

The 3956-CFC includes the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 7122, 92.16 TB SSD Storage
- ▶ FC 7405, Encryption XFC (USB flash drives (four pack))
- ▶ FC 9352, Plant install

3956/3948-XFC SSD Cache Expansion Drawer features

The 3956-XFC includes the following features:

- ▶ FC 0983, TAA Compliance
- ▶ FC 7122, 92.16 TB SSD Storage
- ▶ FC 9354, Plant install
- ▶ FC 9355, Field Merge XFC

Requests for price quotations

This section describes the RPQs that are available for the TS7700 according to component, machine type, and model.

Server VED

The following RPQs are available:

Seven- and eight-cluster hybrid grid configurations are supported by RPQ only.

3952/3948 F07 RPQ

The following RPQs are available:

- ▶ RPQ 8B3670 Top Exit Cabling - TS7700 †
- ▶ RPQ 8B3722 3-Phase WYE Power PDU (only to be used for F07 MES if FC1912 PDU is present)
- ▶ RPQ 8B3723 3 Phase Power Cable
- ▶ RPQ 8B3721 Customer Supplied Rack (HDD Cache)
- ▶ RPQ 8B3749 Customer Supplied Rack (SSD Cache)
- ▶ RPQ 8B3669 Key Locks on Door †
- ▶ RPQ 8B3585 Front/Rear and Side Panel Locking Procedure



IBM TS7700 implementation for IBM z/VM, IBM z/VSE, and IBM z/TPF environments

This appendix describes 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 publications:

- ▶ *z/VM: DFSMS/VM Removable Media Services*, SC24-6278
- ▶ *IBM TS4500 Introduction and Planning Guide*, SC27-5990
- ▶ *z/VSE V6R2.0 Administration*, SC34-2692

This appendix includes the following topics:

- ▶ “Software implementation in z/VM” on page 884
- ▶ “Software implementation in z/VSE (Virtual Storage Extended)” on page 890
- ▶ “Software implementation in z/OS Transaction Processing Facility” on page 894
- ▶ “Implementing Outboard Policy Management for z/TPF” on page 897

Software implementation in z/VM

This section explains how to implement and run the TS7700 under z/VM. It covers the basics for software requirements, implementation, customization, and platform-specific considerations about operations and monitoring. For more information, see *IBM TS4500 Introduction and Planning Guide*, SC27-5990.

Software requirements

The IBM TS7700 is supported by IBM z/VM V6R4 or later. With z/VM, the TS7700 models are transparent to host software. z/VM V6R4, or later, is required for guest and native VM support that provide base CP functions.

TS7700 multi-cluster grid environments disaster recovery

Introduced by VM65789 is the ability for the RMS component of DFSMS/VM to use the COPY EXPORT functions 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, see *z/VM: DFSMS/VM Removable Media Services*, SC24-6278.

For DR tests that involve a TS7700 grid that is connected to hosts that are running z/VM, a keyword in the **DRSETUP** command that is called SELFLIVE is available. This keyword provides a DR host the ability to access its self-created content that was moved into a write-protected category when flash is enabled.

For more information, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide Version 5.4](#).

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 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.

IBM Tape Manager for z/VM(5697-J08) can provide the TMS functions, including the management of one catalog of tape volumes across multiple z/VM LPARs. For more information and documentation, see the [Tape Manager website](#).

For more information about third-party TMSs that support the TS7700 in the z/VM environment, see *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.

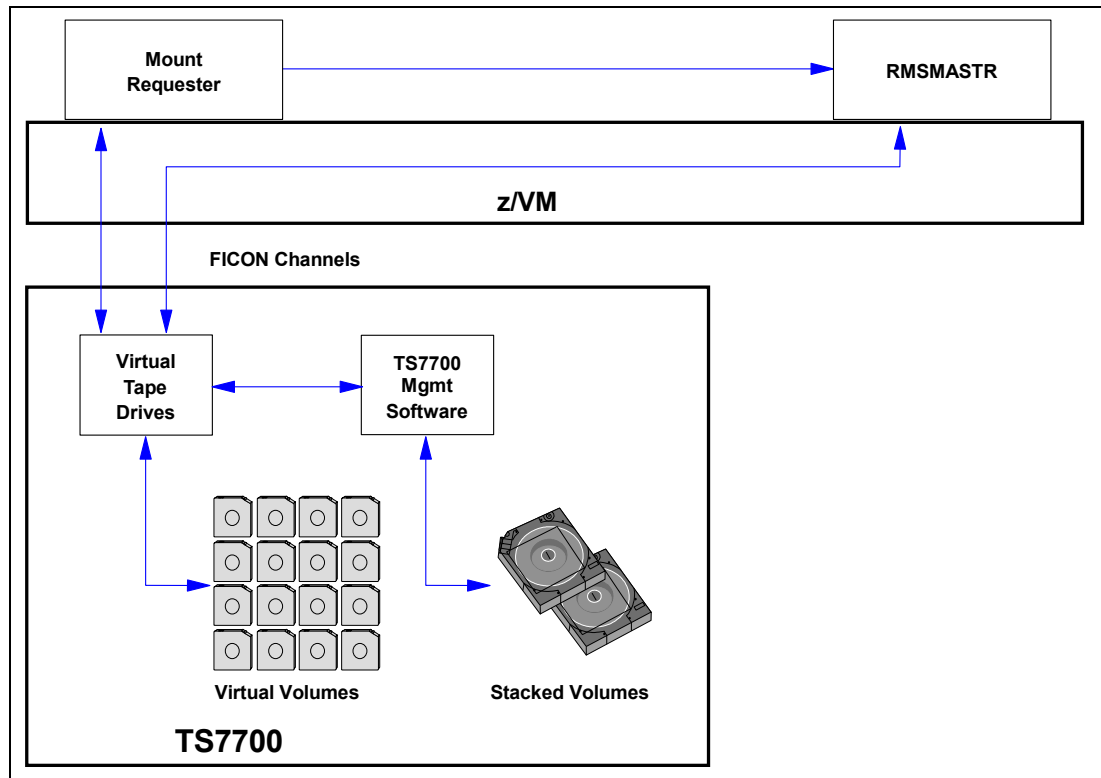


Figure B-1 TS7700 in a native z/VM environment that uses DFSMS/VM

When you use the TS7700C or TS7700T in a z/VM environment, consider that many z/VM applications or system utilities use specific mounts for scratch volumes. With specific mounts, when a mount request is sent from the host, the logical volume might need to be recalled from the stacked cartridge or Object Storage in the cloud if it is not in the Tape Volume Cache (TVC). Instead, you might want to use a larger cache-resident partition on one of more of the clusters in the grid such that it can contain all volumes z/VM can use as workload or other methods to avoid recall of data.

In addition, also consider that your z/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.

Some products were updated to improve scratch management. For example, IBM Tape Manager for z/VM moves free tape volumes to the scratch category, which avoids the recall of unnecessary data on a scratch mount request.

For more information, see [IBM Tape Manager for z/VM and TS7700: Managing Scratch Volumes to Free Cache](#).

DFSMS/VM

After you define the new TS7700 tape library through HCD, you must define the TS7700 to DFSMS/VM if the z/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. For more information, see *DFSMS/VM Removable Media Services*, SC24-6278.

You can access RMS as a component of DFSMS/VM. To enable RMS to run automatic insert bulk processing, you must create the RMBnnnnn data file in the VMSYS:DFSMS CONTROL directory, where nnnnn is the five-character tape library sequence number that is assigned to the TS7700 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.

Implementing Outboard Policy Management for z/VM

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 cannot identify the construct names and cannot change them.

In addition, non-z/OS hosts use multiple Library Manager (LM) categories for scratch volumes. They can also use multiple logical scratch pools on the Library Manager, as listed in Table B-1.

Table B-1 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'

Because the z/VM hosts do not know about constructs, they ignore static construct assignment. The assignment is also kept, even when the logical volume is returned to scratch. *Static assignment* means that logical volumes are assigned construct names when they are inserted. Construct names can also be assigned later.

To implement Outboard Policy Management for z/VM hosts that are 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 11.2.2, “TS7700 definitions” on page 578. You can assign the required static construct names during the insertion, as shown at the bottom part of the window in Figure B-6 on page 898.
3. On the left side of the MI, click **Virtual** → **Virtual Volumes** → **Insert Virtual Volumes**. The window that is shown in Figure B-2 on page 887 opens. Use the window to insert virtual volumes. Select the **Set Constructs** option and enter the construct names.

Figure B-2 Insert logical volumes by assigning static construct names

4. If you want to modify VOLSER ranges and assign the required static construct names to the logical volume ranges through the change logical volume function, select **Logical Volumes** → **Modify Logical Volumes**.

Define groups of logical volumes with the same construct names that are assigned and during insert processing, direct them to separate volume categories so that all volumes in one LM volume category feature 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.

z/VM guest support

This section describes two host environments with which you can 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 z/VM user ID under which the z/OS guest operating system is started for the first time:

```
STDEVOPT LIBRARY CTL
```

z/OS guests

The STDEVOPT statement specifies the optional storage device management functions that are 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 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](#).

z/VSE guests

Some VSE TMSs require VGS support and 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 use 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 that uses 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 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-3 shows the flow and connections of a TS7700 in a z/VSE environment under a VM.

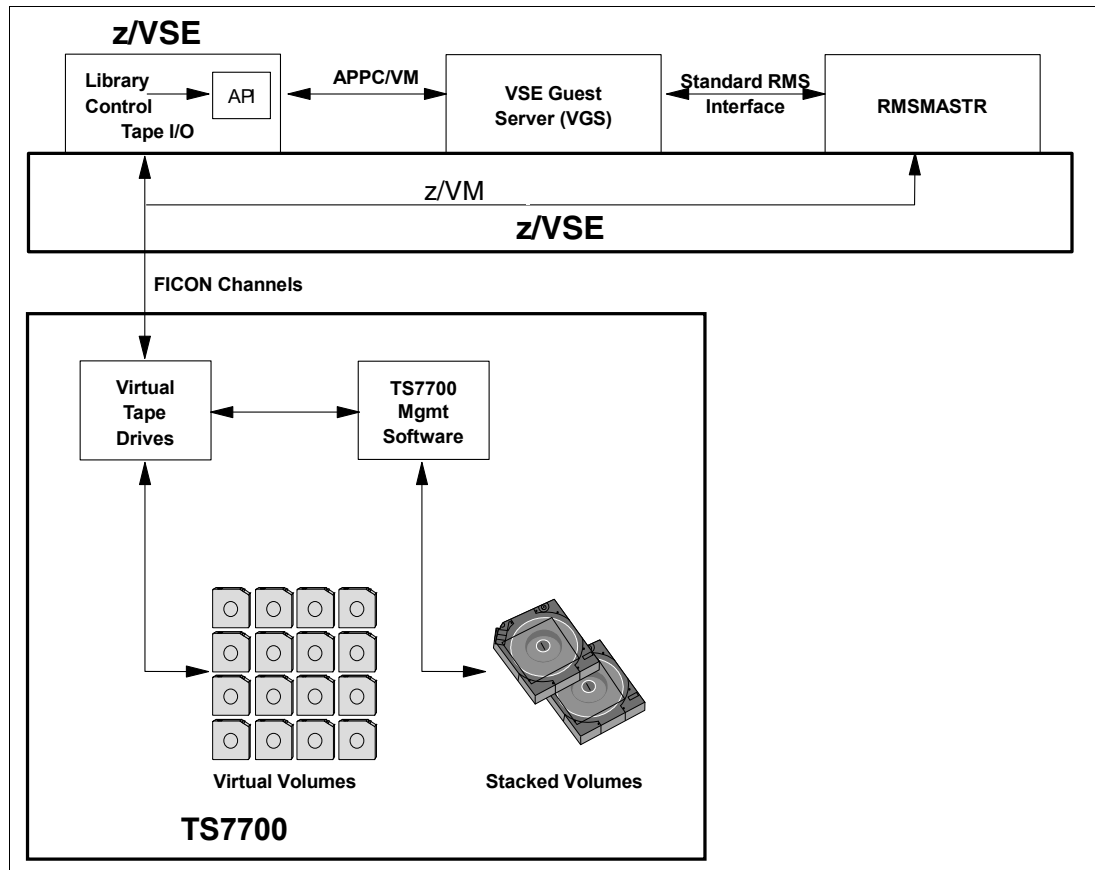


Figure B-3 TS7700 in a z/VSE environment as a VM guest

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-4 on page 890 shows CA-DYNAM/T VSE.

VSE uses original equipment manufacturer (OEM) tape management products that support scratch mounts. Therefore, if you use VSE under VM, you can benefit from the use of the scratch (Fast Ready) attribute for the VSE library's scratch category.

For more information about z/VSE, see *z/VSE V6R2.0 Administration*, SC34-2692.

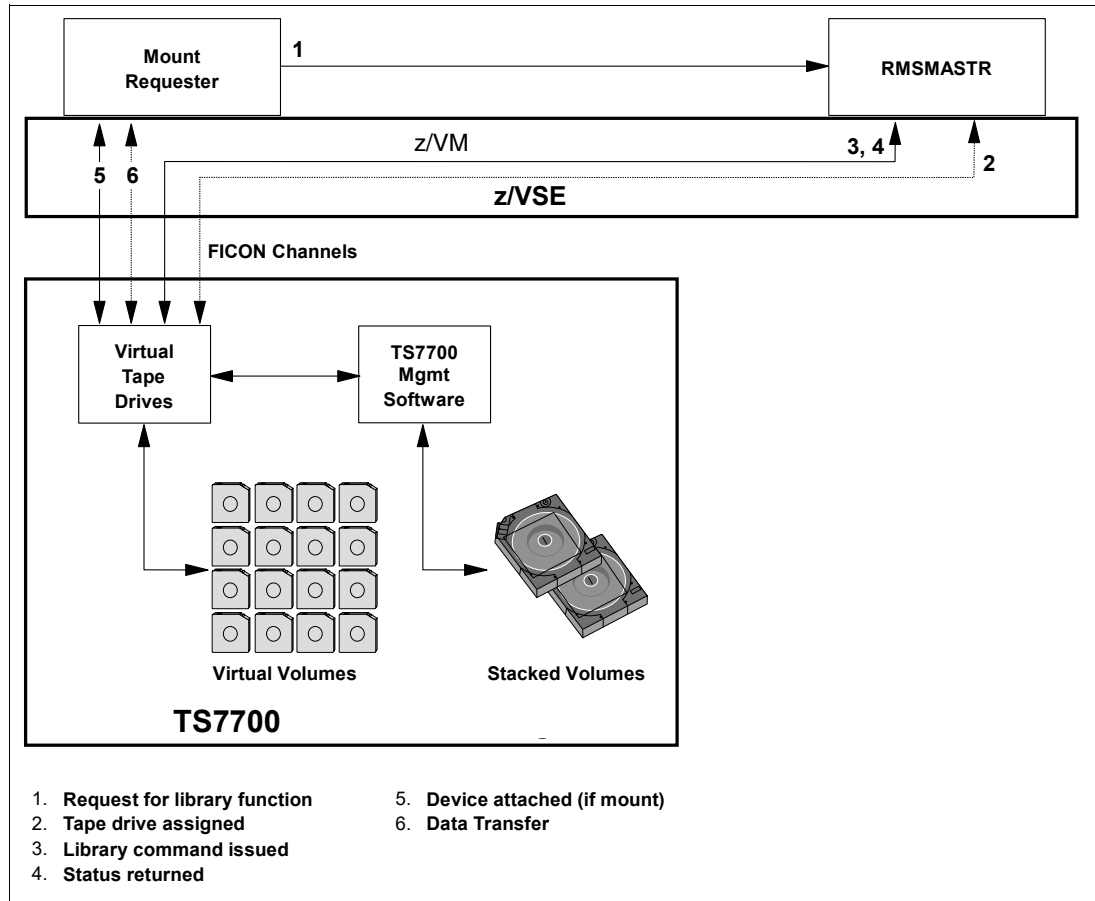


Figure B-4 TS7700 in a z/VSE environment as a VM guest (no VGS)

Software implementation in z/VSE (Virtual Storage Extended)

This section explains how to implement and run the IBM TS7700 under z/VSE. It covers the basics for software requirements, implementation, customization, and platform-specific considerations about operations and monitoring. For more information, see *IBM TS4500 Introduction and Planning Guide, SC27-5990*.

Software requirements

The TS7700 is supported by IBM z/VSE V5.2 or later. With z/VSE, TS7700 is transparent to host software. z/VSE supports the TS7700 as a stand-alone system in transparency mode. z/VSE supports a single node or multi-cluster grid and Copy Export and Logical Write Once Read Many (LWORM).

Native z/VSE

Native support is provided for the stand-alone grid TS7700 configuration in z/VSE Version 5.2 and higher that support all IBM TS1160, TS1150, TS1140, TS1130, TS1120, and 3592-J1A configurations without APARs in all automation offerings. This support includes TS3500 and TS4500 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 was 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 that is used 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 TS1160/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 that uses a VSE Guest Server” on page 888.

For native support under VSE (where TS7700 is used by z/VSE only), select **TLS**. At least one tape drive must be permanently assigned to the VSE.

Defining tape libraries

Next, define your tape library or libraries. This process is done by using 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 on page 892 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
Demount a volume	LIBSERV RELEASE,UNIT=480
Query count of volumes	LIBSERV CQUERY,LIB=TAPELIB1,SRCCAT= SCRATCH01
Query device	LIBSERV DQUERY,UNIT=480
Query inventory of library	LIBSERV IQUERY,LIB=TAPELIB1,SRCCAT=SCRATCH01
Query library	LIBSERV LQUERY,LIB=TAPELIB1
Manage inventory	LIBSERV MINVENT,MEMNAME=ALL,TGTCAT=SCRATCH01
Change category	LIBSERV SETVCAT,VOL=123456,TGTCAT=SCRATCH01
Query library for a volume	LIBSERV SQUERY,VOL=123456,LIB=TAPELIB1
Copy Export	LIBSERV COPYEX,VOL=123456,LIB=TAPELIB1

For more information, see *z/VSE V6R2.0 Administration*, SC34-2692, and *z/VSE System Macros Reference*, SC34-2638.

For DR tests that involve a TS7700 grid that is connected to hosts that are running z/VSE, a keyword on the DRSETUP command is available that is called SELFLIVE. This keyword provides a DR host the ability to access its self-created content that was moved into a write-protected category when flash is enabled.

For more information, see [IBM TS7700 Series z/OS Host Command Line Request User's Guide Version 5.4](#).

Implementing Outboard Policy Management for z/VSE hosts

z/VSE cannot provide SMS constructs to the 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.

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/VSE 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 listed in Table B-2.

Table B-2 Scratch pools and Library Manager volume categories

	Library Manager scratch categories	Number of scratch pools	Library Manager private categories
Native VSE	X'00A0'–X'00BF'	32	X'FFFF'

Because the z/VSE hosts do not know about constructs, they ignore static construct assignment. The assignment is also kept (even when the logical volume is returned to scratch).

To implement Outboard Policy Management for z/VSE hosts that are 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 11.2.2, “TS7700 definitions” on page 578. You can assign the required static construct names during the insertion, as shown at the bottom part of the window in Figure B-6 on page 898.
3. On the left side of the MI, click **Virtual** → **Virtual Volumes** → **Insert Virtual Volumes**. The window that is shown in Figure B-5 opens. Use the window to insert virtual volumes. Select the **Set Constructs** option and enter the construct names.

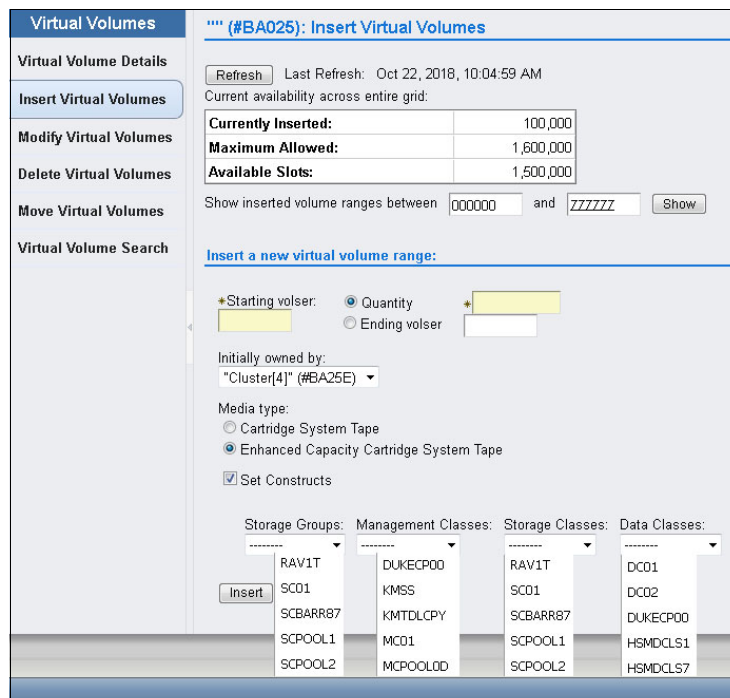


Figure B-5 Insert logical volumes by assigning static construct names

4. If you want to modify VOLSER ranges and assign the required static construct names to the logical volume ranges through the change logical volume function, select **Logical Volumes** → **Modify Logical Volumes**.

Define groups of logical volumes with the same construct names that are assigned and during insert processing. Direct them to separate volume categories so that all volumes in one LM volume category include identical constructs that are 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.

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 TS7700 models. 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 standard z/OS and TS7700 installation processes. For more information about leading practices for implementing the TS7700 with z/TPF, see [IBM Virtualization Engine TS7700 Series Best Practices - TPF Host and TS7700 IBM Virtualization Engine](#).

Software requirements

The IBM TS7700 is supported by IBM z/TPF V1.1 or later. With IBM z/TPF, the TS7700 models are supported in a single node and grid environment with the suitable software maintenance. The category reserve and release functions are not supported by the TS7700.

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 that is obtained from z/TPF. z/TPF features a set of commands (**ztp1f**) that you use to load the volumes in z/TPF-allocated virtual drives.

After a volume is loaded into a z/TPF drive, an automated solution must be in place that passes the volume serial number (VOLSER), tape data set name, and expiration date over to z/OS to process it automatically.

On z/OS, you must update the TMS' catalog and the TCDB so that z/OS can process virtual volumes that are created by z/TPF. Normal expiration processing applies after the z/TPF-written volumes are added to the z/OS TMS catalog and the TCDB. 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 that is 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 to 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 includes 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 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 z/TPF” on page 897.

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 anytime that a volume is loaded into a drive. Then, the user can specify, through the z/TPF user exit, whether the volume inherits 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 features only one operator interface with the TS7700, which is a z/TPF functional message that is 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 that are 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 started by z/TPF require all the clusters in the grid to communicate with one another. Under normal conditions, this communication occurs without delay and has 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 up to 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 elapses, 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 can access 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 (single-cluster or multi-cluster grid). This configuration prevents a single event on a composite library from affecting the primary and secondary devices.
- ▶ If the active and standby z/TPF devices and 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 ensures 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 z/TPF

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/TPF cannot identify the construct names, and cannot change them. 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.

Because the z/TPF hosts do not know about constructs, they ignore static construct assignment. The assignment is kept, even when the logical volume is returned to scratch.

To implement Outboard Policy Management for z/TPF hosts that are 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 11.2.2, "TS7700 definitions" on page 578. You can assign the required static construct names during the insertion as shown at the bottom part of the window in Figure B-6.
3. On the left side of the MI, click **Virtual** → **Virtual Volumes** → **Insert Virtual Volumes**. The window that is shown in Figure B-6 opens. Use the window to insert virtual volumes. Select the **Set Constructs** option and enter the construct names.

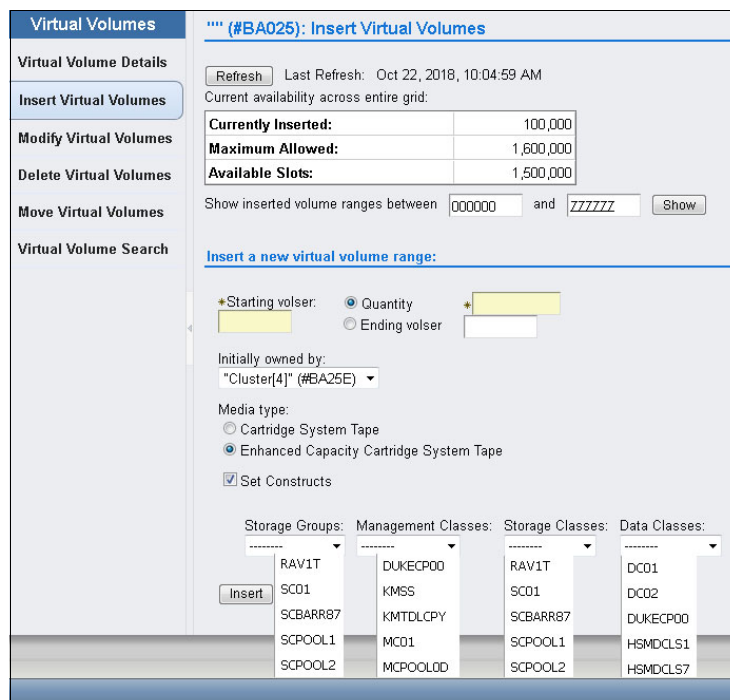


Figure B-6 Insert logical volumes by assigning static construct names

4. If you want to modify VOLSER ranges and assign the required static construct names to the logical volume ranges through the change logical volume function, select **Logical Volumes** → **Modify Logical Volumes**.

Define groups of logical volumes with the same construct names that are assigned and during insert processing. Direct them to separate volume categories so that all volumes in one LM volume category feature identically assigned constructs.

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 in which multiple tape libraries and tape devices, native devices in an IBM System Storage TS3500 tape library, and virtual devices in an IBM TS7700 are used. This appendix also describes the necessary parameters and considerations that must be installed in a job entry subsystem 3 (JES3) environment.

The following examples are used:

- ▶ 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” on page 900
- ▶ “Example with two separate tape libraries” on page 904
- ▶ “Example with three Tape Libraries” on page 906
- ▶ “Processing changes” on page 910

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 more 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.

JES3 or DFSMS cannot verify 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 is described 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 listed 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 listed 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

You can also use it 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 device 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, the 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 being defined 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 the following 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 **XTYPE**s within the referenced library.
- ▶ The complex-wide device type name must be included for all **XTYPE**s 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. 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 names are for 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 that are necessary for the first example

Table C-3 lists the LDG definitions that are needed in HCD. A total of eight esoterics must be defined.

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 (see Figure C-4). A naming convention for XTYPE is not mandatory, but it makes it simpler 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 the following 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 that is 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 that are shown 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 907 shows a JES3 configuration with three TS3500 Tape Libraries that are 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.

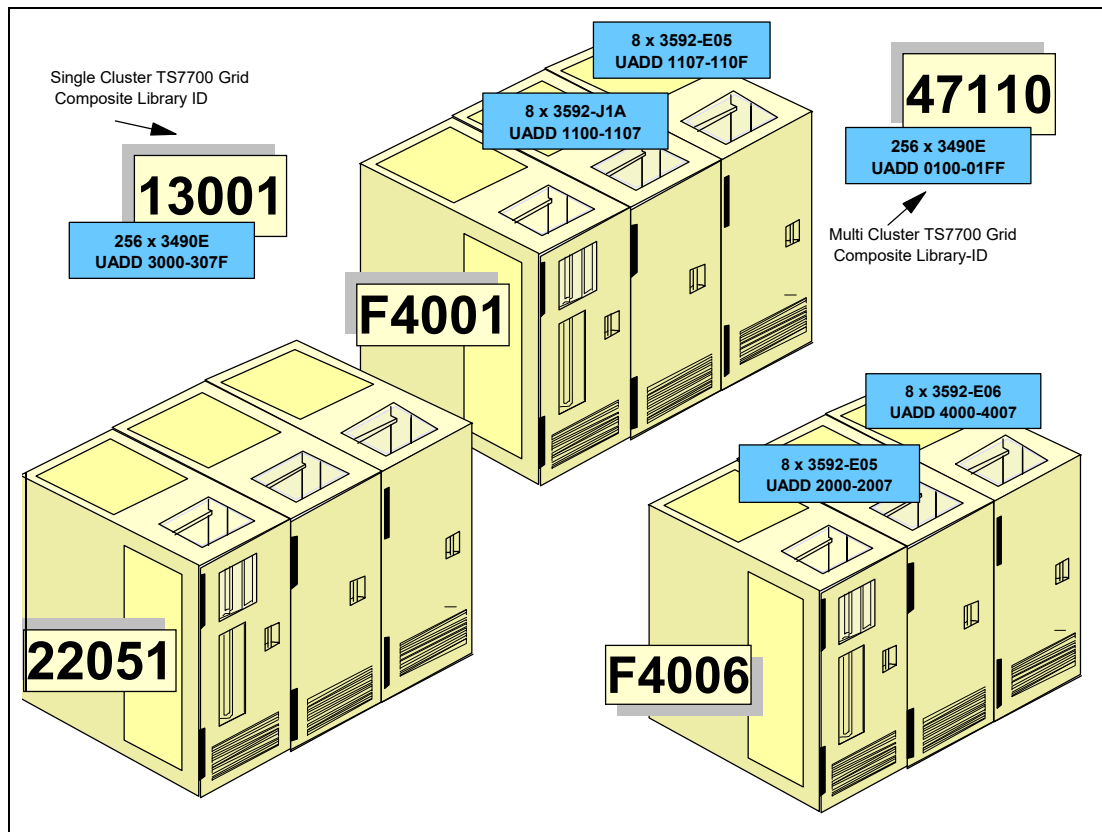


Figure C-7 Second JES3 configuration example

LDG definitions that are needed for the second configuration example

Table C-4 lists 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	Explanation
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 grids, 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	Explanation
Library-specific device type		One definition for each device type in each library, except for the multi-cluster TS7700 grid:
	LDE13001	<ul style="list-style-type: none"> ▶ Represents the virtual drives in the stand-alone cluster TS7700 grid in library 22051 ▶ Represents the virtual drives in the multicluster TS7700 grid in libraries F4001 and F4006
	LDE47110	<ul style="list-style-type: none"> ▶ Represents the 3592-J1A in library F4001 ▶ Represents the 3592-E05 in library F4001
	LDJF4001	<ul style="list-style-type: none"> ▶ Represents the encryption-enabled 3592-E05 in library F4006
	LDKF4001	<ul style="list-style-type: none"> ▶ Represents the 3592-E06 in library F4006
	LDLF4006	
	LDMF4006	

Device statement needed

Figure C-8 lists 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=N0** 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 shows all of 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

More examples

For more examples, see *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

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.
- ▶ 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 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 important 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.

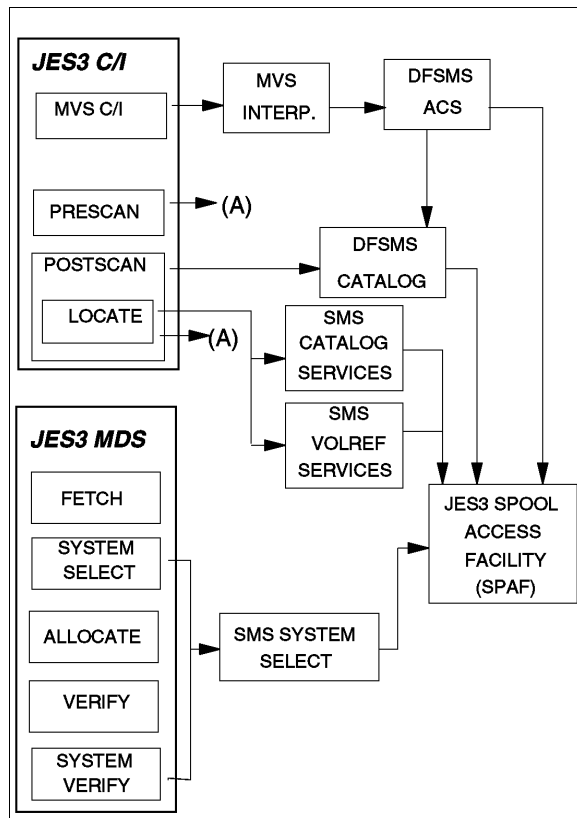


Figure C-11 JES3 CI and MDS processing phases

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 process 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 others) 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 others) 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 then.

MDS allocates an available drive from the available unit addresses for LDGW3495. It passes that device address to the 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 specification 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, "Implementing IBM TS7700" on page 239 and Chapter 13, "Monitoring" on page 667.

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.

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 because pre-execution and job execution can occur on different systems in the JES3plex. However, JES3 itself can be at an earlier 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, this support can be disabled either by a tape hardware specialist (PFE) dialing into the library or through the **MVS LIBRARY REQUEST** command. Verify that the DAA function is enabled at the library by using the following command:

```
LIBRARY REQUEST,composite-libraryname,SETTING
```

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 V2R2 because pre-execution and job execution can occur on different systems in the JES3plex. However, JES3 itself can be at an earlier release level.
2. Make JES3 INISH deck changes as described in the 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 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. 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 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 in this appendix.

For more information about these commands with examples, see *z/OS MVS System Commands*, SA22-7627.

The **DEVSERV QLIB** command can be used to complete the following 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 that are 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 <code>QLIB</code> displays 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. <code>LIST</code> uses the subparameters ACTIVE , INACTIVE , and QUEUE .
LISTALL	Produces a detailed list of all libraries, including the devices and port-IDs within each library. <code>LISTALL</code> uses the subparameters ACTIVE and INACTIVE .
LIBID	Indicates that the request is for a specific library. <code>LIBID</code> uses the subparameters ACTIVE , INACTIVE , VALIDATE , QUEUE , and DELETE .
DDDD	Indicates that the request is either for the library that contains the device <i>dddd</i> , or is for the device <i>dddd</i> itself. A subparameter is required when <code>DDDD</code> is specified. <code>DDDD</code> uses the subparameter SS .
DDR	Displays the limit on storage usage for a tape dynamic device reconfiguration (<code>DDR</code>) swap.
SS	Indicates that <code>QLIB</code> will send a diagnostic state save to the library containing device <code>DDDD</code> . This command is intended to be used at the request of IBM Support Center. For example, <code>SS</code> 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 <code>QLIB SS</code> 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 <code>DEVSUPxx PARMLIB</code> member. If not, an IPL reverts the category codes to what is specified in <code>DEVSUPxx</code> .

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 restart 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 activation. 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 in this appendix. 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**.
- ▶ Creating Volume Maps for logical volumes on tape or in object stores
- ▶ 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, see [IBM Software Tape Tools](#).

For more information about BVIR, see [TS7700 Bulk Volume Information Retrieval Function User's Guide](#).

This appendix includes the following topics:

- ▶ “BVIR jobs to obtain historical data” on page 924
- ▶ “Extra BVIR reporting” on page 929
- ▶ “VEHSTATS reports” on page 937
- ▶ “Creating Volume Maps for logical volumes on tape or in object stores” on page 948

BVIR jobs to obtain historical data

The following JCL can be used to request BVIR data:

- 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-1.
- 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-2 on page 927.

Note: BVIRHSTS can still produce System Measurement Facility (SMF) records from the BVIR historical data, though current VEHSTATS does not accept SMF records as input.

These jobs are also available as members in *userid*.IBMT00LS.JCL after you have installed IBMT00LS.exe on your host.

After you run one of two jobs, you can create various reports by using **VEHSTATS**. For more information, see “VEHSTATS reports” on page 937.

BVIRHSTU

Example E-1 shows the JCL in *userid*.IBMT00LS.JCL member BVIRHSTU.

Example E-1 BVIRHSTU JCL to obtain BVIR historical data

```
//ITS01JOB  CONSOLE,
// MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
// TIME=1440,REGION=2M
//*
/*JOBPARM  SYSAFF=*
//*
/**  BVIR DATA WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL
/**  DISMOUNT AND RE-MOUNT FOR READ.
/**
/**  IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/**  STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO REQUEST STATISTICS
/**  FROM EACH ONE.  USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/**
/**  THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
/**  HISTORICAL STATISTICS FROM THE VTS ASSOCIATED WITH THE VIRTUAL
/**  DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED ON THE
/**  VTS RECEIVING THE REQUEST.  THE FINAL OUTPUT IS WRITTEN TO A
/**  DISK DATA SET AS RECFM=U.
/**  HISTORICAL STATISTICS FOR ALL CLUSTERS IN A GRID ARE RETURNED
/**  FROM A SINGLE BVIR REQUEST TO ANY OF THE CLUSTERS.
/**
/******
/**  CAUTION CAUTION CAUTION CAUTION CAUTION CAUTION  *
/******
/**  IF ONE CLUSTER HAS BEEN DOWN FOR AN EXTENDED PERIOD, YOU NEED  *
/**  TO DIRECT THE BVIRREQ TAPE IN STEP2 TO A CLUSTER THAT HAS BEEN  *
/**  UP THE ENTIRE TIME BY USING A MGMTCLAS VALUE.  IF THE CLUSTER  *
/**  THAT WAS DOWN IS THE ONE PULLING THIS SET OF STATISTICS, ALL  *
```



```

//* CLUSTERS WILL SHOW NO ACTIVITY DURING THAT PERIOD. *
//*****
//*
//* NEXT, RUN VEHSTATS TO GET REPORTS.
//*
//PUTBVIR PROC USERHLQT=USERID, HI-LEVEL FOR TAPE DATA FILES
// TOOLHLQ=TOOLID, HLQ FOR LOAD AND CNTL
// SITE=SITENAME, 2ND LEVEL QUALIFIER
// GRIDID=GRID#, GRID SERIAL NUMBER TO BE PART OF DSN
//*
//* PARAMETER MC COULD BE USED TO ASSING THE TARGET CLUSTER
//* UNIQUE. OTHERWISE YOU MAY RECEIVE THE MESSAGE:
//* "CBR4171I MOUNT FAILED. LVOL=K12345, LIB=LIBVTS,
//* PVOL=?????,RSN=45."
//* THIS IS A KNOW ISSUE DESCRIBED HERE: ->
//* HTTPS://WWW.IBM.COM/SUPPORT/PAGES/NODE/6243136
// MC=#####, DIRECT TO SPECIFIC VTS OR CLUSTER
// UNIT=VTAPE UNITNAME ON THIS GRID
//*
//STEP1 EXEC PGM=IEFBR14
//DEL1 DD UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
// DSN=&USERHLQT..&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=&USERHLQT..&SITE..#&GRIDID..BVIRTAPE,
// MGMTCLAS=&MC,
// 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 IS MANDATORY. IF THE PARAMETER EDATE IS NOT SPECIFIED THEN
* HISTORCAL DATA ONLY FOR 1 DAY WILL BE PULLED OUT
*
*SDATE= 01FEB2009; USE HERE AS DDMONYEAR
*EDATE= 02FEB2009; USE HERE AS DDMONYEAR
SDATE= TODAY- 1; THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*SDATE= LASTWEEK; OR LASTMONTH WITH + OR - OPTIONS ALSO
*EDATE= TODAY;
*
*
//*
//*
//* COPY THE SECOND JOB TO THE MVS INTERNAL READER (SINCE MAY, 2017)
//*
//COPYJOB2 EXEC PGM=IEBGENER
//SYSIN DD DUMMY

```

```

//SYSPRINT DD SYSOUT=*
//SYSUT2 DD SYSOUT=(A,INTRDR) RUN THE JOBS
//SYSUT1 DD DATA,DLM=ZZ
//*JOB1
//*JOB2
//*JOB3
//*JOB4
/*JOBPARM SYSAFF=*
/*
//COPYBVIR PROC USERHLQ=USERID, HI-LEVEL FOR USER DATA FILES
// USERHLQT=USERID, HI-LEVEL FOR TAPE 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=&USERHLQT..&SITE..#&GRIDID..BVIRTAPE,
// DCB=(RECFM=U,BLKSIZE=30000),DISP=(OLD,DELETE)
//SYSUT2 DD DSN=&USERHLQ..&SITE..#&GRIDID..HSTU.D&SDATE..D&EDATE.,
// DCB=(RECFM=U,BLKSIZE=30000),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)
// DD *
*UTCMINUS= 07; ADJUST UTC TO LOCAL TIME WEST OF GREENWICH
*UTCPLUS= 02; ADJUST UTC TO LOCAL TIME EAST OF GREENWICH
*UTCAUTO; ADJUST UTC TO LOCAL TIME AUTOMATICALLY.
*
* UTCMINUS OR UTCPLUS WILL BE BUILT DEPENDING ON
* THE DIFFERENCE BETWEEN UTC AND THE MAINFRAME'S
* LOCAL TIME SETTINGS.
* NOTE. TAKE CARE OF THE DAYLIGHT SAVING TIME PERIOD -
* TIMESTAMPS BEFORE THIS DAY MAY BE OFFSET BY 1 HOUR.
/*
ZZ

```

BVIRHSTV

Example E-2 shows the JCL in *userid*.IBMT00LS.JCL member BVIRHSTV.

Example E-2 BVIRHSTV JCL to obtain BVIR historical data

```
//ITS01JOB CONSOLE,
//MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//TIME=1440,REGION=2M
//*
/*JOBPARM SYSAFF=*
/*
/* THIS IS A TWO JOB MEMBER TO ACCOMODATE EITHER JES2 OR JES3.
/* BVIR DATA WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL
/* DISMOUNT AND RE-MOUNT FOR READ.
/*
/* IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/* STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO REQUEST STATISTICS
/* FROM EACH ONE. USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/*
/* THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
/* HISTORICAL STATISTICS FROM THE VTS ASSOCIATED WITH THE VIRTUAL
/* DRIVE ADDRESS USED. THE BVIR FEATURE MUST BE ACTIVATED ON THE
/* VTS RECEIVING THE REQUEST. THE FINAL OUTPUT IS WRITTEN TO A
/* DISK DATA SET AS RECFM=VB.
/* HISTORICAL STATISTICS FOR ALL CLUSTERS IN A GRID ARE RETURNED
/* FROM A SINGLE BVIR REQUEST TO ANY OF THE CLUSTERS.
/*
/******
/* CAUTION CAUTION CAUTION CAUTION CAUTION CAUTION CAUTION *
/******
/* IF ONE CLUSTER HAS BEEN DOWN FOR AN EXTENDED PERIOD, YOU NEED *
/* TO DIRECT THE BVIRREQ TAPE IN STEP2 TO A CLUSTER THAT HAS BEEN *
/* UP THE ENTIRE TIME BY USING A MGMTCLAS VALUE. IF THE CLUSTER *
/* THAT WAS DOWN IS THE ONE PULLING THIS SET OF STATISTICS, ALL *
/* CLUSTERS WILL SHOW NO ACTIVITY DURING THAT PERIOD. *
/******
/*
/* NEXT, RUN VEHSTATS TO GET REPORTS.
/*
//PUTBVIR PROC USERHLQT=USERID, HI-LEVEL FOR TAPE DATA FILES
// TOOLHLQ=TOOLID, HLQ FOR LOAD AND CNTL
// SITE=SITENAME, 2ND LEVEL QUALIFIER
// GRIDID=GRID#, GRID SERIAL NUMBER TO BE PART OF DSN
/* PARAMETER MC COULD BE USED TO ASSING THE TARGET CLUSTER
/* UNIQUE. OTHERWISE YOU MAY RECEIVE THE MESSAGE:
/* "CBR4171I MOUNT FAILED. LVOL=K12345, LIB=LIBVTS,
/* PVOL=?????,RSN=45."
/* THIS IS A KNOW ISSUE DESCRIBED HERE: ->
/* HTTPS://WWW.IBM.COM/SUPPORT/PAGES/NODE/6243136
/* MC#####, DIRECT TO SPECIFIC VTS OR CLUSTER
// UNIT=VTAPE UNITNAME ON THIS GRID
/*
//STEP1 EXEC PGM=IEFBR14
//DEL1 DD UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
// DSN=&USERHLQT..&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=&USERHLQT..&SITE..#&GRIDID..BVIRTAPE,
// MGMTCLAS=MC,
// 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 IS MANDATORY. IF THE PARAMETER EDATE IS NOT SPECIFIED THEN
* HISTORICAL DATA ONLY FOR 1 DAY WILL BE PULLED OUT
*
*SDATE= 01FEB2009;      USE HERE AS DDMONYEAR
*EDATE= 02FEB2009;      USE HERE AS DDMONYEAR
SDATE= TODAY- 1;        THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*SDATE= LASTWEEK;       OR LASTMONTH WITH + OR - OPTIONS ALSO
*EDATE= TODAY;
*
*
//*
//*
//* COPY THE SECOND JOB TO THE MVS INTERNAL READER (SINCE MAY, 2017)
//*
//COPYJOB2 EXEC PGM=IEBGENER
//SYSIN DD DUMMY
//SYSPRINT DD SYSOUT=*
//SYSUT2 DD SYSOUT=(A,INTRDR) RUN THE JOBS
//SYSUT1 DD DATA,DLM=ZZ
//*JOB1
//*JOB2
//*JOB3
//*JOB4
*JOBPARM SYSAFF=*
//*
//COPYBVIR PROC USERHLQ=USERID, HI-LEVEL FOR USER DATA FILES
// USERHLQT=USERID, HI-LEVEL FOR TAPE 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=IEFBRI4
//DEL2 DD UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
// DSN=&USERHLQ..&SITE..#&GRIDID..HSTV.D&SDATE..D&EDATE
//*
//STEP3 EXEC PGM=CPYHIST
//STEPLIB DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//RECLIST DD SYSOUT=*
//SYSUT1 DD DSN=&USERHLQT..&SITE..#&GRIDID..BVIRTAPE,
// DCB=(RECFM=U,BLKSIZE=30000),DISP=(OLD,DELETE)
//SYSUT2 DD DSN=&USERHLQ..&SITE..#&GRIDID..HSTV.D&SDATE..D&EDATE.,
// DCB=(RECFM=VB,BLKSIZE=30000,LRECL=21996),UNIT=SYSDA,
// DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE)
// PEND
//*
//RUNPROC EXEC COPYBVIR,SDATE=YMMDD,EDATE=YMMDD HERE AS YMMDD
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
// DD *
*UTCMINUS= 07; ADJUST UTC TO LOCAL TIME WEST OF GREENWICH
*UTCPLUS= 02; ADJUST UTC TO LOCAL TIME EAST OF GREENWICH
*UTCAUTO; ADJUST UTC TO LOCAL TIME AUTOMATICALLY.
*
* UTCMINUS OR UTCPLUS WILL BE BUILT DEPENDING ON
* THE DIFFERENCE BETWEEN UTC AND THE MAINFRAME'S
* LOCAL TIME SETTINGS.
*
* NOTE. TAKE CARE OF THE DAYLIGHT SAVING TIME PERIOD -
* TIMESTAMPS BEFORE THIS DAY MAY BE OFFSET BY 1 HOUR.
//*
ZZ

```

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 cluster is not a TS7700T cluster, the following record is returned:

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

Example E-3 shows the JCL to obtain the Volume Map report, which is also contained in the *userid.IBMTTOOLS.JCL* member BVIRVTS.

Example E-3 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-3 on page 929 for the cache report by replacing the last statement (written in bold) with the statement that is shown in Example E-4. This statement creates a report for Cluster 0.

Example E-4 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 grid 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-3 on page 929. Replace the last statement (written in bold) with the statement that is shown in Example E-5, and update the following parameters:

- ▶ VTSID=
- ▶ MC=
- ▶ Also modify the **COPYAUDIT** request at the last line in *userid*.IBMT00LS.CNTL(BVIRAUD) to adjust for your request. Choose to specify **COPYMODE** or not and also choose **INCLUDE** or **EXCLUDE** with their cluster ID(s) in your environment.

Example E-5 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-6 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-6 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=IEFBR14
//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=643)
//SYSUT2 DD DSN=&USERHLQ..&SITE..&VTSID..MESFILE,
//      UNIT=SYSDA,SPACE=(643,(500000,200000),RLSE),
//      DISP=(,CATLG),DCB=(DSORG=PS,RECFM=U,BLKSIZE=643)
//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 shown in Example E-7.

Example E-7 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 cluster is not a TS7700T 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 (see Example E-8). The output of BVIRPHY can then be processed by using BVIRPRPT (see “Physical Volume and Pool Status Report Writer” on page 936) to produce a fully formatted report.

Example E-8 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 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 PHYSICAL VOLUMES BELONGING TO THE VTS ASSOCIATED WITH THE
/** VIRTUAL DRIVE ADDRESS USED. THE BVIR FEATURE MUST BE ACTIVATED
/** ON THE VTS RECEIVING THE REQUEST. THIS IS FOR TS7740 ONLY.
/**
/**BVIRPHY PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
/**      TOOLHLQ=TOOLID,      HLQ FOR LOAD AND CNTL
/**      SITE=SITENAME,      2ND LEVEL QUALIFIER
/**      MC=,      DIRECT TO SPECIFIC VTS OR CLUSTER
/**      VTSID=,      USE CLO, CL1, CL2, ETC TO BE PART OF DSN
/**      UNIT=VTAPE      UNITNAME ON THIS VTS
/**
/**STEP1  EXEC PGM=IEFBR14
/**DEL1   DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
/**          DSN=&USERHLQ..&SITE..&VTSID..BVIRPHY
/**DEL2   DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
/**          DSN=&USERHLQ..&SITE..&VTSID..PHYFILE
/**
/**STEP2  EXEC PGM=IEBGENER
/**SYSPRINT DD  SYSOUT=*
/**SYSUT1  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.CNTL(BVIRPHY)
/**SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..BVIRPHY,MGMTCLAS=&MC,
/**          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
/**          DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
/**SYSIN   DD  DUMMY
/**
/**STEP3  EXEC PGM=IEBGENER
/**SYSPRINT DD  SYSOUT=*
/**SYSUT1  DD  DSN=&USERHLQ..&SITE..&VTSID..BVIRPHY,DISP=(OLD,DELETE),
/**          DCB=(DSORG=PS,RECFM=U,BLKSIZE=420)
/**SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..PHYFILE,
/**          UNIT=SYSDA,SPACE=(420,(500000,200000),RLSE),
/**          DISP=(,CATLG),DCB=(DSORG=PS,RECFM=U,BLKSIZE=420)
/**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 (see Example E-9). The output of BVIRPLNN can then be processed by using BVIRPRPT to produce a fully formatted report. See "Physical Volume and Pool Status Report Writer" on page 936.

Example E-9 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=IEFBR14
//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=420),
//          DISP=(OLD,DELETE)
//SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..POOL&POOL,
//          DCB=(RECFM=U,BLKSIZE=420),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-10 shows the sample JCL.

Example E-10 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 VOLSERS
*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 **BVIRHSTU** and **BVIRHSTV**. The JCL, which was 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 (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-11 shows the sample JCL for VEHSTPO.

Example E-11 VEHSTPO sample JCL

```
//ITS01JOB 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 RETURN CODES (VERSION 16336-07.40 AND LATER):
//* 0 - COMPLETED NORMALLY;
//* 4  COMPLETED OK BUT:
//*     - SMALL ENOUGH NUMBER OF X33 RECORDS WITH IRREGULAR
//*       LAYOUT DETECTED. SEE THE MESSAGE IN JOBL0G;
//*     - WRONG ORDER DETECTED IN THE PROGRAM INPUT;
//*     - CONTROVERSIAL COMBINATION OF REPORT OPTIONS;
//* 8 - COMPLETED BUT BIG NUMBER OF X33 RECORDS WITH IRREGULAR
//*       LAYOUT DETECTED. SEE THE MESSAGE IN JOBL0G.
//* 16 - THE PROGRAM DETECTED CORRUPTED HISTORICAL RECORDS IN THE
//*       INPUT. SEE THE REPORT RECLIST FOR THE DETAILS.
//*     - NO ORDER PARAMETERS FOUND IN SYSCNTL
//*
//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=ORDERXFR,        SHOW JUST CACHE RELATED THROUGHPUT BY CLUSTER
//  ORDER=ORDERALL,        ALL AVAILABLE ORDER STATEMENTS
//  ORDER=ORDERC25,        WHEN THE 4 CLUSTERS ARE CL2, CL3, CL4, CL5
//  ORDER=ORDER6CL,        5 OR 6 CLUSTERS GRIDS (FOR GRAPHING PACKAGE)
//  ORDER=ORDER8CL,        7 OR 8 CLUSTERS GRIDS (FOR GRAPHING PACKAGE)
//  RECL=880,              880 IS WIDE ENOUGH FOR H32PDNN REPORTS
//  610 IS WIDE ENOUGH FOR H30TVCN REPORTS
//  H33GRID 372 IS WIDE ENOUGH FOR REPORT=GRID WITH 4 CLUSTERS
//  507 FOR 5 CLUSTERS, 642 FOR 6 CLUSTERS
//  348 IS WIDE ENOUGH FOR 22 CLUSTERS ON COMPARE REPORT
//  ADD 15 FOR EACH CLUSTER ABOVE 22
//  BLK=0,                 BLKSIZE FOR THE OUTPUT LIBRARY
//  ID=RUN1,                LAST NODE FOR REPORT FILE
//  SDATE=YYMMDD,          YYMMDD BEGINNING DATE  ACTIVATE IF USED
//  EDATE=YYMMDD,          YYMMDD ENDING  DATE  ACTIVATE IF USED
//  GRIDID=GRID#           ID FOR REPORTING SYSTEM
//
//DELETE EXEC PGM=IEFBRI4
//HOURFLAT DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ.&SITE.##&GRIDID..HOURFLAT.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//HOURFCLO DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ.&SITE.##&GRIDID..HOURFCLO.TXT
//HOURFCL1 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ.&SITE.##&GRIDID..HOURFCL1.TXT
//HOURFCL2 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ.&SITE.##&GRIDID..HOURFCL2.TXT
```

```

//HOURFCL3 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL3.TXT
//HOURFCL4 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL4.TXT
//HOURFCL5 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL5.TXT
//HOURFCL6 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL6.TXT
//HOURFCL7 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL7.TXT
//DAYHSMRY DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSMRY.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//DAYHSCL0 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL0.TXT
//DAYHSCL1 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL1.TXT
//DAYHSCL2 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL2.TXT
//DAYHSCL3 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL3.TXT
//DAYHSCL4 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL4.TXT
//DAYHSCL5 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL5.TXT
//DAYHSCL6 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL6.TXT
//DAYHSCL7 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL7.TXT
//WEKHSMRY DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSMRY.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//WEKHSCL0 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL0.TXT
//WEKHSCL1 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL1.TXT
//WEKHSCL2 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL2.TXT
//WEKHSCL3 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL3.TXT
//WEKHSCL4 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL4.TXT
//WEKHSCL5 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL5.TXT
//WEKHSCL6 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL6.TXT
//WEKHSCL7 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL7.TXT
//MNTHSMRY DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..MNTHSMRY.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//MNTHSCL0 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..MNTHSCL0.TXT
//MNTHSCL1 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..MNTHSCL1.TXT
//MNTHSCL2 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..MNTHSCL2.TXT
//MNTHSCL3 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..MNTHSCL3.TXT
//MNTHSCL4 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),

```

```

//          DSN=&USERHLQ. .&SITE. .#&GRIDID. .MNGTHCL4.TXT
//MNGTHCL5 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ. .&SITE. .#&GRIDID. .MNGTHCL5.TXT
//MNGTHCL6 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ. .&SITE. .#&GRIDID. .MNGTHCL6.TXT
//MNGTHCL7 DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ. .&SITE. .#&GRIDID. .MNGTHCL7.TXT
//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=IEFBRI4
//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
//H30TVC2 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC2
//H30TVC3 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC3
//H30TVC4 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC4
//H30TVC5 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC5
//H30TVC6 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC6
//H30TVC7 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC7
//H30TVC8 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30TVC8
//H30COMP DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H30COMP
//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
//H32PD01 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
//          DSN=&&H32PD01

```



```

//H32PD02 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H32PD02
//H32PD03 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H32PD03
//H32PD04 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H32PD04
//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
//H35CLOCL DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H35CLOCL
//H35CLOID DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H35CLOID
//H36OBJSG DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H36OBJSG
//H37CLOSN DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H37CLOSN
//H37OSNCL DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H37OSNCL
//H38OSNPT DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=H38OSNPT
//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
//HOURFLOW DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&HOURFLOW,DCB=(BLKSIZE=0)
//WEKHSMRY DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSMRY.TXT,
// DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL0 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL0.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL1 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL1.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL2 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL2.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL3 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL3.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL4 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL4.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL5 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL5.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL6 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL6.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*WEKHSCL7 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSCL7.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//MNTSMRY DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..MNTSMRY.TXT,
// DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL0 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL0.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL1 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL1.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL2 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL2.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL3 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL3.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL4 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL4.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL5 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL5.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL6 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),
//* DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL6.TXT,
//* DCB=(RECFM=FB,BLKSIZE=0)
//*MNTSCL7 DD UNIT=SYSDA,SPACE=(CYL,(1,1),RLSE),DISP=(MOD,CATLG),

```

```

/**      DSN=&USERHLQ..&SITE..#&GRIDID..MNTSCL7.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)
/**DAYHSCL0 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL0.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL1 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL1.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL2 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL2.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL3 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL3.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL4 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL4.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL5 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL5.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL6 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL6.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**DAYHSCL7 DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSCL7.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)
/**HOURFCL0 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL0.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL1 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL1.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL2 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL2.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL3 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL3.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL4 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL4.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL5 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL5.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL6 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL6.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
/**HOURFCL7 DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
/**      DSN=&USERHLQ..&SITE..#&GRIDID..HOURFCL7.TXT,
/**      DCB=(RECFM=FB,BLKSIZE=0)
//SORTIN  DD UNIT=(SYSDA,1),SPACE=(CYL,(300,100)),
//      DSN=&&SORTIN,DCB=(RECFM=VB,LRECL=30000,BLKSIZE=0)
//SORTOUT DD UNIT=(SYSDA,1),SPACE=(CYL,(300,100)),
//      DSN=&&SORTED,DCB=(RECFM=VB,LRECL=30000,BLKSIZE=0)

```

```

/**
/** COMMENTED ON 31OCT2017. SORT SHOULD WORK NORMALLY WITHOUT
/** THE TEMPORARY DATASETS BELOW. IF NOT - ACTIVATE THE DD-S BELOW
/**
/**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
//RPTSTEP.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
/**
/** THE ORDER STATEMENTS DETERMINE WHICH FIELDS WILL BE REPORTED IN THE
/** DAYSMRY, MONSMRY, HOURFLAT, DAYHSMRY, COMPARE, MNTHSMRY, WEKHSRMRY
/** REPORTS AND WHAT ORDER THEY WILL APPEAR IN.
/**
/** IBMTOOLS.JCL(ORDERV12) IS THE DEFAULT MEMBER
/** IBMTOOLS.JCL(ORDERALL) CONTAINS THE FULL LIST OF VALID ORDERS
/** YOU CAN CREATE YOUR OWN MEMBER WITH YOUR FIELDS AND SEQUENCE.
/** PICK AND CHOOSE FROM THE LIST(S) AND RE-ARRANGE TO FIT YOUR NEEDS.
// DD DISP=SHR,DSN=&USERHLQ..&SITE..IBMTOOLS.JCL(&ORDER)
// DD *
*****
* FILL IN THE FOLLOWING RECORDS AS APPROPRIATE: *
*****
CUSTOMER= TITLENAM VEHSTATS; 1-50 CHAR
*SPLITCLUSTERS; SPLIT FLAT FILES BY CLUSTERS (HOURFLAT, MNTHSMRY,
* DAYHSMRY, WEKHSRMRY)
* YOU SIMPLY NEED TO ACTIVATE THE DD STATEMENTS YOU NEED.
* (THE MESSAGES IEC130I WILL NOT BE ISSUED).
*
*EUROFORMAT; USE COMMA INSTEAD OF PERIOD FOR FRACTIONAL NUMBERS
*DATEFORM= J; DEFAULT FLAT FILE FORMAT IS DDMONYEAR
* USE THIS --> J=JULIAN, A=AMERICAN, E=EUROPEAN, OR I=ISO.
* TO GET ---> YEAR/DDD MM/DD/YEAR DD/MM/YEAR YEAR/MM/DD
*SINGLESPACE; USE SINGLE SPACE BETWEEN FIELDS IN FLAT FILES
*CSVDELIMITER= ','; USE THIS DELIMITER BETWEEN FIELDS IN FLAT FILES
*CSVDELIMITER= '|'; USE THIS DELIMITER BETWEEN FIELDS IN FLAT FILES
*CSVDELIMITER= 'S'; USE THIS TO SPECIFY ; AS DELIMITER IN FLAT FILES
*CSVDELIMITER= 'B'; USE THIS TO SPECIFY SPACE AS DELIMITER IN FLAT
* FILES
*ONEHEADING; ONLY ONE HEADING ON FLAT FILES, NOT BETWEEN CLUSTERS
*NOFILLER; DO NOT WRITE FILLR LINES TO DAYHSMRY
*SHOWVERSION; WRITE ID HEADER TO HOURFLAT FILE
*PRIPPOOL= 1 2 05;
*SECPPOOL= 15 25; DEFINE SECONDARY POOLS SO LVOLS DON'T GET
* COUNTED TWICE FOR ACTIVE_LVOLS FIELD
*
*
QUEAGEMINUTES; REPORT DEF and RUN QUEUE AGE AS MINUTES, NOT SECONDS
USEGB; FOR THE MOST OF THE COUNTERS THAT CONTAIN DATA AMOUNT
* THEIR VALUES ARE CONVERTED TO GIB
*NOINTEGER; ACTIVATE TO SHOW THE SMALL VALUES ABOVE IN GIB AS
* A NUMBER WITH ONE DECIMAL
*

```

```

REPORT= HRS HDSUM COM FLOW; HRS ROLL-UP, COMPARE, AND FLAT FILE SMRY
*   = QTR      REQUEST 15 MINUTE REPORTING AS GENERATED BY TS7740
*   = HRS      REQUEST HOURLY ROLL-UP REPORTING
*   = FLOW     REQUEST DATA FLOW BY CLUSTER - CAN'T USE WITH GRID
*   = GRID     SUMMARIZES ALL CLUSTERS BY GRID - CAN'T USE W/FLOW
*   = SHOP     SUMMARIZES ALL CLUSTERS WITHIN SHOP
*   = COMPARE  REQUEST SIDE BY SIDE CLUSTER COMPARISON
*   = HDSUM    DAILY SUMMARY FLAT FILE - HORIZONTAL 1 DAY/LINE
*   = HXFR     FOR HOURLY ON DEMAND TRANSFER REPORTING
*   = DXFR     FOR DAILY ON DEMAND TRANSFER REPORTING
*UTCAUTO;     ADJUST UTC TO LOCAL TIME AUTOMATICALLY.
*              UTCMINUS OR UTCPLUS WILL BE BUILT DEPENDING ON
*              THE DIFFERENCE BETWEEN UTC AND THE MAINFRAME'S
*              LOCAL TIME SETTINGS.
*              NOTE. TAKE CARE OF THE DAYLIGHT SAVING TIME PERIOD -
*              TIMESTAMPS BEFORE THIS DAY MAY BE OFFSET BY 1 HOUR.
*
*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
*
*   DEFAULT SDATE/EDATE ARE 01JAN1995/01JAN2035
*SDATE= 14JAN2013;  START DATE FOR OUTPUT REPORTING
*SDATE= TODAY- 1;   REPORT JUST YESTERDAY'S DATA
*   ALTERNATIVELY, 'TODAY', 'TODAY- NN', AND 'TODAY+ NN' MAY BE GIVEN
*   NN CAN'T EXCEED 360
*SDATE= LASTWEEK;   REPORT JUST LAST WEEK'S AVTIVITY
*SDATE= LASTMONTH;  REPORT JUST LAST WEEK'S AVTIVITY
*   ALTERNATIVELY, 'LASTMONTH', 'LASTMONTH- NN', 'LASTMONTH+ NN'
*   MAY BE GIVEN. NN CAN'T EXCEED 60
*STIME= 00:00;     START TIME FOR OUTPUT REPORTING
*EDATE= 21JAN2013;  END DATE FOR OUTPUT REPORTING
*EDATE= TODAY- 1;   REPORT JUST YESTERDAY'S DATA
*EDATE= LASTWEEK;   REPORT JUST LAST WEEK'S AVTIVITY
*EDATE= LASTMONTH;  REPORT JUST LAST WEEK'S AVTIVITY
*ETIME= 24:00;     END TIME FOR OUTPUT REPORTING
*
* IF YOU WANT TO LIMIT THE HOURFLAT TO A SUB-SET OF THE ENTIRE PERIOD.
*
*HRSDATE= 14JAN2013;  START DATE FOR HOURFLAT DAYS
*HRSDATE= TODAY- 1;   REPORT JUST YESTERDAY'S DATA
*HRSDATE= LASTWEEK;   REPORT JUST LAST WEEK'S AVTIVITY
*HREDATE= 14JAN2013;  END DATE FOR HOURFLAT DAYS
*HREDATE= TODAY- 1;   REPORT JUST YESTERDAY'S DATA
*HREDATE= LASTWEEK;   REPORT JUST LAST WEEK'S AVTIVITY
*
*SELECTDOW= FRI;     HOURFLAT: LIMITS OUTPUT TO JUST THIS DOW
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
LINES= 58;  LINES= 999 TO PUT DAYSMRY and MONSMRY ON SINGLE PAGE BREAK
*
*VTSNUM = SERNO;  SELECT JUST THIS CLUSTER TO MAKE IT EASIER TO WORK
*              WITH FLAT FILES AND GRAPHING PACKAGE
*              SELECTION IS PERFORMED AFTER THE REPLACEMENT BY
*              DLSER
*
*DLSER= FRSER TOSER;  CHANGE FROM ONE VALUE TO ANOTHER FOR REPORTS
*   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.
*
*GRIDSER= ????? TOSER;  CHANGE BINARY 0 TO NEW GRID SERIAL NUMBER
*      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.
*
SMFNUM = 194;  USER SELECTABLE SMF # FOR STATSMF DATA
*NOHOUR24;  ACTIVATE THIS PARAMETER, IF YOU DO NOT WANT TO CONVERT
*      TIME 00:00 TO 24:00 OF THE PREVIOUS DAY
*      (IF YOU WANT TO USE THE PREVIOUS (OLD) STYLE).
*      THE CONVERSION IS PERFORMED BY DEFAULT -
*      THE REPORTS ARE MORE ACCURATE IN THIS CASE
*      DOES NOT AFFECT ON HOURFLAT
*
* OBSOLETE - VALID FOR THE VERSION CREATED BEFORE 16336-07.40
*      THE VERSIONS BUILT AFTER 01JAN2017 IGNORE DEFDL
*      IF THE CLUSTERS IN YOUR GRID ARE NOT SEQUENTIAL, (0,1,2 ETC),
*      USE THE DEFDL PARAMETER TO DEFINE WHICH ONES ARE ACTUALLY PRESENT.
*      BROWSE THE BVIRHST FILE.
*DEFDL= H2909 0;  CLUSTER SERIAL (23-27 IN RECORD) AND CLUSTER NUMBER
*DEFDL= H2918 2;  CLUSTER SERIAL (23-27 IN RECORD) AND CLUSTER NUMBER
*DEFDL= H2906 3;  CLUSTER SERIAL (23-27 IN RECORD) AND CLUSTER NUMBER
*
/**
/**  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 IF YOU USED BVIRHSTU
/**STATSU DD DISP=SHR,
/**      DSN=&USERHLQ.&SITE.##&GRIDID..HSTU.D&SDATE..D&EDATE.
/**  ACTIVATE THE FOLLOWING IF YOU USED BVIRHSTV
/**STATSVB DD DISP=SHR,
/**      DSN=&USERHLQ.&SITE.##&GRIDID..HSTV.D&SDATE..D&EDATE.
/**  ACTIVATE THE FOLLOWING IF YOU USED BVIRHSTS and PULLED FROM SMF LOG
/**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
/**H30TVC2  EXEC  COPYRPTS,RPT=H30TVC2
/**H30TVC3  EXEC  COPYRPTS,RPT=H30TVC3
/**H30TVC4  EXEC  COPYRPTS,RPT=H30TVC4

```

```
//H30TVC5 EXEC COPYRPTS,RPT=H30TVC5
//H30TVC6 EXEC COPYRPTS,RPT=H30TVC6
//H30TVC7 EXEC COPYRPTS,RPT=H30TVC7
//H30TVC8 EXEC COPYRPTS,RPT=H30TVC8
//H30COMP EXEC COPYRPTS,RPT=H30COMP
//H31IMEX EXEC COPYRPTS,RPT=H31IMEX
//H32TDU12 EXEC COPYRPTS,RPT=H32TDU12
//H32TDU34 EXEC COPYRPTS,RPT=H32TDU34
//H32PD01 EXEC COPYRPTS,RPT=H32PD01
//H32PD02 EXEC COPYRPTS,RPT=H32PD02
//H32PD03 EXEC COPYRPTS,RPT=H32PD03
//H32PD04 EXEC COPYRPTS,RPT=H32PD04
//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
//H35CLOCL EXEC COPYRPTS,RPT=H35CLOCL
//H35CLOID EXEC COPYRPTS,RPT=H35CLOID
//H36OBSJG EXEC COPYRPTS,RPT=H36OBSJG
//H37CLOSN EXEC COPYRPTS,RPT=H37CLOSN
//H37OSNCL EXEC COPYRPTS,RPT=H37OSNCL
//H38OSNPT EXEC COPYRPTS,RPT=H38OSNPT
//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
//HOURFLOW EXEC COPYRPTS,RPT=HOURFLOW
```

Creating Volume Maps for logical volumes on tape or in object stores

This section provides a sample of JCL for a two-step job that can be used to create a map of the logical volumes and what physical stacked volumes they are on and then read it back. Example E-12 shows how to run a job that calls IEBGENER to request and call BVIR to write a volume map to an output data set.

Example E-12 JCL to add Volume Map

```
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=h1q.BVIR.CQUERY,
// UNIT=3490,LABEL=(,SL),
// DISP=(NEW,CATLG),
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
VOLUME MAP
/*
//STEP2 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT1 DD DISP=OLD,DSN=h1q.BVIR.CQUERY
//SYSUT2 DD SYSOUT=A,
// DCB=(DSORG=PS,RECFM=FB,LRECL=80,BLKSIZE=0)
/*
```

Note: for TS7760C use CLOUD VOLUME MAP instead of VOLUME MAP

Several JCL and REXX examples are provided that can help you with the TS7700 migration scenarios that were described earlier in this appendix.

Using EDGUTIL to validate tape configuration database inconsistencies

The JCL that is shown in Example E-13 can help you identify inconsistencies in the Removable Media Management (RMM) control data set (CDS) and the tape configuration database (TCDB).

Example E-13 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.

REXX EXEC to update the library name

Example E-14 shows a sample REXX EXEC program for updating the library name for volumes in the TCDB.

Example E-14 REXX EXEC for updating the library name in the TCDB

```
/* REXX */
/*****/
/* ALTERNVOL */
/*
/* Usage: ALTERNVOL DSN(volserlist) LIB(libname)
/*
/* Before this EXEC is run, you must create the
/* input data set "volserlist". The LISTCAT command
/* can be used to generate the list of volumes
/* to be altered to an output data set.
/*
/* LISTCAT VOLUMEENTRIES(V*)
/* LIBRARY(sourcelib)
/* OUTFILE(ddname)
/*
/* The list generated has the following format:
/* VOLUME-ENTRY----Vvolser
/* For command specifics, 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: ALTERNVOL 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
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

This appendix describes Library Manager volume categories and the platforms that use them. The IBM TS7700 does not use a Library Manager in the IBM 3494-style. Table F-1 lists all the default Library Manager volume categories, platforms that are used, and definitions. 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 0x0001–0xEFFF (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 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 used.
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.
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.

Category (in hex)	Used by	DFSMSrmm	Definition
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.
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 IBM 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.

Category (in hex)	Used by	DFSMSrmm	Definition
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.
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–EFFF	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.
F000–F00D	N/A	N/A	Reserved for future use.
F00E	BTLS	N/A	Indicates a volume in error. Volumes are assigned to the error category during unmount if the volume serial that is specified for unmount does not match the external label of the volume being unmounted.
F00F–FEFF	N/A	N/A	Reserved for future use.

Category (in hex)	Used by	DFSMSrmm	Definition
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	TS7700	N/A	Stacked Volume Insert category for 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 library partition.
FF02	TS7700	N/A	Stacked Volume Scratch category 0 for TS7700. This category is reserved for future use for scratch stacked volumes.
FF03	TS7700	N/A	Stacked Volume Scratch category 1 for TS7700. This category is used by the TS7700 for its scratch stacked volumes.
FF04	TS7700	N/A	Stacked Volume Private category for TS7700. This category includes both scratch and private volumes.
FF05	TS7700	N/A	Stacked Volume disaster recovery category for 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 TS7700 library partition and the Library Manager is in disaster recovery mode.
FF06	TS7700	N/A	This category is used by the TS7700 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	TS7700	N/A	This category is reserved for future hardware functions.
FF08	TS7700	N/A	This category is used by the TS7700 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	TS7700	N/A	Convenience-Eject category. When a tape volume is assigned to the convenience-eject category, it becomes eject pending and the Library 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 inventory. Tip: Logical volumes cannot be ejected from the library. They can be deleted or exported.
FF11	TS7700	N/A	Bulk-Eject category

Category (in hex)	Used by	DFSMSrmm	Definition
FF12	VTS	N/A	Export-Pending category for 3494/VTS, No longer used.
FF13	VTS	N/A	Exported category for 3494/VTS, No longer used.
FF14	VTS	N/A	Import category for 3494/VTS, No longer used.
FF15	VTS	N/A	Import-Pending category for 3494/VTS, No longer used.
FF16	VTS	N/A	Unassigned Category for 3494/VTS, No longer used.
FF17	TS7700	N/A	CopyExport-Hold category. Physical volumes are assigned to this category on completion of processing for a copyexport stacked volume.
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, are assigned to this category.
FF19	TS7700	N/A	Cloud logical Volume Version restore. TS7700 reserve a wanted number of volumes to be used for restore destination from the scratch category. TS7700 selects the specified number of scratch volumes from the scratch category and move them to a reserved category FF19 for restore. The list of the reserved volumes is notified to the host as the output of the LI REQ command.
FF20	TS7700	N/A	Corrupted-Token Volume Category. In a TS7700, volumes are assigned to this category when it is determined that the tokens that are associated with the volume have been corrupted. Prevents 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	TS7700	N/A	3592 Cleaner Volume Category. Cleaner volumes for 3592-type devices in the library are assigned to this category automatically.
FFF5	TS7700	N/A	3592 Service Volume Category. Volumes are assigned to this category by the Library when it detects that a volume has a unique service cartridge VOLSER and a media type compatible with a 3592 device.
FFF6	N/A	N/A	No longer used because the 3590 is not supported by TS7700
FFF7 and FFF8	N/A	N/A	Reserved for library. These categories are reserved for internal library functions.
FFF9	N/A	N/A	No longer used because the 3490 is not supported by TS7700
FFFA	TS7700	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.
FFFB	TS7700	N/A	Purge-Volume Category.

Category (in hex)	Used by	DFSMSrmm	Definition
FFFC	TS7700	N/A	Unexpected-Volume Category. This category is reserved for future use.
FFFD	N/A	N/A	No longer used because the 3590 is not supported by TS7700
FFFE	N/A	N/A	No longer used because the 3490 is not supported by TS7700
FFFF	TS7700	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 describes 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 or 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 that is produced in Deferred mode copy. Therefore, no specific examples need to be added.

This appendix includes the following topics:

- ▶ “General example setup” on page 960
- ▶ “General example setup for tape partitions” on page 975

General example setup

This appendix describes the following examples:

- ▶ 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 includes the following 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 the 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. When 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 situation 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 lists 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	

Feature	Where to find	Comment
Override settings: Force volumes to be mounted on this cluster in a 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. For more information about choosing the best configuration to meet your demands, see 2.5.1, “Homogeneous versus hybrid grid configuration” on page 106.

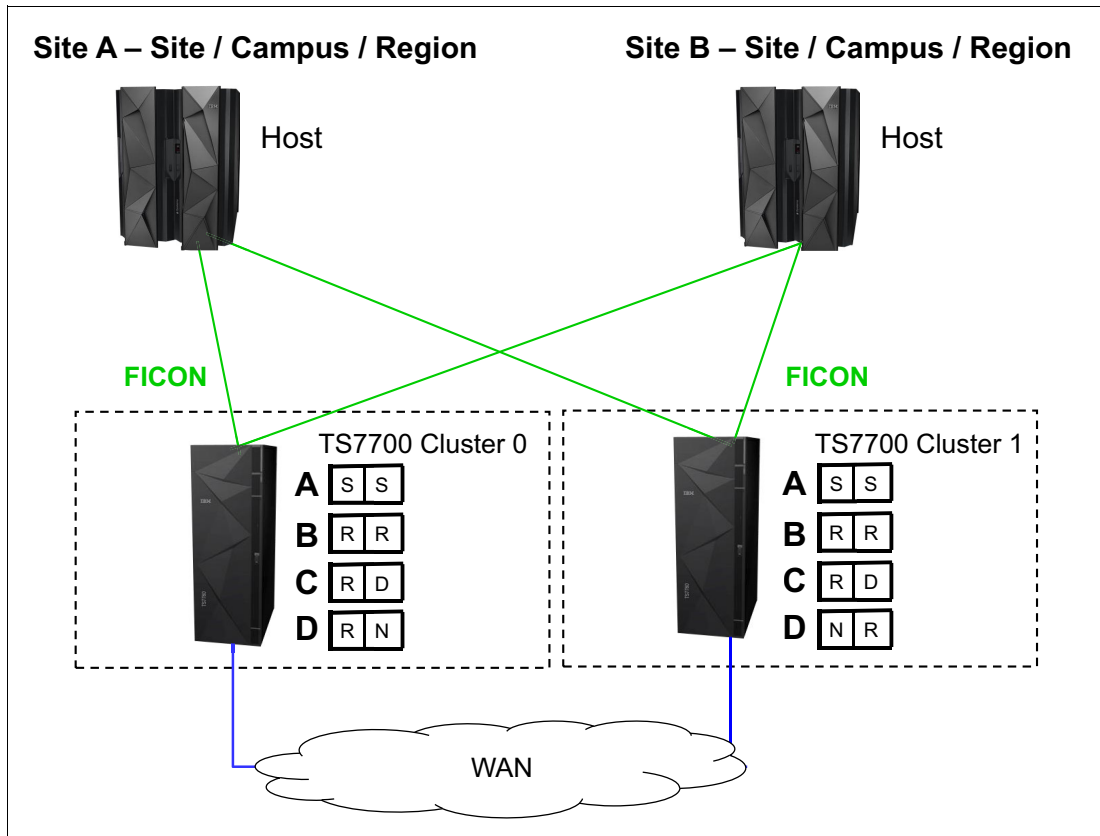


Figure G-1 Example of two-cluster grid for HA limited distance DR

Setting up the configuration

The environment of a two-cluster grid for HA and DR is listed 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 are connected to all clusters.
SAA	Disabled	
DAA for private volumes	Disabled	The default of DAA for private volumes 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 situation 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 fails.
- ▶ 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 likely is selected. However, there are still some cases when the TVC from Cluster 1 is 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 is selected. It 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: A 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 can 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 situation, 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 grid is a two-cluster grid, cluster families do not provide values 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 situation 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 listed 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 are 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	Set on both clusters.
Override settings: Prefer local cache for non-fast ready mounts	On	Set only on Cluster 1.
Override settings: Force volumes to be mounted on this cluster in a 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 the 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 the 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 the 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 the 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, which can result in a remote write to Cluster 1, and the data is 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 is 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 a 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 that is dictated by your company's requirements for DR. In a three-cluster grid configuration, HADR can also be achieved 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 968 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 shows a combined HA and DR solution for a three-cluster grid.

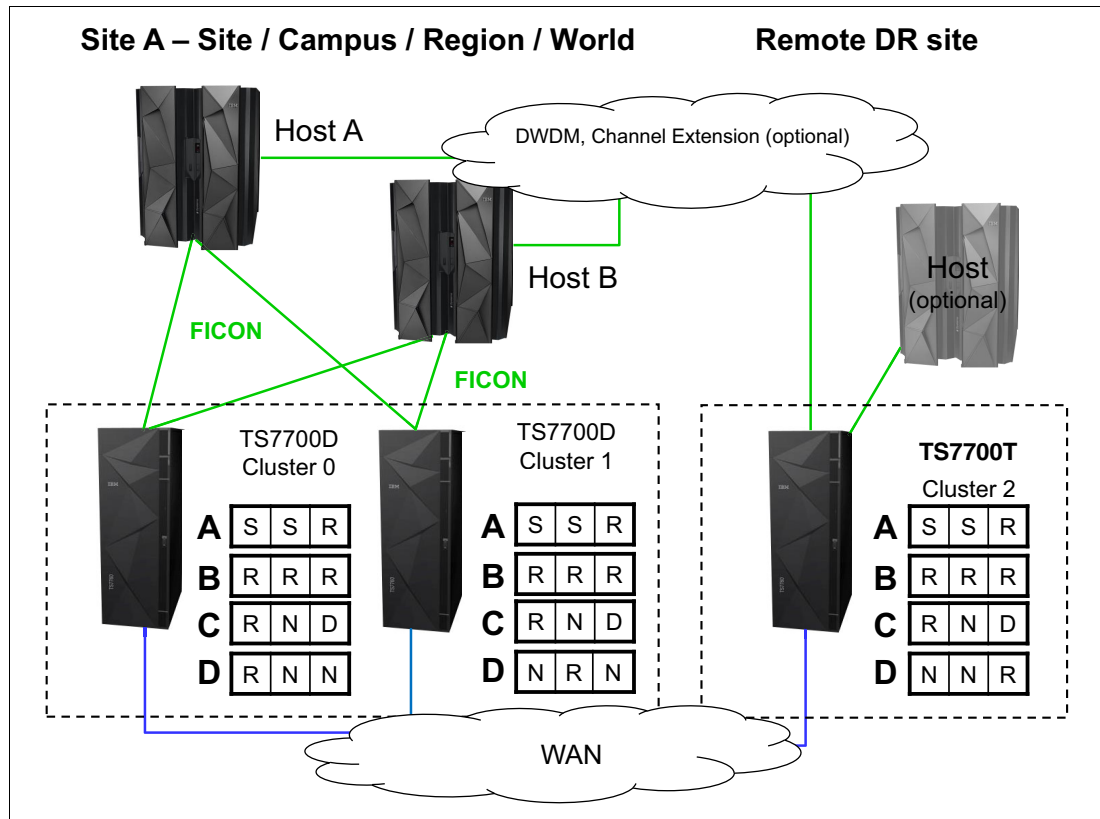


Figure G-3 Three-cluster HA and DR with two TS7700Ds and one TS7700T

Setting up the configuration

The parameter settings are listed 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 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 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 includes 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 is in Cluster 2 only.
- ▶ 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. It 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 964.

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).

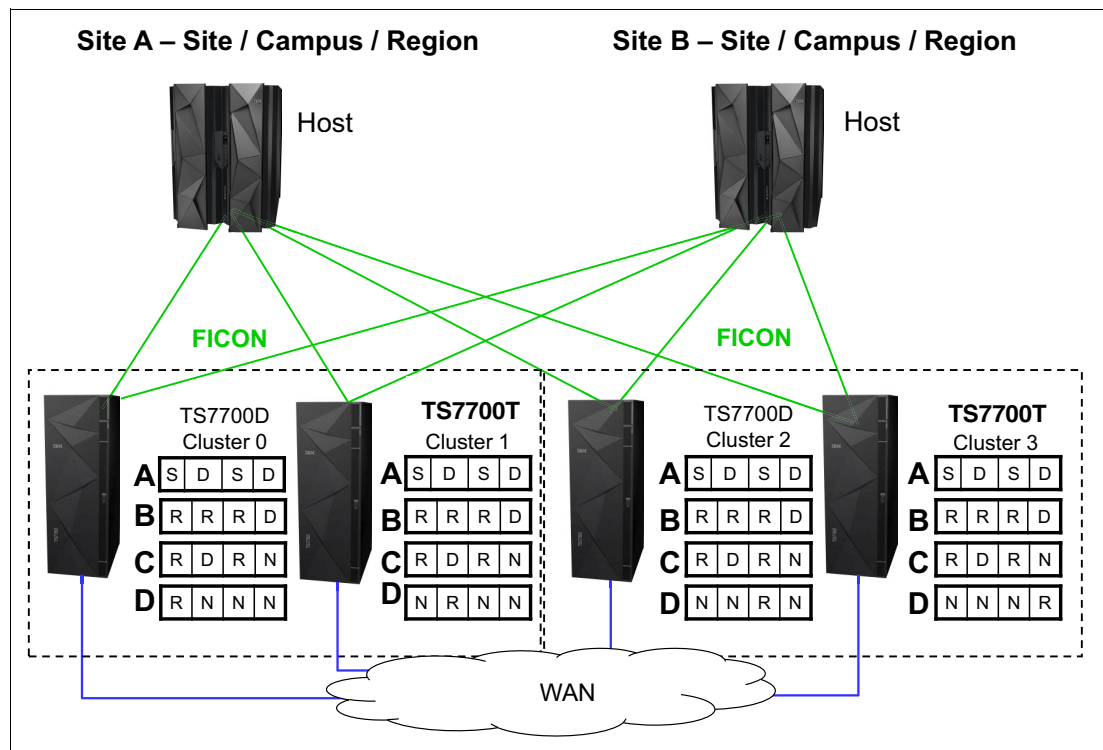


Figure G-4 Examples of a four-cluster grid for HA and DR

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.

Setting up the configuration

The parameter settings are listed 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 the 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. It is a single point of failure, which is not necessary in such a configuration.
- ▶ Cluster 1 has a definition of Deferred copy. 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.

The new configuration is shown in Figure G-5.

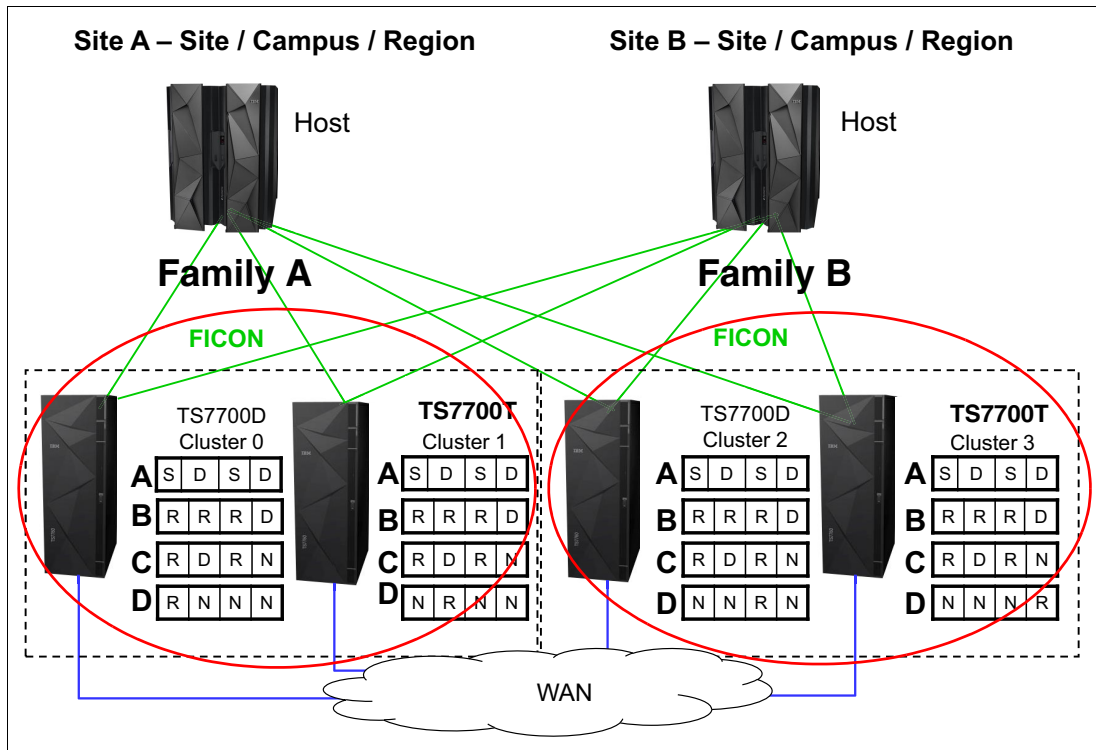


Figure G-5 Examples of a four-cluster grid for HA and DR with cluster families

Setting up the configuration

The parameter settings are listed 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	

Fact	Number/State	Comment
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, it can lead to the 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 lists 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 can be considered:

- ▶ All data in cache: No physical backend.
- ▶ All data is premigrated immediately to physical tape: No use of delay premigration.
- ▶ Only HSM ML2 is kept in cache; all other data is premigrated. Tape partitions are used.
- ▶ Delay premigration is used to expire data in the 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. Therefore, a throughput of 3 TB in an hour is necessary. Assume that this workload is a steady workload, and a compressed host I/O from 875 MBps is expected.

Basic considerations how to find the best configuration and setup

First, the customer defined some requirement, which are 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 whether this option is valid, 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 option is not a feasible solution.

Example 2: All data on physical tape is premigrated now with one tape partition

First, determine what the minimum requirement for the cache is, which should 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, calculate the following information:

- ▶ Db2 archive log= 1 TB a day = 10 days in cache = 10 TB
- ▶ Db2 image copy = 5 TB a day = 3 days in cache = 15 TB
- ▶ Other backup = 15 TB
- ▶ Other data = 0.8 TB
- ▶ HSM = Unpredictable

Adding the numbers $(10 + 15 + 15 + 0.2 + 0.5 + 0.1) = 40.8$ TB

A two-drawer configuration (approx. 63 TB) would also allow to have some HSM data in the cache.

However, this configuration has the following restrictions:

- ▶ It cannot handle the necessary Host I/O sustained throughput.
- ▶ Without Tape partitions, the cache can be controlled only with PG1/PG0, so the Db2 data and SMF data might not be in the cache as requested.

Therefore, this solution is not feasible.

Example 3: HSM ML2 is kept in cache only, all other data is premigrated, and tape partitions are used

To determine the configuration, first calculate the minimum cache capacity:

- ▶ CP0: HSM ML2, 150 TB + 10% free space, and contingency = 165 TB.
- ▶ CP1: Db2: 25 TB + 10% contingency = 28 TB (rounded).
- ▶ CP2: Other data = 15.8 TB. SMF data is treated with PG1. All other data is 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 14, “Performance considerations” on page 747), a seven-drawer configuration could provide the necessary sustained throughput, if no RUN or SYNC copies are produced.

If RUN or SYNC is needed (and both are recommended 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. An issue can result because 15 TB does not fit in the premigration queue. Therefore, do not use this configuration if RUN or SYNC is needed.

Regarding the premigration queue depth, we cannot 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, approximately 15 TB are written, which means that either two TS7700 needs to be installed or the customer must accept that premigration during the peak time is essential to not run in throttling. Also, 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 option might be feasible.

Example 4: Delay premigration is used to expire data in the cache

To determine the configuration, first calculate the minimum cache capacity:

- ▶ CP0: HSM ML2, 150 TB + 10% free space, and contingency = 165 TB.
- ▶ CP1: Db2: 25 TB + 10% contingency = 28 TB (rounded).
- ▶ CP2: Other data (backups with 30 days expires in cache)= 500 TB, + 10 TB for other data. SMF treated with PG1, all other data is treated as PG0.

Therefore, provide 703 TB, which results in 23 drawers. This number of drawers can run the host I/O in sustained mode, and also do some copy activity.

Because all data from the daily backup expires in the cache, the amount of FC 5274 for the premigration queue must 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 installing only 5 TB to cover the Db2 portion only.

In this example, 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 option 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 describes input/output definition file (IODF) and input/output configuration program (IOCP) examples for several different configurations.

Important: These examples are *not* preferred practices or suggested configurations to be adopted. They are for reference only and must be tailored to your environment and needs.

This appendix includes the following topics:

- ▶ “General IODF principles” on page 980
- ▶ “Using switches to connect to the control unit” on page 980
- ▶ “Directly connecting” on page 981
- ▶ “Upgrading to 8-Gb channels” on page 981
- ▶ “Adding more devices” on page 981
- ▶ “Sharing ports” on page 989
- ▶ “LIBPORT-IDs in the MVSCP” on page 990

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 (a path group of up to 8) that are defined to the processor.

A CHPID statement (see Example H-1) specifies the following information:

- ▶ The CSS of the LPARs with access to the CHPID
- ▶ Whether it is shared (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 two- 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 A 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=0
```

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=0
```

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, X
      PATH=(CSS(1),C1,C5,DA,F1), X
      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,
        PATH=(CSS(1),C1,C5,DA,F1),
        LINK=(CSS(1),1D,3D,4F,66),
        UNITADD=((00,16)),UNIT=3490,CUADD=7
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE170 IODEVICE ADDRESS=(E170,16),UNIT=3490,CUNUMBR=(0C18),
        UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/8
    CNTLUNIT CUNUMBR=0C19,
        PATH=(CSS(1),C1,C5,DA,F1),
        LINK=(CSS(1),1D,3D,4F,66),
        UNITADD=((00,16)),UNIT=3490,CUADD=8
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE180 IODEVICE ADDRESS=(E180,16),UNIT=3490,CUNUMBR=(0C19),
        UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/9
    CNTLUNIT CUNUMBR=0C1A,
        PATH=(CSS(1),C1,C5,DA,F1),
        LINK=(CSS(1),1D,3D,4F,66),
        UNITADD=((00,16)),UNIT=3490,CUADD=9
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE190 IODEVICE ADDRESS=(E190,16),UNIT=3490,CUNUMBR=(0C1A),
        UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/A
    CNTLUNIT CUNUMBR=0C1B,
        PATH=(CSS(1),C1,C5,DA,F1),
        LINK=(CSS(1),1D,3D,4F,66),
        UNITADD=((00,16)),UNIT=3490,CUADD=A
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1A0 IODEVICE ADDRESS=(E1A0,16),UNIT=3490,CUNUMBR=(0C1B),
        UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/B
    CNTLUNIT CUNUMBR=0C1C,
        PATH=(CSS(1),C1,C5,DA,F1),
        LINK=(CSS(1),1D,3D,4F,66),
        UNITADD=((00,16)),UNIT=3490,CUADD=B
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1B0 IODEVICE ADDRESS=(E1B0,16),UNIT=3490,CUNUMBR=(0C1C),
        UNITADD=00,PART=(CSS(1),MVSC7,VMT07)
*
*ELWOOD/C
    CNTLUNIT CUNUMBR=0C1D,
        PATH=(CSS(1),C1,C5,DA,F1),
        LINK=(CSS(1),1D,3D,4F,66),
        UNITADD=((00,16)),UNIT=3490,CUADD=C
*$HCDC$    DESC='ELWOOD CLUSTER 0 BARR74'
TAPEE1C0 IODEVICE ADDRESS=(E1C0,16),UNIT=3490,CUNUMBR=(0C1D),
        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 that results from these IOCP statements is shown in Example H-5.

Example H-5 MVSCP with 496 devices that are 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 987 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 maximize your investments. One way to maximize your investment is by using the same hardware for more than one sysplex/host.

Important points to consider are described in this appendix, 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.

This appendix also provides 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 some of the functions that are 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 *Selective Device Access Control* (SDAC). 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, which shield them from unintended or malicious interference between hosts. This process is accomplished by granting access to predetermined ranges of logical volumes by selected groups of devices in a logical control unit (LCU) granularity (also referred to as *LIBPORT-ID*).

It might be possible to use the Tape Management System (TMS) to partition a library by using software, depending on your TMS vendor. This process is referred to as *soft partitioning*, and is accomplished by implementing software controls on all attached z/OS hosts.

If your TMS is DFSMSrmm, the controls are the use of **PARTITION** and **OPENRULE** statements in the EDGRMMxx parmlib member on each LPAR to designate and separate volser ranges. An example is provided that can be used with or without SDAC. All attached hosts are configured correctly for the soft partitioning to work.

An example of a real implementation of this function is provided in this appendix, which describes 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 topics:

- ▶ “Overview of partitioning” on page 992
- ▶ “Definitions and settings in z/OS” on page 993
- ▶ “Definitions on the TS7700 Management Interface” on page 1001
- ▶ “Verification of changes” on page 1010

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. Protect against unauthorized user access for logical volumes on PROD from TEST and vice versa.

The function SDAC, which was 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 you are setting up a multient 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. These naming conventions make it simpler to logically relate all definitions and structures to one host when updates are needed.

Figure I-1 on page 993 shows the setup. Updates are needed on many places. Adapt your current naming standards to your setup.

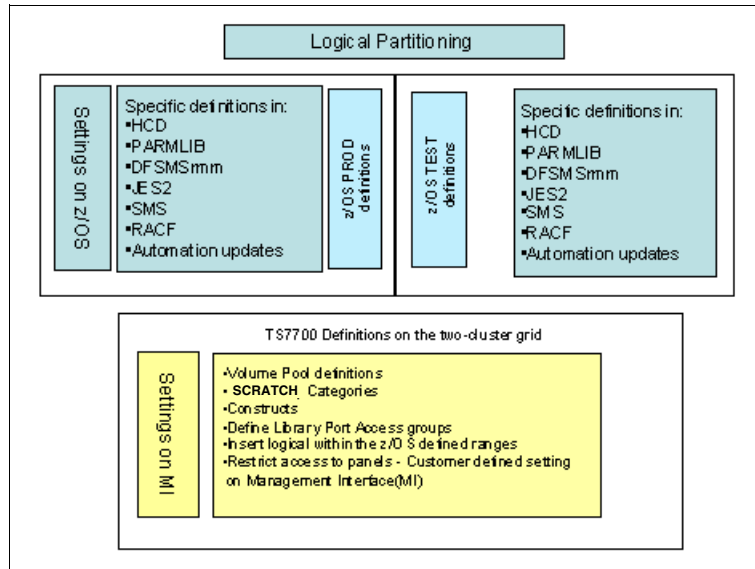


Figure I-1 Logical partitioning overview

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 (HCD) to define:
 - CTLUNITS for the 3490E devices.
 - FICON channels.
 - Esoteric definitions.
- ▶ Parmlib definitions to define for the following 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 function is valid only if a sysplex is used 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 with optional VLPOOL statements.
 - PRTITION and OPENRULE definitions to avoid or customize access and usage of volumes from the other host (replacement of REJECT).
- ▶ JES2 JCL Procedure Library updates
 - JCL definitions to enable start of an OAM-started task.

- ▶ SMS updates for constructs and ACS routines:
 - Definition of libraries and SMS constructs; that is, Storage Class (SC), Management Class (MC), and others.
 - 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 that are needed in the case study that define specific volume ranges, several device addresses, and scratch categories.

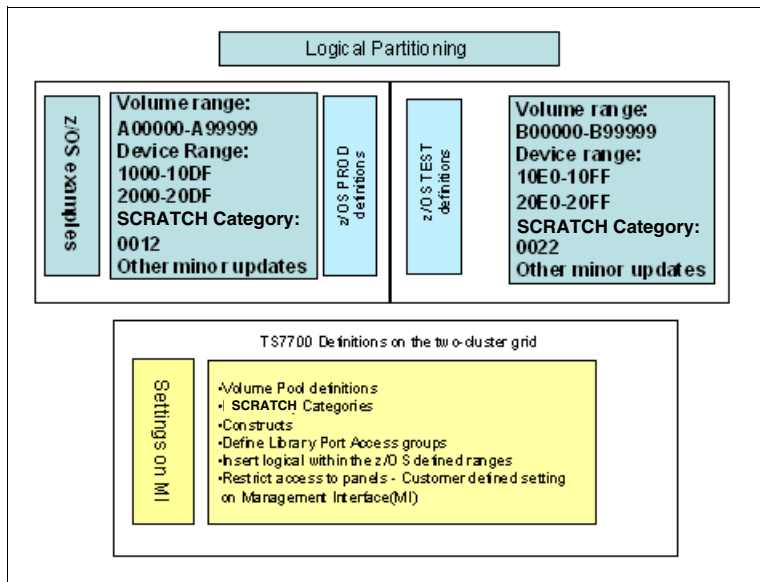


Figure I-2 Software updates with sample values on z/OS

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 features addresses 1000–10FF and for Cluster 1, 2000–20FF (see Table I-1 on page 995).

Table I-1 HCD definitions

Host	Cluster	CTLUNIT definition	LIBPORT-ID
PROD	Cluster 0	CUADD 00-0D (1000-10DF)	01-0E
	Cluster 1	CUADD 00-0D (2000-20DF)	41-4E
TEST	Cluster 0	CUADD 0E-0F (10E0-10FF)	0F-10
	Cluster 1	xUADD 0E-0F (20E0-20FF)	4F-50

More definitions are needed. Consider the following points:

- ▶ The devices must be connected to processors.
- ▶ Devices must be related to esoteric definitions.
- ▶ Definitions regarding the FICON infrastructure and FICON directors are not included.

Normally, you can activate the new definitions dynamically. For more information about HCD definitions, see 6.1, “Hardware configuration definition” on page 240.

Parmlib definitions

Use a 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 of the described parameters can be activated dynamically on the current release of z/OS.

For more information about options in a parmlib, definitions, and commands to activate without IPL, see the following publications:

- ▶ *z/OS MVS System Commands*, SA38-0666
- ▶ *z/OS MVS Initialization and Tuning Reference*, SA23-1380

The updates are within the following members of a 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 that are shown 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, the definitions that are shown 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 includes 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. 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 features the following requirements:

- Devices are defined as AS in HCD.
- Operators or automation products issue **VARY** commands.

Tip: **V 1000,AS,0N** 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.

Enabling AS support is shown in Example I-3.

Example I-3 GRS definition to enable AS support

```

/*-----*/
/* Enable AS support */
/*-----*/
RNLDEF RNL(INCL)
        TYPE(GENERIC)
        QNAME(SYSZVOLS)

```

Note: A Parallel Sysplex normally includes a shared TMS - plex (TMSplex), SMSplex, and TCBD. However, it is possible to define separate TMSplex, SMSplex, and TCDBs if complete separation of volser ranges is wanted

- ▶ IECIOSxx. In this member, you can define specific device ranges and you must separate TEST from the PROD updates:

- TEST updates are shown 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 the 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 shown 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 tape library categories to be used on each system. Use a specific and separate DEVSUPxx parmlib member for TEST that specifies different categories from the PRODDEVSUPxx parmlib member. In this example, you need a unique VOLCAT.VGENERAL that continues the library entries for each LPAR but might share the specifics volcats with the volume entries if you are not using SDAC and want to read the volumes from either LPAR:

- 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,
MEDIA1=0021
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,
MEDIA1=0011,
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, DFSMSrmm is 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 **PRITITION** and **OPENRULE**. **PRITITION** is the preferred method for partitioning. **REJECT** commands, although still supported, are less powerful and work differently than **PRITITION** and **OPENRULE** commands.

If you use **REJECT** commands, convert from the use of **REJECT** commands to use the **PRITITION** and **OPENRULE** commands. You can use this method of soft partitioning with or without SDAC. However, with SDAC, the extra protection of the library that is refusing to mount a volser that is not defined as accessible, which is a complete and thorough method of partitioning that does not rely on correctly configuring the TMS on all sharing hosts.

For more information about options for DFSMSrmm, see *z/OS DFSMSrmm Implementation and Customization Guide*, SC23-6874.

The definitions that are needed in this specific case study are listed in Table I-2.

Table I-2 EDGRMMxx parmlib options for each LPAR

Host	VLPOOL definitions (optional)	PRITITION definitions	OPENRULE definitions
PROD	VLPOOL PREFIX(A*) TYPE(S) DESCRIPTION (PROD DEFAULT') MEDIANAME(*) EXPDTCHECK(N)	PRITITION VOLUME(A*) TYPE(ALL) SMT(ACCEPT) PRITITION VOLUME(B*)TYPE(ALL) SMT(IGNORE) PRITITION VOLUME(*) TYPE(ALL) SMT(IGNORE) NOSMT(IGNORE))	OPENRULE VOLUME(B*) TYPE(RMM) ANYUSE(REJECT) OPENRULE VOLUME(*) TYPE(NORMM) ANYUSE(REJECT)
TEST	VLPOOL PREFIX(B*) TYPE(S) DESCRIPTION (TEST DEFAULT') MEDIANAME(*) EXPDTCHECK(N)	PRITITION VOLUME(B*) TYPE(ALL) SMT(ACCEPT) PRITITION VOLUME(A*) TYPE(NONRMM) SMT(IGNORE) PRITITION VOLUME(*) TYPE(ALL) SMT(IGNORE) NOSMT(IGNORE))	OPENRULE VOLUME(A*) TYPE(RMM) ANYUSE(REJECT) OPENRULE VOLUME(*) TYPE(NORMM) ANYUSE(REJECT)

Define the volume range that is going to be inserted into each host by using the **PRITITION** command. Then, **OPENRULE** with **ANYUSE(REJECT)** is used to limit access to the volumes that are connected to the other host.

The **VLPOOL** statements are optional, but many clients find it convenient to use **EXPDTCHECK(N)** to prevent setting up automation to reply U to IEC507D when mounting a scratch tape with an expiration date.

In the example that is described next, we code explicit statements to accept volumes that are owned by each LPAR, even though the default is **ACCEPT** because we also are coding a **PRITITION** statement with **VOL(*)** to not allow any other volume ranges not explicitly permitted. It is also possible to code the **OPENRULE** statements such that the LPARs can read each other's volumes by using the more granular **OUTPUT** and **INPUT** parameters instead of **ANYUSE**. However, an entry must exist for the volume in the TCDB to be successfully mounted.

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 include how to create ACS routines or ACS constructs. It states the ACS constructs that are needed and suggests a naming convention that helps you understand the constructs relationship:

- ▶ Storage Class (SC) definition:
 - Preference Group 0 assigned to volumes that are unlikely to be accessed.
 - Preference Group 1 for volumes likely to be accessed.
- ▶ Management Class (MC) definition: Not relevant for the partitioning, but you must have separate sets for TEST and PROD.
- ▶ Data Class (DC) definitions: Definition and relation to logical volume size; in this example, 800 MB and 6000 MB.
- ▶ Storage Group (SG) definitions: Pointing the SG to the composite library definition.
- ▶ Automatic Class Selection (ACS) routines and library definitions must also be defined, but they are not described here.

The naming convention in this case defines that all TEST volser definitions are prefixed with TS and PROD with PR.

ACS constructs and definitions for TEST are listed 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

ACS constructs	ACS construct name for TEST/PROD	Important value
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	The library name must match the SMS-defined name for composite library.

a. The 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. However, many installations seem to forget that access to read/write tape volumes and tape data sets is by default not restricted with RACF settings. 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 the following concerns:

- ▶ OAM must start after the IPL.
- ▶ New messages are introduced.
- ▶ Hardware errors and operator interventions occur and must be handled.

For more information, see this [IBM Support web page](#).

Definitions on the TS7700 Management Interface

Updates and definitions that are needed on the windows of the Management Interface (MI) are covered in this section. 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.

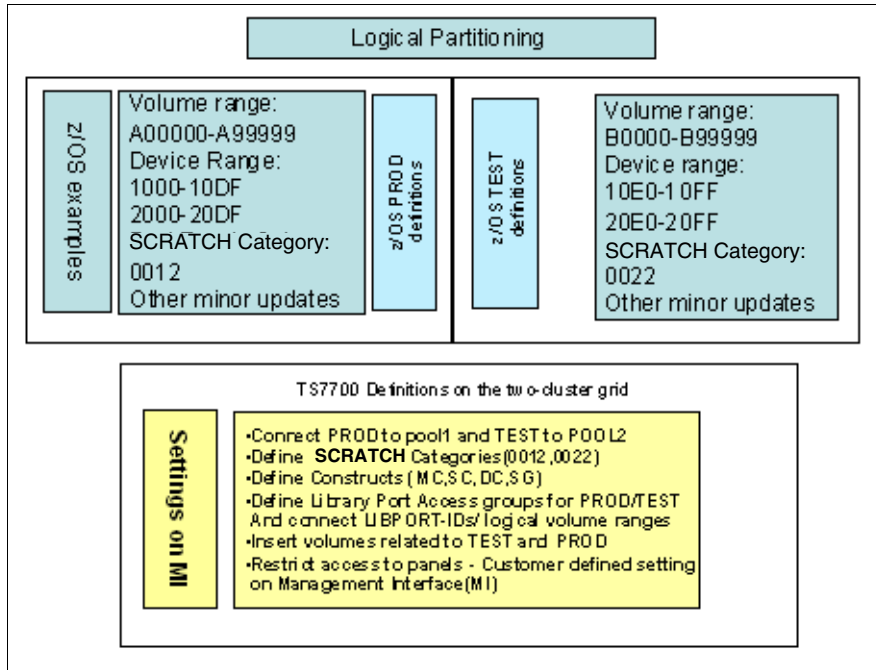


Figure I-3 Management Interface updates for logical partitioning of TEST and PROD

Physical volume pools

Based on your requirements, you must decide 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.

Consider the following points:

- ▶ If sharing the same physical volumes is *acceptable*, all logical volumes from TEST and PROD are mixed on the physical VOLSERS. Separate pools for PROD and TEST are not required.
- ▶ If sharing the 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 1004.




Figure I-4 Volume pools for PROD and TEST

Scratch categories

The TS7700 performs scratch mounts directly to the Tape Volume Cache (TVC) with no physical tape involvement.

Note: At Release 3.0 and higher, all categories that are defined as *scratch* inherit the Fast Ready attribute. In previous 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. Complete the following steps:

1. Define an SC for TEST named TSSCPG0. This SC is for volumes that are 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.

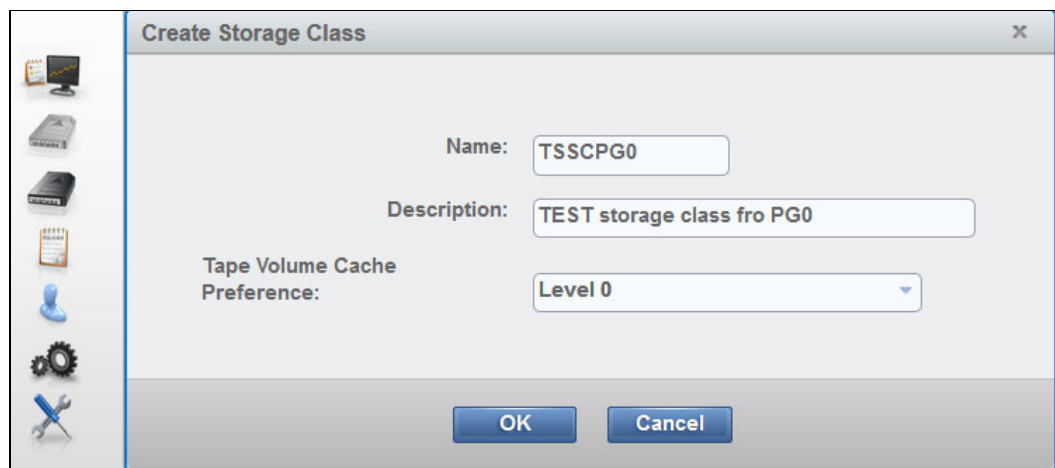


Figure I-5 Define Storage Class for TEST

- Define an MC for TEST that is named TSMCCL0 with only one copy of data and cache residency in cluster0, as shown in Figure I-6. Set the Copy Consistency 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.

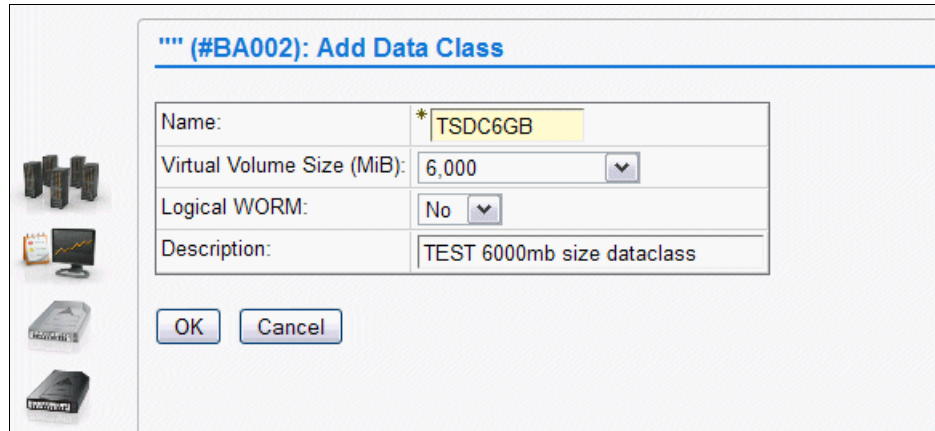
Clusters	Copy Mode	Options
"Yacko[0]" (#BA02A)	Rewind Unload (RUN)	Scratch mount candidate: <input checked="" type="checkbox"/>
"Wacko[1]" (#BA02B)	No Copy	Scratch mount candidate: <input type="checkbox"/>

Figure I-6 Management Class for TEST

- Define a DC for TEST named TSDC800M, as shown in Figure I-7.

Figure I-7 Data Class for TEST

4. Define a DC for TEST that is named TSDC6GB (6000 MB), as shown in Figure I-8.

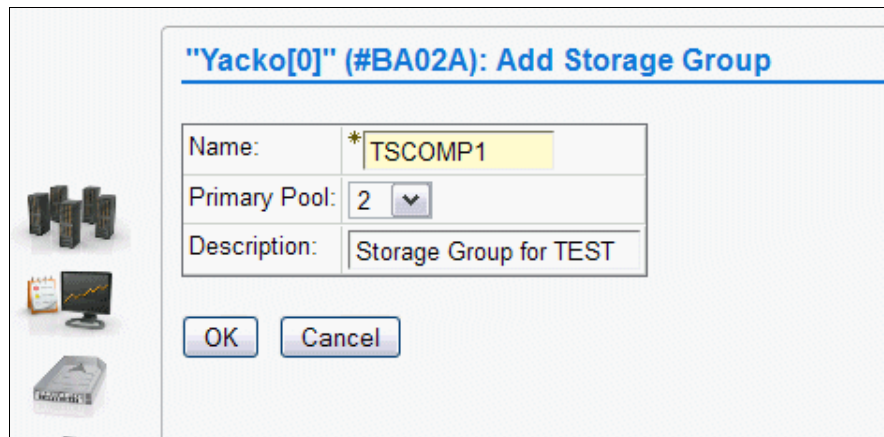


"" (#BA002): Add Data Class	
Name:	*TSDC6GB
Virtual Volume Size (MiB):	6,000
Logical WORM:	No
Description:	TEST 6000mb size dataclass

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 1001. 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.

Complete the following steps for Library Port Access Groups:

1. 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

- 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.

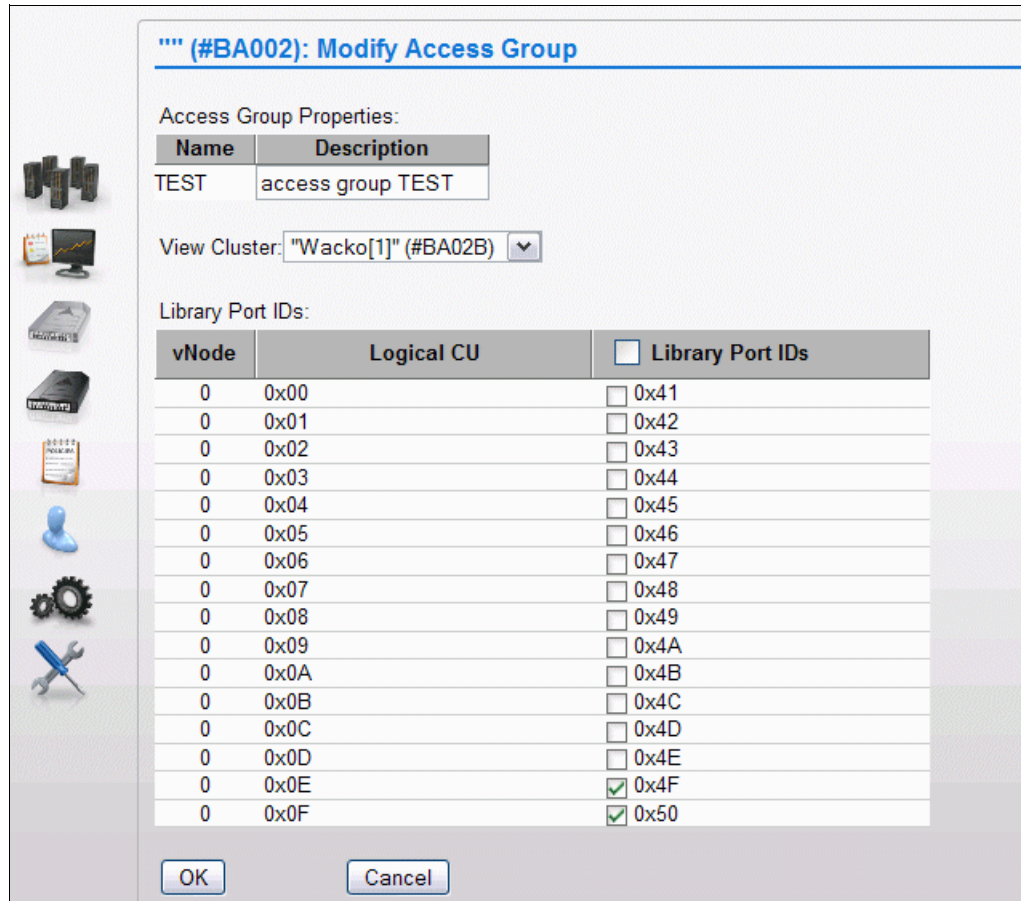


Figure I-11 Modify an access group for TEST LPAR

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

3. Define 41–4E on Cluster 1, as shown in Figure I-13.

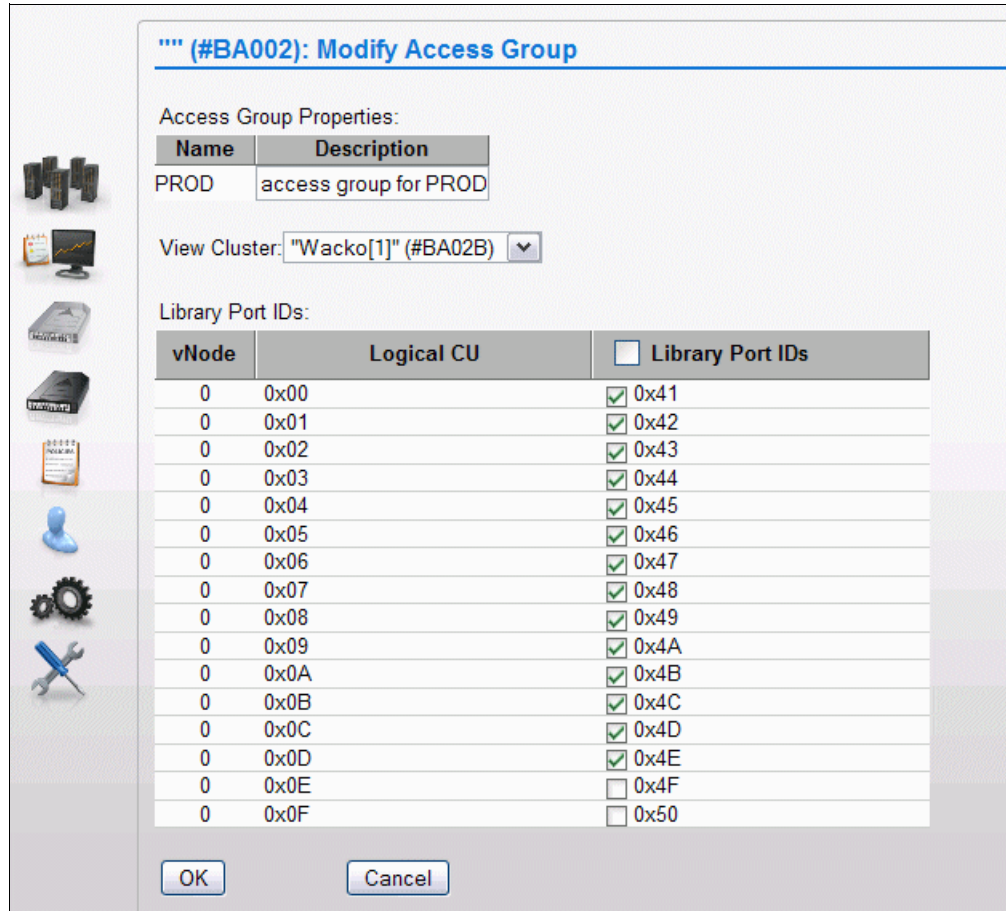


Figure I-13 Modify the access group for PROD LPAR

4. Check to see that each access group's ranges are not overlapping, as shown in Figure I-14.

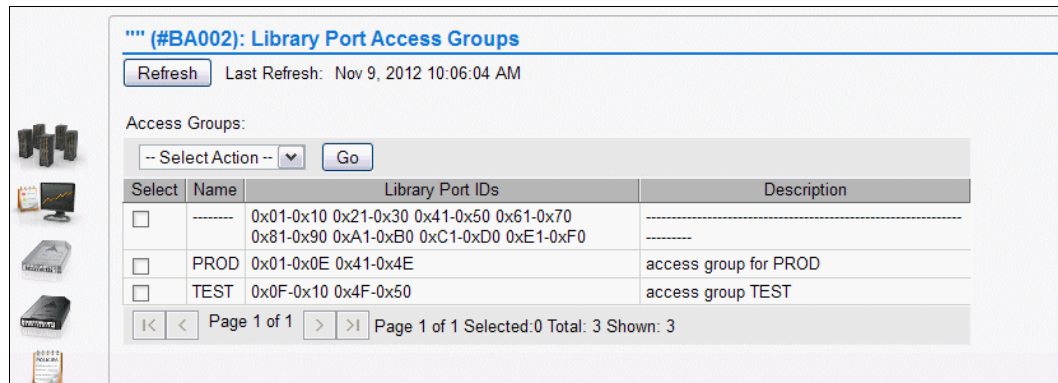


Figure I-14 Summary display of access groups

- Define logical volume ranges to the access groups, as shown in Figure I-15.

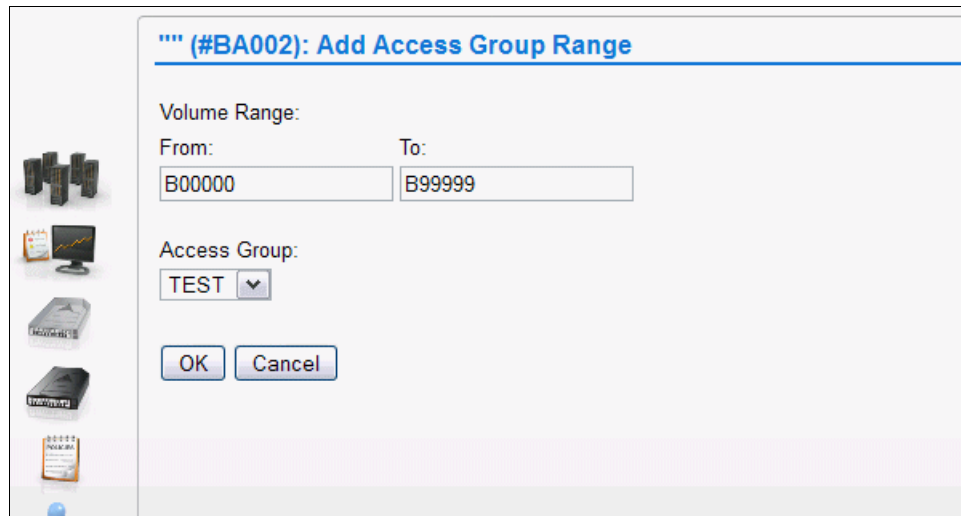


Figure I-15 Add VOLSER range for TEST LPAR

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 that are connected to defined ranges

On the MI window (see 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 shows the 'Virtual Volumes' management interface. On the left is a navigation pane with icons for various actions: Virtual Volume Details, Insert Virtual Volumes (highlighted), Modify Virtual Volumes, Delete Virtual Volumes, Move Virtual Volumes, and Virtual Volume Search. The main area is titled 'Virtual Volumes' and contains the following elements:

- Current availability across entire grid:**

Currently Inserted:	100,000
Maximum Allowed:	1,000,000
Available Slots:	900,000
- Show inserted volume ranges between and
- Insert a new virtual volume range:**
 - *Starting volser:
 - Quantity
 - Ending volser
 - Initially owned by:
 - Media type:
 - Cartridge System Tape (400 MiB)
 - Enhanced Capacity Cartridge System Tape (800 MiB)
 - Set Constructs
 - Storage Groups: Management Classes: Storage Classes: Data Classes:
 -

Figure 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 10.3, "Access icon" on page 492.

Verification of changes

After setting the definitions, evaluate your setup against the one shown in Figure I-18 on page 1011. If you attempt to access volumes from the other host, and SDAC is used, your job fails with this message:

```
CBR4175I VOLUME volser LIBRARY libname ACCESS GROUP DENIES MOUNT
```

If you are not using SDAC, and RMM is your TMS, you see the following message instead:

```
EDG4058I VOLUME 100066 REJECTED BY INSTALLATION OPENRULE COMMAND REJECT ACTION
```

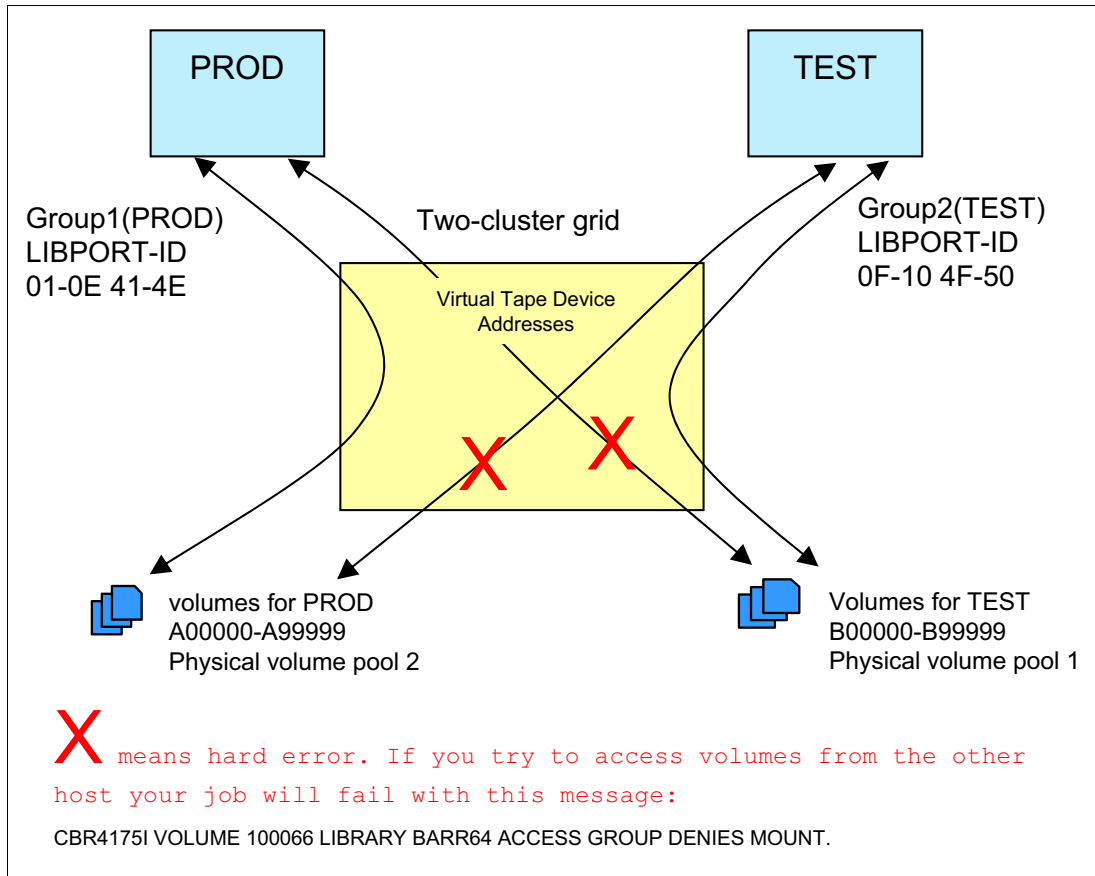


Figure I-18 Result after all definitions are finished

You can run several procedures to ensure that the setup is correct and ready for production. Be sure that you cover the following points:

- ▶ Ensure that all settings are as expected on z/OS and on the TS7700 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 entering the following command:
D SMS,VOL(*volser*)
- ▶ 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,VOL(*volser*)** 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.



Configuring externally managed encryption

The TS7700 currently supports data encryption for data at rest with external key management for two different data storage tiers:

- ▶ Physical tape cartridges
- ▶ Disk cache storage subsystem

In both cases, the TS7700 Server acts as a mediator between the device performing the actual encryption of data (the physical tape drives or disk cache storage subsystem) and an external key server, which is responsible for storing digital keys. The keys are required during the encryption/decryption process. The external key server is reachable through an Ethernet network to provide encryption key services to authorized devices.

The external key server is referred to as the Encryption Key Manager (EKM).

As of TS7700 release 5.3, the only EKM that is supported is the IBM Security Key Lifecycle Manager. IBM Security Key Lifecycle Manager for z/OS support is limited to physical tape drive encryption. IBM Security Key Lifecycle Manager running on open systems hardware supports physical tape drive encryption and disk cache storage encryption within the TS7700. If the TS7700 is to be configured with tape and disk encryption capabilities, the same EKMs must be shared to handle encryption services for both storage tiers.

Starting with Version 4.1 of the IBM Security Key Lifecycle Manager is renamed as *IBM Security Guardium Key Lifecycle Manager 4.1*.

The TS7700 also supports the use of the Container edition for z/OS of the IBM Security Guardium Key Lifecycle Manager.

For planning considerations about data encryption, see section 4.3.18, “Planning for tape encryption in a TS7700T” on page 198, and 4.3.19, “Planning for cache disk encryption in the TS7700” on page 200.

This appendix provides a step-by-step description (from the perspective of a user) of the process to be followed when configuring the TS7700 Virtual Engine to apply external key management when data encryption capabilities are used.

This appendix includes the following topics:

- ▶ “Encrypting physical tape cartridges by using external key management” on page 1014
- ▶ “Disk storage encryption with external key management” on page 1018
- ▶ “Use of digital certificates on TS7700 to EKM connections” on page 1023
- ▶ “Managing IBM Security Guardium Key Lifecycle Manager device groups for TS7700” on page 1033
- ▶ “More information about IBM Security Guardium Key Lifecycle Manager management” on page 1037

Encrypting physical tape cartridges by using external key management

Because of the portable nature of physical tape cartridges, only external key management is supported when physical tape encryption is enabled on the TS7700T.

Complete the following steps to enable external key management for physical tape cartridges. These steps can be completed concurrently while the TS7700 is in an online operational state:

1. Install feature code 9900 (tape encryption configuration) into the TS7700T (if not previously installed). This step can be completed by the customer or an IBM service representative. A list of installed feature codes can be found in the Feature Licenses page by selecting **Cluster Settings** → **Settings** in the Management Interface (MI), as shown in Figure J-1 on page 1015.
2. Enable System-Managed encryption for the tape drives that are assigned to the Logical Library in the TS3500/TS4500 library corresponding to the attached TS7700T:
 - For more information about TS45000 encryption method configuration, see [Tape encryption overview](#).
 - For more information about TS3500 encryption method configuration, see [Setting or changing a drive's method of encryption](#).
3. Configure the TS7700T to see at least one EKM server through the same network infrastructure used by the TS7700 MI. The TS7700T attached 3592 physical tape drives use the IBM Proprietary Protocol (IPP) when communicating with the EKM server.

Cluster common resources:

Cluster-Wide Disk Cache Enabled	40 TB
Cross-Cluster Communication (Grid)	Enabled
FC 5262 Installed Instances	2

Peak data throughput:

vNode	Peak Data Throughput
(v0)	100 MB/s

Currently activated feature licenses:

Select	Feature Code	Feature Desc...	License Key	Node	Node Serial
<input type="radio"/>	5270	200,000 Virtual Volume Increment	e5e6321d04106b8bc6	(h0) hNode	78-8FBC0
<input type="radio"/>	5270	200,000 Virtual Volume Increment	e5e6321d04106b8b4d	(h0) hNode	78-8FBC0
<input type="radio"/>	5270	200,000 Virtual Volume Increment	e5e6321d04106b8b82	(h0) hNode	78-8FBC0
<input type="radio"/>	5270	200,000 Virtual Volume Increment	e5e6321d04106b8ba3	(h0) hNode	78-8FBC0
<input type="radio"/>	5270	200,000 Virtual Volume Increment	e5e6321d04106b8bb8	(h0) hNode	78-8FBC0
<input type="radio"/>	4015	Grid Enablement	e5e6321d04106b8b30	(h0) hNode	78-8FBC0
<input type="radio"/>	9900	Tape Encryption Configuration	e5e6321d04106b8b7d	(h0) hNode	78-8FBC0
<input type="radio"/>	5274	1 TB Active Premigration Queue	e5e6321d04106b8b97	(h0) hNode	78-8FBC0
<input type="radio"/>	9277	External Key Management Certificate	-	(h0) hNode	78-8FBC0
<input type="radio"/>	5273	Enable TS7700 tape attach function	e5e6321d04106b8b28	(h0) hNode	78-8FBC0

Page 2 of 5 | 2 | Go | Total: 47 | Filtered: 47 | Displayed: 10

Figure J-1 Feature code 9900 (tape encryption configuration)

- It is a best practice to have two configured EKMs (primary and secondary), which are replicas of each other for redundancy. For this purpose, the TS7700T must know the IP addresses or named addresses of the target EKMs. This configuration can be done by using the Data at Rest Encryption page by clicking **Settings** → **Cluster Settings** in the MI (see Figure J-2).

Cluster Settings

Cluster Network Settings

Feature Licenses

SNMP

Inhibit Reclaim Schedules

Data at Rest Encryption

Write Protect Mode

Backup Settings

Restore Settings

RSyslog

["0]" (#BA12A): Data at Rest Encryption

Disk encryption: AES256 Enabled local

Tape encryption: Active

Key server type: IBM SKLM (IPP)

Primary key server

Address: 9.47.87.212 Port: 3801 Test connectivity

IPP TLS 1.2

Secondary key server

Address: 9.47.87.171 Port: 3801 Test connectivity

IPP TLS 1.2

Submit Changes

Figure J-2 "Data at rest encryption" configuration page (IPP-only)

Note: Consider the following points:

- ▶ Figure J-2 shows a configuration in which the TS7700 disk cache is not externally key managed. If it is externally key managed, more options might be visible on this page, as shown in Figure J-4 on page 1021. Any displayed KMIP options are not applicable to physical tape encryption. Only IPP options apply to physical tape encryption.

- ▶ Each physical tape drive that is attached to the TS7700 must be authorized by the configured IBM Security Guardium Key Lifecycle Manager to participate in Key Exchanges to perform encryption or decryption.

For more information, see “Managing IBM Security Guardium Key Lifecycle Manager device groups for TS7700” on page 1033.

- ▶ Starting with code level R5.0, TS7700 optionally supports enabling TLS1.2 for the IPP connection for increased security. Consider the following points if this option is activated:
 - The EKM and TS7700 need their individual SSL server certificate installed into each other’s truststore for them to establish a connection. For more information about this configuration, see “Use of digital certificates on TS7700 to EKM connections” on page 1023.
 - The configured port for IPP communication must be changed from the default 3801 to the port in which the target EKM is accepting TLS/SSL traffic. Different EKM setups can be configured differently; therefore, consult with your EKM administrator to ensure that the correct port is used.
 - The configured ports must function within any customer-managed, network-attached firewall infrastructure.

After physical Tape Encryption is enabled, the TS7700T MI can be used to create data management policies that can be used by the host system to control the specific groups of virtual volumes, which are subject to data encryption (it is possible that the same machine can manage encrypted *and* nonencrypted physical volumes). Such data management policies are based on the use of SMS constructs, which are assigned to virtual volumes:

- ▶ **Storage Groups**

In a TS7700T, storage groups determine the physical volume pool in which a virtual volume is stored. Tape encryption is applied based on physical volume pool granularity, and each pool can be configured differently. Therefore, if a virtual volume is required to be encrypted, it should be assigned to a storage group that points to a physical volume pool that has encryption enabled.

- ▶ **Management Classes**

In a TS7700T, the management class can be used to create an optional secondary copy on physical tape. This duplex can be used for extra redundancy or for Copy Export. If a secondary pool is configured, its encryption settings also must be configured and can have encryption settings different than the primary pool.

- ▶ **Storage Classes**

Storage Classes determine the Disk Cache Partition in which a virtual volume is stored. Such partitions are used to control the final storage tier destination of the volume, which can be disk-only, cloud, or physical tape. Selecting a partition that can send data to tape is required in this scenario.

Physical tape cartridge encryption begins after the following prerequisites are met:

- ▶ The TS7700T is configured to communicate with the EKM.

- ▶ The physical tape drives that are installed in the attached TS3500 or TS4500 are configured to use system-managed encryption.
- ▶ One or more physical tape pools have encryption that is enabled.
- ▶ One or more constructs target those physical tape pools.

Note: All data on a physical tape cartridge is encrypted or not based on the pool encryption setting when the tape is moved from an empty to a filling state. Encryption-based pool setting changes are not retroactive and take effect only after a tape is reclaimed, becomes empty and then is reused.

For more information about settings that are related to these construct types and physical volume pool configurations (which include the selection of encryption keys to be used), see the following sources:

- ▶ 11.2.2, “TS7700 definitions” on page 578
- ▶ [TS7700 Encryption Support V2.0](#)

Disk storage encryption with external key management

For disk encryption (tape volume cache encryption), encryption keys can be managed locally within the TS7700 or externally by using an EKM. The external encryption key management method is a best practice for improved security and simplification of the key management process.

The TS7760 and TS7770 support the concurrent update from local key management to external key management. Moving from external key management to local key management is supported through an RPQ request only.

When the external disk encryption is enabled on the TS7700, valid encryption keys must be present on the disk cache storage subsystem when the cache unlocks the drives, or the user generates a new key. If the external key server is enabled in the disk cache storage subsystem, the key is retrieved from the external key server.

The disk cache storage controller requires an encryption key to be present during the following operations:

- ▶ Cache power-on.
- ▶ Cache restart.
- ▶ The user-initiated regeneration of key (rekey) operations.

The disk drive encryption in TS7700 is at the cluster level (storage subsystem cache controller), regardless of whether that cluster is part of a grid or not. The TS7700 architecture does not require members of a grid to have the same configuration for encryption. Therefore, the grid architecture supports any type of combination of TS7700 grid members regardless of the encryption configuration: no encryption, local encryption, external encryption (with key server) or tape encryption can be applied to different members of the same grid.

Enabling internal key management on the TS7760 can be completed only by an IBM service representative. Contact your IBM representative if you want to enable internal key management on a TS7760. This appendix focuses only on external key management for the TS7760 and TS7770.

Encryption of Data at Rest (EDaR) for TS7770

The TS7770 disk cache storage subsystem models support encryption for data at rest. That means encryption is performed by the cache controller system for data that is stored within the entire CSB/XSB and CFC/XFC systems (the controller enclosure and all attached expansion enclosures). So, unlike the CSA model, the CSB and CFC models perform the encryption by controller enclosures, while CSA and prior models used Self-Encrypting Disks.

Data encryption is protected by the Advanced Encryption Standard (AES) algorithm that uses a 256-bit symmetric encryption key in XTS mode, as defined in the IEEE 1619-2007 standard and NIST Special Publication 800-38E as XTS-AES-256. The data encryption key is itself protected by a 256-bit AES key wrap of a key that is derived from the access key that is stored on the EKM server (the same structure applies for internal key management: USB flash drive). The wrapped key is stored in the disk subsystem controllers in non-volatile form.

The data encryption keys are protected by use of NIST's AES key wrap, by using an intermediate 256-bit wrapping key. This intermediate key is in turn protected by using NIST's AES key wrap, by using a 256-bit wrapping key derived from an access key. The access keys are managed using whichever key management methods are enabled on the system. Local key management is supported, in which case the access key is stored on USB flash drives. External key management is also supported, in which case the duty of key management is managed by software on a networked key server (IBM GKLM). The wrapped keys are stored securely on the system in non-volatile memory so that they can be accessed when the system starts up. During normal system operation, all unwrapped keys are stored securely in volatile memory, the contents of which are securely discarded on system shutdown or power loss.

There is no full FIPS 140-2 validation for the current CSB and CFC cache storage subsystem models in TS7770. However, TS7770 does have a NIST Algorithmic Validation for the encryption in CSB and CFC cache models that are used in TS7770. For more information, see the NIST web page [Cryptographic Algorithm Validation Program](#). The algorithmic validation is not a full IPS 140-2 validation, but it is a significant part.

If disk encryption has been enabled for the cache storage subsystem, any data that is written to a disk in the subsystem is also encrypted. When the data is read from the disk, it is decrypted. Encryption for data at rest is full data encryption on data storage (that is, at storage system). Host to storage communication is not encrypted. The data in flight within the grid network is not encrypted unless Secure Data Transfer is enabled. In addition, data movement to back-end physical tape is not encrypted by default. Use physical tape encryption when you offload to physical tape. Data movement to an attached object store supports HTTPS for encryption in flight and requires the object store support encryption for data at rest.

Encryption by using EKM servers for TS7770

Supported EKM Servers (Also, see the whitepaper [IBM TS7700 Disk Encryption for CSB and CFC models: Local, GLM and Thales Support](#)):

- ▶ IBM Security GKLM
- ▶ Thales CipherTrust Manager

EKM servers can have the following characteristics:

- ▶ Physical access to the CSB and CFC cache storage subsystems is not required to process rekeying operation
- ▶ Support for businesses that have requirements not to use USB ports
- ▶ Strong key generation
- ▶ Key self-replication and automatic backups (IBM Security Guardium Key Lifecycle Manager feature functions)
- ▶ Implementation follows an open standard that aids in interoperability
- ▶ Audit detail
- ▶ Ability to administer access to data separately from storage devices

EKM servers create and manage encryption keys that are used by CSB or CFC cache models in TS7700. EKM servers distribute keys remotely without requiring physical access to the disk storage subsystems. For security and simplification of key management, the EKM servers are the best practice method of managing encryption keys on the disk storage subsystems.

An EKM server is a centralized system that generates, stores, and sends encryption keys to the disk storage subsystem. The TS7770 with CSB or CFC supports IBM Security Guardium Key Lifecycle Manager to handle key management on the disk storage subsystem. The IBM GKLM is a management application that creates and manages cryptographic keys for the disk storage subsystem and provide access to these keys through a certificate. Authentication takes place when certificates are exchanged between the disk storage subsystem and the external key server. Certificates must be managed closely because expired certificates can cause disk storage system outages. EKM servers must be installed and configured before they are defined on the CSB or CFC storage subsystems.

IBM GKLM external key server supports Key Management Interoperability Protocol (KMIP), which is a standard for encryption of stored data and management of cryptographic keys. Externally managed encryption stores the lock key to the encryption key inside the IBM GKLM server system. The GKLM server is separate and independent from the TS7700 system. A minimum of one EKM server is required to enable the key server support.

The following are the supported versions of IBM external key servers:

- ▶ IBM Security Key Lifecycle Manager 2.6
- ▶ IBM Security Key Lifecycle Manager 2.7
- ▶ IBM Security Key Lifecycle Manager 3.0
- ▶ IBM Security Key Lifecycle Manager 3.1
- ▶ IBM Security Key Lifecycle Manager 4.0
- ▶ IBM Security Guardium Key Lifecycle Manager 4.1 and later

Encryption keys must be replaced when suspected of compromise. If external encryption is enabled on the system, the TS7700 supports a regeneration of key (rekey) process, where the external key server generates a new key and the previously existing key becomes obsolete.

Note: The TS7700 disk-based encryption is either enabled or disabled at the time of manufacturing. There is no standard method of moving from non-encryption to encryption or vice versa. Therefore, it is a best practice that the TS7770 be configured with disk encryption support at the time of purchase.

Enabling External Key Management

Complete the following steps to enable external key management for disk encryption. These steps are concurrent and can be completed while the TS7700 is operational:

1. If the system is a TS7760, check if the target TS7760 has the external key management certificate that is installed from manufacturing (which was done if the machine was ordered with feature code 9277). You can also have an IBM service representative install the certificate onsite. Feature Code 5277 must be ordered as a miscellaneous equipment specification (MES).

Also, the IBM service representative must install Feature Code 5276 (Disk Encryption with external key management, which is ordered as an MES through your IBM sales representative), if not previously installed. A list of installed feature codes is available in the Feature Licenses window by selecting **Settings** → **Cluster Settings** (see Figure J-3 on page 1021).

Note: All TS7770 solutions include the required certificates.

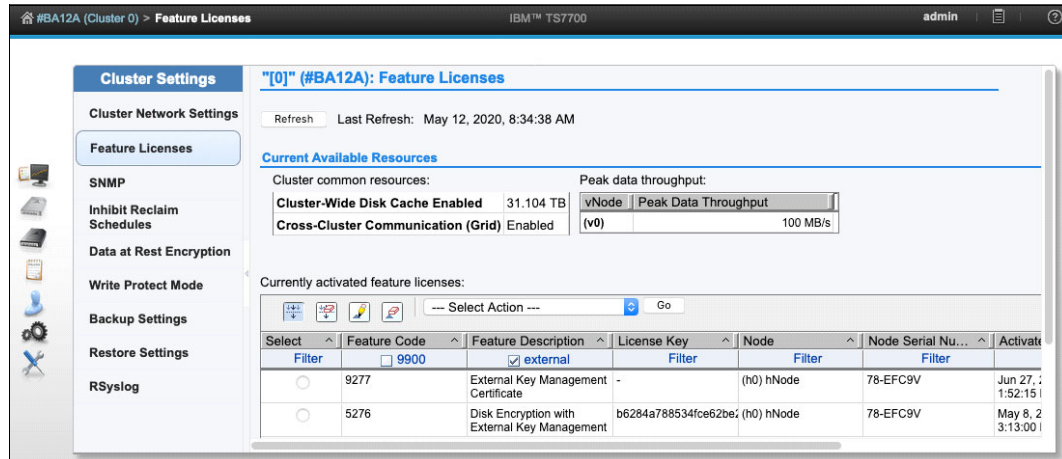


Figure J-3 Feature codes required for disk encryption with external key management

- For the TS7760 and TS7770, configure the TS7700 to see at least one EKM server through the same network infrastructure that is used by the TS7700 MI. The EKM must be a version of IBM Security Guardium Key Lifecycle Manager that runs on an open or distributed systems platform.

Containerized IBM Security Guardium Key Lifecycle Manager instances that are running under zCX are also supported (zOS v2.4 and later). Communication to the IBM Security Guardium Key Lifecycle Manager for disk encryption uses the KMIP protocol for TS7770 systems and the IPP protocol for TS7760 systems.

Physical tape encryption also uses IPP. Therefore, settings might exist for KMIP and IPP, based on your hardware configuration (see Figure J-4).

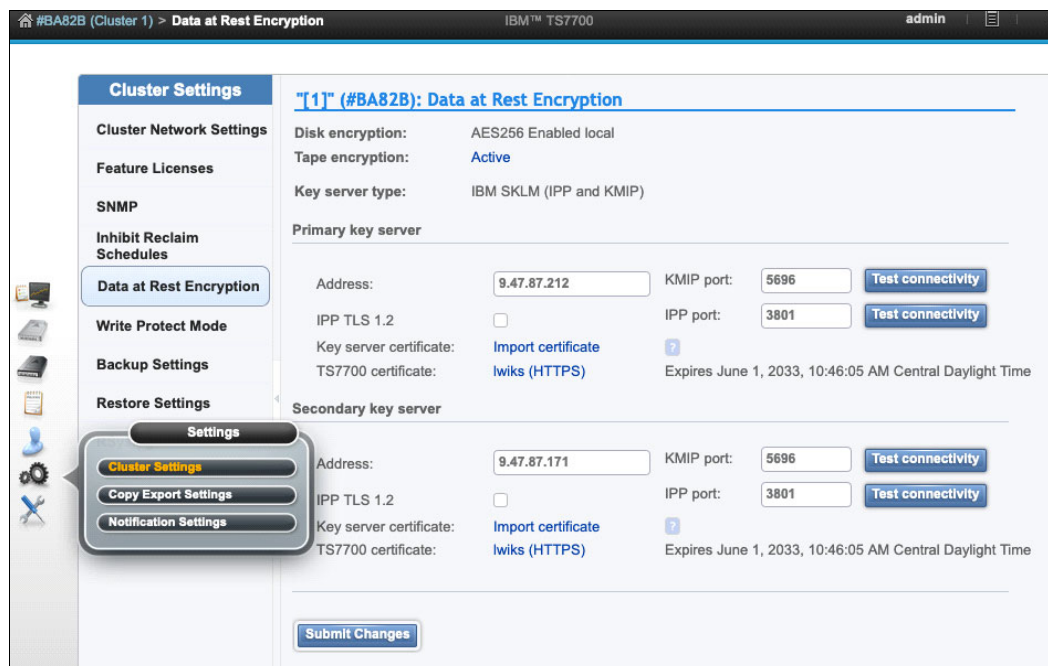


Figure J-4 Data at Rest Encryption window (KMIP and IPP to be configured)

It is a best practice to have two configured EKMs, which are replicas of each other for redundancy. The TS7700 must know the IP addresses or DNS style names of the target EKMs. This configuration can be done by using the Data at Rest Encryption page by selecting **Settings** → **Cluster Settings** (see Figure J-4 on page 1021).

Notes: Consider the following points:

- ▶ Each CSA and CSB disk cache string that is attached to a TS7700 must be authorized by the configured IBM Security Guardium Key Lifecycle Manager to participate in Key Exchanges to provide encryption services. For more information, see “Managing IBM Security Guardium Key Lifecycle Manager device groups for TS7700” on page 1033.
- ▶ TLS 1.2 is “always enabled” for KMIP protocol connections, which are required when enabling external disk encryption on the TS7770-VED.
- ▶ The IPP TLS 1.2 attribute can be optionally activated for the TS7760-VEC when enabling external key management for disk encryption. When physical tape encryption is enabled on a TS7760 or TS7770, the IPP TLS 1.2 attribute is also optional because physical tape encryption relies on IPP.
- ▶ Regarding TLS 1.2 connections:
 - Your EKM and TS7700 need a corresponding SSL server certificate that is installed into each other’s truststore for them to establish a secure connection. For more information about this configuration, see “Use of digital certificates on TS7700 to EKM connections” on page 1023.
 - If the IPP TLS 1.2 attribute is activated, the configured port for IPP communication must be changed from the default 3801 to the port in which the target EKM is accepting TLS/SSL traffic. Different EKM setups can be configured differently; therefore, contact your EKM administrator to ensure that the correct port is used.
 - The configured ports must also be allowed to function within any network-attached firewall infrastructure.

An IBM Service IBM service representative must complete the enablement of external key management on any TS7700.

After disk encryption is enabled in the TS7700, the entire disk subsystem for a specific TS7700 is enabled for encryption by using local or external key management.

Note: No supported method is available to partially encrypt separate portions of the disk subsystem, or to use different key management processes for different portions of the content within the subsystem.

When a disk subsystem is encryption-enabled, a local or external key exchange occurs during the power-on phase of each disk cache string that is attached to the TS7700, or if a rekey is started. A rekey can be started through the MI for the TS7760 with internal or external key management.

For the TS7770, a rekey is possible by using the MI when external key management is enabled. However, the process requires an IBM representative to start a rekey when local key management is enabled.

For more information, see this [IBM Support web page](#).

Use of digital certificates on TS7700 to EKM connections

When EKM exchanges are made through the KMIP or IPP TLS 1.2 protocols, the EKM and the TS7700 are required to authenticate with each other by using identification in the form of a digital certificate. These certificates are used to authenticate and create the secure connection. If only physical tape encryption is used and the IPP TLS 1.2 option is not enabled, you do not need to continue.

To safely identify the validity of the certificates that are exchanged during a KMIP or TLS 1.2 connection, each other's public certificate must be trusted by the other. Meaning, the EKM must have a trust of the TS7700's public certificate and the TS7700 must have a trust of the EKM's public certificate.

A certificate can be trusted by using one of the following methods:

- ▶ The truststore can contain a copy of the exchanged public certificate. This certificate can apply to a self-signed certificate, or a certificate that is signed by a root or intermediate certificate authority (CA).
- ▶ If the public certificate is signed by a root or intermediate CA, you can preferably add the root CA and all relevant intermediate CAs to the truststore. This way, the CA-signed certificate can be verified and trusted by using the CA's public certificate.

It is a best practice to use the CA chain of trust approach because you need to add it to your truststore only once. After it added, any certificates that are signed by the same CA are trusted.

To help simplify this explanation, the TS7700's public certificate is referred to as the *TS7700-PCert*. The EKM's public certificate is referred to as the *EKM-PCert*.

Updating the TS7700 to trust the attached EKMs

Complete one of the following options for each EKM to be attached to the TS7700 when a KMIP or IPP TLS 1.2 is used:

- ▶ Option 1: Upload the EKM-PCert or EKM-PCert's CA chain to the TS7700 truststore.
- ▶ Option 2: Have the TS7700 download the EKM-PCert directly from the EKM and add it to the TS7700 truststore.

These options are described next.

Option 1: Upload the EKM-PCert or EKM-PCert's CA

This option requires you have a copy of the EKM-PCert or the entire chain of trust of the CA or intermediates that are used to sign the EKM-PCert. If you are uploading the CA chain of trust, you can upload each CA and intermediate independently, or you can combine them into one chained certificate and upload that instead. If all EKM-PCerts were signed by the same CA chain, you need only to add the CA chain to the TS7700 once.

Complete the following steps:

1. Log in to the TS7700 MI and open the **Access/SSL** Certificates page.
2. Click **New Certificate**, which opens the Add Certificate window. In the Method to add a certificate tab, select **Upload a certificate file**, as shown in Figure J-5. Click **Next**.

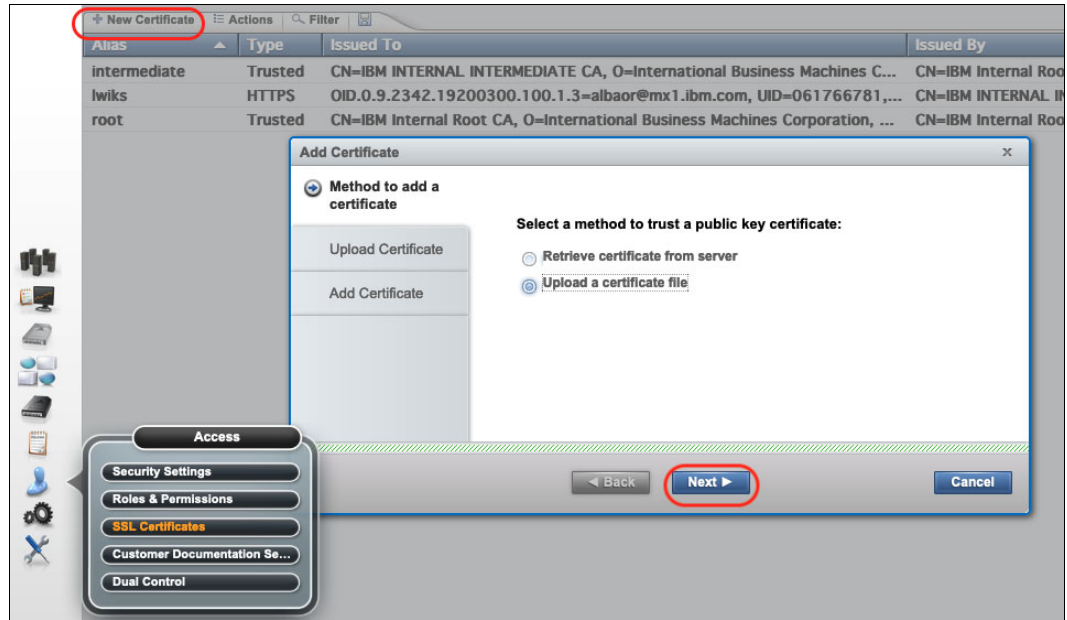


Figure J-5 Open the “Add Certificate” dialog

3. Click **Upload** to open a file manager to select the certificate file to be uploaded from your local workstation (the file manager tool that is used depends on your local operating system), as shown in Figure J-6 on page 1025. This file can be a copy of the EKM-PCert or one or more public certificates from the CA chain of trust. Click **Next** when you are ready to upload the certificate file.
4. Verify that the properties of the uploaded file correspond to the target certificate. Then, assign an alias that identifies it in the TS7700 truststore (as shown in Figure J-7 on page 1025). Click **Finish** to complete the process.

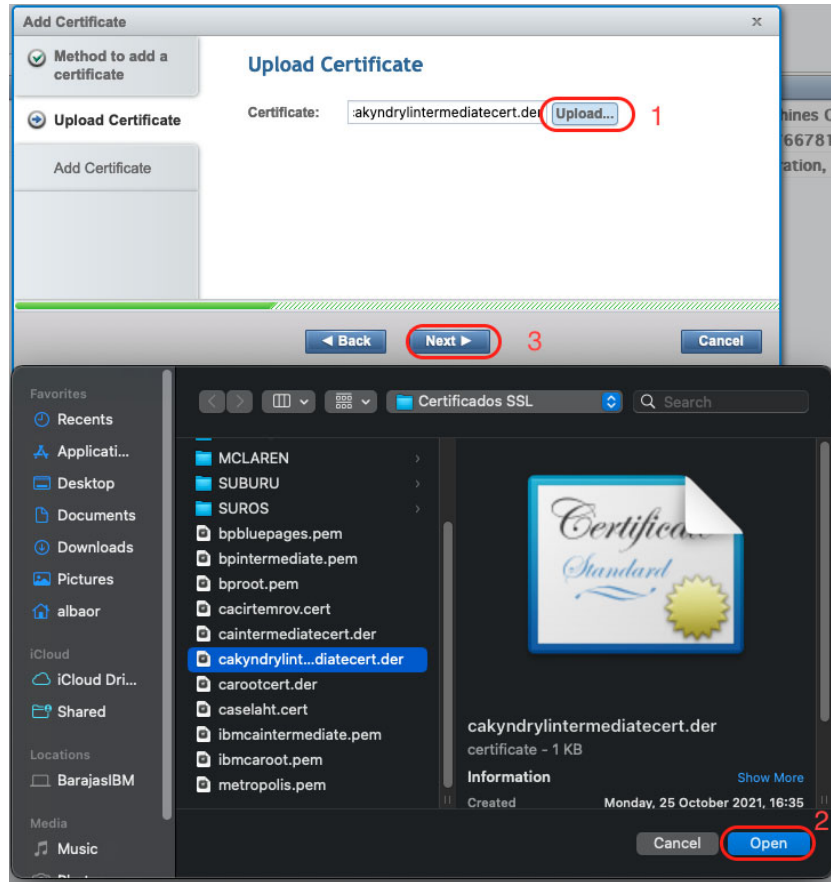


Figure J-6 Selection of SSL certificate file for manual upload

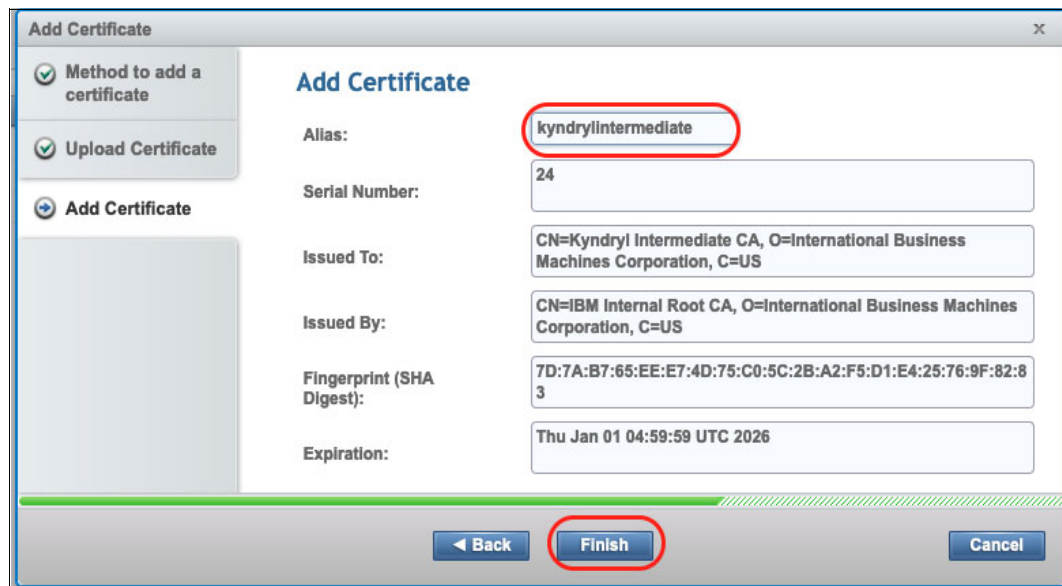


Figure J-7 Assigning an alias for the new certificate and completing registration

5. If you have more than one certificate in the chain, repeat this process for each certificate of the chain.

Option 2: Have the TS7700 download the EKM-PCert directly from the EKM

By using this option, you can download a copy of the EKM-PCert directly from EKM and add it to the TS7700 truststore. This process is essentially the same as uploading the EKM-PCert as described in “Option 1: Upload the EKM-PCert or EKM-PCert’s CA” on page 1023; the only difference is that the EKM-PCert is downloaded from the EKM directly.

Complete the following steps:

1. Log in to the TS7700 MI and open the **Access/SSL** Certificates page.
2. Click **New Certificate**, which opens the Add Certificate window. Then, in the Method to add a certificate tab, select **Retrieve a certificate from server** as the method to trust a public key certificate, as shown in Figure J-8.

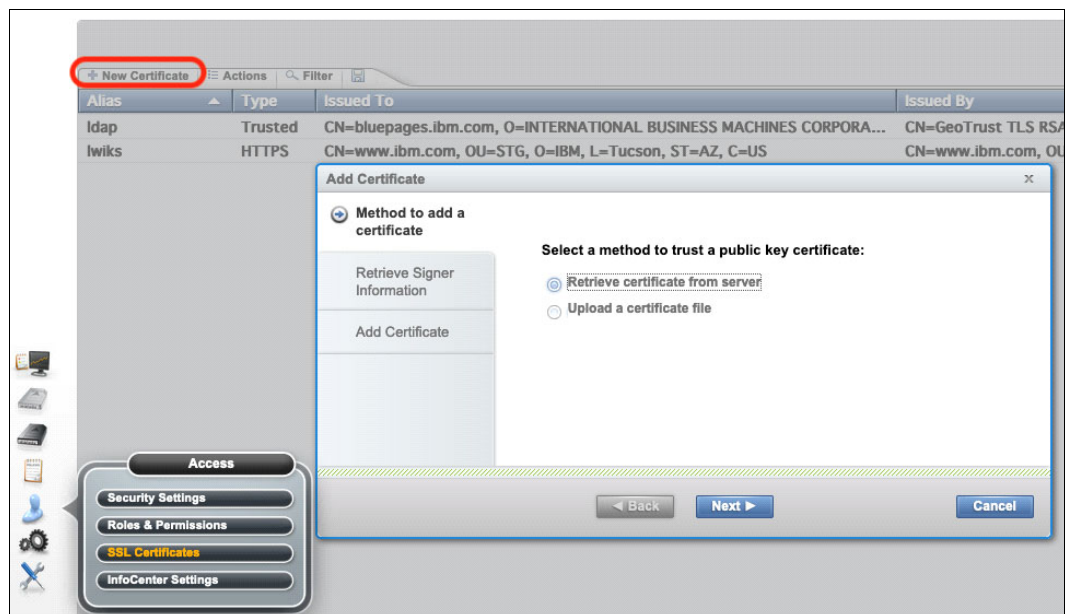


Figure J-8 Open the “Add Certificate” dialog

3. In the **Retrieve Signer Information** tab, enter the IP or DNS style name of the target IBM Security Guardium Key Lifecycle Manager server in the **Host** field, as shown in Figure J-9 on page 1027. Consult with your IBM Security Guardium Key Lifecycle Manager administrator to verify the port in which the target server accepts TLS/SSL traffic (the protocol that is used for this connection). Because different IBM Security Guardium Key Lifecycle Manager setups can be configured differently, the default value (port 443) might not be correct for many cases.

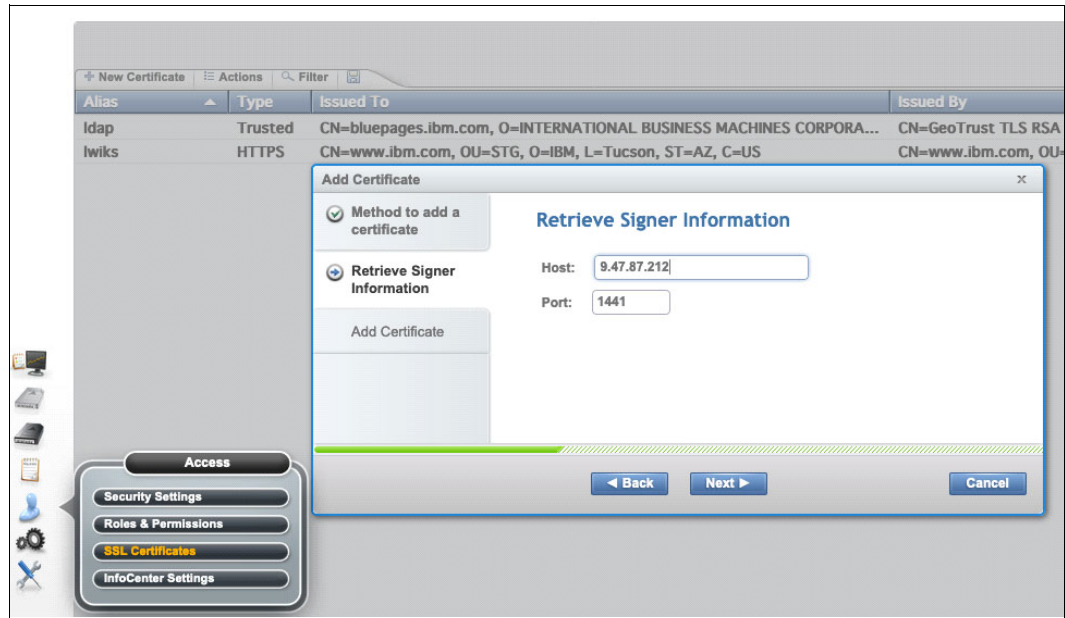


Figure J-9 Retrieving signer information

4. If the connection to the server succeeds by using the provided information, the certificate is retrieved. Next, define an Alias to help you identify the retrieved certificate (see Figure J-10) into the TS7700 truststore. Click **Finish** to complete the process.

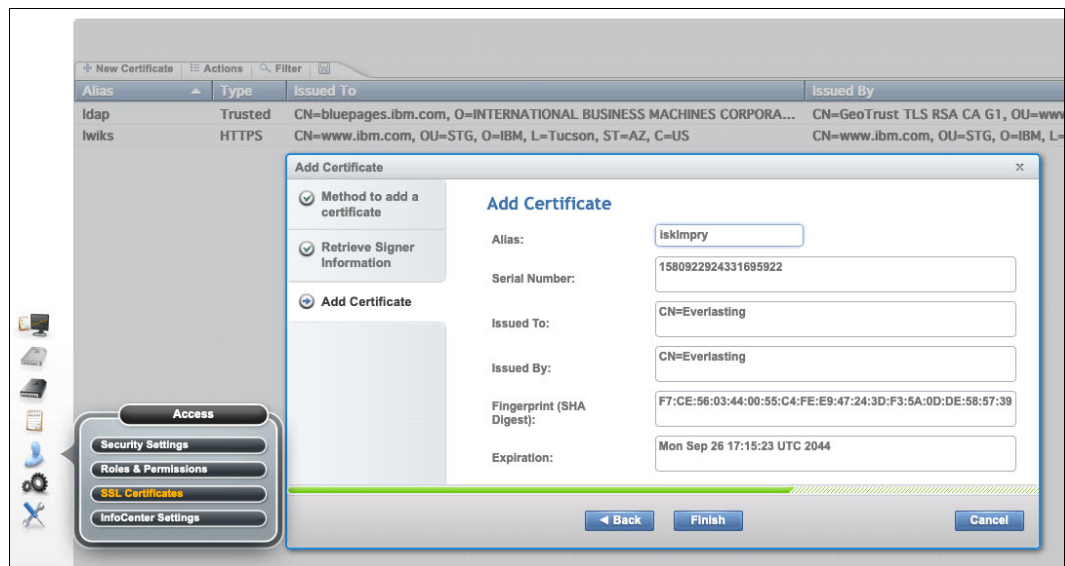


Figure J-10 Choosing an Alias for retrieved certificate

Notes: Because certificates can expire, schedule a reminder to update any trust-based certificates ahead of their expiration. By setting such a reminder, no risk exists of downtime in encryption services.

Updating EKM to trust the TS7700

This section explains how to update the EKM to trust the TS7700-PCert.

TS7700 LWIKS Certificate (TS7700-PCert)

The TS7760 and TS7770 use the LWIKS certificate that is configured in the MI for the TS7700-PCert. LWIKS stands for Lightweight Infrastructure keystore certificate. This certificate is the same certificate that is used by the MI for HTTPS support. Therefore, the EKM must trust the currently configured LWIKS certificate.

The default LWIKS certificate that is installed on all TS7700s is a self-signed certificate. It is a best practice to replace the default LWIKS certificate with your company's CA-signed certificate. For more information, see this [IBM Documentation web page](#).

Consider the following points:

- ▶ If you choose to use the default LWIKS certificate, you must download the LWIKS certificate from the TS7700 and then upload it into your primary and secondary EKM trusts.
- ▶ If you replaced the LWIKS certificate with your company's certificate, you must update the trust of the primary and secondary EKM. Update the trust by adding a copy of the TS7700-PCert certificate or adding the entire CA chain of trust to both EKMs.
- ▶ If CA-signed certificates are used, your CA chain of trust must be updated only once on your EKMs. (This update might be completed already on behalf of other devices in your data center.)

Each TS7700 can have a different LWIKS certificate, or they can be identical. Consider the following points:

- ▶ If one or more of your LWIKS are different and self-signed, you must add each unique LWIKS public certificate to both EKMs.
- ▶ If your LWIKS are the same and self-signed, you must add only a single copy of the LWIKS public certificate to both EKMs. This requirement applies to the default LWIKS certificate.
- ▶ If your LWIKS is any mixture of the same or different but are all signed by the same root CA or intermediate CA, you must add the CA chain of trust only once to each EKM.
- ▶ If your LWIKS are different and signed by different CAs, you must add each unique CA's chain of trust to each EKM.

Downloading a copy of the LWIKS certificate

If you choose to upload a copy of the LWIKS directly to the EKM and you do not possess a copy of the LWIKS public certificate (TS7700-PCert), you can download a copy from the TS7700 by using the process that is described next.

Note: This process is used only if you do not have a CA-signed LWIKS certificate. The LWIKS certificate also is used for HTTPS connectivity into the MI; therefore, you might see that LWIKS are referred to as the *HTTPS certificate* in some circumstances.

Complete the following steps to download the TS7700 LWIKS public certificate (TS7700-PCert):

1. In the TS7700 MI, open the **Settings/Cluster Settings/Data at Rest Encryption** window, which is used to configure the IP addresses of the EKMs for the machine.
2. Click the **lwiks (HTTPS)** hyperlink to download the LWIKS certificate, as shown in Figure J-11. This hyperlink is visible for the following instances only:
 - Target TS7700 is a TS7770-VED, which requires connecting to an EKM that can communicate by using the KMIP protocol when External Disk Encryption is enabled.
 - The IPP TLS 1.2 attribute is activated for the connection to an EKM that can communicate by using that protocol. This customer-selectable option is available for either of the following instances:
 - Target TS7700 is a Tape-Attached machine, and Tape Encryption capability was enabled
 - Target TS7700 is a TS7760-VEC, which requires connecting to an EKM that can communicate by using the IPP protocol when External Disk Encryption is activated.

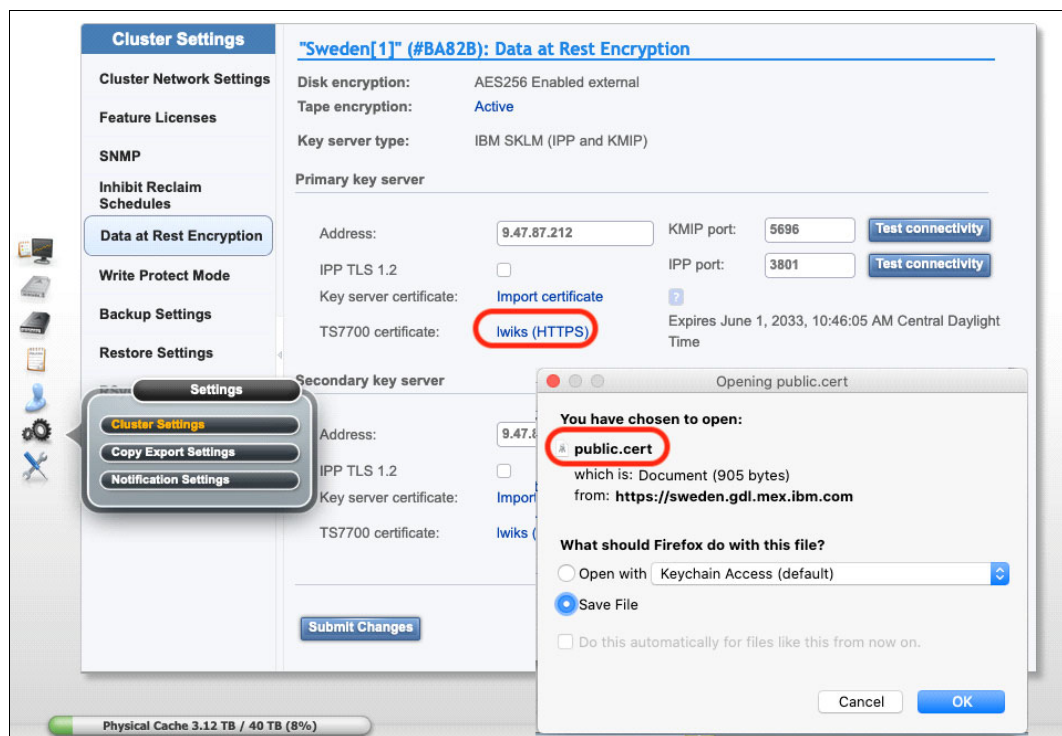


Figure J-11 Downloading the TS7700 HTTPS certificate

Note: If your organization installed a customer-provided certificate on the TS7700 to replace the default LWIKS self-signed certificate and if your organization's CA's trust certificates are available, do *not* download this certificate. Instead, upload the CA chain of trust directly to your EKMs.

Uploading a certificate into an IBM Security Guardium Key Lifecycle Manager trust

Complete the following steps to upload a public certificate into an IBM Security Guardium Key Lifecycle Manager trust:

1. Log in to the GUI of the target IBM Security Guardium Key Lifecycle Manager, and open the **Advanced Configuration/Client Device Certificates** window. Then, click the **Import** control (see Figure J-12).

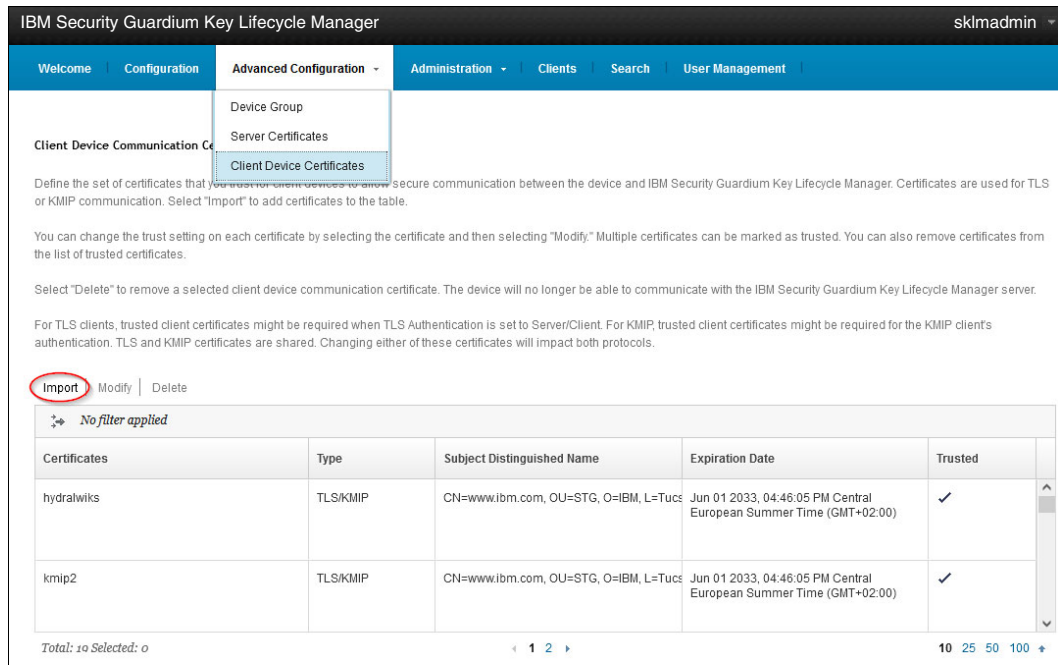


Figure J-12 IBM Security Guardium Key Lifecycle Manager Client Device Certificates

2. The **Import SSL/KMIP Certificate for Clients** opens (see Figure J-13 on page 1031). Complete the following steps:
 - a. Click **Browse** to open the Browse File window.
 - b. Click **Upload** to open a file manager that is directed to your local computer.
 - c. Select the suitable certificate file from your local file system and then click **Open**.

If you obtained the TS7700 certificate by using the method that is described in “Downloading a copy of the LWIKS certificate” on page 1028 (see Figure J-11 on page 1029), the downloaded LWIKS public certificate features a *.cert extension. However, IBM Security Guardium Key Lifecycle Manager can upload certificates with the *.cer extension only.

Because the certificate format is correct, manually rename it to the required file extension.

If the certificate that you want to upload includes a different extension, work with your security team to convert it to a *.cer-based certificate. Often, you need only to rename it.

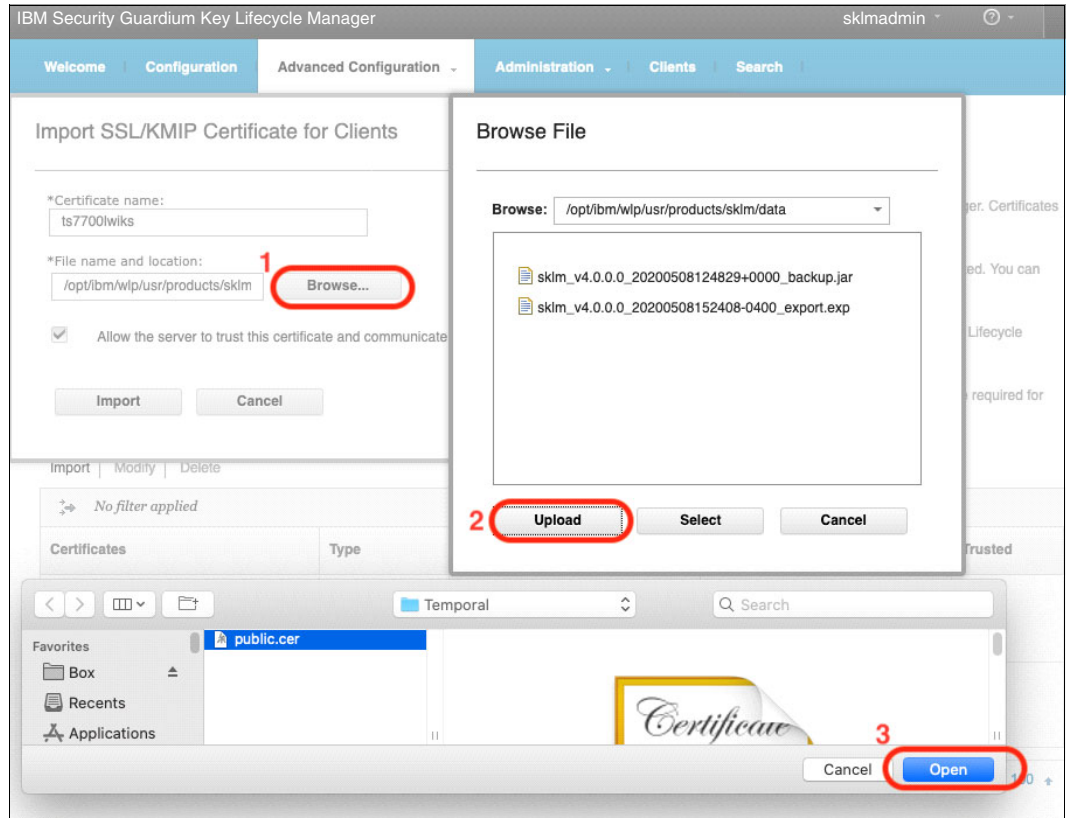


Figure J-13 Uploading TS7700 into the IBM Security Guardium Key Lifecycle Manager

- Now that the certificate file is uploaded from your computer to the IBM Security Guardium Key Lifecycle Manager, select it from the Browse File list (see Figure J-14).

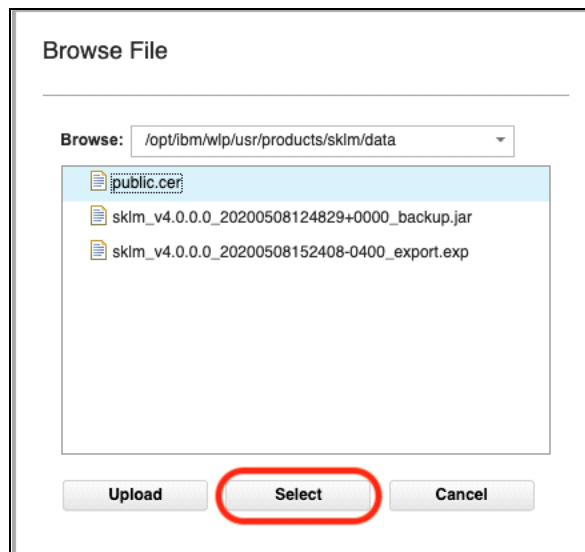


Figure J-14 Selecting uploaded certificate file

4. Enter a name to identify the selected certificate file in the local database (see Figure J-15) and then click **Import**.

Figure J-15 Importing selected certificate file

If the import completes successfully, the certificate is now included in the list of trusted client certificates (see Figure J-16).

Certificates	Type	Subject Distinguished Name	Expiration Date	Trusted
ts7700lwiks	SSL/KMIP	CN=www.ibm.com, OU=STG, O=IBM, I	Jun 01 2033, 10:46:05 AM Central Daylight Time (GMT-05:00)	✓
ts4300_a	SSL/KMIP	CN=KMIPSelfSigned, OU=Storage Pro	Nov 15 2048, 05:56:12 PM Central Standard Time (GMT-06:00)	✓

Figure J-16 Imported client certificate

Note: Repeat these steps for each unique LWIKS certificate when your LWIKS is a self-signed certificate. If CA or Intermediate CA-signed LWIKS certificates are used, each member of the CA chain's public certificate must be uploaded to the EKM by using the method that is described here. If you have a single CA chain of trust certificate that is the combination of all members of the CA chain, you need to complete this process only once.

Managing IBM Security Guardium Key Lifecycle Manager device groups for TS7700

After the TS7700 is configured to connect to IBM Security Guardium Key Lifecycle Manager, which (depending on the requirement of the environment) might include requiring that their individual TLS/SSL certificates are installed into each other's truststore, both machines can establish secure connections for the flow of the encryption key services.

However, the IBM Security Guardium Key Lifecycle Manager does not see the TS7700 Server as an end device to use encryption keys. Instead, the TS7700 Server acts as a gateway between the IBM Security Guardium Key Lifecycle Manager and the end devices that perform the encryption or decryption processes. Therefore, these devices must be authorized by IBM Security Guardium Key Lifecycle Manager to receive and use the encryption keys.

The IBM Security Guardium Key Lifecycle Manager keeps the record of the authorized devices and the corresponding encryption keys to be used by them by mapping them as Device Groups, which are also known as Device Families. For the TS7700, three different applicable IBM Security Guardium Key Lifecycle Manager Device Groups are available (see Table J-1) and displayed in the IBM Security Guardium Key Lifecycle Manager GUI, as shown in Figure J-17 on page 1034.

Table J-1 IBM Security Guardium Key Lifecycle Manager device groups for TS7700

Device group name	TS7700 device type	Key exchange model	IBM Security Key Lifecycle Manager availability
3592	Physical tape drive	Many asymmetric keys to many devices	Available by default
DS5000	CSA Disk Cache	Symmetric keys that are directly tied to a single device	Available by default
SPECTRUM_VIRT	CSB/CFC Disk Cache	Many devices to many keys with access by using certificate	Must be created by IBM Security Key Lifecycle Manager administrator (GPFS family type)

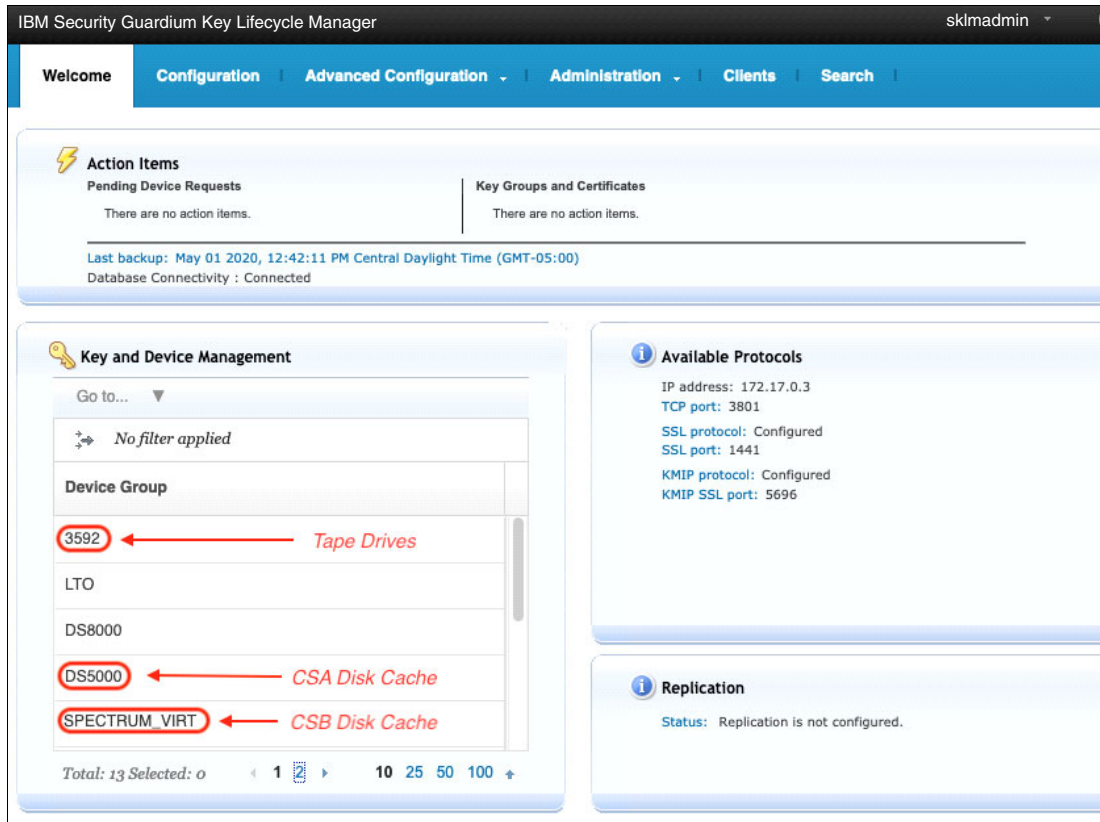


Figure J-17 IBM Security Guardium Key Lifecycle Manager device groups in the GUI

The contents of a specific device group can be displayed by selecting the target group from the list that is displayed in the Key and Device Management section on the Welcome window, and then right-clicking it to select the **Manage keys and devices** option, as shown in Figure J-18.

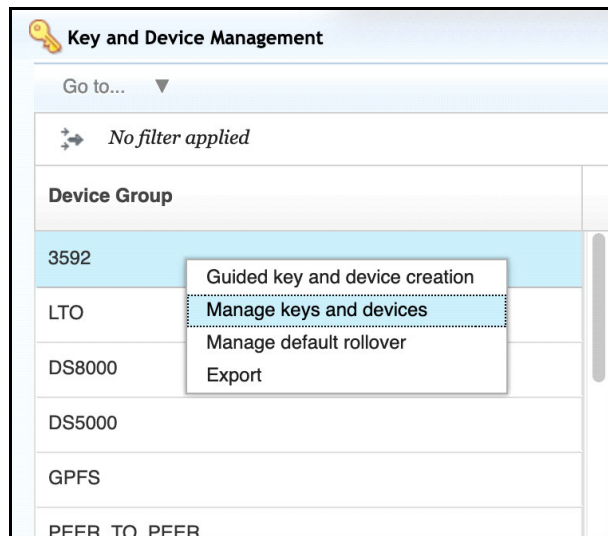


Figure J-18 Opening a selected IBM Security Guardium Key Lifecycle Manager device group

Adding devices to a device group

Each device group can be configured differently regarding policies to authorize new devices. Available options can be selected from a drop-down menu in the window that corresponds to each group, as shown in Figure J-19.

The screenshot shows the IBM Security Guardium Key Lifecycle Manager interface. The top navigation bar includes 'Welcome', 'Configuration', 'Advanced Configuration', 'Administration', 'Clients', and 'Search'. The 'Advanced Configuration' menu is open, showing 'Device Group', 'Server Certificates', and 'Client Device Certificates'. The main content area displays the configuration for device group 3592. It includes a table of certificates and a table of device serial numbers. A dropdown menu at the bottom of the page shows three options for device communication: 'Only accept manually added devices for communication', 'Automatically accept all new device requests for communication', and 'Hold new device requests pending my approval'. The 'Automatically accept all new device requests for communication' option is selected and circled in red.

Certificates	System Default/ Partner	Expiration Date	Private Key
jagkeydef	System Default and Partner	Sep 27 2044, 10:30:46 AM Central Daylight Time (GMT-05:00)	🔑
jagkey1	...	Sep 27 2044, 10:31:37 AM Central Daylight Time (GMT-05:00)	🔑
		Sep 27 2044, 10:32:11 AM	

Device Serial Number	Drive Default Certificate	Drive Partner Certificate
0000013B0071	SYSTEM_DEFAULT	PARTNER_DEFAULT
0000078D8301	SYSTEM_DEFAULT	PARTNER_DEFAULT
0000078D8307	SYSTEM_DEFAULT	PARTNER_DEFAULT
0000078DB657	SYSTEM_DEFAULT	PARTNER_DEFAULT

Figure J-19 The “3592” device group window

In general, the following optional methods are available to add authorized devices to a device group:

- ▶ Accept only manually added devices or certificates for communication
If this option is selected, encryption devices that are attached to the TS7700 must be registered *before* TS7700 is enabled to do encryption. Assistance from an IBM service representative is needed to extract the individual IDs or certificates that are associated to each encryption-capable device (tape drives, CSA, or CSB), which then must be shared to the IBM Security Guardium Key Lifecycle Manager administrator to register them in the IBM Security Guardium Key Lifecycle Manager. For more information, see [Managing and serving keys, certificates, and other cryptographic objects](#).
- ▶ Automatically accept all new device requests for communication
The first time that an encryption-capable device that is attached to the TS7700 requests a key, it is automatically accepted as an authorized device by the IBM Security Guardium Key Lifecycle Manager and the corresponding record is created in its database. This method is the simplest and fastest method that is available from an administration perspective.

- ▶ Hold new device or certificate requests pending my approval

The first time that an encryption-capable device that is attached to the TS7700 requests a key, the IBM Security Guardium Key Lifecycle Manager automatically raises an alert in the Welcome window of the GUI, which includes a link for the Administrator to accept and authorize the new device. The new device cannot perform any encryption or decryption until the administrator grants its permission, which originates errors and warnings on the TS7700 side, depending on the resource delays.

Creating SPECTRUM_VIRT Device Group for CSB/CFC (TS7770)

From the IBM Security Guardium Key Lifecycle Manager perspective, CSB and CFC Disk Cache Storage Subsystems (installed in TS7770 systems) are managed by using a separate device group that is named SPECTRUM_VIRT, which must be associated to the GPFS family type. This device group is *not* created by default when the IBM Security Guardium Key Lifecycle Manager is installed. Creating this IBM Security Guardium Key Lifecycle Manager object requires the EKM administrator to manually add it by using the Device Group GUI page, as shown in Figure J-20.

Note: In IBM Security Guardium Key Lifecycle Manager, device groups that correspond to the GPFS family type (such as the SPECTRUM_VIRT group) do not provide an option to automatically accept all new devices request for communication. Customers that use external key management for encrypted CSB subsystems must choose between the following methods to add devices to this group:

- ▶ Accept only manually added devices or certificates for communication
- ▶ Hold new device or certificate requests pending my approval (best practice setting)

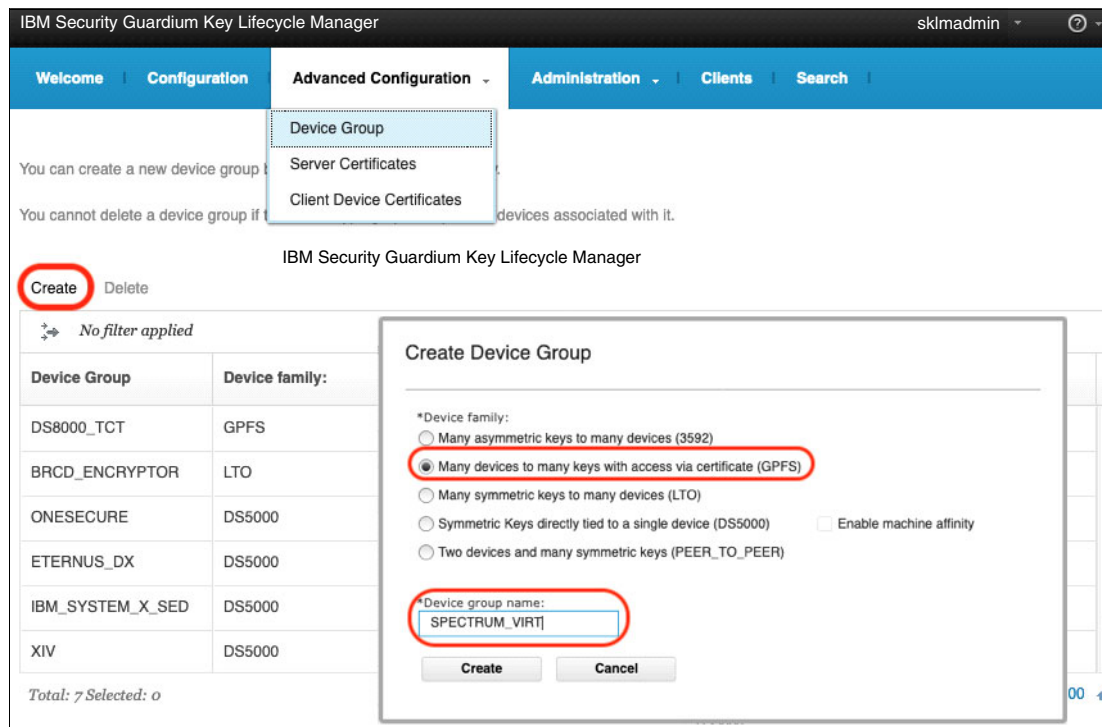


Figure J-20 Creation of the SPECTRUM_VIRT device group for CSB disk cache

More information about IBM Security Guardium Key Lifecycle Manager management

Different versions of IBM Security Guardium Key Lifecycle Manager are available that might require slight variances to the instructions that are provided in this document, which are based on IBM Security Key Lifecycle Manager v4.1 for containerized environments.

For more information, see this [IBM Support web page](#).

If you need to review specific instructions for other IBM Security Key Lifecycle Manager versions, or require more instructions about other IBM Security Key Lifecycle Manager-related tasks (installation, configuration, creating encryption keys, administration, and others), see the following sources:

- ▶ IBM Documentation: [IBM Security Guardium Key Lifecycle Manager 4.2](#)
- ▶ *IBM Security Guardium Key Lifecycle Manager*, SG24-8472

Using Thales CipherTrust Manager

To enable external encryption in CSB or CFC models by using Thales CipherTrust Manager key server, you must create and configure a KMIP interface, and assign trusted local and external CA's, port, mode, and auto registration.

Thales CipherTrust Manager (Thales) Configurations

The TS7700 8.54.0.68 level with 3948-CSB and 3948-CFC supports Thales CipherTrust Manager (CTM) to control encryption key management on the disk storage system.

Thales CipherTrust Key Management Services are a collection of service tiles that allow users to create and manage cryptographic keys and integrate them to external applications. As with GKLM the authentication takes place when certificates are exchanged between the disk storage system and the key server. Certificates must be managed closely because expired certificates can cause disk storage system outages. Key servers must be installed and configured before they are defined on the cache storage subsystems.

Thales CipherTrust Manager key server also supports Key Management Interoperability Protocol (KMIP), which is a standard for encryption of stored data and management of cryptographic keys. Thales creates keys on demand, sharing with the other clustered servers, and provides redundant access. Thales CipherTrust Manager must be configured to use TLS version 1.2.

TS7700 vital product data (VPD) can be analyzed by IBM service representatives to ensure existing hardware, software, and feature codes meet the preceding requirements.

The TS7700 system supports distinct types of configurations on key servers. The following configurations are supported:

- ▶ Thales CipherTrust Manager key servers use an active-active model, where multiple key servers are used to provide redundancy. In these configurations one key server must be specified as the primary key server. The primary key server is the key server that the system uses when you create any new encryption keys. The key is immediately replicated to the other key servers in the cluster. All the key servers that are defined on the system can be used to retrieve keys. Although it is possible to configure a single key server instance, two key servers are recommended to ensure availability of keys if one key server experiences an outage.

- ▶ The TS7700 system supports up to two key servers. If the system is accessing multiple key servers, they must belong to the same cluster of key servers.

General Configuration for External Encryption Key Server (Thales)

Thales CipherTrust Manager does not support self-signed system certificates.

Ensure that the following tasks are completed on the Thales CipherTrust Manager before performing external encryption enablement:

1. KMIP interface must be created by the Thales administrator before you run the External Encryption MES.
2. Get the KMIP interface authentication mode and password if it is required before starting the External Encryption MES.
3. The server certificate that is selected to use with the KMIP interface in Thales CipherTrust Manager should be downloaded to your local machine, ready to upload to the TS7700.
4. The Certificate Authority (CA) that is signing the server certificates must be a trusted entity on the Thales CipherTrust Manager key servers.
 - a. The local CA of the Thales CipherTrust Manager can be used to sign the TS7700 certificate signing request (CSR). Note that this CSR must be created on behalf of the TS7700 using a separate workstation because the TS7700 does not currently support the local creation of CSRs.
 - b. Alternatively, the system certificate can be signed by the CA of your organization. The organization CA must be uploaded as an External CA in the Thales CipherTrust Manager.
5. By default, Thales CipherTrust Manager is configured to require a username in the *common name* field of the client's SSL certificate on TS7700 system. This username is added when the request for certificate is generated by the system.
6. Create a user with the previous username in Thales CipherTrust Manager. This user owns the encryption keys for the system and must be added to the Key Users group.
7. If the TS7700 + CSB or CFC models currently have encryption enabled with USB flash drives, at least one of the USB flash drives must be inserted into the disk storage system before key servers can be configured for managing keys.

Thales CipherTrust Manager Configuration

An administrator of the Thales server is required to enable and setup the Thales for key exchanges to be handled by the Thales CipherTrust Manager. The TS7700 has no control of the Thales configuration. The TS7700 functions as a proxy to pass internal request for keys by the controllers to the Thales server and vice versa. Therefore, the first step is to configure the Thales server for key exchanges. The actual steps to configure the Thales can vary based on the operating system or Thales version. Because of that reason, consult with the Thales administrator for the correct or updated instructions.

The following instructions can be used as a general guideline:

Upload an External CA.

1. Open the Thales management window and log in.
2. Upload any external CA:
 - a. Click **CA** → **External** → **Add external CA** as shown in Figure J-21 on page 1039.




Figure J-21 External certificate authorities

- b. Assign a Name.
- c. Paste the content of the CA file (shown in Figure J-22).



Figure J-22 Creating a KMIP interfaces for Thales

Note: Repeat the process for each certificate authority in the chain of trust.

Download KMIP Certificate.

1. Access the Interface menu and select a KMIP interface:
 - a. Navigate to **Admin Settings** → **Interfaces** → **Select a KMIP interface**. Select "...".
 - b. Click **Download Certificate** as shown in Figure J-23.

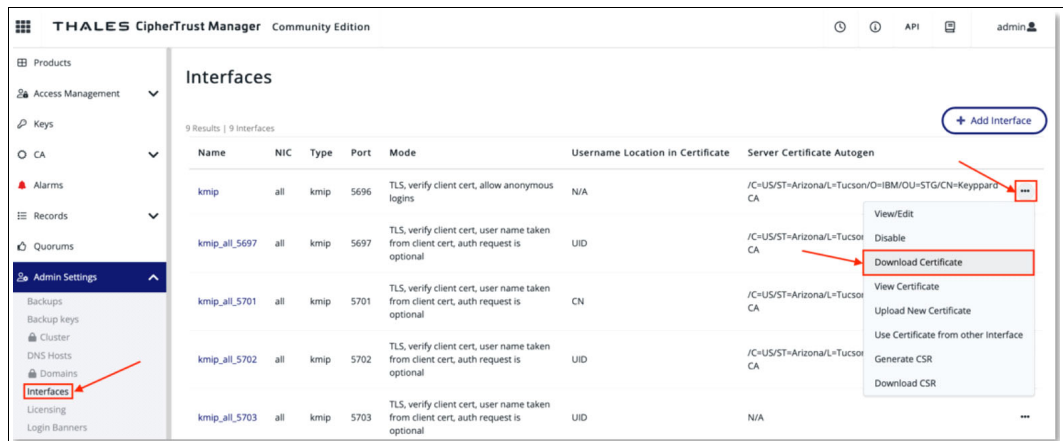


Figure J-23 Interfaces panel

- c. Click **Download Certificate** and **Close** as shown in Figure J-24 on page 1040.

Note: If the certificate has more than one certificate, save each certificate in its own file certificate(s), since each certificate must be uploaded separately.



Figure J-24 Download certificate

- d. Rename the file and add the **.pem** suffix.

Note: If the certificate has more than one certificate, save each certificate in its own **.pem** file. The certificates are uploaded to the TS7700 later.

KMIP Interface Details

1. Access the Interface menu and select a KMIP interface:
 - a. Select **Admin Settings** → **Interfaces** (shown in Figure J-25).

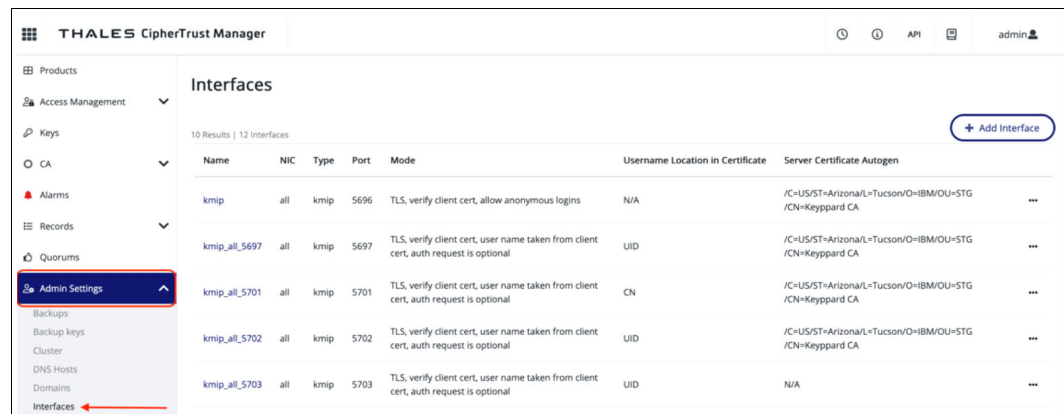


Figure J-25 KMIP interfaces

2. Verify the settings of the selected KMIP interface.
 - a. Access the configuration details shown in Figure J-26 on page 1041.

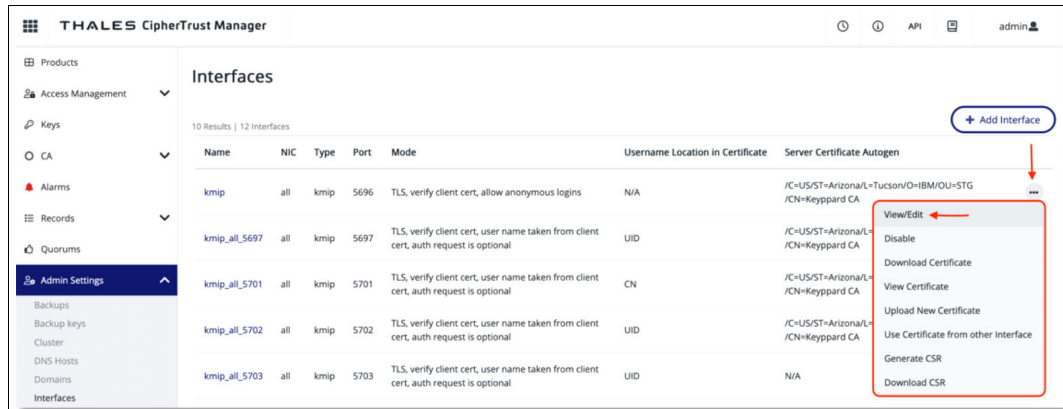


Figure J-26 Interface settings

- b. Verify the interface mode port and verify that the network interface is set to all. See Figure J-27. If a password is needed, ensure that it is defined and available before the defining the encryption MES.

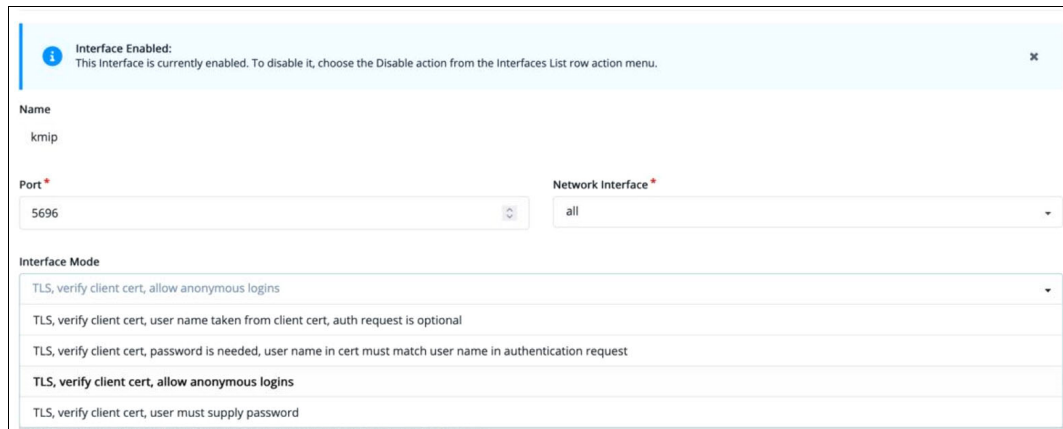


Figure J-27 Interface details

3. Add the Signature Algorithm as needed and enable the auto registration setting. See the example in Figure J-28.

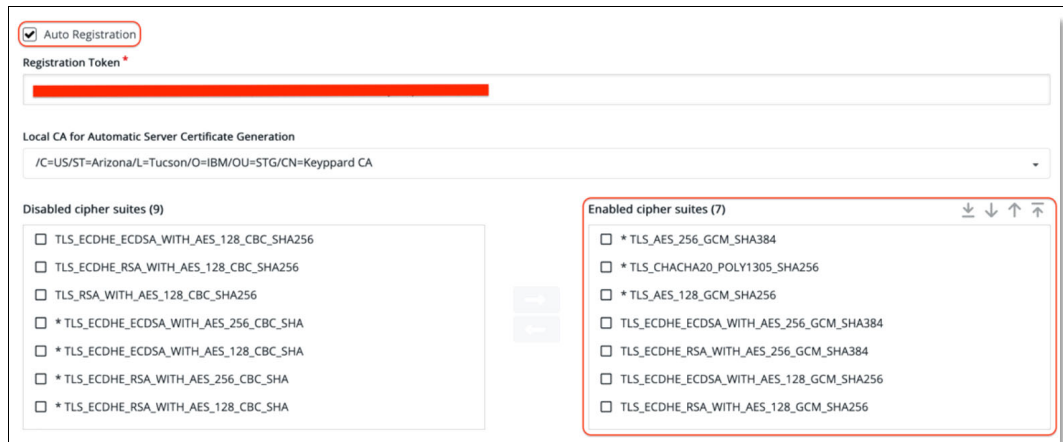


Figure J-28 Signature algorithm

4. If you are using an external CA, add the certificate to the trusted list. Follow these steps and see the example in Figure J-29:
 - a. Go to the drop-down menu.
 - b. Select the certificate → add the certificate.
 - c. When all the changes are made, click **Update**.




Figure J-29 Update details

Create a New User

Follow these steps to create a new user.

1. Go to **Access Management** (Figure J-30).
 - a. Select **Users** → **Add User**




Figure J-30 Users

2. The username must match the CN from the certificate that the user is going to use. Enter the information and click **Add User**. See Figure J-31 on page 1043.

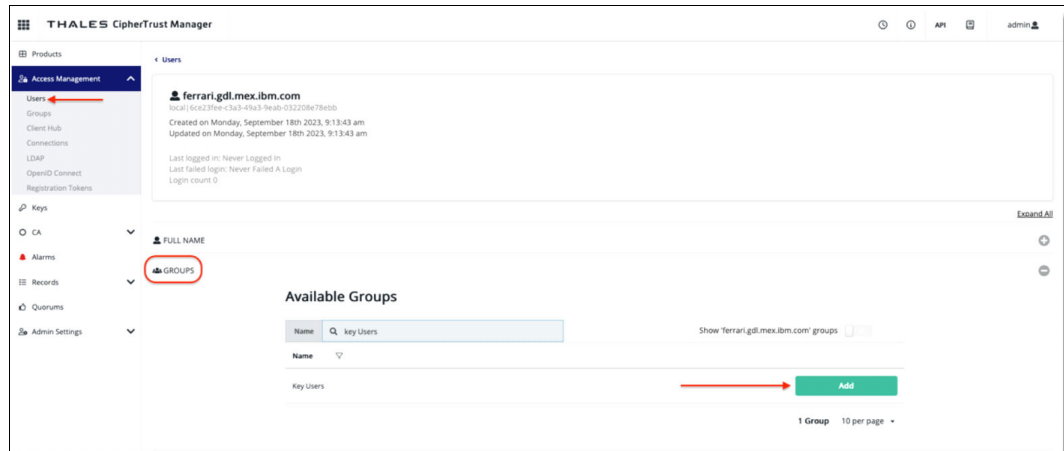


Figure J-33 Adding the key users group

TS7700 Management Interface Configuration

This same section can be used for the TS7700 management interface configuration for Thales CipherTrust Manager: TS7700_MI_configuration

Figure J-34 shows the KMIP ports for the Thales interfaces panel: KMIP 5696.

The screenshot shows the THALES CipherTrust Manager interface with the 'Interfaces' section selected. A table lists 10 results for 13 interfaces. A red arrow points to the 'Port' column, highlighting the value '5696' for the 'kmp' interface.

Name	NIC	Type	Port	Mode	Username Location in Certificate	Server Certificate Autogen
kmp	all	kmp	5696	TLS, verify client cert, allow anonymous logins	N/A	/C=US/ST=Arizona/L=Tucson/O=IBM/OU=STG/CN=keyppard CA
kmp_all_5697	all	kmp	5697	TLS, verify client cert, allow anonymous logins	N/A	/C=US/ST=Arizona/L=Tucson/O=IBM/OU=STG/CN=keyppard CA
kmp_all_5701	all	kmp	5701	TLS, verify client cert, user name taken from client cert, auth request is optional	CN	/C=US/ST=Arizona/L=Tucson/O=IBM/OU=STG/CN=keyppard CA
kmp_all_5702	all	kmp	5702	TLS, verify client cert, user name taken from client cert, auth request is optional	UID	/C=US/ST=Arizona/L=Tucson/O=IBM/OU=STG/CN=keyppard CA
kmp_all_5703	all	kmp	5703	TLS, verify client cert, user name taken from client cert, auth request is optional	UID	N/A
kmp_all_5704	all	kmp	5704	TLS, verify client cert, user name taken from client cert, auth request is optional	SN	/C=US/ST=Arizona/L=Tucson/O=IBM/OU=STG/CN=keyppard CA
kmp_all_5705	all	kmp	5705	TLS, verify client cert, allow anonymous logins	N/A	/C=US/ST=Arizona/L=Tucson/O=IBM/OU=STG/CN=keyppard CA
kmp_all_5706	all	kmp	5706	TLS, verify client cert, user name taken from client cert, auth request is optional	UID	/C=USA/ST=AZ/L=Tucson/O=IBM/OU=Storage/CN=Test Intermediate CA
kmp_all_5707	all	kmp	5707	TLS, verify client cert, allow anonymous logins	N/A	/C=USA/ST=AZ/L=Tucson/O=IBM/OU=Storage/CN=Test Intermediate CA
kmp_all_5898	all	kmp	5898	TLS, verify client cert, user name taken from client cert, auth request is optional	UID	/C=USA/ST=AZ/L=Tucson/O=IBM/OU=Storage/CN=Test Intermediate CA

Figure J-34 KMIP interface ports: 5696

Installing every Certificate into the TS7700 HTTPS

By default, the TS7700 uses a self-signed certificate to manage secure HTTPS connections to the Management Interface identified by the Lightweight Infrastructure Key Store certificate (LWIKS) alias that is included in the TS7700 trust store.

To enable the Thales encryption, you must replace the LWIKS with a non-self-signed certificate that must be provided by the customer and upload every certificate in the chain of trust associated to LWIKS.

Final TS7700 Configuration, completing the External Encryption Enablements

When the Thales server has been configured and when the TS7700 Management Interface has all the required certificates, the last part is that an IBM Service representative must complete the enablement of external key Management using the TS7700 service menu.

More information about Thales CipherTrust Manager (Thales) Configurations

Different versions of Thales CipherTrust Manager are available that might require slight variances to the instructions that are provided in this document, which are based on Thales CipherTrust Manager v2.13.1.

To review specific instructions for other Thales CipherTrust versions, or for more instructions about other Thales CipherTrust Manager-related tasks (installation, configuration, creating encryption keys, administration, and others), see [IBM TS7700 Disk Encryption for CSB and CFC models: Local, GKLM and Thales Support](#).

Abbreviations and acronyms

ABARS	Aggregate Backup And Recovery Support	CL0	CLASSES FOR RNN
ACDS	Active Control Data Set	CL1	CLASSES FOR RNN(CL0),
ACS	Automatic Class Selection	CLI	Command-line Interface
AES	Advanced Encryption Standard	CN	Common Name
ALMS	Advanced Library Management System	CP0	Cache Partition
AMSL	Above Mean Sea Level	CPC	Central Processor Complex
ANSI	American National Standards Institute	CPU	Central Processor Unit
AOS	AOTM, LDAP, Assist On-site	CPYDST	Copy Distribution Report
AOTM	Autonomic Ownership Takeover Manager	CPx	Cache Partition
APARs	Authorized Program Analysis Reports	CRC	Cyclic Redundancy Check
APF	An Authorized Program Facility	CSG	Copy Source Grid
API	Application Programming Interface	CSP	Common Scratch Pool
APIs	Application Programming Interfaces	CSR	Certificate Signing Request
APM	Advanced Policy Management	CSSs	Channel Subsystems
AS	Auto Switchable	CSV	Command-Separated Values
ATAM	Automatic Tape Allocation Manager	CTM	CipherTrust Manager
ATCD	Advanced Type C Data	CU	Control Unit
ATDD	Advanced Type D Data	CUA	Change Use Attribute
ATED	Advanced Type E Data	CUADD	CUs, specifying the logical unit address
ATKE	Advanced Type K Economy	CUIR	Control Unit Initiated Reconfiguration
ATL	Automated Tape Library	CUs	control units
ATLE	Advanced Type L Economy	DAA	Dynamic Allocation Assist
ATLM	Advanced Type M Economy	DASD	Direct Access Storage Device
AWS	Amazon Web Services	DC	Data Class
BOT	Beginning Of Tape	DCT	Deferred Copy Throttling
BTLS	Basic Tape Library Support	DD	Data Definition
BVIR	Bulk Volume Information Retrieval	DDM	Disk Drive Module
CA	Certificate Authority	DDP	Dynamic Disk Pooling
CAP	Cartridge Assignment Policy	DDR	Dynamic Device Reconfiguration
CAPs	Cartridge Assignment Policies	DEF	Deferred
CDS	Control Data Set	DFP	Data Facility Product
CDSs	Control Data Sets	DFSMS	Data Facility Storage Management Subsystem
CE	Channel End	DHCP	Dynamic Host Configuration Protocol
CHPIDs	Css, Which Channel-path Identifiers	DJ	Dump Job
CI	Converter/Interpreter	DN	Distinguished Name
		DNS	Domain Name Server

DOWTO	DD DISP=SHR,DSN=&TOOLHLQ..IB MTOOLS.JCL	GTS	Global Technology Services
DR	Disaster Recovery	GUI	Graphical User-Interface
DRAID	Distributed RAID	Gbps	Gigabit per second
DRI	Data Retention Infrastructure	HA	High availability
DSPs	Dynamic Support Programs	HBA	Host Bus Adapter
DWDM	Directors Or Dense Wave Division Multiplexer	HCD	Hardware Configuration Dialog
ECC	Electronic Customer Care	HCF	Highest Common Format
ECST	Emulated Enhanced Capacity Cartridge System Tape	HCR	Host Console Request
EDT	Eligible Device Table	HDD	Hard Disk Drive
EDaR	Encryption Of Data At Rest	HLQs	High-level Qualifiers
EEH	Extended Error Handling	HMC	Hardware Management Console
EETC	Enterprise Economy Tape Cartridge	HSM	Hierarchical Storage Manager
EK	Encryption Key	HTTPS	Hypertext Transfer Protocol Secure
EKM	Encryption Key Manager	IART	Initial Access Response Time
EKS	Encryption Key Server	ICF	Integrated Catalog Facility
EOF	End Of File	ICL	Integrated Cartridge Loader
EOV	End Of Volume	ICMP	Internet Control Message Protocol
EPO	Emergency Power Off	IDRC	Improved Data-recording Capability
ERA	Error Recovery Action	IETF	Internet Engineering Task Force
ESCON	Enterprise Systems Connection	IOCP	Input/output Configuration Program
ETC	Enterprise Tape Cartridge	IODF	Input/output Definition File
ETCL	Enterprise Extended-length Tape Cartridge	IPL	Initial Program Load
FB	Fixed Block	IPP	Ibm Proprietary Protocol
FC	Feature Code	IPv6	Internet Protocol Version 6
FDE	Full Disk Encryption	ISL	Inter-switch Link
FEOV	Force End-of-volume	ISMF	Interactive Storage Management Facility
FICON	Fibre Channel Connection	ISO	International Organization For Standardization
FIFO	First-in First-out	IT	Information Technology
FL221	Function Level 221	JCL	Job Control Language
FQDN	Fully Qualified Domain Name	JES2	Job Entry Subsystem 2
FST	Field Support Test	JSON	JavaScript Object Notation
GB	Gigabyte	KMP	Key Management Interoperability Protocol
GDGs	Generation Data Groups	KiB	Kibibyte
GDL	Guadalajara Development Lab	LACS	Library Automation Communication Services
GDPS	Geographically Dispersed Parallel Sysplex	LAN	Local Area Network
GGM	Grid-to-Grid Migration	LCS	Library Control System
GKLM	Guardium Key Lifecycle Manager	LCU	Logical Control Unit
GMT	Greenwich Mean Time	LDAP	Lightweight Directory Access Protocol
GRS	Global resource serialization	LDGs	Library device groups
		LIBID	LIBRARY-ID

LIC	Licensed Internal Code	PG0	Preference Group 0
LM	Library Manager	PG1	Preference Group 1
LPAR	Logical Partition	PMR	Problem Management Report
LRU	Least Recently Used	PTFs	Program Temporary Fixes
LTO	Linear Tape-open	PTP	Peer-to-peer
LW	Longwave	QSAM	Queued Sequential Access Method
LWIKS	Lightweight Infrastructure Key Store certificate	QoS	Quality Of Service
LWORM	Logical Write Once Read Many	RACF	Resource Access Control Facility
MAC	Media Access Control	RAM	Random Access Memory
MB	Megabyte	RBAC	Role-based Access Control
MBps	Megabytes per second	RCL	Recall
MC	Management Class	REST	Restful
MDE	Microcode Detected Error	RFC	Request For Comments
MDRs	Miscellaneous Data Records	RMM	Removable Media Manager
MDS	Main Device Scheduler	RMS	Removable Media Services
MDisk	Managed Disks	ROR	Read Only Recovery
MES	Miscellaneous Equipment Specification	ROT	Read Ownership Takeover
MI	Management Interface	RPO	Recovery Point Objective
MIH	Missing Interrupt Handler	RPQ	Request For Price Quotation
MIMs	Media Information Messages	RTO	Recovery Time Objectives
MIs	Management Interfaces	RUN	Rewind Unload
ML2	Migration Level 2	SAA	Scratch Allocation Assistance
MLTCPY	Multi-Copy	SAN	Storage Area Network
MOD	Modify time	SC	Storage Class
MSAD	Microsoft Active Directory	SCDS	Source Control Data Set
MTU	Maximum Transmission Unit	SDAC	Selective Device Access Control
MiB	Mebibytes	SDEPTH	Setup Depth Parameter
MiBps	Measures Throughput	SDT	Secure Data Transfer
NC	No Copy	SDWF	Synchronous Deferred on Write Failure
NTP	Network Time Protocol	SEC	Securities and Exchange Commission
OAM	Object Access Method	SG	Storage Group
OCE	Open Close End-of-volume	SHA	Secure Hash Algorithm
OEM	Original Equipment Manufacturer	SIMs	Service information messages
OS	Operating System	SLA	service level agreement
OSMC	Oam Storage Management Component	SMC	Synchronous mode copy
PB	Petabytes	SMF	System Management Facilities
PCHIDs	Physical Channel Ids	SMS	Storage management subsystem
PCIe	Peripheral Component Interface Express	SNMP	Simple Network Management Protocol
Pds	Partitioned Data Sets	SNS	Surfacing Sense Status Information
Pdse	Partitioned Data Set Extended	SPAF	Spool Access Facility
PFE	Product Field Engineering	SQL	Structured Query Language
		SSA	Serial Storage Architecture

SSD	Sync-Sync Deferred	WTO	Write To Operator
SSD	Solid-state Drive	WTOR	Write To Operator With Reply
SSI	Subsystem Interface	dB	deciBel
SSIC	System Storage Interoperation Center		
SSL	Secure Sockets Layer		
SSPC	System Storage Productivity Center		
SSR	Service Representative		
STA	Simple Tape Attach		
STEP4	Step		
STSM	Senior Technical Staff Member		
SVC	Starts Supervisor Call		
SW	System Shortwave		
SWA	Scheduler Work Area		
TB	Terabytes		
TCDB	Tape Configuration Database		
TCT	Transparent Cloud Tiering		
TD	Time Delayed		
TLS	Tape Library Support		
TLS	Transport Layer Security		
TMC	Tape Management Catalog		
TMM	Tape Mount Management		
TMS	Tape Management System		
TOD	Time Of Day		
TSG	Technical Support Group		
TSO	Time Sharing Option		
TSSC	Total Storage System Console		
TVC	Tape Volume Cache		
UAT	User Acceptance Test		
UDP	User Datagram Protocol		
UIs	User Interfaces		
URI	Uniform Resource Identifier		
UTC	Utc		
UX	Using WC		
VGS	VSE Guest server		
VIO	Virtual I/O		
VOLSER	Volume Serial Number		
VPD	Vital Product Data		
VRSS	Vital Record Specifications		
VSN	Virtual Server Network		
VTS	Virtual Tape Server		
WAN	Wide Area Network		
WORM	Write Once Read Many		
WOT	Write Ownership Takeover		

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

IBM Redbooks publications

The following IBM Redbooks publications provide more 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 Connectivity Handbook*, SG24-5444
- ▶ *IBM z/OS DFSMSHsm Primer*, SG24-5272
- ▶ *Introduction to IBM S/390 FICON*, SG24-5176
- ▶ *Introduction to SAN Distance Solutions*, SG24-6408

For more information about the TS7770 Cloud Object Storage solution and how to implement and integrate this solution into your enterprise, see *IBM TS7700 R5.4 Cloud Storage Tier Guide*, [REDP-5573](#).

For more information about planning and implementing the function of using the TS7700 as an optional target for DS8000 Transparent Cloud Tier that uses DFSMS, see the following publications:

- ▶ *IBM TS7700 R5.4 DS8000 Object Store User's Guide*, [REDP-5583](#)
- ▶ *IBM DS8000 Transparent Cloud Tiering: DS8000 Release 9.3*, [SG24-8381](#)

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, drafts, and other materials, at the following website:

ibm.com/redbooks

Other publications

The following 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 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 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 TS7700 Introduction and Planning Guide*, GA32-0567
- ▶ *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 DFSMSsdfp Storage Administrator*, SC23-6860
- ▶ *z/OS DFSMSsdfp 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

Online resources

The following web pages include more information:

- ▶ Common Information Model (CIM):
<http://www.dmtf.org/standards/cim/>
- ▶ IBM Business Continuity and Recovery Services:
<http://www.ibm.com/services/continuity>
- ▶ IBM Security Guardium Key Lifecycle Manager documentation:
<https://www.ibm.com/docs/en/sgklm>
- ▶ IBM TS3500 tape library at IBM Documentation:
<https://www.ibm.com/docs/en/ts3500-tape-library>
- ▶ IBM TS4500 Tape Library at IBM Documentation:
<https://www.ibm.com/docs/en/ts4500-tape-library>
- ▶ TS7700 documentation at IBM Documentation:
<https://www.ibm.com/docs/en/ts7700-virtual-tape>
- ▶ Web-based Enterprise Management (WBEM):
<http://www.dmtf.org/standards/wbem/>

Technical documents on the IBM Support 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 documents are being created.

To ensure that you reference the newest document, search on the IBM Support website:

<https://www.ibm.com/support/home>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



IBM TS7700 Release 5.4 Guide

SG24-8464-04

ISBN 0738419001



(2.0" spine)
2.0" <-> 2.498"
1052 <-> 1314 pages



SG24-8464-04

ISBN 0738461857

Printed in U.S.A.

Get connected

