



November 2023

**IBM® TS7700 Series
z/OS Host Command Line Request User's
Guide Version 5.4**

Owner: Kousei Kawamura (e35223@jp.ibm.com)
IBM Systems Development, IBM Japan

Yuki Asakura (Yuki.Asakura@ibm.com)
IBM Systems Development, IBM Japan

Takahiro Tsuda (e34096@jp.ibm.com)
IBM Systems Development, IBM Japan

Contents

1	Introduction	5
1.1	<i>Summary of Changes</i>	5
1.2	<i>Configurations Supported</i>	12
1.3	<i>Code Requirements</i>	12
2	Library Request Host Console Command Overview	14
3	Host Command Line Request	15
3.1	<i>Responses to Library Requests</i>	37
3.1.1	CACHE2 Response	38
3.1.2	COPY Response	44
3.1.2.1	COPY, ACTIVITY	44
3.1.2.2	COPY, SUMMARY	48
3.1.2.3	COPY, QUEUE	56
3.1.3	COPYEXP Response	62
3.1.4	GRIDCNTL Response	65
3.1.5	LVOL Response	67
3.1.5.1	LVOL, <volser>, INFO(, FLASH)	67
3.1.5.2	LVOL, <volser>, CLDINFO	81
3.1.5.3	LVOL, <volser>, CLDVERS, <page index>	86
3.1.5.4	LVOL, <volser>, PREFER/MIGRATE	89
3.1.5.5	LVOL, <volser>, REMOVE	90
3.1.5.6	LVOL, <volser>, COPY, KICK/FORCE	94
3.1.5.7	LVOL, <volser>, COPY, CANCEL	96
3.1.6	CLDBKUP Response	98
3.1.6.1	SHOW <page index>	98
3.1.7	PDRIVE Response	100
3.1.8	POOLCNT Response	103
3.1.9	PVOL Response	106
3.1.10	RECALLQ Response	112
3.1.11	RRCLSUN Response	115
3.1.12	SETTING	117
3.1.12.1	Alert Settings	120
3.1.12.2	Cache Settings	134
3.1.12.3	Throttle Settings	140
3.1.12.4	Reclaim Settings	142
3.1.12.5	Device Allocation (Devalloc) Settings	145
3.1.12.6	Copy Settings	147
3.1.12.7	Link Failover Settings	150
3.1.12.8	Delexp Count Settings	151
3.1.12.9	Existdel Settings	152

3.1.13	SETTING2	165
3.1.13.1	Scratch Settings	167
3.1.13.2	Alert Settings	168
3.1.13.3	CACHE Settings	176
3.1.13.4	THROTTLE Settings	179
3.1.13.5	PHYSLIB Settings	181
3.1.13.6	RECALL Settings	187
3.1.13.7	BVIR Settings	189
3.1.13.8	CENCRYPT Settings	189
3.1.13.9	COPY Settings	191
3.1.13.10	VOLINVT Settings	193
3.1.14	STATUS	200
3.1.14.1	GRID Response:	200
3.1.14.2	GRIDLINK Response:	206
3.1.14.3	RDP Response:	216
3.1.14.4	GRLNKACTION Response:	227
3.1.15	COPYRFSH	238
3.1.16	OTCNTL	246
3.1.17	DRSETUP	259
3.1.18	PARTRFSH	260
3.1.19	LOWRANK	264
3.1.20	GGM	268
3.1.21	DIAGDATA	269
3.1.22	OTM	270
3.1.23	FENCE	277
3.1.24	CUIR/LDRIVE	278
3.1.25	CLDSET	279
3.1.25.1	CPMCNTH/CPMCNTL	279
3.1.25.2	CLDPRIOR	280
3.1.25.3	CRCCNT	281
3.1.25.4	CDELCNT	282
3.1.25.5	CPMTOUT	282
3.1.25.6	CRCTOUT	283
3.1.25.7	CDELTOUT	284
3.1.25.8	CENABLMT	285
3.1.25.9	PFRCCPG0	289
3.1.25.10	PFRCCTDL	290
3.1.25.11	PFRCCTRY	290
3.1.25.12	USEHIGH	291
3.1.25.13	USELOW	292
3.1.25.14	USESHOW	293
3.1.26	CLDVR	295
3.1.27	SERVICE	296
3.1.28	LWORMR	305
3.1.29	OCOPY	305

3.1.29.1	OCOPY, SUMMARY	305
3.1.30	OBJSET1	309
3.1.30.1	Alert Settings	310
3.1.31	CRCSET	311
3.1.31.1	GEN Settings	312
3.1.31.2	VERIFY Settings	312
3.1.32	zTape Air-Gap (FC5995)	315
	Error Responses.....	316
4	Appendix – A. Summary table of SETTING/SETTING2/OBJSET1 keywords.....	320
5	Disclaimers:	332

1 Introduction

The IBM TS7700 Series is the latest in the line of tape virtualization products that has revolutionized the way mainframe customers utilize their tape resources. Tape virtualization subsystems have become an essential part of most mainframe customer's operations. Massive amounts of key customer data are placed under the control of the subsystem. The IBM TS7700, with its virtual tape drives, disk cache and integrated hierarchical storage management, is designed to perform tasks with no customer involvement once it has been configured. However, as a customer, you may still need (and sometimes just want) to understand what is going on under the covers to:

- Assess the current utilization of the resources under the management of the TS7700
- Perform problem determination
- Perform proactive management of the resources

This white paper describes a facility of the TS7700 that supports a new z/OS Library Request host console command to allow an operator to request information pertaining to the current operational state of the TS7700, its logical and physical volumes and its physical resources. Although the command is primarily for information retrieval, it can also be used to initiate outboard operations.

It is assumed throughout this white paper that the reader is familiar with the TS7700 and the use of host console commands and responses.

With the 3.2 code release, the TS7700 Management Interface allows an operator to issue a Library Request host console command as if it was issued from the z/OS host. The command and its result are displayed in LI REQ command window. A task is created for each LI REQ request from MI.

1.1 Summary of Changes

Changes for version 1.1

- A new request has been added for control of the copy operations of a TS7700 cluster
- A new request about information about the status of the grid links for a TS7700 cluster has been added.
- Copy status (enabled, disabled) has been added to the Status, Grid response
- Link status has been moved from grid status to a new request that reports on the state of the links between all of the clusters in a grid configuration.
- Updated the definition of the cache preference field in the data returned to the LVOL request.
- Added more error information when an LVOL request fails because information from the owning cluster cannot be obtained.

Changes for version 1.2

- A new request has been added to set and display workflow management and alert threshold settings.

Changes for version 1.3

- A new request has been added to allow a logical volume's cache management preference to be modified.
- An additional column has been added to the grid status response to indicate whether a distributed library has physical tape attached to it.
- An additional column has been added to the grid status response to indicate if a distributed library is a member of a cluster family.

- Added setting reclaim request to be able to limit the number of reclaims.
- Added setting dctavgtd setting request to be able to control deferred copy throttling.
- Group name has been added to the logical volume response.
- Logical volume response has additional distributed library column showing if a copy was removed from that distributed library.
- Logical volume response has additional field displaying how many copies were removed from the grid.
- An additional field has been added to the logical volume response indicating if the volume is Lworm.

Changes for version 1.4

- Added new requests to allow a logical volume to be removed or promoted to the front of removal queue.
- Allow the existing "PREFER" keyword to be applicable to disk-only configuration as well.
- Added a new setting to allow disable/enable removal.
- Added new request to attain LVOL removal information.
- Changed DCTAVGTD description to average compressed host I/O rate vs average host I/O rate.

Changes for version 2.0

- Changed version number to match code release level.
- The changes below require Release 2.0 code on the TS7700.
- Add RRCLSUN – Resident on ReCall SUNset - Useful for accelerating the replacement of older media with newer media
- Add REMOVE INFO keywords to the LVOL command - Used for retrieving the REMOVE status for logical volumes in a TS7720
- Add DEVALLOC, SCRATCH keywords to the SETTING command - Used to enable or disable Scratch Allocation Assist (SAA)
- Add DEVALLOC, PRIVATE keywords to the SETTING command - Used to enable or disable Device Allocation Assist (DAA)
- Add CPYCNT, RUN or DEF keywords to the SETTING command - Used to set the number of concurrent copy threads for processing RUN and Deferred copies for a target cluster
- STATUS GRID now reports which family a cluster belongs to and the number of active RUN and Deferred copies per cluster.
- LVOL response now reports the number of copies removed from TS7720 clusters in the grid, whether the logical volume has been removed from this cluster, and whether the logical volume is a WORM volume or not.
- Changed Default for SETTING RECLAIM RCLMMAX description to 0 instead of -1. This is not related to Release 2.0 but is updated in this release.

Changes for version 2.1

- Changed version number to match code release level.
- Changed every 256KB copy transfer to 32KB for DCOPYT keyword
- The changes below require Release 2.1 code on the TS7700
- Added warning message for CACHE keyword response
- A new request has been added to enable or disable the IP Link Failover for Remote Mounts
- LVOL response now reports Sync Mode Copy
- STATUS GRID now reports Sync-Deferred jobs and the code level for each Distributed Library ID in the grid

- A new request has been added to move logical volumes up front in the copy queue
- A new request has been added to show the Active Copy jobs
- Added REMVTHR setting request to be able to control the removal threshold
- Added DELEXP setting request to be able to edit the Delete Expire field
- LVOL REMOVE now allows to remove volumes on a TS7740 with copy mode of “Exist” or “No-Copy”
- Added Copy Queue Building Alert
- Set the Version to 2 for all the Host Console Request

Changes for version 2.1a, b

- Add CPYRFSH command – Used to refresh copy policy and queue a copy job on the copy target cluster(s) without mounting/demounting a volume.
- Add REMVTHR command – Used to set the threshold where automatic removal is used to remove volumes from the TS7720 cache.

Changes for version 2.1c

- Fix code level typo 8.21.1.x should have been 8.21.0.x

Changes for version 2.1d

- Fix code level for Copy Refresh should be 8.21.0.118 instead of 8.21.0.x.

Changes for version 3.0

- Added the disk encryption status to the CACHE RESPONSE.
- Changed COPYRFSH output version from 1 to 2 because of the newly added error text.
- Added an example of COPYRFSH use case.
- Added LVOL COPY FORCE command - Used to put a copy job on the Distributed Library that already removed the data.
- Changed LVOL COPY KICK output version from 2 to 3 because of the newly added error text.
- Added COPY SUMMARY command – Used to get the summary of copy jobs.
- LVOL REMOVE is capable of removing a copy job on TS7720.
- LVOL REMOVE is allowed when auto removal is disabled.
- Added SETTING ALERT DEFDEG command – Used to prevent a composite library from entering ‘degraded’ state when SyncDeferred or ImmediateDeferred condition occurs.
- Added more explanation of “DELETE EXPIRED” field in LVOL information.
- Added SETTING ALERT REMOVMSG command – Used to suppress the operator intervention event (operator message) when Auto Removal is initiated/stopped on TS7720.
- Added the operator message description generated by SETTING ALERT PCPYLOW/PCPYCRIT.
- Changed SETTING output version from 2 to 3 because of the output format change.
- Changed DCTAVGTD/DCOPYT description to explain the DCT conditions in detail.
- Added Appendix-A SETTING command summary table.

Changes for version 3.0a

- Add column header text to the corresponding field in the response definitions for CACHE, PDRIVE, POOLCNT, RECALLQ, STATUS GRID, and GRIDLINK

Changes for version 3.0b

- Corrected CACHE output for the changes in Release 3.0.

Changes for version 3.0c

- Added SETTING CACHE CPYPRIOR/ CPYPRITH command – Used to limit premigration resources under Grid copy activity.
- Added SETTING DEVALLOC FAMILY command – Used to change DAA (Device Allocation Assistance) for private mount behavior in regard to TVC (Tape Volume Cache) selection algorithm with cluster family configuration.
- Added OTCNTL command – Used to transfer the ownership of logical volumes in the background

Changes for version 3.1

- Introduce a revision number in addition to a version number in the response header.
- CACHE output includes Flash Copy volume information and “ADJUSTED CACHE USAGE” field is added to TS7720 cache state.
- COPY ACTIVITY provides the number of timeout copies.
- COPY SUMMARY supports Time Delayed copy queue status.
- Add LVOL FLASH command – Used to provide Flash copy volume status.
- LVOL provides Time Delayed copy status.
- Change LVOL ‘COPY Q’ output to provide Sync Copy queue correctly.
- Add PVOL DELETE command – Used to eject inaccessible and empty physical volume.
- PVOL provides additional unavailable volume state.
- Add SETTING ALERT LINKDEG command – Used to prevent a composite library from entering ‘link degraded’ state when Grid link degradation occurs.
- Add SETTING CACHE RBPRIOR/ RBTHLVL command – Used to control cache rebuild priority and threshold.
- Add SETTING EXISTDEL command – Used to determine how to handle ‘E’ copy mode volume at mount/demount.
- Add COPYRFSH error message when assigned management class name is missing.
- Add OTCNTL error message when no volumes to be transferred are found.
- Add DRSETUP command – Used for Flash Copy for DR testing function.

Changes for version 3.2

- Add column in the table if the keyword supports TS7720 Tape Attach (TS7720T).
- Add new SETTING keywords:
ALERTS RSDTLOW/RSDTHIGH, COPY SCRATCH, EXISTDEL CRITERIA/WHEN
- Add new SETTING2 keywords:
SCRATCH PFRLOCO, CACHE MAXLGMC.
- Add PARTRFSH command – Used to change cache partition of logical volume on TS7720T.
- CACHE output change (TS7720, 7740 and 7720T).
- Change cluster numbering for COPY ACTIVITY/SUMMAR output.
- LVOL output accommodates 496 devices (RPQ only) and cache partition number.
- Add TS7720T comments for RESDHIGH/RESLOW/PMTHLVL.
- Add OTCNTL CONFIG specific error message.
- Update error messages when invalid key words are issued.

Changes for version 3.3

- Support TS1150 tape drive.
- Add new error message for GRIDLINK when the cluster is under maintenance.

- Add new SETTING2 keyword:
PDRVSLOW, PDRVSCRIT, PHYSLIB
- Add new SETTING keyword RCLMSMAX.
- Add GGM command – Used for GGM function
- Change OTCNTL command and output
- Add DRSETUP, SELFLIVE keyword
- Add LOWRANK command – Used to control the mount TVC or copy source cluster
- Add description on COPYRFSH recall performance

Changes for version 4.0

- Add TS7700 model descriptions (use the new notation of each TS7700 models throughout the document).
- Add more description on PMTHLVL/REMOVTHR/COPYFT to clarify how it works.
- Add SETTING2 keyword OVRSPVL.
- Add explanation of “SVC” value for STATUS, GRID output.
- GRIDCNTL error message enhancement.
- LOWRANK error message enhancement.
- DAA (Device Allocation Assistance) takes account of PHYSLIB, TVCWDEG setting.

Changes for version 4.0.1

- This version includes the updates supported only at the code level of R3.3 PGA2 (8.33.2.9) as of February, in 2017.
- Add DIAGDATA command – Used to get the command handshake elapsed time, timeout count and error count among the clusters in the Grid.
- PARTRFSH command updates cache partition/pool number.
- Use NULL category (0000) to delete the defined category in OTCNTL, CONFIG, CAT command.
- STATUS, GRID output format change.
- Add clearer explanation of MAXLGMC regarding Sync Copy mode case (not specific to R3.3 PGA2).
- Add more description on COPYFT to clarify how it works (not specific to R3.3 PGA2).

Changes for version 4.1.1

- This version includes the updates supported at the code level of R4.1.1 GA (8.41.100.15).
- Add SETTING2, RECALL, BLKRCLSZ keyword to tune the size of data to be recalled for Grid copy.
- Add OTM command.
- Add DRSETUP, LIVEACC keyword to allow the live copy volume access from a DR test host.

Changes for version 4.1.2

- This version includes the updates supported at the code level of R4.1.2 GA (8.41.200.xx).
- Add FENCE keywords for Grid Resiliency function.
- Add CUIR/LDRIVE keywords for C.U.I.R. (Control Unit Initiating Reconfiguration) function.
- LVOL output has the new compression method attributes.
- Add a brief description on what is updated for GGM/DIAGDATA commands.
- Updated EXISTDEL description to explain how volumes are now marked as "removed" when deleted.

- Revise or add a description of COPYFSC/PMMPRIOR/REMOVTHR/DCTAVGTD in SETTING.
- Default value of CPYPRIOR in SETTING is changed.
- Allow PARTRFSH keyword to run on TS7700D.
- Add SETTING, THROTTLE, DCTCURTD keyword.
- Add SETTING2, BVIR, TIMEOUT keyword
- Add SETTING2, THROTTLE, LNKSPED keyword
- Add SETTING2, ALERT, LMTDTHR keyword

Changes for version 4.2

- Describe TS7700C which is TS7700 with cloud enablement feature
- Add TS7700C column to the tables where which model supports the request.
- TS7700C support on the existing LI REQs (CACHE/SETTING/SETTING2/PARTRFSH/COPYRFSH/GGM)
- Add the longest age of Family Deferred copy jobs in COPY, SUMMARY response.
- Add LVOL, <volser>, INFO/CLDINFO and separate LVOL response section based on each keyword.
- Add more Deferred Read Throttling explanation in DCTAVGTD/DCTCURTD.
- Add STATUS, RDP keyword to return information of Read Diagnostic Parameters of FICON port
- Add STATUS, GRLNKACT keyword to return Grid link activity information
- Add more copy source selection behavior in COPYRFSH.
- Add the description about Flash Copy asynchronous operation
- Add the restriction to use GGM when the copy source cluster is at the code level 8.42.x.x or later.
- Add CLDSET to monitor/control cloud data management settings.

Changes for version 5.0

- Add TS7770 (VED) description
- Add new PVOL RELABEL keyword to mark a physical volume for re-labeling
- Add new CACHE2 keyword and output (CACHE is deprecated).
- Add new COPY, QUEUE keyword
- Add new LVOL, COPY, CANCEL keyword
- Add new COPYRFSH, CLDGHOST keyword
- Add new PARTRFSH, MMOUNT keyword
- Add new SERVICE keyword to initiate service-prep/cancel from LI REQ
- Add new SETTING, EXISTDEL, AUTINTVL/AUTCOUNT keyword
- Add new SETTING2, ALERT, RSDOHIGH/ RSDOLOW keyword
- Add new SETTING2, CENCRYPT keyword
- Add new SETTING2, COPY, COPYRFSH/DEF, PAUSE/UNPAUSE keyword
- Add new SETTING2, PHYSLIB, SLDMPRI/ MAINT keyword
- Add new SETTING2, ALERT, CAGAHIGH/CAGALOW/CAGHIGH/CAGLOW keyword
- Add new SETTING2, VOLINVT, LPAGRP keyword
- Add description about Secure Data Transfer setting to STATUS, GRIDLINK output
- Add copy job age to COPY, ACTIVITY output
- Change maximum value of DELEXP count from 2000 to 5000

- Add an attention to specify a distributed (not composite) library for SETTING/SETTING2 command without the second keyword to get the settings from the target library correctly.
- Add more explanation about SETTING, CACHE, PMTHLVL/RBPRIOR/RBTHLVL.
- Add more explanation about SETTING2, PHYSLIB, PRETHDEG.
- Add more explanation about SETTING2, RECALL, BLKRCLSZ.
- Add more explanation about STATUS, RDP.
- Add the description about the variable host performance related to TS7770 with smaller number of CSB/XSB and its mitigation on the concurrent RUN copy counts (adjusted by SETTING, CPYCNT, RUN).

Changes for version 5.0a

- Add OTCNTL MOVE1 keyword
- Add FAILED TRANSFER VOLS field to OTCNTL STAT output
- Add more explanation about LOWRANK
- Add the reference to send the LI REQ command output to a data set.

Changes for version 5.1

- Add new format of LVOL, CLDINFO to show multiple cloud setting per cluster.
- Add new LVOL, CLDVERS keyword.
- Add new CLDBKUP, SHOW keyword.
- Add new CLDSET, PFRCCPG0 keyword.
- Add new CLDSET, PFRCCCTL keyword.
- Add new CLDSET, PFRCCTRY keyword.
- Add new CLDSET, USEHIGH keyword.
- Add new CLDSET, USELOW keyword.
- Add new CLDSET, USESHOW keyword.

Changes for version 5.1a

- Restore changes of version 5.0a missing in version 5.1.
- Add GGM copy and DS8000 Transparent Cloud Tiering statistics to STATUS, GRLNKACT output.
- Add more explanation about SETTING2,VOLINVT.

Changes for version 5.1b

- Add LWORM retention fields in LVOL, INFO output.
- Add new CLDVR keyword.
- Add new LWORMR keyword.

Changes for version 5.11b

- Update STATUS, GRIDLINK percent retransmitted description.

Changes for version 5.22

- Add new OCOPY, SUMMARY keyword.
- New OBJSET1 keyword. Moved RSDOHIGH|RSDOLOW to OBJSET1 from SETTING2.

Changes for version 5.22a

- Add the missing state in PVOL response (VOLUME STATE).

Changes for version 5.3

- Change LI REQ keyword table format.
- Add new drive format modes for TS1160 support.
- Add new CRCSET keyword.
- Add 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 (FC5995).

Changes for version 5.4

- Add 25GB scaling for SETTING COPY TIMEOUT.

1.2 Configurations Supported

The Library Request host console command function is supported on all configurations of the TS7700, including Grid configurations, that meet minimum code level requirements.

1.3 Code Requirements

Library Request host console command support was introduced with TS7700 code level 8.3.x.x. Although there are no specific library manager code changes to support the function, there are other library manager functions required for other functions that are part of the 8.3.x.x level of the TS7700 code.

For the latest level of support added with version 1.1, the TS7700 must be at a code level 8.4.1.x or later.

For the latest level of support added with version 1.2, the TS7700 must be at a code level of 8.5.x.x or later.

For the latest level of support added with version 1.3, the TS7700 must be at a code level of 8.6.x.x or later.

For the latest level of support added with version 1.4 of this document, the TS7700 must be at a code level of 8.7.x.x or later.

For the latest level of support added with version 2.0 and 2.0a of this document, the TS7700 must be at a code level of 8.20.x.x or later.

For the latest level of support added with version 2.1 of this document, the TS7700 must be at a code level of 8.21.x.x or later.

For the latest level of support added with version 2.1a and b of this document, the TS7700 must be at a code level of 8.21.0.118 or later.

For the latest level of support added with version 3.0, 3.0a and 3.0b of this document, the TS7700 must be at a code level of 8.30.x.x or later.

For the latest level of support added with version 3.0c of this document, the TS7700 must be at a code level of 8.30.1.x (R3.0 PGAx) or later.

For the latest level of support added with version 3.1 of this document, the TS7700 must be at a code level of 8.31.x.x or later.

For the latest level of support added with version 3.2 of this document, the TS7700 must be at a code level of 8.32.x.x or later.

For the latest level of support added with version 3.3 of this document, the TS7700 must be at a code level of 8.33.x.x or later.

For the latest level of support added with version 4.0 of this document, the TS7700 must be at a code level of 8.40.x.x or later.

For the latest level of support added with version 4.0.1 of this document, the updates included in the version is contained at a code level of 8.33.2.9 (3.3 PGA2) only as of January, 2017.

For the latest level of support added with version 4.1.1 of this document, the updates included in the version is contained at a code level of 8.41.100.15 (4.1.1 GA).

Updates contained within version 4.1.2 of this document are associated with the TS7700 release level 8.41.200.XX (4.1.2 GA).

Updates contained within version 4.2 of this document are associated with the TS7700 release level 8.42.x.x (4.2 GA).

Updates contained within version 5.0 of this document are associated with the TS7700 release level 8.50.x.x (5.0 GA).

Updates contained within version 5.0a of this document are associated with the TS7700 release level 8.50.1.25 (5.0 PGA1).

Updates contained within version 5.1 and 5.1a of this document are associated with the TS7700 release level 8.51.0.x (5.1 GA).

Updates contained within version 5.1b and 5.11b of this document are associated with the TS7700 release level 8.51.1.x (5.1 PGA1).

Updates contained within version 5.22 of this document are associated with the TS7700 release level 8.52.200.x (5.22 GA).

Updates contained within version 5.3 of this document are associated with the TS7700 release level 8.53.x.x (5.3 GA).

z/OS support is also required for this function and is provided at z/OS V1R6 and above. Refer to OAM APAR OA20065 and device services APARs OA20066, OA20067 and OA20313.

2 Library Request Host Console Command Overview

The Library Request host console command provides a simple way for an operator to determine the status of the TS7700, to obtain information about the resources of the TS7700 or to perform an operation in the TS7700. It could also be used with automation software to obtain and analyze operational information that could then be used to alert a storage administrator that there is something that should be examined further. The command requires a library name to be specified which could be a composite or a distributed library and it will also allow 1 to 4 keywords to be specified, with each keyword being a maximum of 8 characters. The specified keywords will be passed to the TS7700 identified by the library name to instruct it on what type of information is being requested or which operation is to be performed. Based on the operation requested through the command, the TS7700 will then return information to the host that will be displayed as a multi-line Write To Operator (WTO) message.

A TS7700 library is made up of a composite and one or more distributed libraries. The composite library represents the logical view of the aggregate of the underlying physical aspects of the library. In essence it is virtualizing the underlying physical TS7700s so they look like one library. From a host job processing perspective, the virtual tape device addresses and logical volumes it uses are part of the composite library although the actual resources used is on one of the underlying distributed libraries. A distributed library represents the view of the resources owned and managed by a specific TS7700 and associated physical library. For a single TS7700, it has a composite library view and a distributed library view. In a Grid configuration, there is a composite view for the entire configuration as well as individual distributed library views for each of the TS7700s. Most of the host requests are related to the physical resources of a TS7700 and as such are directed to the distributed library name for a specific TS7700. Logical volumes, however, have a view that spans all of the distributed libraries and is directed to the composite library name.

This white paper is organized to first explain the format of the Library Request host console command and supported keywords, followed by format descriptions for the responses.

3 Host Command Line Request

From the z/OS host console, an operator will enter the following command:

Command Syntax:

```
>> _LIBRARY_ _REQUEST_ _ , _library_name_ >
    | _LI_ | | _REQ_ |
> _ , keyword1
    | _ , keyword2 | | _ , L= _a_ |
    | _ , keyword3 | | _name_ |
    | _ , keyword4 | | _name-a_ | ><
```

Required Parameters

REQUEST | REQ

Specifies a request to obtain information from the TS7700 or to perform an outboard operation.

library_name

Specifies the library name associated with the TS7700 to which the request should be directed. The library name specified may be a composite or a distributed library and which library is applicable depends on the other keywords specified.

keyword1

Specifies which operation is to be performed at the TS7700.

Optional Parameters

The optional keyword parameters are dependent on the first keyword specified. Based on the first keyword specified, zero or more of the additional keywords may be appropriate.

keyword2

Specifies additional information in support of the operation specified with the first keyword.

keyword3

Specifies additional information in support of the operation specified with the first keyword.

keyword4

Specifies additional information in support of the operation specified with the first keyword.

L={a | name | name-a}

Specifies where to display the results of the inquiry: the display area (L=a), the console name (L=name), or both the console name and the display area (L=name-a). The name parameter can be an alphanumeric character string.

Note: the keywords specified must be from one to eight characters in length and can consist of alphanumeric (A-Z and 0-9) and the national character set (\$*#@#%). A keyword cannot contain any blanks. The only checking performed by the host is to verify that the specified keywords conform to the supported character set. The validity of the keywords themselves and the keywords in combination with each other is verified by the TS7700 processing the command.

If the request is specific to the composite library, then the composite library name must be specified. If the request is specific to a distributed library, then the distributed library name must be used. If a request for a distributed library is received on a virtual drive address on a TS7700 cluster of a different distributed library, the request is routed to the appropriate cluster for handling and the response is routed back through the requesting device address.

When the command is entered, z/OS will generate the following console message to indicate that the command has been initiated:

CBR1020I Processing LIBRARY command:

Request,library_name,keyword1,keyword2,keyword3,keyword4

When the TS7700 completes processing the request, it will return a set of formatted text line that will be displayed to the console or display area specified in the request. See the section on *Responses to Library Requests*.

If a request specifies an invalid keyword or the request cannot be routed to the appropriate TS7700 cluster for processing, a formatted error message is returned. See the section on *Error Responses*.

The following table lists the keyword requests supported, the type of information the request returns or the action taken and whether the request is applicable to a composite and/or distributed library.

At the code level 8.32.x.x, TS7720T which has a resident only partition and up to seven tape attached partitions is supported. The table has columns to provide if the keyword is supported on TS7740, TS7720 and TS7720T. Regarding the keyword supported on TS7720T, if it's not supported on TS7720 (TS7740), the command is applicable to the attributes of the tape attached (resident only) partition only.

At the code level 8.40.x.x, the new model TS7760 is introduced. At the code level 8.42.x.x, a cloud attachment feature can be applied to TS7760. At the code level 8.50.x.x, the new model TS7770 is introduced. The following model notations are used throughout this document:

- TS7740: V06 and V07
- TS7720: VEA and VEB with no tape library.
- TS7720T: VEB with tape library
- TS7760: VEC with no tape library
- TS7760T: VEC with tape library
- TS7760C: VEC with cloud (with no tape library)
- TS7770: VED with no tape library
- TS7770T: VED with tape library
- TS7770C: VED with cloud (with no tape library)
- TS7700: All models (TS77X0, TS77X0T and TS77X0C) are included.
- TS7700D: TS7700 Disk only models (TS7720, TS7760 and TS7770) are included.
- TS7700T: TS7700 Tape attach models (TS7720T, TS7760T and TS7770T) are included (TS7740 is NOT included in the white paper).
- TS7700C: TS7700 with cloud enablement feature (TS7760C and TS7770C).

- TS7700O: TS7700 with advanced object store feature for DS8000. TS7700O is not considered as a “model”. The feature can be applied to all models of TS7700D/TS7700T/TS7700C.



November 2023

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
CACHE2				Requests information about the current state of the cache and the data managed within it associated with the specified distributed library. “CACHE” was deprecated at R5.0.	N	Y	Y	Y	Y	Y	Y
COPY	ACTIVITY			Requests information about Active Copy jobs	N	Y	Y	Y	Y	Y	Y
		RUN									
		SDEF									
		IDEF									
		DEF									
	SUMMARY			Requests summary information about all the copy jobs.	N	Y	Y	Y	Y	Y	Y
	QUEUE	RUN	NEXT/ OLDEST/ ACTIVE	Requests information about individual copy job	N	Y	NA	Y	Y	Y	Y
		DEF									
		RFSH									
		FDEF									
		TDEL									
		RCAL									
COPYEXP	<volser>	RECLAIM		Requests that the specified physical volume that has been exported	N	Y	Y	N	Y	N	Y

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				previously in a copy export operation, be made eligible for priority reclaim.							
		DELETE		Requests that the specified physical volume that has been exported previously in a copy export operation, be deleted from the TS7700 database. The volume must be empty.	N	Y	Y	N	Y	N	Y
GRIDCNTL	COPY	ENABLE/ DISABLE		Requests that copy operations for the specified distributed library be enabled/disabled. Enabled copy operations can again use the specified distributed library as the source or target for copies. Disabled copies that are in progress are allowed to complete, but no new copies using the specified distributed library as the source or target are initiated.	N	Y	Y	Y	Y	Y	Y

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
LVOL	<volser>	INFO		Requests information about a specific logical volume.	Y	N	Y	Y	Y	Y	Y
			FLASH								
		FLASH									
		CLDINFO	<page index>	“LVOL, <volser> with no 3 rd KW” was deprecated at R4.2 and replaced to the request with 3 rd KW “INFO”.	N	Y	N	N	N	Y	Y
		CLDVERS									
		PREFER									
		MIGRATE		Requests a change in the cache management for a logical volume.	N	Y	Y	Y	Y	Y	Y
		REMOVE									
		REMOVE	PROMOTE								
		REMOVE	INFO								
		COPY	KICK	Kick requests to move a logical volume to the front of the copy queue FORCE is used to put a copy of the removed logical volume into the queue and move it to the front of the copy queue. CANCEL is used to cancel an in-flight copy job and delay it to be processed later.	N	Y	Y	Y	Y	Y	Y
			FORCE								
			CANCEL								
CLDBKUP	SHOW			Requests information about cloud backups.	N	Y	N	N	N	Y	Y
		<page index>									

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
PDRIVE				Requests information about the physical drives and their current usage associated with the specified distributed library.	N	Y	Y	N	Y	N	Y
POOLCNT				Requests information about the media types and counts, associated with a specified distributed library, for volume pools beginning with the value in keyword 2.	N	Y	Y	N	Y	N	Y
	<pool index>										
PVOL	<volser>			Requests information about a specific physical volume.	N	Y	Y	N	Y	N	Y
		DELETE		Requests the specified physical volume record to be deleted from the TS7700 database. The specified physical volume must be empty and not physically in the library.	N	Y	Y	N	Y	N	Y
		RELABEL	YES/NO	YES – Instructs the TS7700 to mark this physical volume for re-labeling the next time the cartridge	N	Y	NA	N	Y	N	Y

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				is written from the beginning of tape. NO – Cancels the re-labeling of the physical volume.							
RECALLQ	<volser>			Requests the content of the recall queue starting with the specified logical volume.	N	Y	Y	N	Y	N	Y
		PROMOTE		Requests that the specified logical volume be promoted to the top of the recall queue.	N	Y	Y	N	Y	N	Y
RRCLSUN	ENABLE/ DISABLE			Requests that force residency on recall for sunsetting function be enabled/disabled.	N	Y	Y	N	Y	N	Y
	STATUS			Requests to display residency on recall for sunsetting function setting.							
SETTING	See description section			Requests to change various settings.	See description section						
SETTING2	See description section			Requests to change various settings.	See description section						
STATUS	GRID			Requests information about the copy, reconcile and ownership takeover status of	Y	N	Y	Y	Y	Y	Y

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				the libraries in a Grid configuration							
	GRIDLINK			Requests information about the status and performance of the links between the TS7700s in the Grid configuration	N	Y	Y	Y	Y	Y	Y
	RDP	<FxPyz>		Requests Read Diagnostic Parameters (RDP) information of a specified FICON port.	N	Y	Y	Y	Y	Y	Y
			REFRESH								
	GRLNKACT			Requests information about Grid link point-in-time activity.	N	Y	Y	Y	Y	Y	Y
COPYRFSH	<volser>			Refresh copy policy and queue a copy job on the copy target cluster(s) without mounting/demounting a volume.	N	Y	Y	Y	Y	Y	Y
		<source ID>									
			NORECALL								
		CLDGHOST									
OTCNTL	START	<source ID>		Request to transfer the ownership of logical volumes in the background.	N	Y	Y	Y	Y	Y	Y
			<max vol num>								
	STOP										
	MOVE1	<volser>									
	DIST										
		<starting category>									
	STAT										

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
DRSETUP	CONFIG	CAT	<category>								
		DELAY	<delay>								
	<DR family name>	ADD	<cluster ID>	Add a cluster to the DR family	Y	N	Y	Y	Y	Y	Y
		REMOVE		Remove a cluster from the DR family							
		WP	ENABLE/ DISABLE	Enable/disable write protect mode within the DR family							
		FLASH	ENABLE/ DISABLE	Enable/disable Flash Copy within the DR family							
		DOALL	ENABLE/ DISABLE	Enable/disable write protect mode and flash copy with a single command							
		LIVECOPY	FAMILY/ NONE	Allow or disallow usage of a live copy within the DR family							
		SELFIVE	ENABLE/ DISABLE	Enable/disable accessing live copy created after time zero	N	Y	Y	Y	Y	Y	Y
	SHOW			View information about the DR family	Y	N	Y	Y	Y	Y	Y
		<DR family name>									
	LIVEACC	<volser>	<category>	Change category and allow live copy volume access from a DR host	N	Y	Y	Y	Y	Y	Y

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
PARTRFSH	<volser>			Refreshes the storage class setting.	N	Y	N	Y	Y	Y	Y
		MMOUNT									
LOWRANK	TVC	SHOW		View and control the mount TVC	N	Y	Y	Y	Y	Y	Y
		<cluster ID mask>									
	COPY	SHOW		View and control the copy source cluster							
		<cluster ID mask>									
GGM	CTG	ENABLE/ DISABLE		Enable/Disable GGM function on the CTG.	N	Y	Y	Y	Y	Y	Y
		START		Start GGM function on the CTG. 4 th KW IDLE is the option.							
			IDLE								
		STOP		STOP GGM function on the CTG. 4 th KW CLEAR is optional.							
			CLEAR								
		RECOVER		Recover the error status when CTG communication fails during the online phase.							
	MC/ SC/ DC/ SG	INHERIT		Set storage construct name assigned to the replicated logical volume in the CTG.							
		DEFAULT									
		FIXED	<storage construct name>								
	TGTCAT	INHERIT		Set the category assigned to the replicated logical							
		FIXED	<category>								

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				volume in the CTG.							
	SRCCAT	<category>		Set the category assigned to the copy source logical volume in the CSG.							
	QUEUE	<volser>	<CSG composite library sequence number>	Queue the specified volume for GGM copy into the CTG.							
	REMOVEQ	<volser>		Remove the specified volume from the CTG GGM copy queue							
	CPYCNT	<count>		Set the number of dedicated GGM copy tasks.							
	FORCECMP	<volser>		Force the GGM completion of a volume unable to complete CTG copies.							
	SRCVLCHG	<volser>	<category>	Change the copy source volume's category in the CSG.							
	BVIRHIST	CLEAR	<volser>	Clear BVIR history record of the <volser>.							
			ALL	If ALL is specified in the 4 th KW, clear all the BVIR history.							

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
	OPMSG	ENABLE/ DISABLE		Set if the operator message is surfaced upon the GGM operation.							
	COPY	SUMMARY		Display the summary of the current GGM status.							
		DETAIL	<volser>	Display the current GGM activity/status of the <volser>.							
DIAGDATA	SHOW			Requests information about the command handshake elapsed time between two clusters in the Grid.	Y	N	Y	Y	Y	Y	Y
		AVG									
		MAX									
		MIN									
		TMO									
		ERR									
	SHOW			Requests information about the command handshake elapsed time from the target cluster to peer cluster(s) in the Grid.	N	Y	Y	Y	Y	Y	Y
	RESET			Reset information about the command handshake elapsed time on the target cluster.							
OTM	SET	DISABLE	<cluster ID>	Disable ownership takeover mode.	N	Y	Y	Y	Y	Y	Y

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
		ROT		Enable Read-Only ownership takeover mode.							
		WOT		Enable Read-Write ownership takeover mode.							
	SHOW			Display the current ownership takeover status							
FENCE	ENABLE/DISABLE			Enable/disable remote cluster fence function.	Y	N	Y	Y	Y	Y	Y
	THRESHLD	SCRVOAVG	<value>	Set the thresholds to determine when the remote cluster fence action is triggered.	Y	N	Y	Y	Y	Y	Y
		PRIVOAVG									
		VCAVG									
		TOKAVG									
		TMO									
		ERR									
		EVALWIN									
	TIME	DELAY	<value>	Adjust the timing requirements used to determine when a remote cluster fence action is applied.	N	Y	Y	Y	Y	Y	Y
		CONSCNT									
	SHOW			Request information about the remote cluster fence function.							
	ACTION	PRI	NONE	Configure what primary action takes place against a remote cluster							
			ALERT								
			OFFLINE								

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
			REBOOT	when a fence action is applied.							
			REBOFF								
		SEC	ENABLE/ DISABLE	Configure whether the secondary fence action (grid network isolation) is initiated against a remote cluster when the primary action is unsuccessful.							
		AIXDUMP	ENABLE/ DISABLE	Configure whether an AIX system dump is automatically initiated as part of the fencing action when REBOOT/REBOFF is initiated.							
CUIR	SETTING			Display the option to automatically vary offline z/OS connected devices using CUIR functions as part of the service-preparation process.	Y	N	Y	Y	Y	Y	Y
		SERVICE	ENABLE/ DISABLE	Enable/disable the option to automatically vary offline z/OS connected devices using CUIR functions as part of the service-							
		FENCE									
		ALL									

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				preparation process.							
	AONLINE			Display the option to automatically vary online z/OS connected devices, which were previously varied offline through CUIR during service-preparation, as part of service-cancel operations.							
	SERVICE	ENABLE/ DISABLE		Enable/disable the option to automatically vary online z/OS connected devices, which were previously varied offline through CUIR during service-preparation, as part of service-cancel operations.							
	FENCE										
	ALL										
LDRIVE				Request summary information about connected host systems and their associated logical device pathing states for the target cluster	Y	Y	Y	Y	Y	Y	

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				(distributed) or grid (composite).							
	GROUP	<PGID index>		Request detailed information about a particular connected host system's pathing state by path group id (PGID) for its connected logical devices to the target cluster.	N	Y	Y	Y	Y	Y	Y
CLDSET				Display the current values of all CLDSET cloud data handling settings	N	Y	N	N	N	Y	Y
	CPMCNTH	<value>		Set the highest number of concurrent data premigration tasks to cloud							
	CPMCNTL	<value>		Set the lowest number of concurrent data premigration tasks to cloud							
	CLDPRIOR	<value>		Set the thresholds to begin increasing the number of cloud pre-migration tasks corresponding to the size of unpemigrated cloud data.							
	CRCCNT	<value>		Set the number of concurrent data							

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				recall tasks from cloud							
	CDELCNT	<value>		Set the number of concurrent data deletion tasks on cloud							
	CPMTOUT	<value>		Set the data premigration timeout to cloud							
	CRCTOUT	<value>		Set the data recall timeout from cloud							
	CDELTOU	<value>		Set the data deletion timeout on cloud							
	CENABLM	ALL	ENABLE/DI SABLE	Enable/disable the options of data premigration to cloud, data recall from cloud and data deletion on cloud at a time							
	PFRCCPG0	ENABLE/DI SABLE		Enable/disable ghost copy of PG0 replication.							
	PFRCTDL	ENABLE/DI SABLE		Enable/disable ghost copy of time delayed replication.							
	PFRCCTRY	<value>		Set the number of copy retries before ghost copy is done.							
	USEHIGH	<cloud nickname>	<value>	Set the high warning level of amount of data stored in a cloud that the cloud							

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				nickname points to.							
	USELOW	<cloud nickname>	<value>	Set the low warning level of amount of data stored in a cloud that the cloud nickname points to.							
	USESHOW			Show USEHIGH and USELOW settings.							
		<page index>									
SERVICE	ENTER			Initiate service-prep.	N	Y	NA	Y	Y	Y	Y
		FORCE		Initiate service-prep forcibly.							
	CANCEL			Cancel service-prep.							
	SHOW			Show service-prep status.							
CLDVR	RESERVE	<number>	<category>	Reserve scratch volumes to FF19 category as candidates of volume version restore destination.	N	Y	N	N	N	Y	Y
	DESTCAT	<category>		Set restore destination category.							
	DESTVOL	<volser>		Set restore destination volume.							
	<volser>	<insert version>	<data level>	Perform volume version restore.							

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
LWORMR	SHOW			Show LWORM Retention special settings configured by IBM support	N	Y	Y	Y	Y	Y	Y
OCOPY	SUMMARY			Requests summary information about all the object copy jobs.	N	Y	N	N	N	N	Y
OBJSET1	See description section			Requests to change various object control settings.	See description section						
CRCSET	GEN	LVOL	DISABLE/ CRC32	Request to change CRC generation settings.	Y	N	N	Y	Y	Y	Y
	VERIFY	LVOLWRR	DISABLE/ ENABLE	Request to change CRC verification settings.	Y	N	N	Y	Y	Y	Y
		LVOLWRL									
		LVOLCPYS									
		LVOLCPYT									
		LVOLTRCL									
		LVOLTPMG									
		LVOLRCLM									
STABACK	<pvol>	VOL	<lvol>	Request zTape Air-Gap function to backup to <pvol> of logical volumes identified with KW3 & KW4 See description section for zTape Air-Gap (FC5995)	N	Y	N	Y	N	N	N
		ALL									
		STG	<storage group name>								
STAREST	<pvol>			Request zTape Air-Gap function to restore from <pvol> logical	N	Y	N	Y	N	N	N

KW1	KW2	KW3	KW4	Description	Comp	Dist	7740	7700D	7700T	7700C	7700O
				volumes previously stored by STABACK See description section for zTape Air-Gap (FC5995)							
STASTAT				Request zTape Air-Gap function to report status of library (Physical volumes in storage or export station) Or current or prior STABACK STAREST task. See description section for zTape Air-Gap (FC5995)	N	Y	N	Y	N	N	N
	TASK										
		<page number>									
STACANC				Request zTape Air-Gapfunction to cancel current STABACK or STAREST in progress at next interruptible point See description section for zTape Air-Gap (FC5995)	N	Y	N	Y	N	N	N



November 2023

3.1 Responses to Library Requests

In response to a recognized request using the library request function, the TS7700 will perform the request and return data that is formatted as 70 character length text lines.

The z/OS host receiving the data will generate a multi-line console output with the data. The first three lines are generated by the host software to identify the request. The rest of the lines are taken from the response data. z/OS does not format or interpret the response data. Each line to display is already formatted by the TS7700.

For example:

```
CBR1280I Library library_name request.
Keywords: keyword1,keyword2,keyword3,keyword4
-----
---
data line 1
data line 2
...
data line N
```

The data returned for each of the requests are described in the following sections.

Prior to 8.31.x.x, the version is provided in the response header and the number of the version is increased when the response format is changed. At the code level 8.31.x.x or above, the TS7700 starts providing a revision as well as a version number in the response header. They are updated when the response format is changed in the following manner:

- The revision is increased when the new change is compatible.
- The version is increased when the new change is incompatible.

The “compatible change” means:

- Adding new lines at the end of the response.
- Adding new columns at the right end of the response.
- New words are used (format itself is not changed).

The “incompatible change” means:

- Adding new lines in the middle of the response.
- Adding new columns in the middle of the response.
- Shifting words in the line of the response.

This concept may be useful for the application which parses the response data to determine if the application needs to be updated to handle the new changes.

Please refer to the IBM Support document (<https://www.ibm.com/support/pages/node/6193419> (“How can I send the output of a TS7700 LI REQ command issued on a z/OS host to a data set?”)) if the LI REQ response needs to be sent to a data set instead of the system console.

3.1.1 CACHE2 Response

In response to the CACHE2 request, the cluster associated with the distributed library in the request will examine the current state of its tape volume cache.

Code level 8.50.x.y deprecates the CACHE request and introduces this request at version 5.0 in which each space quantity is expressed as a one to four digit number followed by a unit symbol composed of either a space or decimal metric prefix symbol (e.g. K, M, G, T, etc. where a G is 10⁹, a T is 10¹², etc.) then a B indicating bytes.

(Note) When “CACHE” instead of “CACHE2” is still issued to 8.50.x.x cluster, the following text shows up at the end of the output to encourage to use “CACHE2”:

```
"CACHE REQUEST DEPRECATED; SUPERSEDED BY CACHE2."
```

The premigration throttle value applies to the tape attach partitions of a TS7700T. The copy throttle value applies to all cache partitions.

If the request is issued when the cache is in the process of initialization, the following warning message will be displayed with the partial information of the cluster's tape volume cache. It is recommended to try again after the initialization is finished.

WARNING: CACHE INITIALIZATION IS STILL IN PROGRESS.

INFORMATION DISPLAYED MAY NOT REFLECT THE ACTUAL STATUS OF CACHE.

The response lines are formatted as follows:

(TS7740 (only “CACHE” is supported because TS7740 cannot install 8.50.x.x level))

The cache content of the tape managed cache partition 0 is provided.

```
TAPE VOLUME CACHE STATE V4 .0
TS7700 TIERING CAPABILITIES: TAPE
TOTAL INSTALLED/ENABLED GBS: 28604 / 9000
TOTAL ADJUSTED CACHE USED GBS: 24
CACHE ENCRYPTION STATUS: CAPABLE
TIERED CACHE PARTITIONS
CP   ALLOC   USED   PG0   PG1  PMIGR D_PMIGR   COPY   PMT
CPYT
  0    9000    12    0    12    0      0      0      0
```

VERSION 5.0 & UP - ALL TS7700 MODELS

In the “TIERED CACHE PARTITIONS” part, “TCO” (Tape/Cloud/Object) column next to “CP” indicates the partition type information. The meaning of the column is detailed in the table below.

```
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

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

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	0
2	NNY	4000GB	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	0
4	YNN	4000GB	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	0
6	NNN	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	0

Line	Bytes	Name	Description
1	0:24	Header Info	'TAPE VOLUME CACHE STATE V'
	25:26	Version	The version number for the response. The number is left justified and padded with blanks. Starts from 5 for CACHE2 (R5.0).
	27	Dot	'.'
	28:29	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts from 0 for CACHE2 (R5.0).
	30:69	Blanks	
2	0:31	Header Info	'TOTAL SPACE INSTALLED/ENABLED: '
	32:37	Cache Size	The quantity of cache currently installed. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	38	Separator	'/'
	39:44	Enabled Size	The quantity of cache enabled through the cache enablement features. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	45:69	Blanks	
3	0:33	Header Info	'TOTAL ADJUSTED CACHE SPACE USED: '
	34:39	Total adjusted used cache size.	The TS7700 tape volume cache is divided into two filesystems, namely /smfs0 and /smfs1. Logical volumes written to the TS7700 cache are balanced between the two filesystems. When either one filesystem gets

Line	Bytes	Name	Description
			full, the cache is full. The total adjusted used cache size represents the used cache space derived by doubling the higher cache usage of the two filesystems. This value is used in the algorithm to trigger host write throttling due to running low on free cache space. It is also used to control data migration. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	40:69	Blanks	
4	0:23	Header Info	'CACHE ENCRYPTION STATUS: '
	24-25	Blanks	
	26:41	Disk Encryption Status	The disk encryption status is shown as a string. The status string is right justified and padded with blanks. Encryption Status: • 'CAPABLE' – The disk controller is capable but not enabled for encryption. • 'NOT CAPABLE' – The disk controller is not capable of encryption. • 'ENABLED-INTERNAL' – The disk controller is capable and enabled for internal or local encryption. • 'ENABLED-EXTERNAL' – The disk controller is capable and enabled for external encryption.
	42:69	Blanks	
5	0:32	Header Info	'OVERCOMMITTED CACHE PARTITIONS: '
	33-(33+2*x)	Overcommitted cache partitions	Overcommitted partitions are listed in a comma delimited format. For example, if CP1, CP3, and CP7 are overcommitted, this line will show "1, 3, 7". If no overcommitted cache partition exists, "NONE" is shown.
	(33+2*x):69	Blanks	
6	0:38	Header Info	'CACHE RESIDENT ONLY PARTITION'
	39:69	Blanks	
7	0	Blank	
	1:27	Header Info	'PRIVATE CACHE SPACE USED: '
	28:33	Private Volume Data Quantity	The quantity of logical volume data assigned to non-scratch categories in the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	34:69	Blanks	
8	0	Blank	
	1:27	Header Info	'SCRATCH CACHE SPACE USED: '
	28:33	Scratch Volume Data Quantity	The number of logical volume data assigned to scratch categories in the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	34:69	Blanks	

Line	Bytes	Name	Description
9	0	Blank	
	1:58	Header Info	'CP ALLOC USED PIN PKP PRM COPY CPYT'
	59:69	Blanks	
10	0:1	Blanks	
	2	Cache Partition (CP)	The cache partition number.
	3:4	Blanks	
	5:10	Space Allocated (ALLOC)	The quantity of the cache that has been allocated. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	11:12	Blanks	
	13:18	Space Used (USED)	The TS7700's Tape Volume Cache is divided into two filesystems. This value represents the total space used by both filesystems. This includes not only logical volumes but also overhead files like throttling files, etc. The nominal cache usage is the sum of PIN, PKP, and PRM Data Quantities. The Space Used and the sum of PIN, PKP and PRM Data Quantities may slightly differ because of the internal calculation methods used. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	19:20	Blanks	
	21:26	PIN Data Quantity (PIN)	The quantity of the cache that contains pinned volume copy retention group data. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	27:28	Blanks	
	29:34	PKP Data Quantity (PKP)	The quantity of the cache that contains prefer keep volume copy retention group data. The count expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), is right justified and padded with blanks.
	35:36	Blanks	
	37:42	PRM Data Quantity (PRM)	The quantity of the cache that contains prefer remove volume copy retention group data. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	43:44	Blanks	
	45:50	Data quantity requiring Copy (COPY)	The quantity of the cache that contains data that needs to be copied to another cluster. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	51:52	Blanks	
	53:58	Copy Throttling (CPYT)	This field indicates whether copy throttling is active. If throttling is active, a value in milliseconds is indicated. The value is right justified and padded with blanks. If throttling is not active, 0 is indicated.
	59:69	Blanks	

Line	Bytes	Name	Description
11	0	Blank	
	1:22	Header Info	'FLASH COPY INFORMATION'
	23:69	Blanks	
12	0	Blank	
	1:21	Header Info	'INDEX ENABLED SIZE'
	22:69	Blanks	
13:20	0:4	Blanks	
	5	Flash family index (INDEX)	DR family index from 1 to 8.
	6:10	Blanks	
	11:13	Enabled (ENABLED)	'YES' if Flash Copy is enabled for the DR family. 'NO' if Flash Copy is not enabled for the DR family
	14:21	Size (SIZE)	The quantity of the cache consumed by Flash Copy volumes for the DR family. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. 'B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	22:69	Blanks	
21	0:30	Header Info	'TIERED CACHE PARTITIONS'
	31:69	Blanks	
22	0	Blank	
	1:66	Header Info	' CP TCO ALLOC USED PG0 PG1 PMIGR D_PMIGR COPY PMT CPYT'
	67:69	Blanks	
23:29	0	Blanks	
	1:2	Cache Partition (CP)	The cache partition number for the tiered cache partitions.
	3	Blank	
	4:6	Tape/Cloud/Object (TCO)	Column to indicate whether tape or cloud is attached to the partition and whether the partition is for object data or not. T means Tape attached partition. C means Cloud attached partition. O means the partition is not for logical volume but for object data. The character 'Y' under each letter indicates "Yes", and 'N' indicates "No".
	7:13	Space Allocated (ALLOC)	The quantity of the cache that has been allocated to the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. 'B', KB, MB, GB, TB, etc.), right justified and padded with blanks. The amount of the cache allocated to partition 0 is the enabled size less the amounts allocated to other partitions.
	14	Blanks	
	15:20	Space Used (USED)	The TS7700's Tape Volume Cache is divided into two filesystems. This value represents the total space used by both filesystems. This includes not only logical volumes but also overhead files like throttling files, etc. The nominal cache usage is the sum of PG0 and PG1 Data Quantities. The Space Used and the sum of PG0 and PG1 Data Quantities may slightly differ because of the internal calculation methods used.

Line	Bytes	Name	Description
			The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	21	Blank	
	22:27	PG0 Data Quantity (PG0)	The quantity of the cache that contains preference group 0 data in the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	28	Blank	
	29:34	PG1 Data Quantity (PG1)	The quantity of the cache that contains preference group 1 data in the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks. In the Object Cache Partition case, object data is always shown as PG1 data. Copying object data to tape or cloud is not supported so in fact, object data is data that always exists only in the cache.
	35	Blank	
	36:41	Data Quantity requiring Premigration (PMIGR)	The quantity of the cache that contains data that needs to be copied to a physical volume for the partition. Amount of delayed premigration data waiting for the delay time is not included. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	42:43	Blanks	
	44:49	Data Quantity Delaying Premigration GBs (D_PMIGR)	The quantity of the cache that contains data that has been delayed to be copied to a physical volume for the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	50	Blank	
	51:56	Data Quantity requiring Copy (COPY)	The quantity of the cache that contains data that needs to be copied to another cluster for the partition. The count is expressed in decimal with up to four digits followed by a unit symbol (e.g. ' B', KB, MB, GB, TB, etc.), right justified and padded with blanks.
	57	Blank	
	58:61	Premigration Throttling (PMT)	This field indicates whether premigration throttling is active for the partition. If throttling is active, a value in milliseconds is indicated. The value is right justified and padded with blanks. If throttling is not active, 0 is indicated. Note: The same throttling value is applied to all tape partitions.
	62	Blank	
	63:66	Copy Throttling (CPYT)	This field indicates whether copy throttling is active for the partition. If throttling is active, a value in milliseconds is indicated. The value is right justified and padded with blanks. If throttling is not active, 0 is indicated. Note: The same throttling value is applied to all partitions.
	67:69	Blanks	

3.1.2 COPY Response

In response to the COPY request, the cluster associated with the distributed library in the request will collect the information on the copy jobs based on the second keyword specified in the request. The information contained in this response pertains to tape logical volume copies and not to object copies associated with transparent cloud tiering.

If ACTIVITY is specified in the second keyword, the information on the active copy jobs associated with the distributed library in the request is collected.

If SUMMARY is specified in the second keyword, the information on all the copy jobs associated with the distributed library in the request is collected.

If QUEUE is specified in the second keyword, the information on the copy jobs associated with the distributed library in the request is collected. This keyword is supported at the code level of 8.50.x.x (R5.0) or above.

3.1.2.1 COPY, ACTIVITY

If the third keyword associated with the request (RUN, SDEF, IDEF, DEF) is specified, only the information on the active copy jobs associated with the third keyword is collected. If the third keyword is not specified the receiving cluster will show all active copy jobs. The response lines are formatted as follows:

```
COPY ACTIVITY V3 .1
  TOTAL ACTIVE COPIES:  1
                        RUN:  1
  SYNC-DEFERRED (SDEF) :  0
  IMM-DEFERRED (IDEF) :  0
  DEFERRED (DEF) :  0
```

```
VOLSER TYPE          SOURCE_CLUSTER SIZE (MB)  COPY_TIME
-----
```

ZS0000	RUN	Palomino (c7)	752	0000:03:21
--------	-----	---------------	-----	------------

TIMED-OUT	COPIES	DURING THE LAST HOUR (06 MINUTES SINCE LAST RESET)						
CLUSTER	0	1	2	3	4	5	6	7
TOTAL	NA	0	NC	NC	NC	NC	NC	NC
0								

At 8.31.x.x, the timeout copy job status is added at the bottom of the output. TS7700 monitors how many copy jobs have timed out and keeps its count. The count is reset once an hour.

At 8.32.x.x, the cluster number starts from 0 to 7 which used to be from 1 to 8 till 8.31.x.x.

At 8.50.x.x, the copy_time (elapsed time how long the copy job has been in flight) is provided.

Note: Up to 40 active copy jobs status is provided prior to 8.31.x.x. At 8.31.x.x, up to only 30 active copy jobs status is provided to get the space for the timeout copy job status.

If the search didn't match any records, the following message is displayed

No record matched the searching criteria

Line	Bytes	Name	Description
1	0:14	Header Info	'COPY ACTIVITY V'
	15:1	Version	The version number for the response. The number is left justified and padded with blanks. Start with 2 (at the code level 8.210.x.x). The version is incremented to: - 3 at 8.32.x.x.
	16	Dot	'.'
	17:18	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0. The revision is incremented to: - 1 at 8.50.x.x.
	19:69	Blanks	
2	0	Blank	
	1:20	Header Info	'TOTAL ACTIVE COPIES: '
	21	Blank	
	22:23	Total Active Jobs	The total number of active copies for the response. The number if right justified and padded with blanks. Without activity the field will be set to Zero
	24:69	Blanks	
3	0:16	Blanks	
	17:20	Header Info	'RUN: '
	21	Blank	
	22:23	"RUN"Active Jobs	The total number of active "RUN" copies for the response. The number if right justified and padded with blanks. Without activity the field will be set to Zero
	24:69	Blanks	
4	0	Blank	
	1:20	Header Info	'SYNC-DEFERRED (SDEF) : '
	21	Blank	
	22:23	"SDEF"Active Jobs	The total number of active "SDEF" copies for the response. The number if right justified and padded with blanks. Without activity the field will be set to Zero
	24:69	Blanks	
5	0:1	Blanks	
	2:20	Header Info	'IMM-DEFERRED (IDEF) : '
	21	Blank	
	22:23	"IDEF"Active Jobs	The total number of active "IDEF" copies for the response. The number if right justified and padded with blanks. Without activity the field will be set to Zero
	24:69	Blanks	

6	0:6	Blanks	
	7:20	Header Info	'DEFERRED (DEF) : '
	21	Blank	
	22:23	"DEF"Active Jobs	The total number of active "DEF" copies for the response. The number is right justified and padded with blanks. Without activity the field will be set to Zero
	24:69	Blanks	
7	0	Blank	
	1:39	Separator	All dash '-' characters
	40:69	Blanks	
8	0	Blank	
	1:50	Header Info	'VOLSER TYPE SOURCE_CLUSTER SIZE (MB) COPY_TIME'
	51:69	Blanks	
9	0	Blank	
	1:39	Separator	All dash '-' characters
	40:69	Blanks	
When no record matched the searching criteria			
10	0:39	Header Info	'No record matched the searching criteria'
	40:69	Blanks	
When at least one record was found. Up to 40 active copy jobs status is provided prior to 8.31.x.x. At 8.31.x.x or later, up to only 30 active copy jobs status is provided.			
N	0	Blank	
	1:6	Volser	Active copy task job
	7	Blank	
	8:11	Job Type	It could either be: ' RUN' 'SDEF' 'IDEF' ' DEF'
	12:25	Library name	
	26	Blank	
	27:30	Header Info	(Cx) Where X specifies the Cluster ID selected as the source for the copy
	31:39		This field indicates the size in MBs for each volume listed, If there are more than 99,999,999 volumes, the display stays at 99,999,999. The number is right justified and padded with blanks.
	40	Blank	
	41:50	Elapsed time of the copy job (COPY_TIME)	The elapsed time since the copy job has become in-flight. It is formatted as follows: HHHH:MM:SS for example: 0000:03:12 (3 minutes 12 seconds)
	51:69	Blanks	

M	0:69	Separator	All dash '-' characters. If no active copy job exists, this line locates at the line 11. If N active copies exist, the line locates at the line M (= 10+N).
M+1	0:38	Header Info	'TIMED-OUT COPIES DURING THE LAST HOUR ('
	39:40	Passed minutes since last reset	The count of timeout copy job is reset once in an hour. This field shows how many minutes it has passed since the last reset.
	41:66	Header Info	'MINUTES SINCE LAST RESET) '
	67:69	Blanks	
M+2	0:69	Header Info	'CLUSTER 0 1 2 3 4 5 6 7 TOTAL' The cluster ID starts from 0 to 7 in the response.
M+3	0:7	Blanks	
	8:12	Count of time out copy pulling from cluster 1 in last T minutes	The count of time out copy job trying to pull from cluster1 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster1, 'NA' is displayed. If cluster1 is not configured, 'NC' is displayed
	13:14	Blanks	
	15:19	Count of time out copy pulling from cluster 2 in last T minutes	The count of time out copy job trying to pull from cluster2 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster2, 'NA' is displayed. If cluster2 is not configured, 'NC' is displayed
	20:21	Blanks	
	22:26	Count of time out copy pulling from cluster 3 in last T minutes	The count of time out copy job trying to pull from cluster3 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster3, 'NA' is displayed. If cluster3 is not configured, 'NC' is displayed
	27:28	Blanks	
	29:33	Count of time out copy pulling from cluster 4 in last T minutes	The count of time out copy job trying to pull from cluster4 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster4, 'NA' is displayed. If cluster4 is not configured, 'NC' is displayed
	34:35	Blanks	
	36:40	Count of time out copy pulling from cluster 5 in last T minutes	The count of time out copy job trying to pull from cluster5 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster5, 'NA' is displayed. If cluster5 is not configured, 'NC' is displayed
	41:42	Blanks	
	43:47	Count of time out copy pulling from	The count of time out copy job trying to pull from cluster6 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster6, 'NA' is displayed.

		cluster 6 in last T minutes	If cluster6 is not configured, 'NC' is displayed
	48:49	Blanks	
	50:54	Count of time out copy pulling from cluster 7 in last T minutes	The count of time out copy job trying to pull from cluster7 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster7, 'NA' is displayed. If cluster7 is not configured, 'NC' is displayed
	55:56	Blanks	
	57:61	Count of time out copy pulling from cluster 8 in last T minutes	The count of time out copy job trying to pull from cluster8 in last T minutes. T is "Passed minutes since last reset". If a local cluster is cluster8, 'NA' is displayed. If cluster8 is not configured, 'NC' is displayed
	62:64	Blanks	
	65:69	Total count of time out copy n last T minutes	The total count of time out copy job on the local cluster in last T minutes. T is "Passed minutes since last reset".

3.1.2.2 COPY, SUMMARY

No third and fourth keywords are applicable. The response lines are formatted as follows:

COPY SUMMARY V3 .1

DISTRIBUTED LIBRARY Lipizzan CLUSTER 1

COPY COUNT BY ORIGINATING CLUSTER BY COPY QUEUE

CLUSTER	0	1	2	3	4	5	6	7	*
RUN	NA	0	0	0	NC	NC	NC	NC	0
SyncDef	NA	0	0	0	NC	NC	NC	NC	0
ImedDef	NA	0	0	0	NC	NC	NC	NC	0
Def	NA	10	0	0	NC	NC	NC	NC	0
CpyRfsh	NA	0	0	0	NC	NC	NC	NC	15
FamDef	NA	0	0	0	NC	NC	NC	NC	0
TDel	NA	7	0	0	NC	NC	NC	NC	0

COPY SIZE GB BY ORIGINATING CLUSTER BY COPY QUEUE

CLUSTER	0	1	2	3	4	5	6	7	*
RUN	NA	0	0	0	NC	NC	NC	NC	0
SyncDef	NA	0	0	0	NC	NC	NC	NC	0
ImedDef	NA	0	0	0	NC	NC	NC	NC	0
Def	NA	7	0	0	NC	NC	NC	NC	0
CpyRfsh	NA	0	0	0	NC	NC	NC	NC	11
FamDef	NA	0	0	0	NC	NC	NC	NC	0
TDel	NA	5	0	0	NC	NC	NC	NC	0

COPY COUNT BY COPY TASK BY COPY QUEUE

TASK	COPY	RECALL	DELAY	TDEL
------	------	--------	-------	------

RUN	0	0	0	0	
SyncDef	0	0	0	0	
ImedDef	0	0	0	0	
Def	0	0	10	0	
CpyRfsh	0	0	15	0	
FamDef	0	0	0	0	
TDel	0	0	0	7	

COPY SIZE GB BY COPY TASK BY COPY QUEUE					
TASK	COPY	RECALL	DELAY	TDEL	

RUN	2	0	0	0	
SyncDef	0	0	0	0	
ImedDef	0	0	0	0	
Def	0	0	7	0	
CpyRfsh	0	0	11	0	
FamDef	0	0	0	0	
TDel	0	0	0	5	

LONGEST COPY QUEUE AGE :			4399		
LONGEST TIME DELAYED TASK :			25049		
LONGEST FAM DEF QUEUE AGE :			0"		

The second line shows the distributed library name and ID in the request.

The first and second matrices show the total copy count and size grouped by the copy queue (per each line) and originating cluster (per each column) on the distributed library in the request.

Each line shows the copy queue type.

- RUN : A copy job with Rewind Unload copy mode
- SyncDef : A copy job with Synchronous copy mode. A copy job is posted when any event occurs which prevents the synchronous updates to either or both clusters assigned to synchronous copy mode.
- ImedDef : A copy job with Rewind Unload copy mode, but downgraded to immediate deferred copy due to any event which prevents Rewind Unload copy from completing.
- Def : A copy job with Deferred copy mode
- CpyRfsh : A copy job posted by "COPYRFSH" Library Request command
- FamDef : A copy job with Deferred copy mode, but downgraded to family deferred copy because of cluster family setting.
- TDel : A copy job with Time Delayed copy mode.

Each column shows the originating cluster. The originating cluster is the distributed library ID from which the copy target cluster may copy the data and it is recorded on each copy job when the copy job is initially created. The actual copy source will be determined again when the copy is attempted, but in most cases, the originating cluster would be the actual copy source.

If no originating cluster is recorded when the copy job is initially created, it is displayed in the most right column '*'. The typical case is when a copy job is created by "COPYRFSH" Library Request command. If no copy source (keyword 3) is specified in COPYRFSH command, no originating cluster is recorded on its copy job and it's shown in this column.

The column where the originating cluster is equal to the distributed library in the request always shows 'NA' because the copy source should reside outside of the local cluster. 'NC' means Not Configured. The cluster which is not configured in this grid always shows 'NC'.

The third and fourth matrices show the total copy count and size which are grouped by the copy queue (per each line) and copy task (per each column) on the distributed library in the request.

Each column shows the copy task.

- COPY : A copy job which is ready to be executed or already in progress.
- RECALL : A copy job whose copy source data was already migrated. It is waiting for the copy source to recall data into cache.
- DELAY : A copy job which was delayed for some reason.
- TDEL : A copy job with Time Delayed copy mode which is still delayed until the delay is expired. Once the delay is expired, the task of the copy job is changed from TDEL to COPY.

The above example shows:

- The distributed library "Lipizzan" (distributed library ID = 1) returns the copy job status.
- There are 10 Deferred copies with DELAY task whose originating cluster is 2. Their total size is 7GB.
- There are 15 COPYRFSH copies with DELAY task which have no originating cluster (this indicates COPYRFSH Library Request was issued to the volumes with no second keyword (copy source). Their total size is 11GB.
- There are 7 Time Delayed copies with TDEL task whose originating cluster is also 2. Their total size is 5GB.
- The longest copy queue age except Time Delayed copy is 4399 seconds. The longest Time Delayed copy queue age is 25049 seconds.
- The distributed library IDs 5 to 8 are not configured in this Grid.

At 8.31.x.x, Time Delayed copy mode is introduced. When this copy mode is assigned to the logical volume, its copy job created on a copy target cluster won't start until the delay defined in the management class option is expired. Once it's expired, the copy job may start. While the copy job is

delayed, it's counted in the column of TDEL task and its age is used to calculate LONGEST TIME DELAYED TASK. Once the delay is expired, its task is moved from TDEL to COPY and its age is used to calculate LONGEST COPY QUEUE AGE.

At 8.32.x.x, the cluster number starts from 0 to 7 which used to be from 1 to 8 till 8.31.x.x.

At 8.42.x.x, the longest copy queue age among the family deferred copy jobs is supported and provided at the end of the response.

If any unexpected error occurs and the request is not completed, the following error text is returned:

AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = XX

Line	Bytes	Name	Description
1	0:13	Header Info	`COPY SUMMARY V`
	14:15	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.30.x.x). The version is incremented to: - 2 at 8.31.x.x - 3 at 8.32.x.x.
	16	Dot	`.`
	17:18	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0. - 1 at 8.42.x.x.
	19:69	Blanks	
2	0	Blank	
	1:19	Distributed Library Info	`DISTRIBUTED LIBRARY`
	20	Blank	
	21:28	Distributed Library Name	The distributed library name which responds to the request. The name is left justified and padded with blanks.
	29	Blank	
	30:36	Cluster Info	`CLUSTER`
	37	Blank	
	38	Distributed Library ID	Distributed Library ID (1-8)
	39:69	Blanks	
3	0:69	Separator	All dash '-' characters
4	0:46	Header Info	`COPY COUNT BY ORIGINATING CLUSTER BY COPY QUEUE`
	47:69	Blanks	
5	0:60	Header Info	`CLUSTER        0        1        2        3        4        5        6 7        *`
	61:69	Blanks	
6	0:60	Separator	All dash '-' characters
	61:69	Blank	

7	0:3	Blanks	
	4:6	“RUN”	`RUN`
	7:60	RUN copy count	The line shows the RUN copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
8	0:6	“SyncDef”	`SyncDef`
	7:60	Sync Deferred copy count	The line shows the Sync Deferred copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
9	0:6	“ImedDef”	`ImedDef`
	7:60	Immediate Deferred copy count	The line shows the Immediate Deferred copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
10	0:3	Blanks	
	4:6	“Def”	`Def`
	7:60	Deferred copy count	The line shows the Deferred copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
11	0:6	“CpyRfsh”	`CpyRfsh`
	7:60	Copy Refresh copy count	The line shows the Copy Refresh copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
12	0	Blanks	
	1:6	“FamDef”	`FamDef`
	7:60	Family Deferred copy count	The line shows the Family Deferred copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
13	0:2	Blanks	
	3:6	“TDel”	`TDel`
	7:60	Time Delayed copy count	The line shows the Time Delayed copy count grouped by each originating cluster. Each field has “1 blank + copy count (from 0 to 99999)”. The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
14	0:69	Separator	All dash '-' characters
15	0:48	Header Info	`COPY SIZE GB BY ORIGINATING CLUSTER BY COPY QUEUE`
	49:69	Blanks	
16	0:60	Header Info	`CLUSTER      0      1      2      3      4      5      6 7      *`

	61:69	Blanks	
17	0:60	Separator	All dash '-' characters
	61:69	Blank	
18	0:3	Blanks	
	4:6	"RUN"	'RUN'
	7:60	RUN copy size	The line shows the total size of RUN copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
19	0:6	"SyncDef"	'SyncDef'
	7:60	Sync Deferred copy size	The line shows the total size of Sync Deferred copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
20	0:6	"ImedDef"	'ImedDef'
	7:60	Immediate Deferred copy size	The line shows the total size of Immediate Deferred copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
21	0:3	Blanks	
	4:6	"Def"	'Def'
	7:60	Deferred copy size	The line shows the total size of Deferred copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
22	0:6	"CpyRfsh"	'CpyRfsh'
	7:60	Copy Refresh copy size	The line shows the total size of Copy Refresh copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
23	0	Blanks	
	1:6	"FamDef"	'FamDef'
	7:60	Family Deferred copy size	The line shows the total size of Family Deferred copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
24	0:2	Blanks	
	3:6	"TDel"	'TDel'
	7:60	Time Delayed copy size	The line shows the total size of Time Delayed copy in GB grouped by each originating cluster. Each field has "1 blank + copy size (from 0 to 99999)". The line includes total 9 fields for each originating cluster.
	61:69	Blanks	
25	0:69	Separator	All dash '-' characters
26	0:36	Header Info	'COPY COUNT BY COPY TASK BY COPY QUEUE'

	37:69	Blanks	
27	0:2	Blanks	
	3:46	Header Info	`TASK            COPY            RECALL            DELAY            TDEL`
	47:69	Blanks	
28	0:36	Separator	All dash '-' characters
	37:69	Blank	
29	0:3	Blanks	
	4:6	"RUN"	`RUN`
	7:36	RUN copy count	The line shows the RUN copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
30	0:6	"SyncDef"	`SyncDef`
	7:36	Sync Deferred copy count	The line shows the Sync Deferred copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
31	0:6	"ImedDef"	`ImedDef`
	7:36	Immediate Deferred copy count	The line shows the Immediate Deferred copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
32	0:3	Blanks	
	4:6	"Def"	`Def`
	7:36	Deferred copy count	The line shows the Deferred copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
33	0:6	"CpyRfsh"	`CpyRfsh`
	7:36	Copy Refresh copy count	The line shows the Copy Refresh copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
34	0	Blanks	
	1:6	"FamDef"	`FamDef`
	7:36	Family Deferred copy count	The line shows the Family Deferred copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
35	0:2	Blanks	
	3:6	"TDel"	`TDel`
	7:36	Time Delayed copy count	The line shows the Time Delayed copy count grouped by each copy task. Each field has "5 blank + copy count (from 0 to 99999)". The line includes total 3 fields for each copy task.

	37:69	Blanks	
36	0:69	Separator	All dash '-' characters
37	0:38	Header Info	'COPY SIZE GB BY COPY TASK BY COPY QUEUE'
	39:69	Blanks	
38	0:2	Blanks	
	3:36	Header Info	'TASK COPY RECALL DELAY TDEL'
	37:69	Blanks	
39	0:36	Separator	All dash '-' characters
	37:69	Blank	
40	0:3	Blanks	
	4:6	"RUN"	'RUN'
	7:36	RUN copy size	The line shows the total size of RUN copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
41	0:6	"SyncDef"	'SyncDef'
	7:36	Sync Deferred copy size	The line shows the total size of Sync Deferred copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
42	0:6	"ImedDef"	'ImedDef'
	7:36	Immediate Deferred copy size	The line shows the total size of Immediate Deferred copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
43	0:3	Blanks	
	4:6	"Def"	'Def'
	7:36	Deferred copy size	The line shows the total size of Deferred copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
44	0:6	"CpyRfsh"	'CpyRfsh'
	7:36	Copy Refresh copy size	The line shows the total size of Copy Refresh copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
45	0	Blanks	
	1:6	"FamDef"	'FamDef'
	7:36	Family Deferred copy size	The line shows the total size of Family Deferred copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
46	0:2	Blanks	

	3:6	"TDel"	`TDel`
	7:36	Time Delayed copy size	The line shows the total size of Time Delayed copy in GB grouped by each copy task. Each field has "5 blank + copy size (from 0 to 99999)". The line includes total 3 fields for each copy task.
	37:69	Blanks	
47	0:69	Separator	All dash '-' characters
48	0:23	Header Info	`LONGEST COPY QUEUE AGE :`
	24:27	Blanks	
	28:35	Longest copy queue age	The longest copy queue age in second on the distributed library. The number is left justified and padded with blanks The age is calculated from all copy jobs except Time Delayed copy jobs which are still delayed (i.e. included in TDEL task).
	36:69	Blanks	
49	0:26	Header Info	`LONGEST TIME DELAYED TASK :`
	27	Blank	
	28:35	Longest Time Delayed copy queue age	The longest Time Delayed copy queue age in second on the distributed library. The number is left justified and padded with blanks The age is calculated only from Time Delayed copy jobs which are still delayed (i.e. the copy jobs only in TDEL task are used). Once they exit the delay state and move to COPY task, they are no longer used to calculate this age. They are used to calculate "Longest copy queue age".
	36:69	Blanks	
50	0:26	Header Info	`LONGEST FAM DEF QUEUE AGE :`
	27	Blank	
	28:35	Longest Family Deferred copy queue age	The longest Family Deferred copy queue age in second on the distributed library. The number is left justified and padded with blanks The age is calculated only from Family Deferred copy jobs in any cluster family members in the Grid. If Family Deferred copy jobs are expedited by LI REQ, LVOL, COPY, KICK and exit the family deferred queue, they are no longer used to calculate this age. They are used to calculate "Longest copy queue age".
	36:69	Blanks	

3.1.2.3 COPY, QUEUE

The specified target distributed library returns the status of the pending copy jobs which meet the criteria specified by the third and fourth keyword. The following third and fourth keywords are supported:

[Third keyword]

This determines the type of the copy jobs to provide. The following type can be used:

- RUN: Provide the status of the pending RUN copy jobs only
- DEF: Provide the status of the pending Deferred copy jobs only (CopyRefresh/Family-Deferred/Time Delayed copy jobs are excluded)
- RFSH: Provide the status of the pending CopyRefresh copy jobs only
- FDEF: Provide the status of the pending Family-Deferred copy jobs only

- TDEL: Provide the status of the pending Time Delayed copy jobs which already expire the defined time delay only
- RCAL: Provide the status of the pending copy jobs which require the recall from the copy source cluster

Note: Time Delayed copy jobs which still stay within the defined time delay cannot be queried.

[Fourth keyword]

This determines what status of the copy jobs is provided. The following state can be used:

- NEXT: Provide the status of the pending copy jobs which are waiting for the execution. They are from the oldest to youngest in order (i.e. the copy job listed on the top of the output is the one to be executed first).
- OLDEST: Provide the status of the pending copy jobs which are waiting for the execution. They are from the youngest to oldest in order (i.e. the copy job listed on the top of the output is the one to be executed last).
- ACTIVE: Provide the status of the pending copy jobs which are currently in-flight. They are from the oldest to youngest in order (i.e. the copy job listed on the top of the output is the one which has the longest in-flight time).

Note: Fourth keyword can be omitted. If it's omitted, "NEXT" is automatically used.

For example,

- ✓ If the in-flight RUN copy jobs are required, RUN/ACTIVE should be specified.
- ✓ If the list of the pending Deferred copy jobs which will be executed next is required, DEF/NEXT should be specified.
- ✓ If the list of the pending Family Deferred copy jobs which will be executed last is required, FDEF/OLDEST should be specified.
- ✓ If the list of the in-flight copy jobs which require the recall, RCAL/ACTIVE should be specified.

There are 2 cases where TS7700 internally changes (upgrades) copy job priority:

- When the family-deferred copy job has stayed for more than 12 hours, its queue priority is upgraded from the family-deferred to normal deferred.
- When the time-delayed copy job had already expired its time delay and has stayed for more than 12 hours, its queue priority is upgraded from time delayed to normal deferred.

On both cases, LI REQ still categorizes the upgraded copy job in the original type (i.e. the copy job can be found by using the third keyword FDEF or TDEL) but the TYPE field in the LI REQ output shows DEF. The example of the output is:

```
COPY QUEUE V1 .0
```

```
SEARCH RESULT: COPY, QUEUE, FDEF , NEXT
```

```
-----
VOLSER QUEUE  SIZE          AGE          POSTED_TIME  RET D_RSN
```

```

-----
ZL0002   DEF      0 9999:59:59 2019-03-05 08:52:50      0
ZL0001  FDEF      0 0000:01:18 2019-03-06 08:53:09      0

```

(*) ZL0002 was upgraded from family-deferred to deferred internally. The output with the third keyword "FDEF" includes ZL0002 but TYPE field shows "DEF". Also AGE shows 9999:59:59. They mean its priority was upgraded.

COPY QUEUE V1 .0

```

SEARCH RESULT: COPY, QUEUE, TDEL      , NEXT

```

```

-----
VOLSER QUEUE  SIZE          AGE          POSTED_TIME  RET D_RSN
-----
ZL0005   DEF      0 9999:59:59 2019-03-05 08:52:50      0
ZL0004   TDEL      0 0000:03:47 2019-03-06 08:53:29      0

```

(*) ZL0005 was upgraded from time delayed to deferred internally. The output with the third keyword "TDEL" includes ZL0005 but TYPE field shows "DEF". Also AGE shows 9999:59:59. They mean its priority was upgraded.

If LI REQ, LVOL, <volser>, COPY, KICK (FORCE) is issued to a copy job, it's not categorized in the original type when it's in FDEF/TDEL/FRSH, then it will be categorized to DEF.

The example of the output is:

COPY QUEUE V1 .0

```

SEARCH RESULT: COPY, QUEUE, DEF      , NEXT

```

```

-----
VOLSER QUEUE  SIZE          AGE          POSTED_TIME  RET D_RSN
-----
ZL0001   DEF      0 9999:59:59 2019-03-06 08:53:09      0

```

(*) ZL0001 was originally a family-deferred copy job and it's upgraded by LI REQ, LVOL, <volser>, COPY, KICK command. It then appears in the output with the third keyword "DEF" and TYPE field shows "DEF" as well. AGE shows 9999:59:59 because it's priority was upgraded.

Up to 40 copy jobs in the specified criteria can be provided in the response and the response lines are formatted as follows:

COPY QUEUE V1 .0

```

SEARCH RESULT: COPY, QUEUE, DEF      , NEXT

```

```

-----
VOLSER QUEUE  SIZE          AGE          POSTED_TIME  RET D_RSN
-----

```

```

ZA0003   DEF      3 0000:08:38 2019-02-08 06:15:19      0
ZA0002   DEF      3 0000:06:35 2019-02-08 06:17:22      0
ZA0001   DEF      3 0000:06:00 2019-02-08 06:17:57      0
ZA0039   DEF      3 0000:04:05 2019-02-08 06:19:52      0
ZA0034   DEF      3 0000:01:40 2019-02-08 06:22:17      0
ZA0033   DEF      3 0000:01:11 2019-02-08 06:22:46      0
ZA0032   DEF      3 0000:00:42 2019-02-08 06:23:15      0
ZA0031   DEF      3 0000:00:13 2019-02-08 06:23:44      0
ZA0000   DEF      3 0000:03:43 2019-02-08 06:16:16      1 S-DIS

```

Line	Bytes	Name	Description
1	0:11	Header Info	`COPY QUEUE V`
	12:13	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.50.x.x).
	14	Dot	`.`
	15:16	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.50.x.x).
	17:69	Blanks	
2	0	Blank	
	1:28	Header Info	`SEARCH RESULT: COPY, QUEUE, `
	29:36	Third Keyword	The third keyword specified in the LI REQ command is provided. The keyword is left justified and padded with blanks.
	37:38	Header Info	`,`
	39:46	Fourth Keyword	The fourth keyword specified in the LI REQ command is provided. The keyword is left justified and padded with blanks. If no fourth keyword is specified, `NEXT` is provided.
	47:69	Blanks	
3	1	Blank	
	2:61	Separator	All dash `-` characters
	62:69	Blanks	
4	0	Blank	
	1:61	Header Info	`VOLSER QUEUE SIZE AGE POSTED_TIME RET D_RSN`
	62:69	Blanks	
5	0	Blank	
	2:61	Separator	All dash `-` characters
	62:69	Blanks	
6	If no copy jobs in the specified criterion (based on the third and fourth keyword) exist, this line is just provided.		

	0	Blank	
	1:40		'No record matched the searching criteria'
	41:69	Blanks	
6-45	If one or more copy jobs exist, up to 40 copy jobs can be provided.		
	0	Blank	
	1:6	Copy job volser (VOLSER)	The volser of the copy job
	7:8	Blanks	
	9:12	Copy queue (QUEUE)	The current copy queue of the copy job. The value is right justified and padded with blanks. The possible values are: RUN: RUN copy DEF: Deferred copy (except CopyRefresh/Family-Deferred/Time Delayed) RFSH: Copy Refresh copy FDEF: Family Deferred copy TDEL: Time Delayed copy This value may be different with the third keyword (for example, if "RCAL" is used as the third keyword to get the recalling copy job list, the current copy job queue such DEF (deferred copy job) can be provided).
	13	Blank	
	14:18	Copy job size (SIZE)	The copy job size in MiB. The size is converted to MiB and it's round-down. If the size is smaller than 1 MiB, 0 is provided. The size is right justified and padded with blanks.
	19	Blank	
	20:29	Copy job age (AGE)	Age of the copy job since the last copy job state change. For example, a deferred copy job is posted with 1 minute delay first. Then, once it exits the delay, it changes the state internally and ready to be executed. This Age is calculated from the last state change and provides how long the copy job stays in the current state. The age is formatted as follows: HHHH:MM:SS If it has larger age than 10000 hours. 9999:59:59 is provided (note: when the copy job is upgraded internally or LI REQ, LVOL, <volser>, COPY, KICK command, it shows 9999:59:59 too. Note: From TS7700 Management Interface, the incoming queued volumes can be downloaded in a spreadsheet. "Age in Queue" in the spreadsheet is calculated from the original creation time which is equivalent to POSTED_TIME in this LI REQ output.
	30	Blank	
	31:49	Copy job posted time (POSTED_ TIME)	This is the timestamp in UTC when the copy job was originally posted (created). The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS
	50	Blank	

	51:54	Retry count (RET)	This is the retry count how many times the copy job has been attempted and delayed. The count is right justified and padded with blanks.
	55	Blank	
	56:60	Delay reason (D_RSN)	This provides the reason why the copy job attempt has been delayed. The possible values are as follows: ● NONE: The copy is not delayed. ● NOSRC: No available copy source was found. ● MNTED: The volume was still mounted. ● R-ERR: Recall failed in all possible copy sources. ● S-DIS: Due to the host copy disabled status, all copy sources are unavailable. ● HOT: The local cluster showed the volume was hot. ● IMDEF: Immediate deferred. The immediate copy was monitored. ● NOMON: Immediate deferred. The immediate copy was not monitored. ● TKMIS: The token showed a mismatch. A retry was required. ● TKERR: An unexpected token error occurred. A retry was required. ● LATER: The copy attempt to the last valid copy source failed. The copy will be performed again. ● C-CAN: The copy was cancelled. ● W-RUN: The non-immediate copy is waiting for the pending immediate copy on the same volser. ● FAMDL: Another local family member is copying the data. ● EXACS: Parallel copy is supported, but the volume was updated or had a scratch mount while the current copy was in progress. ● EXP-C: Copy failed due to the expired server certificate. Certificate must be replaced from MI SSL certificates page. ● GHOST: The copy on the cloud is not accessible from the local cluster yet. A retry will be performed.
	61:69	Blanks	
46	If the distributed library has more than 40 copy jobs, this line is provided to indicate that.		
	1	Blank	
	2:50	Message	'More records matched the searching criteria exist'
	51:69	Blanks	

3.1.3 COPYEXP Response

In response to the COPYEXP request, the cluster associated with the distributed library in the request will take actions against the volser specified based on the other keywords specified with the request.

If RECLAIM is specified, the cluster will check to see that the volume has a state of copy exported. If it does, it will make the volume eligible for reclaim. The reclaim will be performed at a higher priority than normal reclaim processing and will ignore the inhibit reclaim schedule.

If DELETE is specified, the cluster will check to see that the volume has a state of copy exported and that it is empty. If it meets those criteria, it is removed from the database of the cluster.

When the RECLAIM keyword is specified:

If the volume has a state of copy exported and contains active data, the following text is returned:

```
COPY EXPORT VOLUME RECLAIM V2 .0
VOLUME SCHEDULED FOR RECLAIM
```

If the volume does not have a state of copy exported, the following text is returned:

```
COPY EXPORT VOLUME RECLAIM V2 .0
VOLUME IS NOT AN EXPORT VOLUME
```

If the volume has a state of copy exported but is empty, the following text is returned:

```
COPY EXPORT VOLUME RECLAIM V2 .0
VOLUME IS EMPTY
```

If the volume specified is not in the library, the following text is returned:

```
COPY EXPORT VOLUME RECLAIM V2 .0
VOLUME NOT IN LIBRARY SPECIFIED
```

Line	Bytes	Name	Description
1	0:27	Header Info	'COPY EXPORT VOLUME RECLAIM V'
	28:29	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	30	Dot	'.'
	31:32		The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	33:69	Blanks	
If the volume has a state of copy exported and contains active data, the following line is returned			
2	0	Blank	
	1:28	Header Info	'VOLUME SCHEDULED FOR RECLAIM'
	29:69	Blanks	
If the volume does not have a state of copy exported, the following line is returned			
2	0	Blank	
	1:30	Header Info	'VOLUME IS NOT AN EXPORT VOLUME'
	31:69	Blanks	

If the volume has a state of copy exported but is empty, the following line is returned			
2	0	Blank	
	1:15	Header Info	'VOLUME IS EMPTY'
	16:69	Blanks	
If the volume specified is not in the library, the following line is returned			
2	0	Blank	
	1:31	Header Info	'VOLUME NOT IN LIBRARY SPECIFIED'
	32:69	Blanks	

When the DELETE keyword is specified.

If the volume has a state of copy exported and is empty, the following text is returned:

```
COPY EXPORT VOLUME DELETE V2 .0
VOLUME DELETED
```

If the volume does not have a state of copy exported, the following text is returned:

```
COPY EXPORT VOLUME DELETE V2 .0
VOLUME IS NOT AN EXPORT VOLUME
```

If the volume has a state of copy exported, but is not empty, the following text is returned:

```
COPY EXPORT VOLUME DELETE V2 .0
VOLUME IS NOT EMPTY
```

If the volume has already been deleted, the following text is returned:

```
COPY EXPORT VOLUME DELETE V2 .0
VOLUME NOT IN LIBRARY SPECIFIED
```

Line	Bytes	Name	Description
1	0:26	Header Info	'COPY EXPORT VOLUME DELETE V'
	27:28	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	29	Dot	'.'
	30:31		The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	32:69	Blanks	
If the volume has a state of copy exported and is empty, the following line is returned			
2	0	Blank	
	1:14	Header Info	'VOLUME DELETED'
	15:69	Blanks	
If the volume does not have a state of copy exported, the following line is returned			
2	0	Blank	
	1:30	Header Info	'VOLUME IS NOT AN EXPORT VOLUME'

Line	Bytes	Name	Description
	31:69	Blanks	
If the volume has a state of copy exported, but is not empty, the following line is returned			
2	0	Blank	
	1:19	Header Info	'VOLUME IS NOT EMPTY'
	20:69	Blanks	
If the volume has already been deleted, the following line is returned			
2	0	Blank	
	1:31	Header Info	'VOLUME NOT IN LIBRARY SPECIFIED'
	32:69	Blanks	

3.1.4 GRIDCNTL Response

In response to the GRIDCNTL request, the cluster associated with the distributed library in the request will take actions based on the additional keywords specified.

If a second keyword of COPY is specified, the actions the cluster will take have to do with its copy capabilities. The third keyword specifies what action is to be taken.

When the DISABLE keyword is specified:

If the cluster's copy capabilities are currently enabled, the cluster's copy capabilities will become disabled and the cluster will no longer be a candidate as a source for copies and will not initiate any copies where it is the target. Any copies already in execution will continue. All copies which are currently queued will remain queued. Future operations which result in new copies will also be queued but not processed. This may result in an immediate-deferred state being entered. Volumes where the disabled site has the only consistent copy will generate a copy backlog at peer clusters. Copy operations between other members of the Grid configuration are not affected. The following text is returned:

```
GRID COPY CAPABILITIES V2 .1
SUCCESSFUL DISABLE FOR SOURCE AND TARGET COPIES
```

If the cluster's copy capabilities are currently disabled, the following text is returned:

```
GRID COPY CAPABILITIES V2 .1
ALREADY DISABLED FOR SOURCE AND TARGET COPIES
```

When the ENABLE keyword is specified:

If the cluster's copy capabilities are currently disabled, the cluster's copy capabilities will become enabled and the cluster is again a candidate as a source for copies and will initiate any copies where it is the target. Any queued copy operations will proceed. Copy operations for other members of the Grid configuration are not affected. The following text is returned:

```
GRID COPY CAPABILITIES V2 .1
SUCCESSFUL ENABLE FOR SOURCE AND TARGET COPIES
```

If the cluster's copy capabilities are currently enabled, the following text is returned:

```
GRID COPY CAPABILITIES V2 .1
ALREADY ENABLED FOR SOURCE AND TARGET COPIES
```

Note: When a cluster is powered on or restarts, its copy capabilities default to enabled regardless of whether it had previously been disabled using this request.

If the request is issued to a cluster that is not configured in a Grid:

```
DISTRIBUTED LIBRARY DOES NOT EXIST
```

Line	Bytes	Name	Description
1	0:23	Header Info	'GRID COPY CAPABILITIES V'
	24:25	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.

	26	Dot	`.'`
	27:28	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0. The revision is incremented to 1 at 8.40.x.x.
	29:69	Blanks	
If the library the request is sent to is not in a Grid configuration, the following line is returned			
2	0	Blank	
	1:40		`CLUSTER NOT PART OF A GRID CONFIGURATION`
	41:69	Blanks	
If the cluster is enabled for copy and the keyword specified DISABLE, the following line is returned.			
2	0	Blank	
	1:37	Header Info	`DISABLED FOR SOURCE AND TARGET COPIES`
	38:69	Blanks	
If the cluster has already been disabled for copy and the keyword specified DISABLE, the following line is returned.			
2	0	Blank	
	1:45	Header Info	`ALREADY DISABLED FOR SOURCE AND TARGET COPIES`
	46:69	Blanks	
If the cluster is disabled for copy and the keyword specified ENABLE, the following line is returned.			
2	0	Blank	
	1:36	Header Info	`ENABLED FOR SOURCE AND TARGET COPIES`
	37:69	Blanks	
If the cluster has already been enabled for copy and the keyword specified ENABLE, the following line is returned.			
2	0	Blank	
	1:44	Header Info	`ALREADY ENABLED FOR SOURCE AND TARGET COPIES`
	45:69	Blanks	

3.1.5 LVOL Response

3.1.5.1 LVOL, <volser>, INFO(, FLASH)

Prior to R4.2, in response to the LVOL request for logical volume zzzzzz, when no other keywords are specified, the cluster receiving the request will collect the information on the volume and create the response. In a Grid configuration, the response indicates if each distributed library is to have a copy of the volume and the current validity state for the volume. For the distributed libraries that have a valid copy of the volume, the location of the copy, in cache or on a physical volume in a primary pool, is indicated. The response lines are formatted as follows:

If the logical volume exists:

```
LOGICAL VOLUME INFORMATION V5 .1
LOGICAL VOLUME:                Y03469
MEDIA TYPE:                     ECST
COMPRESSED SIZE (MB):           469
MAXIMUM VOLUME CAPACITY (MB):   832
CURRENT OWNER:                  TS001A
MOUNTED LIBRARY:                TS001A
MOUNTED VNODE:                  00
MOUNTED DEVICE:                 1E
TVC LIBRARY:                    TS001A
MOUNT STATE:                    RECALLING FROM AB0203
CACHE PREFERENCE:               PG0
CATEGORY:                       000F
LAST MOUNTED (UCT):             2006-05-23 19:34:23
LAST MODIFIED (UTC):            2006-05-23 19:35:54
LAST MOUNTED VNODE:             00
LAST MODIFIED DEVICE:           0000
TOTAL REQUIRED COPIES:           2
KNOWN CONSISTENT COPIES:        2
KNOWN REMOVED COPIES:           1
IMMEDIATE-DEFERRED:             N
DELETE EXPIRED:                 N
RECONCILIATION REQUIRED:         N
LWORM VOLUME:                   N
FLASH COPY:                     NOT ACTIVE
FORMAT ID:                      6
COMPRESSION METHOD:              FICON
3490 COUNTERS HANDLING:         SURFACE EOT
```

LIBRARY	RQ	CACHE	PRI	PVOL	SEC	PVOL	COPY	ST	COPY	Q	COPY	MD	REM
TS001A	N	N		AB0203		AD3056		CMPT		-		RUN	Y
TS001B	N	N		PA9405				REQR		DEF		DEF	N

LIBRARY	CP												
TS001A	1												
TS001B	0												

Starting from R4.2 (8.42.x.x), this LI REQ, LVOL, <volser> with no third keyword is no longer maintained due to the limitation of the maximum lines of the output and it is replaced by the new third keyword "INFO". If the third keyword "INFO" is not used, the following 2 lines are always provided to prompt the user to use the new keyword:

```
>>> THIS LI REQ IS NO LONGER MAINTAINED. PLEASE ADD 3RD KW 'INFO'
>>> TO GET LVOL INFORMATION WITH NEW FORMAT.
```

With the third keyword “INFO” supported in R4.2, the response lines are formatted as follows:

```
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 CPYS, REQ, REMOVED  : 1, 1, 0 (N)
DEL EXP, WHEN (UTC)       : N, -
HOT, FLASH COPY           : 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 CC
Lipizzan  N  Y ----- ----- CMP - DEF  N  2  0  C
Arabian   N  N ----- ----- NOR - NOC  N  0  0  A
cluster2  N  N ----- ----- NOR - NOC  N  0  0  N
Palomino  N  N ----- ----- NOR - NOC  N  1  0  N
```

Note: At 8.31.x.x, Flash Copy for DR testing function is supported. LVOL response includes the current Flash Copy status of each logical volume in the line 25 “FLASH COPY:”.

Note: At 8.32.x.x, the cache partition information is added at the end of output. Also the device number is provided with 4-digits to cover up to 496 (x1EF) devices. 496 devices feature is supported via RPQ only.

Note: As part of the 8.40.200.X release, different compression methods are supported. The chosen compression format for a given volume is displayed in the LVOL response.

Note: From R5.1, the cloud consistency (CC) column is introduced in the LVOL INFO output. If all the clusters in the domain are at the code level 8.51.x.x or above, the CC column is displayed. Otherwise the output is shown in V1.0 format (not having CC column).

Note: At 8.51.1.x, LWORM retention function is supported. The logical volume's LWORM retention state and time are added. This line is provided when all clusters in the domain are at the code level of 8.51.1.x or above. Otherwise, the output is provided in V2.0 format (this line is not provided).

If the third keyword ‘FLASH’ is specified, LVOL response provides the status of Flash Copy volume itself and its response lines are formatted as follows (line 25 “FLASH COPY:” isn't provided because the response is for Flash Copy volume status itself):

```

LOGICAL VOLUME INFORMATION V5 .0
FLASH COPY VOLUME:           Z10000
MEDIA TYPE:                   ECST
COMPRESSED SIZE (MB):         0
MAXIMUM VOLUME CAPACITY (MB): 800
CURRENT OWNER:                stardust
MOUNTED LIBRARY:
MOUNTED VNODE:
MOUNTED DEVICE:
TVC LIBRARY:                  elm
MOUNT STATE:
CACHE PREFERENCE:             ---
CATEGORY:                     100F
LAST MOUNTED (UCT):           1970-01-01 00:00:00
LAST MODIFIED (UTC):          2013-10-01 23:58:03
LAST MOUNTED VNODE:           00
LAST MODIFIED DEVICE:         0000
TOTAL REQUIRED COPIES:         -
KNOWN CONSISTENT COPIES:      -
KNOWN REMOVED COPIES:         -
IMMEDIATE-DEFERRED:           -
DELETE EXPIRED:               N
RECONCILIATION REQUIRED:       N
LWORM VOLUME:                 -
FORMAT ID:                    6
COMPRESSION METHOD:            FICON
3490 COUNTERS HANDLING:       SURFACE EOT

```

```

-----
  LIBRARY  RQ  CACHE  PRI  PVOL  SEC  PVOL  COPY  ST  COPY  Q  COPY  CP  REM
stardust   N      Y    -----  -----  CMPT      -      DEF    N
-----

```

```

  LIBRARY  CP
stardust   1

```

Starting from R8.42 (8.42.x.x), this LI REQ LVOL, <volser>, FLASH with no fourth keyword is no longer maintained due to the limitation of the maximum lines of the output and it is replaced by the new third keyword “INFO”. If the third keyword “INFO” is not used, the following 2 lines are always provided to prompt the user to use the new keyword:

```

>>> THIS LI REQ IS NO LONGER MAINTAINED. PLEASE ADD 3RD KW 'INFO'
>>> TO GET LVOL INFORMATION WITH NEW FORMAT.

```

With the third and fourth keyword “INFO, FLASH” supported in R4.2, the response lines are formatted as follows:

```

LOGICAL VOLUME INFO V3 .0
FLASH COPY 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, FLASH COPY                : 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 CC
Lipizzan N Y ----- ----- CMP - DEF N 2 0 C

```

Note: The Maximum Volume Capacity and the Cache Preference fields are not reporting correctly prior to the R1.4 (8.4.0.32) code release.

Note: From R5.1, the copy consistency (CC) column is introduced in the LVOL INFO output. If all the clusters in the domain are at the code level 8.51.x.x or above, the CC column is displayed. Otherwise the output is shown in V1.0 format (not having CC column).

Note: At 8.51.1.x, LWORM retention function is supported. The logical volume's LWORM retention state and time are added. This line is provided when all clusters in the domain are at the code level of 8.51.1.x or above. Otherwise, the output is provided in V2.0 format (this line is not provided).

The cluster processing the request must obtain information from the owning cluster to perform the request. If the owning cluster is not available, the following is included in the response for the library:

```
AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 1083
```

If information is needed from a distributed library and that library is unavailable, the following is included in the response for the library:

```
TS001C    DISTRIBUTED LIBRARY UNAVAILABLE
```

If the logical volume does not exist:

```
LOGICAL VOLUME                zzzzzz
VOLUME NOT IN LIBRARY SPECIFIED
```

If the third keyword 'FLASH' is specified but no Flash Copy is enabled, the following is included in the response for the library:

```
THERE IS NO FLASH COPY AVAILABLE FOR VOLSER zzzzzz
```

Line	Bytes	Name	Description
1	0:20	Header Info	` LOGICAL VOLUME INFO V '`
	21:22	Version	The version number for the response. The number if left justified and padded with blanks. Starts with 1. The version is

Line	Bytes	Name	Description
			incremented to 2 if all the clusters in the domain are at the code level 8.51.x.x or above. The version is incremented to 3 if all clusters in the domain are at the code level 8.51.1.x or above.
	23	Dot	`.'`
	24:25	Revision	The revision number for the response. The number is left justified and padded with blanks.
	26:69	Blanks	
2	0	Blank	For logical volume response
	1:32	Header Info	`LOGICAL VOLUME : `
	33:38	Volser	The logical volume information is being returned for. The volser is left justified and padded with blanks.
	39:69	Blanks	
2	0	Blank	For Flash Copy volume response
	1:32	Header Info	`FLASH COPY VOLUME : `
	33:38	Volser	The Flash Copy volume information is being returned for. The volser is left justified and padded with blanks.
	39:69	Blanks	
If the volume does not exist in the composite library the request is sent to, the following line is returned			
			`LOGICAL VOLUME VVVVVV` ` VOLUME NOT IN LIBRARY SPECIFIED`
If the volume exists in the composite library the request is sent to, the following lines are returned			
3	0	Blank	
	1:32	Header Info	`MEDIA, FMT, MAX(MB), CWRAP : `
	33:69	Media Type (MEDIA), Format ID of the logical volume (FMT), Maximum Volume Capacity (MAX), 3490 counters handling (CWRAP)	Media Type, Format ID of the logical volume, Maximum Volume Capacity and 3490 counters handling are provided with `,'` separator. The entire line is left justified and padded with a blank. Media Type: The media type assigned to the logical volume when it was inserted. Valid media types are CST and ECST. Format ID of the logical volume: The logical volume's Format ID which indicates the format of the data written in the logical volume. -2 Default value - may not reflect the actual level of a volume if that volume was created prior to R3.0 and has not yet been accessed. The ID is reverted to -2 when data of the volume is removed/deleted. 1/2/3 Data was imported from a B10/B20 VTS to a TS7700. 4 Data was written on a TS7700 at code level prior to R1.6. 5 Data was written on a TS7700 at code level R1.6 through R4.1.1 or at code level R4.1.2 or above in a mixed Grid code configuration.

Line	Bytes	Name	Description
			on. If the volume is not currently mounted or does not have a mount in progress, ',' is provided. (*) Mounted Library and Device are separated by '/'. Mount State: If the volume is mounted or a mount is in progress for the volume, this field indicates the mount state. If the volume is not currently mounted or does not have a mount in progress, ',' is provided. The following are the mount states that can be indicated: 'MOUNTING' A mount for the volume has been received and is in execution. 'RECALLREQ' A mount for the volume has been received and requires a recall. 'RECALLING FROM zzzzzz' A recall is in progress from physical volume zzzzzz. 'MOUNTED' A mount for the volume has been completed.
7	0	Blank	
	1:32	Header Info	'CACHE PREFERENCE, CATEGORY : '
	33:69	Cache Preference Setting (CACHE PREFERENCE), Category (CATEGORY)	Cache Preference Setting and Category are provided with ',' separator. The entire line is left justified and padded with a blank. Cache Preference Setting: This field indicates the cache preference associated with the volume at the tape volume cache library. The cache preference for a volume is determined when the volume was last closed. Cache preference settings returned while a volume is mounted are invalid. The following are the cache preference levels that can be indicated: 'PG0' The volume is preferred to be removed from cache after volume close, and in a Grid configuration, after copies have been performed. 'PG1' The volume is preferred to remain in cache after volume close. '***' The cache preference level is currently unknown. Flash Copy volume response always shows '---'. Category: The number, in hexadecimal, of the category that the volume is assigned to. The category's attribute is also provided: '(INSERT)' Insert category '(SCRATCH)' Scratch category '(PRIVATE)' Private category
8	0	Blank	
	1:32	Header Info	'LAST MOUNTED (UTC) : '

Line	Bytes	Name	Description
	33:51	Last Mount Activity Timestamp (LAST MOUNTED)	The timestamp in UTC when the volume was last mounted or unloaded by a host. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2006-05-23 19:34:23 If the volume has not yet been mounted, this field is set to 1970-01-01 00:00:00
	52:69	Blanks	
9	0	Blank	
	1:32	Header Info	`LAST MODIFIED LIB/DV, UTC (UTC) : '`
	33:69	Last Modified Library (LAST MODIFIED LIB)/Last Modified Device (DV), Last Modified Activity Timestamp (UTC)	Last Modified Library, Device and Activity Timestamp are provided with ',' or '/' separator. The entire line is left justified and padded with a blank. Last Modified Library: The name of the distributed library when the volume was mounted and modified. If the volume is not yet updated, ',' is provided. Last Modified Device: The number, in hexadecimal, of the virtual device that the volume was mounted on when it was last modified. If the volume has never been mounted and written to, ',' is provided. Last Modified Activity Timestamp: The timestamp in UTC when the contents of the volume was last modified by a host. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2006-05-23 19:34:23 If the volume has not yet been mounted, this field is set to 1970-01-01 00:00:00
10	0	Blank	
	1:32	Header Info	`KNOWN CPYS, REQ, REMOVED : '`
	33:69	Known Consistent Copies (KNOWN CPYS), Required Copies (REQ), Known Removed Copies (REMOVED) (Immediate deferred flag)	Known Consistent Copies, Required Copies, Known Removed Copies and Immediate deferred flag are provided with ',' separator. The entire line is left justified and padded with a blank. Known Consistent Copies: The total number of copies for the volume that are known to be consistent. If a cluster in a Grid configuration is unavailable, whether it has a consistent copy of the volume is unknown and would not be included in this number. The total includes the volume created or modified by the host. If the volume has never been mounted and written to, this field is 0. Flash Copy volume response always shows '-'.

Line	Bytes	Name	Description
			Required Copies: The total number of copies for the volume required based on the copy policies in place at the time the volume was last unloaded from a device. The total includes the volume created or modified by the host. If the volume has never been mounted and written to, this field is 0. Flash Copy volume response always shows '-'. Known Removed Copies: The total number of copies for the volume that are known to have been removed. If a cluster in a Grid configuration is unavailable, whether it has a copy that was removed is unknown and would not be included in this number. Flash Copy volume response always shows '-'. Immediate deferred flag: Indicates whether the volume is currently in the immediate copy deferred state. 'Y' The volume is in the immediate copy deferred state. 'N' The volume is not in the immediate copy deferred state. If the volume has never been mounted or written to, this flag is set to 'N'. Flash Copy volume response always shows '-'.
11	0	Blank	
	1:32	Header Info	'DEL EXP, WHEN (UTC) : '
	33:69	Delete expired flag (DEL EXP), Delete expired Timestamp (WHEN)	Delete expired flag and Delete expired Timestamp are provided with ',' separator. The entire line is left justified and padded with a blank. Delete expired flag: Indicates whether the contents of the volume have been deleted from the database through the delete expired volume data function. When a volume's contents have been deleted, the volume can only be re-used as a scratch volume. 'Y' The volume's data has been deleted. 'N' The volume's data has not been deleted. If the volume has never been mounted or written to, this flag is set to 'N'. Note : The delete expire function runs on each distributed library independently. It could be possible that one distributed library has already deleted volume data, but the other does not yet. LVOL response is a composite view and this delete expired flag is set to 'Y' only when all the available clusters have deleted the contents of the volume.

Line	Bytes	Name	Description
			Delete expired Timestamp: The timestamp in UTC when the contents of the volume was delete expired. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2006-05-23 19:34:23 If the volume has not been delete expired, this field is set to ‘,’.
12	0	Blank	
	1:32	Header Info	`HOT, FLASH COPY : '`
	33:69	Reconciliation flag (HOT), Flash Copy status of the logical volume (FLASH COPY)	Reconciliation flag and Flash Copy status of the logical volume are provided with ‘,’ separator. The entire line is left justified and padded with a blank. Delete expired flag: Indicates whether the volume's status on all cluster's in a Grid configuration is consistent or if status changes are needed or in progress. This flag is primarily the result of one of the clusters in a Grid being unavailable when a status change occurs. When the cluster is again available, the volumes status will be updated as a background task. `Y' The volume's status is not consistent across all of the clusters.</p> <p>`N' The volume's status is consistent across all of the clusters. Flash Copy status of the logical volume: The logical volume's Flash Copy status. The status is left justified and padded with blanks. `NOT ACTIVE' Flash Copy is not enabled.</p> <p>`ACTIVE' Flash Copy is enabled and the content or attribute of the Flash Copy volume is still the same with its logical volume. `CREATED' Flash Copy is enabled and the content or attribute of the Flash Copy volume is already modified since the Flash Copy is enabled.
13	0	Blank	
	1:32	Header Info	`LWORM RET STATE, TIME (UTC) : '`
	33	Retention state	LWORM retention state. The possible values are: `N': The volume is not LWORM-retained.</p> <p>`D': The volume is date based LWORM-retained. `F': The volume is forever LWORM-retained. Please refer to “LWORM Retention User's Guide” too. This line is provided only when the Grid is at 8.51.1.x or above.

Line	Bytes	Name	Description
	34:35	Header Info (Comma separator)	` , `
	36:54	Retention time	LWORM retention time. The timestamp in UTC until when the volume is LWORM-retained. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2021-04-13 00:00:00 If the volume is not LWORM-retained (i.e. the retention state is 'N'), 'NA' is provided. If the volume is forever LWORM-retained (i.e. the retention state is 'F'), '- ' is provided. Please refer to “LWORM Retention User’s Guide” too. This line is provided only when the Grid is at 8.51.1.x or above.
	55:69	Blanks	
14	0:69	Separator	All dash '-' characters
15	0:69	Header Info	` LIBRARY RQ CA P-PVOL S-PVOL CPS CPQ CPP RM CP CD CC `
For each configured distributed library, a line is formatted as follows.			
For Flash Copy volume status, the information of the distributed library in DR family is only provided.			
If the distributed library is not available			
15+N	0	Blank	
	1:8	Library	The name of the distributed library.
	9:11	Blanks	
	12:42	Message	`DISTRIBUTED LIBRARY UNAVAILABLE`
	43:69	Blanks	
If the distributed library is available (N clusters are configured)			
15+N	0	Blank	
	1:8	Library	The name of the distributed library.
	9:10	Blanks	
	11	Reconciliation Required (RQ)	Indicates whether the status of the volume on this library current or not and needs to be reconciled with the rest of the Grid. `Y` Reconcile is required. The volume’s status is not consistent with the rest of the Grid. `N` The volume’s status is consistent with the rest of the Grid.
	12:13	Blanks	
	14	Cache (CA)	Indicates whether a valid copy of the volume is resident in the cache of the distributed library. `Y` The valid copy of the volume is resident.

Line	Bytes	Name	Description
			`N' The volume is not resident or is not valid.
	15	Blank	
	16:21	Primary Pool Volume (P-PVOL)	If a valid copy of the logical volume resides on a primary pool volume, this field contains the volser of the physical volume. If a valid copy is not resident on a primary pool physical volume, this fields is set to all dashes ('-').
	22	Blank	
	23:28	Secondary Pool Volume (S-PVOL)	If a valid copy of the logical volume resides on a secondary pool volume, this field contains the volser of the physical volume. If a valid copy is not resident on a primary pool physical volume, and a secondary volume pool is associated with the logical volume, this field is set to all dashes ('-'). If a secondary volume pool is not associated with the logical volume, this field contains all blanks.
	29	Blank	
	30:32	Copy State (CPS)	This field indicates the state of the copy for the logical volume to the distributed library. The following are the copy states that can be indicated: `CMP' : The copy of the volume is complete.</p> <p>`REQ' : A copy of the volume is required and has not yet begun. `INP' : A copy of the volume is required and is in progress.</p> <p>`REC' : An update to the volume's copy locations is being processed. Copy state will be updated when reconcile is complete. `TDL' : A copy of the volume is required and is still delayed based off of Time Delayed copy mode and delay option. The value is right justified and padded with blanks.
	33	Blank	
	34:36	Copy Queue (CPYQ)	This field indicates the queue that the copy is in on the distributed library. The following are the copy queues that can be indicated: `RUN' : The copy operation for the volume is in the Rewind/Unload copy consistency point queue.</p> <p>`SYD' : The copy operation for the volume is in the Sync copy consistency point queue. `DEF' : The copy operation for the volume is in the Deferred copy consistency point queue.</p> <p>`IMD' : The policy for the copy operation for the volume is the Rewind/Unload copy consistency point, but it could not be completed and has been placed in the immediate copy deferred queue. `TDL' : The copy operation for the volume is in the Time Delayed copy consistency point queue.

Line	Bytes	Name	Description
			If no copy queue exists, '-' is provided.
	37	Blank	
	38:40	Copy Policy (CPP)	This field indicates the copy consistency policy for the volume on the distributed library. The following are the copy policies that can be indicated: 'RUN' : The distributed library is to have a copy consistent with the rewind/unload operation for the volume. 'SYN' : Two different distributed library are to have a copy consistent with the rewind/unload operation for the volume. 'TDL' : The distributed library is to have a copy since the delay defined in the management class is expired after the rewind/unload operation for the volume. 'NOC' : The distributed library is not to have a copy of the volume. 'EXS' : The copy consistency point for the volume on the distributed library has changed from a valid copy mode to No Copy mode. The contents of the logical volume have not changed so a valid copy still exists.
	41:42	Blanks	
	43	Volume Removed (RM)	This field indicates that the volume has been removed from this distributed library. Auto removal as well as LVOL REMOVE request could remove a volume from the distributed library. 'Y' : The volume has been removed. 'N' : The volume has not been removed.
	44:45	Blanks	
	46	Cache Partition (CP)	The cache partition assigned to the logical volume on each distributed library. If the distributed library is TS7700D or TS7740, it's always 0. The possible value is -1 or 0-7. When the logical volume is just inserted then no mount/demount is executed to the volume yet, no cache partition is assigned to the volume and -1 will be provided. When TS7700T is upgraded from TS7700D, the cache partition 0 is automatically assigned during the upgrade process for all volumes residing in the TS7700D cache.
	47:48	Blanks	
	49	Cloud Data (CD)	This field indicates the number of the logical volume on cloud (object data) which is accessible from the distributed library.
	50:51	Blanks	
	52	Cloud Consistency (CC)	Indicates whether the cluster is aware of the existence of the volume in the cloud. 'N' : The cluster is not aware of the volume in the cloud 'A' : The cluster is aware of the volume in the cloud, but yet to check the existence. 'C' : The cluster has already checked the existence of the volume in the cloud.

Line	Bytes	Name	Description
	53:69	Blanks	

3.1.5.2 LVOL, <volser>, CLDINFO

Starting from R4.2, LI REQ, LVOL, <volser>, CLDINFO is supported to provide the logical volume status on cloud (i.e. object data). If the logical volume is already premigrated to cloud and latest data is accessible from a cluster in the Grid, the response lines are formatted as follows:

If all the clusters in the domain are at the code level 8.51.x.x or above:

```
LOGICAL VOLUME CLOUD INFORMATION V2 .0
```

```
LOGICAL VOLUME:                Z99999
```

```
-----
LIBRARY ST RK      POOL/ ACCOUNT/CL  CNT      SIZE (MiB)
Lipizzan  M   1  MYPOOL01/MYACCT01/2    12         61250
           2  MYPOOL02/MYACCT02/2    40        700000
           3  MYPOOL03/MYACCT01/2    99       1000000
           4  MYPOOL04/MYACCT02/2    10       1800000
Arabian   - NO DATA IS PREMIGRATED TO CLOUD
Oak       M   1  MYPOOL01/MYACCT01/0    12         61250
Palomino  - NO DATA IS PREMIGRATED TO CLOUD
```

If the logical volume does not exist:

```
LOGICAL VOLUME                ZZZZZZ
VOLUME NOT IN LIBRARY SPECIFIED
```

Line	Bytes	Name	Description
1	0:33	Header Info	'LOGICAL VOLUME CLOUD INFORMATION V'
	34:35	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.51.x.x.
	36	Dot	'.'
	37:38	Revision	The revision number for the response. The number is left justified and padded with blanks.
	39:69	Blanks	
2	0	Blank	For logical volume response
	1:30	Header Info	'LOGICAL VOLUME: '
	31:36	Volser	The logical volume information is being returned for. The volser is left justified and padded with blanks.
	37:69	Blanks	

Line	Bytes	Name	Description
3	0:69	Separator	All dash '-' characters
4	0:54	Header Info	` LIBRARY ST RK POOL/ ACCOUNT/CL CNT SIZE (MiB) '
	55:69	Blanks	
For each configured distributed library (N clusters are configured), a line is formatted as follows. For Flash Copy volume status, the information of the distributed library in DR family is only provided.			
If the distributed library is not available			
-	0	Blank	
	1:8	Library	The name of the distributed library.
	9:10	Blanks	
	11:41	Message	'DISTRIBUTED LIBRARY UNAVAILABLE'
	42:69	Blanks	
If the distributed library is available			
Each distributed library could premigrate the logical volume to up to 4 different clouds. Each cloud data configured to the library is displayed in separate line.			
-	0	Blank	
	1:8	Library	The name of the distributed library. This item is displayed only once per library at the first line.
	9:10	Blanks	
	11	Volume Data Status (ST)	Indicates whether the volume on this library current is premigrated to cloud and accessible from the distributed library or not. 'P' : The volume is premigrated to cloud (i.e. the volume is also on cache) 'M' : The volume is migrated to cloud (i.e. the volume is no longer on cache) 'A' : The volume is available in the cloud (i.e. the volume is not on cache but can be recalled from cloud as another cluster migrated it) '-' : The volume is not on cloud (this does NOT mean no consistent volume exists on cache or tape; there is no consistent volume on cloud, or the cluster has no access to it). If no volume data exists on cloud, additional text is provided: 'NO DATA IS PREMIGRATED TO CLOUD' This item is displayed only once per library at the first line.
	12:13	Blanks	
	14	Cloud Premigration Rank (RK)	The cloud premigration rank.
	15	Blank	
	16:23	Cloud Pool Nickname (POOL)	The cloud pool nickname defined by the customer where the logical volume is premigrated to, right justified and padded with blanks.
	24	Separator	`/'

Line	Bytes	Name	Description
	31:36	Volser	The logical volume information is being returned for. The volser is left justified and padded with blanks.
	37:69	Blanks	
3	0:69	Separator	All dash '-' characters
4	0:53	Header Info	` LIBRARY ST POOL/ ACCOUNT/CL POOL/ ACCOUNT/CL`
	54:69	Blanks	
For each configured distributed library, a line is formatted as follows.			
For Flash Copy volume status, the information of the distributed library in DR family is only provided.			
If the distributed library is not available			
4+N	0	Blank	
	1:8	Library	The name of the distributed library.
	9:10	Blanks	
	11:41	Message	`DISTRIBUTED LIBRARY UNAVAILABLE`
	42:69	Blanks	
If the distributed library is available (N clusters are configured)			
Each distributed library could premigrate the logical volume to up to 4 different clouds. Each line could provide up to 2 cloud data. If the distributed library has premigrated the logical volume to more than 2 different clouds (pools), the line 4+N+1 provides additional 1 or 2 cloud data status. At 8.42.x.x, only 1 cloud pool per each volume is supported.			
4+N	0	Blank	
	1:8	Library	The name of the distributed library.
	9:10	Blanks	
	11	Volume Data Status (ST)	Indicates whether the volume on this library current is premigrated to cloud and accessible from the distributed library or not. `P` : The volume is premigrated to cloud (i.e. the volume is also on cache) `M` : The volume is migrated to cloud (i.e. the volume is no longer on cache) `A` : The volume is available in the cloud (i.e. the volume is not on cache but can be recalled from cloud as another cluster migrated it) `-` : The volume is not on cloud (this does NOT mean no consistent volume exists on cache or tape). If no data exists on cloud, additional text is provided: `NO DATA IS PREMIGRATED TO CLOUD`
	12	Blank	
	13:20	Cloud Pool Nickname (POOL)	The cloud pool nickname defined by the customer where the logical volume is premigrated to, right justified and padded with blanks.
	21	Separator	`/`

Line	Bytes	Name	Description
	22:29	Cloud Account Nickname (ACCOUNT)	The cloud account ID nickname defined by the customer, right justified and padded with blanks.
	30	Separator	` / '
	31	Cluster ID to premigrate data to cloud (CL)	The cluster ID which has premigrated the logical volume to cloud.
	32:33	Blanks	
	34:41	Cloud Pool Nickname (POOL)	The cloud pool nickname defined by the customer where the logical volume is premigrated to, right justified and padded with blanks.
	42	Separator	` / '
	43:50	Cloud Account Nickname (ACCOUNT)	The cloud account ID nickname defined by the customer, right justified and padded with blanks.
	51	Separator	` / '
	52	Cluster ID to premigrate data to cloud (CL)	The cluster ID which has premigrated the logical volume to cloud.
	53:69	Blanks	
4+N+1	0:12	Blanks	
	13:20	Cloud Pool Nickname (POOL)	The cloud pool nickname defined by the customer where the logical volume is premigrated to, right justified and padded with blanks.
	21	Separator	` / '
	22:29	Cloud Account Nickname (ACCOUNT)	The cloud account ID nickname defined by the customer, right justified and padded with blanks.
	30	Separator	` / '
	31	Cluster ID to premigrate data to cloud (CL)	The cluster ID which has premigrated the logical volume to cloud.
	32:33	Blanks	
	34:41	Cloud Pool Nickname (POOL)	The cloud pool nickname defined by the customer where the logical volume is premigrated to, right justified and padded with blanks.
	42	Separator	` / '
	43:50	Cloud Account Nickname (ACCOUNT)	The cloud account ID nickname defined by the customer, right justified and padded with blanks.
	51	Separator	` / '
	52	Cluster ID to premigrate data to cloud (CL)	The cluster ID which has premigrated the logical volume to cloud.

Line	Bytes	Name	Description
	53:69	Blanks	

3.1.5.3 LVOL, <volser>, CLDVERS, <page index>

Starting from R5.1, LI REQ, LVOL, <volser>, CLDVERS, <page index> is supported to provide the list of volume versions stored on the cloud. The list is separated into pages made up of 7 records per each. Only one page can be displayed at a time, and the page index can be specified by the fourth keyword. The default page index 1 is applied when the fourth keyword is not set.

If the volume versions are already created for the specified volume, the response lines are formatted as follows:

```
LOGICAL VOLUME CLOUD VERSION INFORMATION V1 .0
```

```
LOGICAL VOLUME: ZCLD00 STATUS: NORMAL
```

```
-----
0: myverylong40charactersgpprefixname0123456/7683/BA030/BA87A/3957T/
   Z00000/10700000000000000000/1020000000000000000/20191007212440
IV,DL,SZ: 0000000000000000278,0000000000000000150,1208
PL,ACT   : MYPOOL,MYACCT
D,P-TIME: 2020-05-14 01:59:21, 2020-05-14 01:57:08
-----
```

```
1:      lipizzan/41c6/BA092/BA92C/CF55V/ZCLD00/278/147/20200509004039
IV,DL,SZ: 0000000000000000278,0000000000000000147,1208
PL,ACT   : MYPOOL,MYACCT
D,P-TIME: 2020-05-09 01:01:10, 2020-05-09 00:43:08
-----
```

```
2:      lipizzan/41be/BA092/BA92A/35B30/ZCLD00/278/145/20200509003102
IV,DL,SZ: 0000000000000000278,0000000000000000145,1208
PL,ACT   : MYPOOL,MYACCT
D,P-TIME: 2020-05-09 00:34:48, 2020-05-09 00:31:03
-----
```

```
3:      lipizzan/43fb/BA092/BA92A/35B30/ZCLD00/278/144/20200509001658
IV,DL,SZ: 0000000000000000278,0000000000000000144,1208
PL,ACT   : MYPOOL,MYACCT
D,P-TIME: 2020-05-09 00:30:57, 2020-05-09 00:16:59
-----
```

```
4:      lipizzan/41c6/BA092/BA92A/35B30/ZCLD00/278/143/20200509001331
-----
```

IV,DL,SZ: 00000000000000000278,00000000000000000143,1208

PL,ACT : MYPOOL,MYACCT

D,P-TIME: 2020-05-09 00:16:40, 2020-05-09 00:13:32

5: lipizzan/41c6/BA092/BA92C/CF55V/ZCLD00/278/142/20200508054359

IV,DL,SZ: 00000000000000000278,00000000000000000142,820358

PL,ACT : MYPOOL,MYACCT

D,P-TIME: 2020-05-09 00:13:28, 2020-05-08 05:44:01

6: lipizzan/0f30/BA092/BA92C/CF55V/ZCLD00/278/139/20200507090256

IV,DL,SZ: 00000000000000000278,00000000000000000139,1208

PL,ACT : MYPOOL,MYACCT

D,P-TIME: 2020-05-07 09:38:13, 2020-05-07 09:02:57

8 MORE VERSIONS EXIST...

If no volume version exists for the specified volume:

NO RECORDS FOUND

The maximum index of version which can be displayed in the output is limited to 3-digit size, and the number of versions which can be displayed per one page is limited to 7, so only 142 page indexes are allowed. If the specified page index is too large, the following error is displayed:

THE INPUT IS TOO LARGE: 143

Line	Bytes	Name	Description
1	0:41	Header Info	'LOGICAL VOLUME CLOUD VERSION INFORMATION V'
	42:43	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1.
	44	Dot	'.'
	45:46	Revision	The revision number for the response. The number is left justified and padded with blanks.
	47:69	Blanks	
2	0:16	Header Info	' LOGICAL VOLUME: '
	17:22	Volser	The volume version information is being returned for.
	23:31	Header Info	' STATUS: '
	31:69	Status	The status of the most recent version of the volume. The status string is left justified and padded with blanks. 'NORMAL' is displayed.
3	0:69	Separator	All dash '-' characters

Line	Bytes	Name	Description
For each volume versions, output is sorted by newest to oldest and formatted as follows.			
3+6N	0:2	Version Index	The index of the cloud backup. Starts with 0.
	3:4	Header Info	` : '`
	5:68	Object Name	The name of the object exists on the cloud. If the name is longer than 64 characters, the exceeded part is wrapped to the next line. The name is left justified and padded with blanks.
3+6N+1	5:68	Object Name	The name of the object exists on the cloud. The name is left justified and padded with blanks.
3+6N+2	0:11	Header Info	` IV, DL, SZ : '`
	12:30	Insert Version	The insert version of the volume.
	31	Comma	` , '`
	32:50	Data Level	The data level of the volume.
	51	Comma	` , '`
	52:69	Size	The size of the version (byte). The size string is left justified and padded with blanks.
3+6N+3	0:11	Header Info	` PL, ACT : '`
	12:69	Cloud Info	The cloud pool nickname and the cloud account, separated by comma. The cloud info string is left justified and padded with blanks.
3+6N+4	0:11	Header Info	` D, P-TIME : '`
	12:30	Deleted Time Stamp	The time stamp the version was deleted. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS For example: 2020-05-14 01:59:21 If the volume has not yet been deleted, this field is set to 1970-01-01 00:00:00
	31:32	Comma + Blank	` , '`
	33:51	Premigrated Time Stamp	The time stamp the version was premigrated. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS For example: 2020-05-14 01:59:21
	52:69	Blanks	
3+6N+5	0:69	Separator	All dash '-' characters
If there are still more versions to show, the following message is displayed in the end of the output.			
		Message	` X MORE VERSIONS EXIST...` Or if the page index is 142: ` X MORE VERSIONS EXIST BUT REACHED MAXIMUM DISPLAYABLE VALUE` where 'X' is the number of remaining versions. The message string is left justified and padded with blanks.

3.1.5.4 LVOL, <volser>, PREFER/MIGRATE

In response to the LVOL request when the PREFER or MIGRATE keyword is specified, the cluster associated with the distributed library in the request will first determine if the specified logical volume is in the cache. The MIGRATE keyword is not applicable in a disk-only configuration without cloud tier (7700C). If the volume is in the cache, the MIGRATE keyword will modify the volume's cache management position in the cache. A volume's cache management position determines the order in which volumes are removed from the cache when space is needed and if the volume is subject to removal even if space is not needed. In the terms that have been used to describe the management of volumes in the TS7740/7700T(CPx)/7700C(CPx), volumes that have been assigned to preference group 1 (PG1) are removed from cache when space is needed in Least Recently Accessed order and are never subject to removal unless space is needed. Volumes that have been assigned to preference group 0 (PG0) are most likely to be removed from cache when space is needed, in largest in size first order. PG0 volumes are also subject to be removed from cache even when space is not needed and in this case are removed in smallest in size first order.

If the PREFER keyword is specified, the volume's cache management position is modified as if it was the most recently accessed volume. If the cluster is not a disk-only configuration, the volume is managed as if it had been assigned to PG1. If the cluster is a disk-only configuration, the timestamp when the volume becomes eligible for removal is refreshed thus emulating a mount access for read. The assigned removal policy is not modified.

The following text is returned when the PREFER keyword is specified:

```
LOGICAL VOLUME CONTROL V2 .0
LOGICAL VOLUME zzzzzz UPDATED FOR PREFERRED RETENTION IN CACHE
```

Line	Bytes	Name	Description
1	0:23	Header Info	'LOGICAL VOLUME CONTROL V'
	24:25	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	26	Dot	'.'
	27:28	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	29:69	Blanks	
2	0	Blank	
	1:14	Header Info	'LOGICAL VOLUME'
	15	Blank	
	16:21	Volser	The logical volume the information is being returned for. The volser is left justified and padded with blanks.
	22	Blank	
	23:62	Header Info	'UPDATED FOR PREFERRED RETENTION IN CACHE'
	63:69	Blanks	

If the MIGRATE keyword is specified, the volume's cache management position is modified as a PGO volume that is the largest in size (preferred for removal first)

The following text is returned when the MIGRATE keyword is specified:

```
LOGICAL VOLUME CONTROL V2 .0
LOGICAL VOLUME zzzzzz UPDATED FOR PREFERRED REMOVAL FROM CACHE
```

Line	Bytes	Name	Description
1	0:23	Header Info	'LOGICAL VOLUME CONTROL V'
	24:25	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1.
	26	Dot	'.'
	27:28	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	29:69	Blanks	
2	0	Blank	
	1:14	Header Info	'LOGICAL VOLUME'
	15	Blank	
	16:21	Volser	The logical volume the information is being returned for. The volser is left justified and padded with blanks.
	22	Blank	
	23:62	Header Info	'UPDATED FOR PREFERRED REMOVAL FROM CACHE'
	63:69	Blanks	

If the MIGRATE keyword is specified for a volume that is in a cluster in a disk-only configuration, the following error text is returned (See 3.12 Error Responses):

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

3.1.5.5 LVOL, <volser>, REMOVE

If the REMOVE keyword is specified, then an immediate attempt will be made to remove this volume from cache assuming at least one copy exists on another cluster. Pinned volumes and volumes that are still retained due to the minimum retention time can also immediately be removed. The REMOVE request behavior is dependent on the type of cluster and code level installed. With 8.7.x.x or older, only a TS7700D supports the REMOVE request. Beginning with code level 8.21.x.x or later, the REMOVE request is supported in the TS7740 when the copy policy of the volume is set to 'N' (No Copy) or 'E' (Exist) at the removal site. As of 8.32.x.x, all cluster models support the REMOVE request, though TS7740, TS7700T tape managed partitions (1 through 7) will only allow removal of volumes set to 'N' (No Copy) or 'E' (Exist). At 8.42.x.x, TS7700C also supports the REMOVE request as same as TS7700T.

Two functional enhancements are added for the REMOVE request at the code level 8.30.x.x:

1. Prior to the code level of 8.30.x.x, if no data exists in the cache, the REMOVE request always returns "NOT IN CACHE" regardless of the current in-flight copy job state. There was a case that the wrong copy mode is set to the volumes and copy jobs were queued on the cluster which is not supposed to have the copies, then eventually the cluster's cache was filled with the unnecessary copies. At the code level 8.30.x.x or above, the REMOVE request to such an in-flight copy can be used to remove the copy jobs to prevent the cache from being filled up by unnecessary copies. Only the TS7700D or TS7700T with the resident cache partition removes an in-flight copy. Once the copy job is removed successfully, the volume token is updated as if the volume is removed from the cache and the following text is returned:

'LOGICAL VOLUME XXXXXX ONLY COPY QUEUE REMOVED'

2. Prior to the code level of 8.30.x.x, the REMOVE request is rejected if automatic removal is disabled by "SETTING CACHE REMOVE DISABLE" on the cluster which receives the request and the error text 'NOT REMOVED FROM CACHE' is always returned. It is changed at 8.30.x.x that REMOVE request is accepted even if automatic removal is disabled.

Note: The automatic removal does not work if it is disabled. Only the removal request through the host command line request works.

If the optional keyword PROMOTE is supplied with the REMOVE keyword, then the LRU value for the logical volume will be updated to move that logical volume to the front of removal queue.

Volumes which are pinned or are in fast-ready categories are not candidates for promotion. The removal threshold must still be crossed before removal takes place. In addition, volumes in fast-ready categories will be removed first before attempting to remove the promoted volumes.

If the REMOVE was specified and the logical volume was successfully removed from cache:

Line	Bytes	Name	Description
1	0:23	Header Info	'LOGICAL VOLUME CONTROL V'
	24:25	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	26	Dot	'.'
	27:28	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	29:69	Blanks	
2	0	Blank	
	1:14	Header Info	'LOGICAL VOLUME'
	15	Blank	
	16:21	Volser	The logical volume the information is being returned for. The volser is left justified and padded with blanks.
	22	Blank	
	23:40	Header Info	'REMOVED FROM CACHE'
	41:69	Blanks	

If REMOVE was specified and the volume could not be removed:

Line	Bytes	Name	Description
2	0	Blank	
	1:14	Header Info	'LOGICAL VOLUME'
	15	Blank	
	16:21	Volser	The logical volume the information is being returned for. The volser is left justified and padded with blanks.
	22	Blank	
	23:44	Header Info	'NOT REMOVED FROM CACHE'
	45:69	Blanks	

If REMOVE PROMOTE was specified:

Line	Bytes	Name	Description
2	0	Blank	
	1:14	Header Info	'LOGICAL VOLUME'
	15	Blank	
	16:21	Volser	The logical volume the information is being returned for. The volser is left justified and padded with blanks.
	22	Blank	
	23:62	Header Info	'UPDATED FOR PREFERRED REMOVAL FROM CACHE'
	63:69	Blanks	

If REMOVE INFO was specified:

Line	Bytes	Name	Description
1	0:27	Header Info	'LOGICAL VOLUME INFORMATION V'
	28:29	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	30	Dot	'.'
	31:32	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	33:69	Blanks	
2	0	Blank	
	1:15	Header Info	'LOGICAL VOLUME: '
	16:22	Blanks	
	23:28	Volser	The logical volume the information is being returned for. The volser is left justified and padded with blanks.
	29:69	Blanks	
3	0:69	Separator	All dash '-' characters
4	0	Blank	
	1:8	Library	The name of the distributed library.
	9:69	Blanks	

Line	Bytes	Name	Description
5	0:1	Blanks	
	2:16	Header Info	`REMOVAL POLICY: '
	17:22	Blanks	
	23:24	Removal Policy	The removal policy of the volume the information is being returned for. The following are the removal policies that may be returned: ` 0' = (Prefer Remove) The copy of a private volume is removed as long as an appropriate number of copies exists on peer clusters, the pinning duration (in X hours) has elapsed since last access and the available free space on the cluster has fallen below the removal threshold. The order of which volumes are removed under this policy is based on their least recently used (LRU) access times. Volumes in Group 0 are removed prior to the removal of volumes in Group 1 except for any volumes in fast-ready categories which are always removed first. Archive and backup data would be a good candidate for this removal group since it won't likely be accessed once written. ` 1' = (Prefer Keep) The copy of a private volume is removed as long as an appropriate number of copies exists on peer clusters, the pinning duration (in X hours) has elapsed since last access, the available free space on the cluster has fallen below a threshold and LRU group 0 has been exhausted. The order of which volumes are removed under this policy is based on their least recently used (LRU) access times. Volumes in Group 0 are removed prior to the removal of volumes in Group 1 except for any volumes in fast-ready categories which are always removed first. ` 4' = (Pinned) The copy of the volume is never removed from this TS7700D or TS7700T (resident partition) cluster. The pinning duration is not applicable and is implied as infinite. Once a pinned volume is moved to scratch, it becomes a priority candidate for removal similarly to the next two policies. This feature must be used cautiously to prevent TS7700D or TS7700T CP0 cache full condition.
	25:69	Blanks	
6	0:1	Blanks	
	2:15	Header Info	`REMOVAL STATE: '
	16:22	Blanks	
	23:30	Removal State	The current residency of the volume the information is being returned for. `                ' Volume has not been removed.</p> <p>`PINNED ' Removal Policy Pinned or Disabled. `HOLD          ' Volume is not a valid removal candidate yet.</p> <p>`RESIDENT' Volume has not been removed. `RETAINED'     Volume set to never be removed.</p> <p>`DEFERRED' Removal of volume has been deferred. `REMOVED ' Volume has been removed.

Line	Bytes	Name	Description
	31:69	Blanks	
7	0:1	Blanks	
	2:20	Header Info	'REMOVAL TIME (UTC) : '
	21:22	Blanks	
	23:52	Removal Time	The timestamp of when the volume the information being requested was removed. If the Removal State specified HOLD then this timestamp denotes the time when this volume becomes eligible to be removed.
	53:69	Blanks	

If the volume is not in cache and no in-flight copy job exists on the TS7700D and TS7700T (resident cache partition), the following text is returned:

'LOGICAL VOLUME zzzzzz NOT IN CACHE'

If the REMOVE keyword is specified for a volume that is in a cluster that is not a disk-only configuration and the volume selected for removal has a copy mode different to "Exist" or "No copy", the following text is returned:

'LOGICAL VOLUME xxxxxx NOT REMOVED FROM CACHE'

If the REMOVE keyword is specified for a volume that is in the tape partition on TS7700T, the following text is returned:

'VOLUME xxxxxx IN PARTITION Y COULD NOT BE REMOVED'

3.1.5.6 LVOL, <volser>, COPY, KICK/FORCE

If the COPY KICK or FORCE keyword is specified for a logical volume an attempt will be done to promote it to the front of the copy queue.

When a logical volume is removed, the COPY KICK command does not put a copy job on the distributed library and the text "VOLUME IS NOT QUEUED" is returned. When it is required to get the removed volume back by copying it from another consistent distributed library, the COPY FORCE command needs to be used. The COPY FORCE command puts a copy job against a removed volume and promotes it to the front of the copy queue as well as the COPY KICK command.

The following text is returned when the Logical Volume exists in the queue:

```
LOGICAL VOLUME COPY KICK V3 .0
VOLUME MOVED TO FRONT OF QUEUE CLASS
```

Line	Bytes	Name	Description
1	0:25	Header Info	'LOGICAL VOLUME COPY KICK V'

	26:27	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x and 3 at 8.30.x.x.
	28	Dot	`.`
	29:30	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	31:69	Blanks	
2	0	Blank	
	1:36	Header Info	`VOLUME MOVED TO FRONT OF QUEUE CLASS`
	37:69	Blanks	

When this keyword is used against a logical volume in the recall required queue it will just speed up the recall request. Once recalled, it will again go to the back of the active list. A second kick would be needed to move it up front there too.

The following text is returned when the Logical Volume is not in the copy queue:

```
LOGICAL VOLUME COPY KICK V3 .0
VOLUME IS NOT QUEUED
```

Line	Bytes	Name	Description
2	0	Blank	
	1:20	Header Info	`VOLUME IS NOT QUEUED`
	21:69	Blanks	

The following text is returned when the Logical Volume is already active:

```
LOGICAL VOLUME COPY KICK V3 .0
VOLUME IS ALREADY ACTIVE
```

Line	Bytes	Name	Description
2	0	Blank	
	1:24	Header Info	`VOLUME IS ALREADY ACTIVE`
	25:69	Blanks	

The following text is returned when the Logical Volume is already consistent and no copy is required:

```
LOGICAL VOLUME COPY KICK V3 .0
VOLUME IS ALREADY CONSISTENT
```

Line	Bytes	Name	Description
2	0	Blank	
	1:28	Header Info	'VOLUME IS ALREADY CONSISTENT'
	29:69	Blanks	

An error response is formatted as follows:

```
AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 3
```

Line	Bytes	Name	Description
1	0:59	Header Info	'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = '
	60:61	Return Code	Indicates the area that detected the error. The return code is right justified and padded with blanks.
	62:69	Blanks	

3.1.5.7 LVOL, <volser>, COPY, CANCEL

If the COPY CANCEL keyword is specified for a logical volume, an attempt will be done to cancel the in-flight copy and delay the copy job. This command may be useful when the in-flight copy needs to be cancelled immediately, or the copy job may get stuck in “in-flight” state and it needs to be cancelled (reset). This command does not delete a copy job.

The command can be issue to the copy job which is not in “in-flight” state and the copy job is just delayed and retried later. But the command fails if the copy job is currently recalling the copy source volume.

The following text is returned when the Logical Volume exists in the queue:

```
LOGICAL VOLUME COPY CANCEL V1 .0
VOLUME ZA0040 COPY JOB WAS CANCELLED SUCCESSFULLY
```

Line	Bytes	Name	Description
1	0:27	Header Info	'LOGICAL VOLUME COPY CANCEL V'
	28:29	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.50.x.x.
	30	Dot	'.'
	31:32	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision starts with 0 at 8.50.x.x.
	33:69	Blanks	
2	0	Blank	

	1:49	Header Info with volser	'VOLUME ZA0040 COPY JOB WAS CANCELLED SUCCESSFULLY' (byte 8-13 contains the volser)
	50:69	Blanks	

The following text is returned when the specified volser is not in the copy queue:

```
VOLUME VVVVVV COPY JOB WAS NOT FOUND
```

The following text is returned when the specified volser is recalling the copy source volume:

```
VOLUME VVVVVV COPY JOB IS IN RECALL TASK
```

The following text is returned when an unexpected error occurs:

```
VOLUME VVVVVV COPY JOB CANCEL REPORTED INTERNAL ERROR. RC 1
```

3.1.6 CLDBKUP Response

3.1.6.1 SHOW <page index>

Starting from R5.1, CLDBKUP SHOW <page index> is supported to provide the list of cloud pool database backups ever created on the cloud. The list is separated into pages made up of 49 records per each. Only one page can be displayed at a time, and the page index can be specified by the third keyword. The default page index 1 is applied when the third keyword is not set.

If the pool backups are already created on the cloud, the response lines are formatted as follows:

```
CLDBKUP SHOW V1 .0
POOLNAME BACKUPID BACKUPNAME DESCRIPTION
MYPOOL01 10000001 S1234.C25B1.B25B1.20200121123456.001 DESCRIPT001
MYPOOL01 10000002 S1234.C25B1.B25B1.20200121123457.002 DESCRIPT002
MYPOOL01 10000003 S1234.C25B1.B25B1.20200121123458.003 DESCRIPT003
[...]
MYPOOL01 10000048 S1234.C25B1.B25B1.20200121123458.048 DESCRIPT048
MYPOOL01 10000049 S1234.C25B1.B25B1.20200121123458.049 DESCRIPT049
128 MORE BACKUPS EXIST...
```

If no pool backups exist on the cloud:

```
NO RECORDS FOUND
```

Line	Bytes	Name	Description
1	0:13	Header Info	`CLDBKUP SHOW V'
	14:15	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1.
	16	Dot	`.'
	17:18	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	19:69	Blanks	
2	0	Blank	
	1:8	Header Info	`POOLNAME'
	9	Blank	
	10:17	Header Info	`BACKUPID'
	18	Blank	
	19:28	Header Info	`BACKUPNAME'
	29:55	Blanks	
	56:66	Header Info	`DESCRIPTION'
	67:69	Blanks	
For each pool backups on the cloud, output is sorted by newest to oldest and formatted as follows.			
3+N	0	Blank	

Line	Bytes	Name	Description
	1:8	Pool Nickname	The nickname of the cloud pool the backup was created for. Left justified.
	9	Blank	
	10:17	Backup ID	The backup ID automatically generated when the backup was created. Left justified.
	18	Blank	
	19:54	Backup Name	The name of the backup. Left justified.
	55	Blank	
	56:67	Description	The user provided description when the cloud export was initiated. Left justified.
	68:69	Blanks	
If there are still more backups to show, the following message is displayed in the end of the output.			
		Message	' X MORE BACKUPS EXIST...' where 'X' is the number of remaining backups. The message string is left justified and padded with blanks.

Note: Pool Nickname shows the pool nickname having been assigned when the backup was created. Even if the nickname has been changed after created the backup, the Pool Nickname still shows the old pool nickname.

3.1.7 PDRIVE Response

In response to the PDRIVE request, the cluster associated with the distributed library in the request will examine the state and usage of its physical drives. A response line is provided for every drive that is installed. The type is based on the physical machine type and model of the drive and its mode further defines the current mode the drive is in. A drive is available if the cluster has access to the drive. A drive's role indicates its current usage. A drive is idle if it is not currently in use for another role or it is not mounted. Once a task has been completed, the physical volume may remain mounted on the drive for up to ten minutes waiting for another role to be assigned to it. A drive is being used for migration if it is currently copying data from the tape volume cache to a physical tape. A drive is being used for recall if it is currently copying data from a physical tape to the tape volume cache. A drive is being used for reclaim if it is the source or target of a reclaim operation. A drive is being used for secure data erasure if it is currently erasing the physical volume. If a drive is mounted, the physical volume mounted and the pool associated with the physical volume is indicated. If a logical volume is being transferred to or from the physical volume, it is indicated. The response lines are formatted as follows:

```
PHYSICAL DRIVES V2 .1
```

SERIAL NUM	TYPE	MODE	AVAIL	ROLE	POOL	PVOL	LVOL
000007861241	3592E05	E05	Y	MIGR	01	AB3945	L03013
000007834290	3592E05	E05	Y	MIGR	01	AC0419	V04019
000007838493	3592E05	E05	Y	RECA	03	AA4019	T10349
000007825193	3592E05	E05	Y	IDLE	00		
000007868475	3592E05	E05	Y	MIGR	01	AC5019	L03987
000007854964	3592E05E	E05	Y	IDLE	02	AA9201	
000007855978	3592E05E	E05	Y	RCRT	02	AA9501	V05678
000007832984	3592E05E	E05	Y	IDLE	00		
000007827381	3592E05	E05	Y	SECE	04	AF0718	
000007861543	3592E05E	E05	Y	RCRS	02	AD7581	V05678

Line	Bytes	Name	Description
1	0:16	Header Info	'PHYSICAL DRIVES V'
	17:18	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	19	Dot	'.'
	20:21	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0. The revision is incremented to 1 at 8.33.x.x.
	22:69	Blanks	
2	0:2	Blanks	
	3:12	Header Info	'SERIAL NUM'
	13:17	Blanks	
	18:21	Header Info	'TYPE'
	22:23	Blanks	

Line	Bytes	Name	Description
	24:27	Header Info	'MODE'
	28:29	Blanks	
	30:34	Header Info	'AVAIL'
	35:36	Blanks	
	37:40	Header Info	'ROLE'
	41:42	Blanks	
	43:46	Header Info	'POOL'
	47:50	Blanks	
	51:54	Header Info	'PVOL'
	55:58	Blanks	
	59:62	Header Info	'LVOL'
	63:69	Blanks	
For each drive installed in the cluster, a line is formatted as follows			
N	0	Blank	
	1:12	Drive Serial Number (SERIAL NUM)	This field indicates the serial number of the physical drive.
	13	Blank	
	14:21	Drive Type and Model (TYPE)	This field indicates the physical drive machine type and model installed. The following are the values that can be indicated: '3592J1A' The installed drive is a 3592 model J1A. '3592E05' The installed drive is a 3592 model E05 (TS1120). '3592E05E' The installed drive is a 3592 model E05 (TS1120) is encryption capable. '3592E06' The installed drive is a 3592 model E06 (TS1130). '3592E08' The installed drive is a 3592 model E08 (TS1150). '359260F' The installed drive is a 3592 model E08 (TS1160). The value is right justified and padded with blanks.
	22:23	Blanks	
	24:27	Drive Format Mode (MODE)	This field indicates the format the physical drive is operating in. The following are the values that can be indicated: 'J1A' The drive is operating in J1A format. 'E05' The drive is operating in E05 (TS1120 native) format. 'E05E' The drive is operating in E05 (TS1120 native) encrypting format. 'E06' The drive is operating in E06 (TS1130 native) format. 'E06E' The drive is operating in E06 (TS1130 native) encrypting format. 'E07' The drive is operating in E07 (TS1140 native) format. 'E07E' The drive is operating in E07 (TS1140 native) encrypting format. 'E08' The drive is operating in E08 (TS1150 native) format.

Line	Bytes	Name	Description
			'E08E' The drive is operating in E08 (TS1150 native) encrypting format. '55F' The drive is operating in 55F (TS1155 native) format. '55FE' The drive is operating in 55F (TS1155 native) encrypting format. '60F' The drive is operating in 60F (TS1160 native) format. '60FE' The drive is operating in 60F (TS1160 native) encrypting format. The value is right justified and padded with blanks. Note: TS7700 does not support TS1155 drive but TS1160 drive supports 55F/55FE format for JD/JL media.
	28:33	Blanks	
	34	Availability (AVAIL)	This field indicates whether the drive is available for use by the TS7740. The following are the values that can be indicated: 'Y' The drive is available for use. 'N' The drive is unavailable for use.
	35:36	Blanks	
	37:40	Role (ROLE)	The current role the drive is performing. The following are the values that can be indicated: 'IDLE' The drive is currently not in use for another role or is not mounted. 'SECE' The drive is currently being used to erase a physical volume. 'MIGR' The drive is being used to copy a logical volume from the tape volume cache to a physical volume. 'RECA' The drive is being used to recall a logical volume from a physical volume to the tape volume cache. '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.
	41:44	Blanks	
	45:46	Physical Pool (POOL)	If the drive is in use (not idle), the pool of the mounted physical volume. If the drive is idle, this field contains '00'.
	47:48	Blanks	
	49:54	Physical Volume (PVOL)	If the drive is in use (not idle), the volser of the mounted physical volume. If the drive is idle, this field contains all blanks.
	55:56	Blanks	
	57:62	Logical Volume (LVOL)	If the drive is in use (not idle), the volser of the logical volume being processed. If the drive is idle, this field contains all blanks.
	63:69	Blanks	

If the distributed library specified is for a cluster in a disk-only configuration, the following error text is returned (See 3.12 Error Responses):

'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'

3.1.8 POOLCNT Response

In response to the POOLCNT request for pool PP, the cluster associated with the distributed library in the request will examine the contents of its physical pools and generate a formatted response. The response provides a count of the physical volumes in each pool that are empty, are in the process of being filled or have been marked as full as well as volumes that are marked as read only, unavailable or copy-exported. The count for empty includes physical volumes that have been assigned to the pool as well as volumes that were borrowed from the common scratch pool but have not yet been returned. For pool 0 (the common scratch pool), only the empty counts are provided.

The response begins with the pool specified in keyword 2. If it is not provided, the response data begins with pool 0. If there is more pool information than can be provided in the response lines, the last line will indicate that more data is available. If all of the data lines for a pool will not fit in the response, none of them are included. To obtain more data, issue the request with a starting pool number of the last pool in the prior response.

Counts are provided for each media type associated with the pools (up to a maximum of 8) and the response lines are formatted as follows:

```
PHYSICAL MEDIA COUNTS V2 .0
POOL MEDIA  EMPTY  FILLING    FULL  ERASE      ROR  UNAVAIL  CXPT
   0      J      340
   0      K      523
   0     JA       83
   0     JB       28
   0     JJ        2
   1     JA     134        3    609        0        0        0        0
   2     JJ      23        2    325        0        0        0        0
```

Last line if more data is available:

```
MORE POOL COUNT DATA IS AVAILABLE
```

If there are no volumes associated with the pool specified or that follow it, the following text is returned.

```
PHYSICAL MEDIA COUNTS V2 .0
POOL MEDIA  EMPTY  FILLING    FULL  ERASE      ROR  UNAVAIL  CXPT
NO VOLUMES IN SPECIFIED POOLS
```

If the pool specified is greater than the number of pool supported by the TS7700, the following text is returned.

```
PHYSICAL MEDIA COUNTS V2 .0
POOL MEDIA  EMPTY  FILLING    FULL  ERASE      ROR  UNAVAIL  CXPT
LIBRARY SUPPORTS 32 POOLS
```

Line	Bytes	Name	Description
1	0:22	Header Info	'PHYSICAL MEDIA COUNTS V'
	23:24	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	25	Dot	'.'

Line	Bytes	Name	Description
	26:27	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	28:69	Blanks	
2	0	Blank	
	1-4	Header Info	'POOL'
	5	Blank	
	6:10	Header Info	'MEDIA'
	11:12	Blanks	
	13:17	Header Info	'EMPTY'
	18:19	Blanks	
	20:26	Header Info	'FILLING'
	27:29	Blanks	
	30:33	Header Info	'FULL'
	34:35	Blanks	
	36:40	Header Info	'ERASE'
	41:44	Blanks	
	45:47	Header Info	'ROR'
	48:49	Blanks	
	50:56	Header Info	'UNAVAIL'
	57:58	Blanks	
	59:62	Header Info	'CXPT'
	63:69	Blanks	
For each media type defined for a pool, a line is formatted as follows			
N	0:2	Blanks	
	3:4	Pool (POOL)	Pool number. The pool number is right justified and padded with a leading blank.
	5:8	Blanks	
	9:10	Media Type (MEDIA)	The media type defined for the pool. Media types with a single character are padded with a leading blank.
	11:12	Blanks	
	13:17	Empty Count (EMPTY)	The count of the physical volumes that are empty for the media type. The count is right justified and padded with leading blanks.
	18:21	Blanks	
	22:26	Filling Count (FILLING)	The count of the physical volumes that are in the filling state for the media type. The count is right justified and padded with leading blanks. This field is all blanks for pool 0.
	27:28	Blanks	
	29:33	Full Count (FULL)	The count of the physical volumes that have been marked full for the media type. The count is right justified and padded with leading blanks. This field is all blanks for pool 0.
	34:35	Blanks	

Line	Bytes	Name	Description
	36:40	Erase Count (ERASE)	The count of the physical volumes that have been reclaimed, but need to be erased before they become empty. The count is right justified and padded with leading blanks. This field is all blanks for pool 0.
	41:42	Blanks	
	43:47	ROR Count (ROR)	The count of the physical volumes that are in the read only recovery state. The count is right justified and padded with leading blanks.
	48:51	Blanks	
	52:56	Unavailable (UNAVAIL)	The count of the physical volumes that are in the unavailable or destroyed state. The count is right justified and padded with leading blanks.
	57:58	Blanks	
	59:62	Copy Exported (CXPT)	The count of the physical volumes that are in the copy exported state. The count is right justified and padded with leading blanks.
	63:69	Blanks	
If there are more complete pools to report then will fit in the response, the following line is provided			
M	0:32	Notice	'MORE POOL COUNT DATA IS AVAILABLE'
	33:69	Blanks	
If there are no volumes assigned to the pool specified or after it, the following line is provided			
3	1:29	Notice	'NO VOLUMES IN SPECIFIED POOLS'
	30:69	Blanks	
If the pool specified is greater than the pools supported, the following line is provided			
3	1:24	Notice	'LIBRARY SUPPORTS 32 POOLS'
	24:69	Blanks	

If the distributed library specified is for a cluster in a disk-only configuration, the following error text is returned (See 3.12 Error Responses):

'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'

3.1.9 PVOL Response

In response to the PVOL request for physical volume xxxxxx, the cluster associated with the distributed library in the request will extract data for the volume out of its database. The lines are formatted as follows:

If the physical volume exists:

```
PHYSICAL VOLUME INFORMATION V2 .0
  PHYSICAL VOLUME:  AB0203
  MEDIA TYPE:        JA
  DRIVE MODE:        E05
  FORMAT:            TS7700
  VOLUME STATE:      READ-WRITE
  CAPACITY STATE:    FULL
  CURRENT POOL:      1
  MBYTES WRITTEN:    497413
  % ACTIVE DATA:    85.2
  LAST INSERTED:     2006-12-23 19:34:27
  WHEN EXPORTED:     N/A
  MOUNTS:            34
  LOGICAL VOLUMES:   254
  ENCRYPTED:          N
```

If the physical volume is encrypted, the key labels associated with the volume are included in the output:

```
.
.
.
  ENCRYPTED:          Y
  "FINANCIAL ACCOUNTING RECORDS"
  "SEC INTERCHANGE AUDIT RECORDS"
```

If the physical volume does not exist in the library:

```
PHYSICAL VOLUME INFORMATION V2 .0
  PHYSICAL VOLUME:  ABC123
  VOLUME NOT IN LIBRARY SPECIFIED
```

Line	Bytes	Name	Description
1	0:28	Header Info	'PHYSICAL VOLUME INFORMATION V'
	29:30	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	31	Dot	'.'
	32:33	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	34:69	Blanks	
2	1:16	Header Info	'PHYSICAL VOLUME: '
	17:18	Blank	

Line	Bytes	Name	Description
	19:24	Volser	The physical volume the information is being returned for. The volser is right justified and padded with blanks.
	25:69	Blanks	
If the volume does not exist in the distributed library the request is sent to, the following line is returned			
3	0	Blank	
	1:31		'VOLUME NOT IN LIBRARY SPECIFIED'
	32:69	Blanks	
If the volume exists in the distributed library the request is sent to, the following lines are returned			
3	0	Blank	
	1:11	Header Info	'MEDIA TYPE: '
	12:18		
	19:20	Media Type	The media type of the volume. Media types with a single character are padded on the left with a blank.
	21:69	Blanks	
4	0	Blank	
	1:11	Header Info	'DRIVE MODE: '
	12:18		
	19:22	Drive Mode	The drive mode the media was last written under. The value is left justified and padded with a blank.
	23:69	Blanks	
5	0	Blank	
	1:7	Header Info	'FORMAT: '
	8:18	Blanks	
	19:24	Format	The current state of the physical volume. The following are the values that can be indicated: 'TS7700' The volume is written in native TS7700 format. 'VTS' The volume is written in the legacy VTS format. The value is left justified and padded with blanks.
	25:69	Blanks	
6	0	Blank	
	1:13	Header Info	'VOLUME STATE: '
	13:18	Blanks	
	19:38	Volume State	The current state of the physical volume. The following are the values that can be indicated: 'READ-WRITE' The volume can be read from or written to. 'READ-ONLY' The volume contains valid data, but had encountered read and or write errors that exceeded an allowable threshold. It is only accessible for read. 'MANUALLY EJECTED' The volume is manually ejected and unavailable to the TS7740/7700T.

Line	Bytes	Name	Description
			'MISPLACED' The volume is unavailable to the TS7740/7700T and can't be located. 'INACCESSIBLE' The volume is unavailable to the TS7740/7700T and currently in a location that the cartridge accessor cannot access. 'COPY_EXPORTED' The volume has been removed from the library through the copy export operation. 'COPY_EXPORT_PENDING' The volume is in the process of copy export operation and it doesn't finish yet. 'CE_RECLAIM' The volume has been copy exported and received the library request command COPYEXP,<physical volume>, RECLAIM to make the volume eligible for reclamation. Please refer to the section of COPYEXP Response for the details. The value is left justified and padded with blanks.
	39:69	Blanks	
7	0	Blank	
	1:15	Header Info	'CAPACITY STATE: '
	15:18	Blanks	
	19:25	Capacity State	The current capacity state of the physical volume. The following are the values that can be indicated: 'EMPTY' The volume contains no data and is available for use as a physical scratch volume. 'FILLING' The volume contains valid data, but is not yet full. It is available for additional data to be added to it. 'FULL' The volume contains valid data. As some point it was marked as full and additional data cannot be added to it. A volume can be marked full in some cases short of the volume capacity limit. 'UNKNOWN' The volume's capacity state is unknown. The value is left justified and padded with blanks.
	26:69	Blanks	
8	0	Blank	
	1:13	Header Info	'CURRENT POOL: '
	14:18	Blanks	
	19:20	Current Pool	The pool number the volume is currently assigned to. The pool number is left justified and padded with a blank.
	21:69	Blanks	
9	0	Blank	
	1: 15	Header Info	'MBYTES WRITTEN: '
	16:18	Blanks	

Line	Bytes	Name	Description
	19:26	MB Written	The number of MB that had been written to the physical volume at the time it was marked as full. The count is left justified and padded with blanks.
	27:69	Blanks	
10	0	Blank	
	1: 14	Header Info	'% ACTIVE DATA: '
	15:18	Blanks	
	19:23	% Active Data	The percent of the data that was written to the physical volume when it was marked full, that is currently still valid. The number is shown to 1/10th of a percent with a decimal point separating the 10ths. The value is left justified and padded with blanks.
	24:69	Blanks	
11	0	Blank	
	1: 14	Header Info	'LAST INSERTED: '
	15:18	Blanks	
	19:37	Last Inserted Timestamp	The field indicates the date and time that the physical volume was last inserted into the library. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2006-05-23 19:34:23 If the volume has not yet been mounted, this field is set to 1970-01-01 00:00:00
	38:69	Blanks	
12	0	Blank	
	1: 14	Header Info	'WHEN EXPORTED: '
	15:18	Blanks	
	19:37	Last Exported Timestamp	The field indicates the date and time that the physical volume was last exported from the library. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2006-05-23 19:34:23 If the volume has not been exported, this field is set to 'N/A' and padded with blanks.
	38:69	Blanks	
13	0	Blank	
	1: 7	Header Info	'MOUNTS: '
	8:18	Blanks	
	19:23	Mount Count	The number of times that the physical volume has been mounted since it was inserted into the library. The count is right justified and padded with leading blanks.
	24:69	Blanks	
14	0	Blank	
	1:16	Header Info	'LOGICAL VOLUMES: '

Line	Bytes	Name	Description
	17:18	Blanks	
	19:24	Logical Volume Count	The number of active logical volume resident on the physical volume. The value is left justified and padded with blanks.
	25:69	Blanks	
15	0	Blank	
	1:10	Header Info	'ENCRYPTED:'
	11:18	Blanks	
	19	Encryption Flag	This field indicates whether the physical volume's data has been encrypted. A volume can change between being encrypted and non-encrypted only when re-used as a scratch volume. The following are the values that can be indicated: 'Y' The volume's data has been encrypted. 'N' The volume's data has not been encrypted.
	20:69	Blanks	
If the encryption flag is 'Y', the key labels assigned to the volume are added to the output. Up to two key labels can be assigned to a volume.			
N	0:65	Key label	A key label is between 1 and 64 characters in length. For this field, the key label is enclosed in double quotes. For key labels that are less than 64 characters in length, the label (with surrounding quotes) is left justified and padded with blanks.
	66:69	Blanks	

If the distributed library specified is for a cluster in a disk-only configuration, the following error text is returned (See 3.12 Error Responses):

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

At the code level 8.31.0.x or above, the third keyword "DELETE" is supported. The keyword provides a method to delete the specified physical volume record from TS7700 database. The specified physical volume must be empty and not physically in the library in order to be deleted successfully.

When the specified physical volume is deleted successfully, the following text is returned:

```
COPY EXPORT VOLUME DELETE V2 .0
VOLUME DELETED
```

Line	Bytes	Name	Description
1	0:26	Header Info	'COPY EXPORT VOLUME DELETE V'
	27:28	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 2.
	29	Dot	'.'
	30:31	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	32:69	Blanks	
2	0	Blank	

Line	Bytes	Name	Description
	1:14	Blank	'VOLUME DELETED'
	15:69	Blanks	

If the specified physical volume is not an empty volume, the following error text is returned:

```
VOLUME NOT EMPTY
```

If the specified physical volume still physically exists in the library, the following error text is returned:

```
VOLUME IS NOT AN UNAVAILABLE VOLUME
```

If the specified physical volume doesn't exist in the library, the following error text is returned:

```
VOLUME NOT IN LIBRARY SPECIFIED
```

If any unexpected error occurs and the request is not completed, the following error text is returned:

```
AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = XX
```

Starting with the code level 8.50.0.x, the third keyword "RELABEL" is supported. This keyword provides a method to instruct the TS7700 to write a new internal volser label to the physical volume the next time it is written from the beginning of tape. Once the physical volume has been re-labeled, this flag will be disabled. This method allows the operator to insert a tape cartridge into the TS7700 where the external volser label does not match the internal volser label.

For the relabel command, the usage is expected to be the following:

1. The physical tape is inserted into the physical library and is known to have an external label that is different than the internally written label.
2. Before the TS7700 selects the inserted physical tape for writing, the customer should issue a library request command to relabel this tape. In the case where the TS7700 had already attempted to use this tape for writing and encountered the label mismatch error, this tape may be put in read-only recovery where the tape will be ejected. If the tape is already in read-only, issuing a library request command to relabel this tape has no effect since the tape will be ejected by the read-only processing. In such a case, you must wait for the eject completion, then reinsert the tape back into the library and then resend the library request to relabel this tape.

3.1.10 RECALLQ Response

In response to the RECALLQ request, the cluster associated with the distributed library will examine the current contents of its recall queue and provide a formatted response.

Unless keyword 3 specifies 'PROMOTE', the response begins with the logical volume specified in keyword 2. If keyword 2 is not provided, the response data begins with the first volume in the queue. If keyword 3 specifies 'PROMOTE', and the logical volume specified in keyword 2 is in the recall queue and not currently in progress or scheduled, it will be moved to the top of the queue for volumes to be scheduled. After the move is complete, the response data begins with the first volume in the queue. If the specified volume is already in progress or is scheduled, its position in the queue is not changed and the response data begins with the first volume in the queue.

If there are more queue entries then can be provided in the response lines, the last line will indicate that more data is available. To obtain more data, issue the request with the last logical volume in the prior response.

There are three kinds of entries for the logical volumes in the queue, those that are in progress, those that have been scheduled to be executed next and those that remain to be scheduled.

The response lines are formatted as follows:

```
RECALL QUEUE V2 .0
  POS      LVOL      PVOL1      PVOL2      TIME
  IP  L00121  AB0456
  IP  L99356  AA0350          201
  SC  Y30458  AB0456          148
  SC  L54019  AA0350          145
   1  L67304  AC0101          135
   2  T09356  P00010  P00167   126
   3  Y35728  AD5901          102
```

Last line if more data is available:

```
MORE RECALL QUEUE DATA IS AVAILABLE
```

If the request specifies a volume that is not in the recall queue, the following text is returned:

```
RECALL QUEUE V2 .0
  POS      LVOL      PVOL1      PVOL2      TIME
VOLUME NOT IN RECALL QUEUE
```

If there are no volumes in the recall queue, the following text is returned:

```
RECALL QUEUE V2 .0
  POS      LVOL      PVOL1      PVOL2      TIME
RECALL QUEUE EMPTY
```

Line	Bytes	Name	Description
1	0:13	Header Info	'RECALL QUEUE V'
	14:15	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.

Line	Bytes	Name	Description
	16	Dot	`.'`
	17:18	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	19:69	Blanks	
2	0:1	Blanks	
	2:4	Header Info	`POS`
	5:8	Blanks	
	9:12	Header Info	`LVOL`
	13:15	Blanks	
	16:20	Header Info	`PVOL1`
	21:23	Blanks	
	24:28	Header Info	`PVOL2`
	29:31	Blanks	
	32:35	Header Info	`TIME`
	36:69	Blanks	
If a specific volume is specified and it is not in the queue, the following line is returned.			
3	0	Blank	
	1:26	Header Info	`VOLUME NOT IN RECALL QUEUE`
	27:69	Blanks	
If there are no volumes in the recall queue, the following line is returned.			
3	0	Blank	
	1:18	Header Info	`RECALL QUEUE EMPTY`
	19:69	Blanks	
For each recall queue entry, a line is formatted as follows			
3:N	0:1	Blanks	
	2:4	Position (POS)	Position of the logical volume in the recall queue. The following are the values that can be indicated: `IP` A recall is in progress for the volume. `SC` The volume has been scheduled for a recall. If optimization is enabled, the TS7740 will schedule recalls so that multiple recalls are processed from the same physical volume. Position The volume's current position in the list of volumes that have not yet been scheduled. The position is a three digit number. The values are right justified and padded with blanks.
	5:6	Blanks	
	7:12	Logical Volume (LVOL)	The logical volume to be recalled. The volser is right justified and padded with blanks.
	13:14	Blanks	

Line	Bytes	Name	Description
	15:20	Physical Volume 1 (PVOL1)	The physical volume the logical volume resides on. The volser is right justified and padded with blanks.
	21:22	Blanks	
	23:28	Physical Volume 2 (PVOL2)	If the logical volume spans a physical volume, the second physical volume the volumes resides on. The volser is right justified and padded with blanks.
	29:30	Blanks	
	31:35	Time in Queue (TIME)	The time that the logical volume has been in the recall queue, in seconds. The value is right justified and padded with leading blanks. If the time is greater than 9999 seconds, 9999 is indicated.
	36:69	Blanks	
If there are more queue entries than can fit in the response, the following line is provided			
M	0:34	Notice	'MORE RECALL QUEUE DATA IS AVAILABLE'
	35:69	Blanks	

If the distributed library specified is for a cluster in a disk-only configuration, the following error text is returned (See 3.12 Error Responses):

'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'

3.1.11 RRCLSUN Response

In response to the RRCLSUN request, the cluster associated with the distributed library will enable, disable or display the current status of the force residency on recall feature. The TS7700 has a resident on recall function to accelerate format change to newest one. The resident on recall function makes the logical volume(LVOL) on the older formatted physical volume(PVOL) invalidated just after recall. As a result, the recalled LVOL becomes “resident” then it will premigrate/migrate to a newer formatted PVOL.

There are three valid options for keyword 2 of a RRCLSUN request.

If a keyword 2 of ENABLE is specified the force residency on recall feature will be enabled.

If a keyword 2 of DISABLE is specified the force residency on recall feature will be disabled.

If a keyword 2 of STATUS is specified the current setting of the force residency on recall feature will be displayed in the response.

The response lines are formatted as follows:

```
RESIDENT ON RECALL V2 .0
RRCLSUN:      ENABLED
```

Line	Bytes	Name	Description
1	0:19	Header Info	'RESIDENT ON RECALL V'
	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	22	Dot	'.'
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	25:69	Blanks	
2	0:1	Blanks	
	2:9	Header Info	'RRCLSUN: '
	10:14	Blanks	
	15:22	Resident on Recall for Sunset	The currently configured value for Resident on Recall for Sunset Physical Volumes. The following values can be configured and are left justified and padded with blanks: 'ENABLED' The TS7740 resident on recall function to accelerate format change to newest one is enabled. The resident on recall function makes the lvol on the older formatted pvol invalid just after recall. As a result, the recalled lvol becomes “resident” then it will premigrate/migrate to a newer formatted pvol. 'DISABLED' The TS7740 resident on recall function to accelerate format change to newest one is disabled. This is the DEFAULT setting unless modified by the customer.
	23:69	Blanks	

If the distributed library specified is for a cluster in a disk-only configuration, the following error text is returned (See 3.12 Error Responses):

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

3.1.12 SETTING

The SETTING request provides information about many of the current workflow and management settings of the cluster specified in the request as well as the ability to modify the settings. It also allows alerts to be set for many of the resources managed by the cluster.

In response to the SETTING request, the cluster associated with the distributed library in the request will modify its settings based on the additional keywords specified. If no additional keywords are specified, the request will just return the current settings. See the example of the data returned after the rest of the keyword descriptions.

When a value is specified, lead blanks or zeroes are ignored.

Note: All settings are persistent across machine restarts, service actions or code updates. The settings are not carried forward as part of Disaster Recovery from Copy Exported tapes or the recovery of a system.

The following table lists the keywords supported for the SETTING request, whether the request is applicable to a composite and/or distributed library, and whether the request is applicable on a cluster with physical tape and/or a disk-only configuration.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
SETTING	ALERT	COPYLOW	Value	N/A	Y	Y	Y	Y	Y
SETTING	ALERT	COPYHIGH	Value	N/A	Y	Y	Y	Y	Y
SETTING	ALERT	PDRVLOW	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	ALERT	PDRVCRIT	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	ALERT	PSCRLOW	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	ALERT	PSCRCRIT	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	ALERT	RESLOW	Value	N/A	Y	Y	Y	Y	Y
SETTING	ALERT	RESHIGH	Value	N/A	Y	Y	Y	Y	Y
SETTING	ALERT	RSDTLOW	Value	N/A	Y	N/A	N/A	Y	Y
SETTING	ALERT	RSDTHIGH	Value	N/A	Y	N/A	N/A	Y	Y
SETTING	ALERT	PCPYLOW	Value	N/A	Y	Y	Y	Y	Y
SETTING	ALERT	PCPYCRIT	Value	N/A	Y	Y	Y	Y	Y
SETTING	ALERT	DEFDEG	ENABLE/ DISABLE	Y	N/A	Y	Y	Y	Y
SETTING	ALERT	LINKDEG	ENABLE/ DISABLE	Y	N/A	Y	Y	Y	Y
SETTING	ALERT	REMOVMSG	ENABLE/ DISABLE	N/A	Y	N/A	Y	Y	Y
SETTING	CACHE	COPYFSC	ENABLE/ DISABLE	N/A	Y	Y	N/A	Y	Y
SETTING	CACHE	RECLPG0	ENABLE/ DISABLE	N/A	Y	Y	N/A	Y	Y
SETTING	CACHE	PMPRIOR	Value	N/A	Y	Y	N/A	Y	N/A

SETTING	CACHE	PMTHLVL	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	CACHE	REMOVE	ENABLE/ DISABLE	N/A	Y	N/A	Y	Y	Y
SETTING	CACHE	REMOVTHR	Value	N/A	Y	N/A	Y	Y	Y
SETTING	CACHE	CPYPRIOR	ENABLE/ DISABLE	N/A	Y	Y	N/A	Y	N/A
SETTING	CACHE	CPYPRITH	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	CACHE	RBPRIOR	Value	N/A	Y	Y	Y	Y	Y
SETTING	CACHE	RBTHLVL	Value	N/A	Y	Y	Y	Y	Y
SETTING	THROTTLE	COPYFT	ENABLE/ DISABLE	N/A	Y	Y	Y	Y	Y
SETTING	THROTTLE	ICOPYT	ENABLE/ DISABLE	N/A	Y	Y	Y	Y	Y
SETTING	THROTTLE	DCOPYT	Value	N/A	Y	Y	Y	Y	Y
SETTING	THROTTLE	DCTAVGTD	Value	N/A	Y	Y	Y	Y	Y
SETTING	THROTTLE	DCTCURTD	Value	N/A	Y	Y	Y	Y	Y
SETTING	RECLAIM	RCLMMAX	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	RECLAIM	RCLMSMAX	Value	N/A	Y	Y	N/A	Y	N/A
SETTING	DEVALLOC	SCRATCH	ENABLE/ DISABLE	Y	N/A	Y	Y	Y	Y
SETTING	DEVALLOC	PRIVATE	ENABLE/ DISABLE	Y	N/A	Y	Y	Y	Y
SETTING	DEVALLOC	FAMILY	RESPECT /IGNORE	N/A	Y	Y	Y	Y	Y
SETTING	CPYCNT	RUN	Value	N/A	Y	Y	Y	Y	Y
SETTING	CPYCNT	DEF	Value	N/A	Y	Y	Y	Y	Y
SETTING	COPY	IMMSNS	ALL/UNE XP/NONE	N/A	Y	Y	Y	Y	Y
SETTING	COPY	TIMEOUT	Value	N/A	Y	Y	Y	Y	Y
SETTING	COPY	SCRATCH	ALWAYS / NEVER/ NONTDL Y	Y	N/A	Y	Y	Y	Y
SETTING	LINK	FAILOVER	ENABLE/ DISABLE	N/A	Y	Y	Y	Y	Y
SETTING	DELEXP	COUNT	Value	Y	N/A	Y	Y	Y	Y
SETTING	EXISTDEL	CRITERIA	STALE/A LL/NONE	N/A	Y	Y	Y	Y	Y
SETTING	EXISTDEL	WHEN	ATCLOSE /AUTO	N/A	Y	Y	Y	Y	Y
SETTING	EXISTDEL	AUTINTVL	Value	N/A	Y	N/A (*1)	Y	Y	Y

SETTING	EXISTDEL	AUTCOUNT	Value	N/A	Y	N/A (*1)	Y	Y	Y
---------	----------	----------	-------	-----	---	-------------	---	---	---

(*1) This command is not supported on TS7740 since it is new to R5.0 and R5.0 is not supported on TS7740.

When the request has the first keyword “SETTING” only, the response just provides the current settings on the target distributed library specified in the target library name. If the composite library is specified in the target library name, the returned setting values depend on the distributed library which receives the request then the unexpected setting values could be returned because it may vary which distributed library receives the request. The distributed (not composite) library name must be always specified when it is required to get the setting values from the specific distributed library.

If the distributed library specified is for a TS7700D, and the request is not valid in a TS7700D based on the table above the following error text is returned (See 3.1.18 Error Responses):

```
'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

If the distributed library specified is for a TS7740, and the request is not valid with a TS7740 based on the table above the following error text is returned (See 3.1.18 Error Responses):

```
'NOT SUPPORTED IN A NON DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
```

If the distributed library specified is for a cluster in a non TS7700T configuration, and the request is not valid in a non TS7700T configuration based on the table above, the following error text is returned (See 3.1.18 Error Responses):

```
'ONLY SUPPORTED IN A TS7700 TAPE ATTACHED VIRTUALIZATION ENGINE'
```

3.1.12.1 Alert Settings

If a second keyword of ALERT is specified, the cluster will set different thresholds at which a message is sent to all hosts attached to the cluster and in a Grid configuration, all hosts attached to all clusters. The third keyword specifies which alert threshold is to be set and the fourth the threshold value. All messages indicate the distributed library and will result in the following z/OS host console message:

CBR3750I Message from library *distributed library name*: Message Text

Thresholds can be set for many of the resources managed by the cluster. For each resource, two settings are provided, one that is warning that the resource is approaching a value that may result in an impact to the operations of the attached hosts. A second is to provide a warning that the resource has exceeded a value that may result in an impact to the operations of the attached hosts. When the second warning is reached, the warning message is repeated every 15 minutes.

The message text includes a message identifier that may be used to automate the capture and routing of these messages. To assist in routing messages to appropriate individuals for handling, the messages that are indicating that a resource is approaching an impact value will use message identifiers in a range of AL0000-AL4999. Message identifiers in a range of AL5000-AL9999 will be used for messages that are indicating that a resource has exceeded an impact value.

Note: For the messages where a variable is included (the setting), the value returned is left-justified without leading zeroes or right padding. For example: AL5000 Uncopied data of 1450 GB above high warning limit of 1050 GB.”

The following ALERT thresholds are supported:

Keyword 3	Keyword 4	Description
COPYHIGH	Value	Uncopied Data High Warning Limit This is the threshold, in GBs of data in cache that needs to be copied to other TS7700s in a Grid configuration, at which the TS7700 will generate a message indicating that the amount of uncopied data has exceeded a high warning limit. For this message to be generated, the amount of uncopied data must be above the value specified in keyword 4 for 150 seconds. As long as the amount of uncopied data remains above the threshold, the above warning limit message is repeated every 15 minutes. If the message had been generated and the amount of uncopied data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of uncopied data is below the high warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than the enabled cache size less 500GB cannot be set. If a value is specified that is not 500GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500GB. A value less than 100GB cannot be set. If a value is specified that is not 100GB or greater, the threshold is set to 100GB. If the uncopied data low warning limit is not zero and a value is specified that is not 100GB greater than the uncopied data low warning limit, the uncopied data low warning limit will be changed so that it is 100GB less than the value specified (but never lower than 0). Message Text (xxxxxxx is the amount of uncopied data, yyyyyyyy is the threshold): When above the threshold: “AL5000 Uncopied data of xxxxxxxx GB above high warning limit of yyyyyyyy GB.” When below the threshold: “AL5001 No longer above uncopied data high warning limit of yyyyyyyy GB.”
COPYLOW	Value	Uncopied Data Low Warning Limit This is the threshold, in GBs of data in cache that needs to be copied to other TS7700s in a Grid configuration, at which the TS7700 will generate a message indicating that the amount of uncopied data has exceeded a low warning limit. For this message to be generated, the amount of uncopied data must be above the value specified in keyword 4 for 150 seconds. If the message had been generated and the amount of uncopied data is at or falls below

Keyword 3	Keyword 4	Description
		that threshold for 150 seconds, a message is generated indicating that the amount of uncopied data is below the low warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. If the uncopied data high warning limit is set to 0, then the uncopied data low warning limit cannot be set to a value greater than the enabled cache size less 500GB or to a value less than 100GB. If a value is specified that is not 500GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500GB. If a value is specified that is less than 100GB, the threshold is set to 100GB. If the uncopied data high warning limit is not zero, the uncopied data low warning limit cannot be set to a value greater than the uncopied data high warning limit less 100GB. If a value is specified that is not 100GB lower than the uncopied data high warning limit, the threshold is set to the uncopied data high warning limit less 100GB (but never lower than 0). Message Text (xxxxxxx is the amount of uncopied data, yyyyyyy is the threshold): When above the threshold: “AL0000 Uncopied data of xxxxxxxx GB above low warning limit of yyyyyyy GB.” When below the threshold: “AL0001 No longer above uncopied data low warning limit of yyyyyyy GB.”
PDRVCRIT	Value	Available Physical Drive Critical Warning Limit This is the threshold, in number of physical drives, at which the TS7700 will generate a message indicating that the number of available physical drives has fallen below the critical warning limit. For this message to be generated, the number of available physical drives must be below the value specified in keyword 4 for 15 minutes. As long as the number of available physical drives is below the threshold, the below the threshold message is repeated every 15 minutes. If the message had been generated and the number of available physical drives is at or has risen above that threshold for 15 minutes, a message is generated indicating that the number of available physical drives is above the critical warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than the number of installed physical drives minus 1 cannot be set. If a value is specified that is greater than the

Keyword 3	Keyword 4	Description
		number of installed physical drives minus 1, the threshold is set to the number of installed physical drives minus 1. A value less than 3 cannot be set. If a value is specified that is less than 3, the threshold is set to 3. If the available physical drive low warning limit is not zero and a value is specified that is not 1 less than the available physical drive low warning limit, the available physical drive low warning limit will be changed so that it is 1 more than the value specified (but never more than the number of installed physical drives). Message Text (xx is the number of available drives, yy is the threshold): When fallen below the threshold: “AL5004 Available physical drives xx below critical limit of yy.” When risen above the threshold: “AL5005 Available physical drives no longer below critical limit of yy.” This setting only applied to the active read-write drives. If the cluster has the sunset drives for the media compatibility, it is required to configure SETTING2, PDRVSCRT to set the alert to the sunset drives too.
PDRVLOW	Value	Available Physical Drive Low Warning Limit This is the threshold, in number of physical drives, at which the TS7700 will generate a message indicating that the number of available physical drives has fallen below the low warning limit. For this message to be generated, the number of available physical drives must be below the value specified in keyword 4 for 15 minutes. If the message had been generated and the number of available physical drives is at or has risen above that threshold for 15 minutes, a message is generated indicating that the number of available physical drives is above the low warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. If the available physical drive critical warning limit is set to 0, then the available physical drive low warning limit cannot be set to a value greater than the number of installed physical drives or less than 3. If a value is specified that is greater than the number of installed physical drives, the threshold is set to the number of installed physical drives. If a value is specified that is less than 3, the threshold is set to 3. If the available physical drive critical warning limit is not zero and a value is specified that is not 1 greater than the available physical

Keyword 3	Keyword 4	Description
		drive critical warning limit, the available physical drive low warning limit will be changed so that it is 1 more than the available physical drive critical warning limit (but never greater than the number of installed physical drives). Message Text (xx is the number of available drives, yy is the threshold): When fallen below the threshold: “AL0004 Available physical drives xx below low limit of yy.” When risen above the threshold: “AL0005 Available physical drives no longer below low limit of yy.” This setting only applied to the active read-write drives. If the cluster has the sunset drives for the media compatibility, it is required to configure SETTING2, PDRVSLOW to set the alert to the sunset drives too.
PSCRCRIT	Value	Physical Scratch Volume Critical Warning Limit This is the threshold, in number of scratch physical volumes, at which the TS7700 will generate a message indicating that the number of available scratch physical volumes has fallen below the critical warning limit. For this message to be generated, the number of available physical scratch volumes for one of the active general pools (pools 1-32) must be below the value specified in keyword 4 for 16 minutes. Available volumes include those in pool 0 if borrowing is allowed for the pool. All media types allowed for borrowing are considered. As long as the number of scratch physical volumes is below the threshold, the below the threshold message is repeated every 16 minutes. If the message had been generated and the number of available physical scratch volumes for the pool is at or has risen above that threshold for 16 minutes, a message is generated indicating that the number of available physical scratch volumes is above the critical warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than 190 cannot be set. If a value is specified that is greater than 190, the threshold is set to 190. A value less than 5 cannot be set. If a value is specified that is less than 5, the threshold is set to 5. If the physical scratch volume low warning limit is not zero and a value is specified that is not 10 less than the physical scratch volume low warning limit, the physical scratch volume low

Keyword 3	Keyword 4	Description
		warning limit will be changed so that it is 10 more than the value specified (but never greater than 200). Message Text (xxx is the number of available physical scratch volumes, yyy is the threshold, zz is the pool number): When fallen below the threshold: “AL5006 Available physical scratch volumes of xxx below critical limit of yyy for pool zz.” When risen above the threshold: “AL5007 Available physical scratch volumes no longer below critical limit of yyy for pool zz.” Note: The TS7700 will enter panic reclaim if the number of scratch volumes available to a defined pool is less than 2, including ones that it could borrow from pool 0.
PSCRLOW	Value	Physical Scratch Volume Low Warning Limit This is the threshold, in number of scratch physical volumes, at which the TS7700 will generate a message indicating that the number of available scratch physical volumes has fallen below the low warning limit. For this message to be generated, the number of available physical scratch volumes for one of the active general pools (pools 1-32) must be below the value specified in keyword 4 for 16 minutes. Available volumes include those in pool 0 if borrowing is allowed for the pool. All media types allowed for borrowing are considered. If the message had been generated and the number of available physical scratch volumes for the pool and media type is at or has risen above that threshold for 16 minutes, a message is generated indicating that the number of available physical scratch volumes is above the low warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. If the physical scratch volume critical warning limit is set to 0, then the physical scratch volume low warning limit cannot be set to a value greater than 200 or less than 5. If a value is specified that is greater than 200, the threshold is set to 200. If a value is specified that is less than 5, the threshold is set to 5. If the physical scratch volume critical warning limit is not zero and a value is specified that is not 10 greater than the physical scratch volume critical warning limit, the physical scratch volume low warning limit will be changed so that it is 10 more than the physical scratch volume critical warning limit (but never greater than 200). Message Text (xxx is the number of available physical scratch volumes, yyy is the threshold, zz is the pool number):

Keyword 3	Keyword 4	Description
		When fallen below the threshold: “AL0006 Available physical scratch volumes of xxx below low limit of yyy for pool zz.” When risen above the threshold: “AL0007 Available physical scratch volumes no longer below low limit of yyy for pool zz.” Note: The TS7700 will enter panic reclaim if the number of scratch volumes available to a defined pool is less than 2, including ones that it could borrow from pool 0.
RESHIGH	Value	Resident Data High Warning Limit This is the threshold, in GBs of resident data, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a high warning limit. For this message to be generated, the amount of resident data must be above the value specified in keyword 4 for 150 seconds. As long as the amount of resident data remains above the threshold, the above warning limit message is repeated every 15 minutes. If the message had been generated and the amount of resident data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of resident data is below the high warning limit. This is not to say that this alert will be present 150 seconds after either exceeding this limit or falling below this limit. This alert will be presented after at least 150 seconds. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than the enabled cache size less 500GB cannot be set. If a value is specified that is not 500GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500GB. A value less than 100GB cannot be set. If a value is specified that is not 100GB or greater, the threshold is set to 100GB. If the resident data low warning limit is not zero and a value is specified that is not 100GB greater than the resident data low warning limit, the resident data low warning limit will be changed so that it is 100GB less than the value specified (but never lower than 0). Message Text (xxxxxxx is the amount of resident data, yyyyyyy is the threshold): When above the threshold: “AL5008 Resident data of xxxxxx GB above high warning limit of yyyyyyy GB.”

Keyword 3	Keyword 4	Description
		When below the threshold: “AL5009 No longer above resident data high warning limit of xxxxxxxx GB.” Note: For a TS7740, resident data is the data that has not been copied to back end physical volumes. For a TS7700D, resident data is the data that occupies the cache. Note: For a TS7700T/7700C, this alert is applied to the resident data in the resident-only cache partition (cache partition 0).
RESLOW	Value	Resident Data Low Warning Limit This is the threshold, in GBs of resident data, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a low warning limit. For this message to be generated, the amount of resident data must be above the value specified in keyword 4 for 150 seconds. If the message had been generated and the amount of resident data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of resident data is below the low warning limit. This is not to say that this alert will be present 150 seconds after either exceeding this limit or falling below this limit. This alert will be presented after at least 150 seconds. The default value is 0, which indicates no warning limit is set and no messages are to be generated. If the resident data high warning limit is set to 0, then the resident data low warning limit cannot be set to a value greater than the enabled cache size less 500GB or to a value less than 100GB. If a value is specified that is not 500GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500GB. If a value is specified that is less than 100GB, the threshold is set to 100GB. If the resident data high warning limit is not zero, the resident data low warning limit cannot be set to a value greater than the resident data high warning limit less 100GB. If a value is specified that is not 100GB lower than the resident data high warning limit, the threshold is set to the resident data high warning limit less 100GB (but never lower than 0). Message Text (xxxxxxx is the amount of resident data, xxxxxxxx is the threshold): When above the threshold: “AL0008 Resident data of xxxxxxxx GB above low warning limit of xxxxxxxx GB.” When below the threshold:


Keyword 3	Keyword 4	Description
		“AL0009 No longer above resident data low warning limit of yyyyyyyy GB.” Note: For a TS7740, resident data is the data that has not been copied to back end physical volumes. For a TS7700D, resident data is the data that occupies the cache. Note: For a TS7700T/7700C, this alert is applied to the resident data in the resident-only cache partition (cache partition 0).
RSDTHIGH	Value	Resident Data High Warning Limit for TS7700T/7700C This is the same threshold with RESDHIGH, but only applicable to the total TS7700T tape/cloud attached cache partitions (cache partitions from 1 to 7). Message Text (xxxxxxx is the amount of resident data, yyyyyyyy is the threshold): When above the threshold: “AL5015 Sum of resident data in tape partitions of xxxxxxxx GB above high warning limit of yyyyyyyy GB.” When below the threshold: “AL5016 Sum of resident data in tape partitions no longer above resident data high warning limit of yyyyyyyy GB.” This is only applicable to TS7700T/7700C. To set the same alert to TS7700T/7700C resident-only partition (cache partition 0), RESDHIGH needs to be used. Note: Resident data is the data that has not been copied to back end physical volumes in the tape attached partitions. It does not include the data that Time Delayed Premigration is configured and Premigration Delay Time has not passed yet.
RSDTLOW	Value	Resident Data Low Warning Limit for TS7700T/7700C This is the same threshold with RESDLOW, but only applicable to the total TS7700T tape/cloud attached cache partitions (cache partitions from 1 to 7). Message Text (xxxxxxx is the amount of resident data, yyyyyyyy is the threshold): When above the threshold: “AL0012 Sum of resident data in tape partitions of xxxxxxxx GB above low warning limit of yyyyyyyy GB.” When below the threshold:

Keyword 3	Keyword 4	Description
		“AL0013 Sum of resident data in tape partition no longer above low warning limit of yyyyyyyy GB.” This is only applicable to TS7700T/7700C. To set the same alert to TS7700T/7700C resident-only partition (cache partition 0), RESDLOW needs to be used. Note: Resident data is the data that has not been copied to back end physical volumes in the tape attached partitions. It does not include the data that Time Delayed Premigration is configured and Premigration Delay Time has not passed yet.

Keyword 3	Keyword 4	Description
PCPYLOW	Value	Pending Copy Low This is the threshold in GB of volumes in the copy queue, following are some limitations for this alarm to be set The difference between the crit value (described below) and the low value must be equal or bigger than 500. The lower value for this field must be equal or bigger than 500GB. When adjusting the lower / high value, if the difference is lower than 500, the lower value is decreased automatically to accommodate the difference. Message Text (xx is the distributed library name, yyyy is total pending copy GB and zzzz is the threshold of the alert : “AL0011 Distributed Library xx has successfully fallen below the inbound copy backlog low warning limit of zzzz GB” When risen above the threshold: “AL0010 Distributed Library xx has a total pending inbound copy backlog of yyyy GB which is above the low warning limit of zzzz GB”
PCPYCRIT	Value	Pending Copy Crit This is the upper limit in GB for volumes in the copy queue. The difference between the crit value (described above) and the low value must be equal or bigger than 500. When adjusting the lower / high value, if the difference is lower than 500, the lower val is decreased automatically to accommodate the difference. Message Text (xx is the distributed library name, yyyy is total pending copy GB and zzzz is the threshold of the alert : When fallen below the threshold: “AL5011 Distributed Library xx has successfully fallen below the inbound copy backlog critical warning limit of zzzz GB.” When risen above the threshold: “AL5010 Distributed Library xx has a total pending inbound copy backlog of yyyy GB which is above the critical warning limit of zzzz GB”
DEFDEG	ENABLE, DISABLE	May be used to prevent a composite library from entering the ‘degraded’ state when the ‘SyncDeferred’ or the ‘ImmediateDeferred’ condition occurs. When the ENABLE keyword is specified, the ‘degraded’ state is reported to the host through the operational state change attention.

Keyword 3	Keyword 4	Description
		When the DISABLE keyword is specified, the 'degraded' state is NOT reported to the host through the operational state change attention. The default is enabled. This command is acceptable only when all the clusters in the grid are at the code level 8.30.xx.xx or above. The corresponding operator messages below are still generated regardless of the DEFDEG setting: "G0005 Distributed library xx has entered the immediate deferred state." "G0032 Distributed Library xx has entered the synchronous deferred state due to volser yyyyyy."
LINKDEG	ENABLE, DISABLE	May be used to prevent a composite library from entering the 'link degraded' state when Grid link degradation occurs. When the ENABLE keyword is specified, the 'degraded' state is reported to the host through the operational state change attention. When the DISABLE keyword is specified, the 'degraded' state is NOT reported to the host through the operational state change attention. The default is enabled. This command is acceptable only when all clusters in the domain are at the code level 8.31.xx.xx or above. The corresponding operator messages below are still generated regardless of the LINKDEG setting: "G0030 Library XXXXX, PPP, AAA Grid Link are degraded." "G0031All grid links for this cluster have left the degraded state"

Keyword 3	Keyword 4	Description
REMOVMSG	ENABLE, DISABLE	May be used to prevent a distributed library (except TS7740) from reporting Automatic Removal start and stop events to the MI and operator messages to the host when Auto Removal occurs. When the ENABLE keyword is specified, an event will be posted to the TS7700 MI when automatic removal of logical volumes starts and another event will be posted when the TS7700 stops automatic removal of logical volumes. An associated host message will be sent for each event as well. When the DISABLE keyword is specified, no event associated with the logical volume removals will be posted and no host message will be sent. The default is enabled. This command is acceptable only on TS7700D with the code level 8.30.x.x or 8.21.0.144 (R2.1 PGA4) and also TS7700T/7700C. Note: If this request is sent to pre-8.32.x.x cluster and the command target cluster is TS7700T, the command fails. It is because pre-8.32.x.x cluster can see the remote TS7700T as TS7740 as it has a physical tape library and this command is rejected by the cluster. In that case, this command needs to be issued to the cluster with the code level 8.32.x.x, or MI panel on the cluster with the code level 8.32.x.x can be used.



Alert Setting Diagram

3.1.12.2 Cache Settings

If a second keyword of CACHE is specified, the cluster will modify how it controls the workflow and content of the tape volume cache.

The following CACHE settings are supported:

Keyword 3	Keyword 4	Description
COPYFSC	ENABLE, DISABLE	Copies To Follow Storage Class Preference When the ENABLE keyword is specified, logical volumes copied into this cluster's disk cache from a peer TS7700 are managed using the preference group actions defined within this target cluster's Storage Class construct definition. When the DISABLE keyword is specified, logical volumes copied into this cluster's disk cache from a peer TS7700 are always managed as PG0 (prefer to be migrated from cache). This setting only applies to preference group (e.g. PG0, PG1). All other behaviors applied to copied logical volumes utilize the Storage Class construct settings of this target cluster. The default is disabled.
PMPRIOR	Value	Premigration Priority Threshold This is the threshold, in GBs of unpremigrated data, at which the TS7700 will begin increasing the number of premigration tasks that will be allowed to compete with host I/O for cache and processor cycles. The amount of unpremigrated data must be above the value specified in keyword 4 for 150 seconds before the additional premigration tasks are added. As the amount of data to premigrate continues to grow above this threshold setting, so do the number of enabled premigration tasks until the maximum is reached. If the amount of unpremigrated data subsequently falls below this threshold for at least 150 seconds, the number of premigration tasks may be reduced depending on host I/O demand. If I/O host demand is high, the number of premigration tasks will eventually be reduced to a minimum of one. The default value is 1600. If a value is specified that is higher than the premigration throttling threshold value, it is set to the premigration throttling threshold value. Note: If TS7700T has only one FC 5274 (i.e. maximum amount of queued premigration content is 1TB), an internal PMPRIOR of 600 is used even though the SETTING command will continue to show 1600 (default value). Only when the PMPRIOR value is explicitly set, will it show a value other than the default of 1600.
PMTHLVL	Value	Premigration Throttling Threshold

Keyword 3	Keyword 4	Description
		This setting is applicable to TS7700T only. The equivalent setting for TS7700C is fixed at a value of the size of the premigration queue. This is the threshold, in GBs of unpremiered data, at which the TS7700 will begin introducing a delay in responding to host write operations on all virtual tape device addresses of the TS7700. The amount of unpremiered data must be above the value specified in keyword 4 for 150 seconds before the delay is introduced. The amount of delay begins with a few milliseconds and increases to 2000 milliseconds as the amount of unpremiered data approaches the total cache capacity. The amount of unpremiered data must fall below this threshold for 150 seconds for all delay to be removed. The default value is 2000. A value greater than the enabled cache size less 500GB cannot be set. If a value is specified that is not 500GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500GB. For the TS7700T, a value greater than the total size of the active premigration queue cannot be set. If a value is specified that is not lower than that, the threshold is set to the total amount of the active premigration queue. The size of the active premigration queue is set to total amount of FC 5274 in R3.2 GA (8.32.0.84) and (total amount of FC 5274 + 5TB * total amount of FC 5279) in R4.2 GA (8.42.0.71) or above. Note: If TS7700T has only one FC 5274 (i.e. maximum amount of queued premigration content is 1TB), an internal PMTHLVL of 1000 is used even though the SETTING command will continue to show 2000 (default value). Only when the PMTHLVL value is explicitly set, will it show a value other than the default of 2000.
RECLPG0	ENABLE, DISABLE	Recalls Preferred to be Removed from Cache When the ENABLE keyword is specified, logical volumes that are recalled into cache are managed as PG0 (prefer to be removed from cache). This overrides the actions defined for the Storage Class associated with the recalled volume. When the DISABLE keyword is specified, logical volume that are recalled into cache are managed using the actions defined for the Storage Class construct associated with the volume as defined at the TS7700. The default is disabled.
REMOVE	ENABLE,	Enable or disable automatic removal when the cache usage size crosses the removal threshold

Keyword 3	Keyword 4	Description
	DISABLE	When the ENABLE keyword is specified, the automatic removal will be enabled on this disk-only cluster. When the DISABLE keyword is specified, the automatic removal will be disabled on this disk-only cluster. Default value is enabled. Note: Only TS7700D, TS7700T and TS7700C clusters support this command. But if this request is sent to pre-8.32.x.x cluster and the command target cluster is TS7700T, the command fails. It is because pre-8.32.x.x cluster can see the remote TS7700T as TS7740 as it has a physical tape library and this command is rejected by the cluster. In that case, this command needs to be issued to the cluster with the code level 8.32.x.x, or MI panel on the cluster with the code level 8.32.x.x can be used.
REMOVTHR	Value	Automatic removal starts when the cache usage size crosses the removal threshold When automatic removal is enabled on this disk-only cluster (TS7700D) or a tape/cloud-attached cluster (TS7700T/7700C), logical volume removal will start when the free cache space is below the removal threshold plus 1000GB (1000GB is the out-of-cache warning threshold). For a TS7700T/7700C, the free space is relative to the resident only cache partition (CP0). The automatic removal state will exit when the free space goes above this threshold plus 500GB. Event message OP0164 will surface when the removal begins and AL5012 will surface when the cluster exits the removal state. The default value is 3000 (in GB). A value of 2000 to 10000 can be set for all TS7700D, TS7700T and TS7700C model clusters. As of 8.40.x.x, TS7700D models support a maximum value equal to the total size of the disk cache minus 2000GB. For example, if the configured disk cache is 400TB in size, the maximum value for the removal threshold on a TS7720D with 8.40.x.x or later code installed would be 398000. If a value outside the supported range is provided, it will be automatically set to the minimum or maximum supported value depending on whether the provided value was too small or too large. Note: Only TS7700D, TS7700T and TS7700C clusters support this command. TS7700T clusters are only supported with code levels 8.32.x.x or higher. TS7720T clusters that exist within a mixed code level grid may result in command failures when the LI REQ command is issued through a cluster running a code level prior to 8.32.x.x. If this occurs, the command must be issued through a host

Keyword 3	Keyword 4	Description		
		attached cluster running 8.32.x.x or later or via the TS7720T's Management Interface LI REQ GUI panel.		
CPYPRIOR	ENABLE, DISABLE	Limit premigration resources under Grid copy activity When the ENABLE keyword is specified, the number of premigration tasks will start decreasing when Grid copy activity crosses the threshold defined by CPYPRITH command. When the DISABLE keyword is specified, the number of premigration tasks will not decrease. Default value is ENABLE. Note: the default value of CPYPRIOR in pre-8.41.x.x cluster was DISABLE. However, new default value ENABLE is applied in 8.41.x.x or later even if the cluster is upgrade from pre-8.41.x.x.		
CPYPRITH	Value	Premigration resources start being limited when Grid copy activity crosses the threshold When CPYPRIOR is enabled, the number of premigration tasks will start decreasing when the total Grid copy activity (the data transfer rate to read/write the data from/into the local cluster's cache) is above the threshold. The default value is 100 (in MB/sec). A value of 10 to 1000 can be set. If a value less than 10 or more than 1000 is given, an error is returned and the value is not changed. PMPRIOR overrides CPYPRITH threshold when both conditions are satisfied. For example, - When unpremigrated data has already crossed PMPRIOR threshold, the number of premigrated tasks is not reduced although it then crosses CPYPRITH threshold. - When unpremigrated data has already crossed CPYPRITH threshold, the number of premigration tasks is till increased if it then crosses PMPRIOR threshold.		
RBPRIOR	Value	Cache Rebuild Priority After the replacement of a failed disk drive, the cache controller must rebuild the array. The cache rebuild and copyback affects the cache performance. There are five cache rebuild priority levels, 1 - 5. A value of 1 is the lowest, which means the cache controller will rebuild the cache array at the slowest pace but will have the least impact on cache performance. A value of 5 is the highest rebuild priority but will have the most impact to cache performance. The RBPRIOR value is a setting that instructs how the TS7700 will manage the cache rebuild, if one is in progress. A value of 0 - 5 may be set. This value is defined as the following: <table><tr><td>Setting</td><td>Cache Controller's Rebuild Priority Level</td></tr></table>	Setting	Cache Controller's Rebuild Priority Level
Setting	Cache Controller's Rebuild Priority Level			

Keyword 3	Keyword 4	Description												
		<table><tr><td>0</td><td>The TS7700 will dynamically manage the cache controller's rebuild priority level based on the TS7700's resource usage, setting the cache controller's rebuild priority to either 1 or 4 for the CSA cache model and either 1 or 5 for the CSB cache model.</td></tr><tr><td>1</td><td>Lowest priority level</td></tr><tr><td>2</td><td>Low priority level</td></tr><tr><td>3</td><td>Medium priority level</td></tr><tr><td>4</td><td>High priority level</td></tr><tr><td>5</td><td>Highest priority level</td></tr></table> If a value less than 0 or more than 5 is given, an error is returned and the value is not changed.	0	The TS7700 will dynamically manage the cache controller's rebuild priority level based on the TS7700's resource usage, setting the cache controller's rebuild priority to either 1 or 4 for the CSA cache model and either 1 or 5 for the CSB cache model.	1	Lowest priority level	2	Low priority level	3	Medium priority level	4	High priority level	5	Highest priority level
0	The TS7700 will dynamically manage the cache controller's rebuild priority level based on the TS7700's resource usage, setting the cache controller's rebuild priority to either 1 or 4 for the CSA cache model and either 1 or 5 for the CSB cache model.													
1	Lowest priority level													
2	Low priority level													
3	Medium priority level													
4	High priority level													
5	Highest priority level													
RBTHLVL	Value	Cache Rebuild Priority Threshold The RBTHLVL value is only applicable when dynamic cache rebuild prioritization is used (RBPRIOR is set to 0). The RBTHLVL keyword represents the upper average disk I/O threshold in which the TS7700 will allow the cache rebuild and copyback to run at the HIGH level. The average disk I/O includes disk reads and writes from host, replication, premigration and recall activity, and this rate is in megabyte per second (MB/s). The average disk I/O rate is an average of the disk I/O rate over a twenty-minute period, computed using samples taken every two minutes. At R5.0 PGA1, R5.1 or later code level, default value is 3500 MB/sec. At R5.0 GA, default value was 2500 MB/sec. Prior to R5.0, default value was 200 MB/sec. If R5.0 GA or earlier code level machine who has default RBTHLVL setting (2500 or 200) is upgraded to R5.0 PGA1, R5.1 or later code level, new default value 3500 is set to RBTHLVL automatically. A value of 100 to 5000 can be set. If a value less than 100 or more than 5000 is given, the default value (2500) is automatically set. The value of 2500 MB/s is set by default to favor host, copy, and premigration disk I/O over a faster cache rebuild or copyback. In general, when there is disk activity from host, copy, and/or premigration where the average disk I/O rate is in the range of 100 MB/s to RBTHLVL, the TS7700 will favor the host, copy, and/or premigration over a faster cache rebuild and copyback. As an example, if RBPRIOR is set to 0 and RBTHLVL is set to 3500 MB/s, when the average disk I/O rate is equal to or greater than 100MB/s and below the RBTHLVL value of 3500 MB/s, the TS7700 will set the cache rebuild and copyback to run at the LOWEST level to allow resources to the host, copy, and premigration activity. If the average disk I/O rate is less than 100MB/s, the TS7700 will set the cache rebuild and copyback to run at the HIGH level since there is low disk activity from host, copy, and premigration activity. If the average disk I/O rate is greater than or equal to the RBTHLVL value of 3500												

Keyword 3	Keyword 4	Description
		MB/s, then the TS7700 will set the cache rebuild and copyback to the HIGH level because the average disk I/O rate meets or exceeds the upper average disk I/O threshold.

3.1.12.3 Throttle Settings

If a second keyword of THROTTLE is specified, the cluster will modify how it controls the different data flow rates into and out of the cluster.

The following THROTTLE settings are supported:

Keyword 3	Keyword 4	Description
COPYFT	ENABLE, DISABLE	Full Cache Copy Throttling When the ENABLE keyword is specified, throttling when the cache is full of uncopied data to other TS7700s is enabled. Meaning, if the TS7700 cache (for TS7700T, this includes all cache partitions combined) is full of data waiting to be copied out to other clusters, host writes and inbound copies will be throttled to allow priority for outbound copies. When the DISABLE keyword is specified, throttle when the cache is full of uncopied data to other TS7700s is disabled. TS7700T/7700C always works as if the setting is set to ENABLE regardless of the actual setting. On TS7700D, Full Cache Copy Throttling actually never works because it always becomes “out of cache” state prior to reaching Full Cache Copy Throttling threshold then this setting is not effective. Therefore, this setting is actually applicable to TS7740 only. The default is enabled. Note: Full Cache Copy Throttling is applied when COPYFT is enabled, amount of data to be copied out to other clusters exceeds (cache size – 200GB) and other TS7700 clusters have been up more than 24 hours.
DCOPYT	Value	Deferred Copy Throttle This is a delay, in milliseconds, that the TS7700 will introduce for every 32KB copy transfer for all deferred mode copy tasks in which it is the source. Refer to the description for the Deferred Copy Throttling Average Threshold (DCTAVGTD) for an explanation as to how deferred copy throttling is applied. A value of 0 to 250 milliseconds can be set. The default value is 125 milliseconds.
ICOPYT	ENABLE, DISABLE	Immediate Copy Throttling When the ENABLE keyword is specified, immediate copy throttling is enabled. The depth of the immediate copy queue is examined as well as the amount of time copies have been in the queue to determine if the throttle should be applied. When the DISABLE keyword is specified, immediate copy throttling is disabled. The default is enabled.

Keyword 3	Keyword 4	Description
DCTAVGTD	Value	Deferred Copy Throttling Average Threshold This value represents the average compressed Host I/O rate over a twenty minute period. The average compressed host I/O rate and other system utilization factors described below are used to determine when deferred copy throttling is applied. This type of throttling will be activated when any of the three cluster busy conditions are true and either of the DCTAVGTD or DCTCURTD thresholds have been exceeded. • Busy Condition 1: CPU percent busy is greater than 85% • Busy Condition 2: Disk percent busy is greater than 85% • Busy Condition 3: There is more than 150MB/s of system disk write activity, which includes local host writes, remote writes into this cluster, inbound copy activity or recalls from back end tape. DCTAVGTD is compared to the average host read+write rate over the last 20 minutes, including remote writes into this cluster. DCTCURTD is compared to the average host read+write rate over the last 30 seconds, including remote writes into this cluster. When deferred copy throttling is applied, the deferred copies are delayed as specified by the DCOPYT value. The default value is 100 in MB/sec. A value of 0 will set the threshold to the default. A value of 1 to 500 can be set.
DCTCURTD	Value	Deferred Copy Throttling Current Threshold This value represents the average compressed Host I/O rate over a thirty seconds period. See description of DCTAVGTD in details. The default value is 100 in MB/sec. A value of 0 will set the threshold to the default. A value of 1 to 500 can be set.

3.1.12.4 Reclaim Settings

If a second keyword of RECLAIM is specified, the cluster will modify how the Reclaim background tasks controls the workflow and content of the tape volume cache.

Also note that if a valid RECLAIM request is received while reclaims are inhibited, that request will take effect as soon as reclaims are no longer inhibited by the Inhibit Reclaim schedule.

The following RECLAIM settings are supported:

Keyword 3	Keyword 4	Description																														
RCLMMAX	Value	Reclaim Maximum Tasks Limit Occasionally, it is desired to have fewer reclaims running and use the service resource for other activities in the cluster. If keyword 3 is specified as RCLMMAX, the cluster can be directed to limit the maximum number of reclaims to a certain value below that which the cluster would normally limit as given by Table 1: <table><tr><th>Number of Available Drives</th><th>Maximum Number of Reclaims</th></tr><tr><td>3</td><td>1</td></tr><tr><td>4</td><td>1</td></tr><tr><td>5</td><td>1</td></tr><tr><td>6</td><td>2</td></tr><tr><td>7</td><td>2</td></tr><tr><td>8</td><td>3</td></tr><tr><td>9</td><td>3</td></tr><tr><td>10</td><td>4</td></tr><tr><td>11</td><td>4</td></tr><tr><td>12</td><td>5</td></tr><tr><td>13</td><td>5</td></tr><tr><td>14</td><td>6</td></tr><tr><td>15</td><td>6</td></tr><tr><td>16</td><td>7</td></tr></table> Table 1 - Cluster Calculated Maximum Number of Reclaims Based on the Number of Available Drives Note that the Inhibit Reclaim schedule as set through the MI has precedence over the limiting of reclaims. That is, if the cluster is in a period of time where reclaims are inhibited, reclaims will not be started even if a valid RCLMMAX has been received. The limiting of reclaims applies when beginning new reclaims. Currently running reclaims are not cancelled or stopped even if the number of reclaims already running is greater than this ‘value’. The minimum valid ‘value’ is 1.	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
Number of Available Drives	Maximum Number of Reclaims																															
3	1																															
4	1																															
5	1																															
6	2																															
7	2																															
8	3																															
9	3																															
10	4																															
11	4																															
12	5																															
13	5																															
14	6																															
15	6																															
16	7																															

		The maximum valid 'value' is one less than the maximum number of reclaims that are calculated by the cluster based on the number of available drives. See Table 1 above. If a value of 0 is given, reclaim limiting is disabled and the cluster sets the maximum number of reclaims to the value calculated from Table 1 based on the number of available drives. Note: The limiting of reclaims will remain in effect until a value of '0' is received for keyword 4. EXAMPLE Assume that: 1. 2 reclaims are currently running before a request is received. 2. A RCLMMAX request is received with the value of 3. 3. There are 13 available drives. The default maximum number of reclaims that can be running for normal reclaim processing on this cluster is 5 (from Table 1 on page 73.) Currently, 2 reclaims are already running and the new maximum is 3. So 1 more reclaim may be started while the original 2 are running. If there were 4 reclaims running when this same example started, then since the new maximum is 3, 2 of the 4 running reclaims would have to finish before the cluster could start a new reclaim.
RCLMSMAX	Value	Reclaim Maximum Tasks Limit For Sunset Media When the resident on recall function is enabled by RRCLSUN command, TS7700 accelerates to change the format of the sunset media to the newest one. This function provides a method to limit the maximum number of concurrent reclamation tasks that run against sunset media using read-only 3592 tape drives in heterogeneous configuration. The default absolute maximum number of concurrent tasks is the lesser of the available sunset drives and non-sunset drives minus one: Maximum number of the concurrent reclaim for the sunset media = (the lesser of "Number of the available non-sunset drives" and "Number of the available Sunset drives") - 1

		The RCLMSMAX value must be between 1 and the absolute maximum. Or, disable any limits by setting the value to 0 which will result in the absolute maximum being used. EXAMPLE - The cluster has 6 non-sunset drives and 4 sunset drives. - All 10 drives are available. - The maximum number of the concurrent reclaim for the sunset media is 3 ((the lesser of 6 and 4 (i.e. 4)) - 1). - RCLMSMAX can set the value from 0 – 3. If it sets 2, the maximum number of the concurrent reclaim for the sunset media is limited to 2. - If RCLMSMAX sets the value 0, it changes the setting back to the default. The default is 0 which means no additional limitation for the maximum number of the concurrent reclaim for the sunset media is set. Note that this command is restricted to heterogeneous drive configuration and does not limit the number of concurrent reclamation tasks with homogeneous drive configuration.
--	--	---

3.1.12.5 Device Allocation (Devalloc) Settings

If a second keyword of DEVALLOC is specified, the grid will modify how Device Allocation Assist requests are handled.

The following DEVALLOC settings are supported:

Keyword 3	Keyword 4	Description
SCRATCH	ENABLE/ DISABLE	Device Allocation Assist for Scratch Volumes If a keyword 4 of ENABLE is specified then a grid at code level 8.20.x.x or greater shall process Device Allocation Assist commands for Scratch volumes. A keyword 4 of DISABLE shall remove the capability to process Device Allocation Assist commands for Scratch volumes. Note that the default behavior is for this feature to be DISABLED.
PRIVATE	ENABLE/ DISABLE	Device Allocation Assist for Private Volumes If a keyword 4 of DISABLE is specified then a grid at code level 8.20.x.x or greater shall inhibit processing Device Allocation Assist commands for Private volumes. A keyword 4 of ENABLE shall allow the capability to process Device Allocation Assist commands for Private volumes. Note that the default behavior is for this feature to be ENABLED.
FAMILY	RESPECT/ IGNORE	Cluster family configuration is taken into account when Device Allocation Assist for Private Mount returns the optimal cluster list If a keyword 4 of RESPECT is specified then Device Allocation Assist commands for Private Mount take into account the cluster family configuration which the distributed library is included in order to determine the optimal cluster order. If a keyword 4 of IGNORE is specified then the cluster family configuration is ignored. Note that the default behavior is for this feature to be RESPECT. Example of the use-case where this option is helpful: - 4 clusters (C0-C3) exist in the grid. C0/C2 in family A and C1/C3 in family B. - The host workload is balanced across the grid. Each family has at least one consistent data. - DAA to C0/C2 always ranks C0/C2 higher than C1/C3 and vice versa because of the family configuration. If a host tends to send DAAs to a specific cluster, the next private mount tends to be directed to the cluster(s) in the same family. Therefore it causes the imbalance of the private mounts.

		- If IGNORE is set, the family configuration is ignored then the private mounts will be distributed across the grid (between the families).
--	--	---

3.1.12.6 Copy Settings

If a second keyword of CPYCNT or COPY is specified, the cluster will modify how it controls the volume copies on the specified cluster.

If a second keyword of CPYCNT is specified, the cluster will modify how many concurrent threads are allowed to process either RUN or Deferred copies.

If a second keyword of COPY is specified, the cluster will modify the reporting method of immediate-deferred state in the CCW (RUN) ERA35 sense data with a third keyword of IMMSNS. And the cluster will modify the volume copy timeout value with a third keyword of TIMEOUT.

In R5.0 (8.50.x.x) or above on a TS7770 with a smaller number of CSB/XSB cache drawers (2 or 4 drawers) and 10Gb interconnect, host write throughput can be impacted with high RUN copy thread counts. This is more likely to occur when minimal latency exists between clusters which allows RUN copies to use a high amount of disk cache resource. If variable host performance is observed, reducing the number of concurrent RUN copy threads from 40 or above to 20 can help to stabilize the throughput by giving more priority to host writes.

The following CPYCNT settings are supported:

Keyword 3	Keyword 4	Description
RUN	Number of Concurrent RUN Copy Threads	The number of concurrent copy threads for processing RUN copies The allowed values for copy thread counts is from 5 to 128. The default value is 20 for clusters with two 1GB Ethernet links and 40 for clusters with four 1GB Ethernet links or two 10GB Ethernet links. "Volume Copy Timeout Time" may need to be increased when the task count is increased.
DEF	Number of Concurrent Deferred Copy Threads	The number of concurrent copy threads for processing Deferred copies The allowed values for copy thread counts is from 5 to 128. The default value is 20 for clusters with two 1GB Ethernet links and 40 for clusters with four 1GB Ethernet links or two 10GB Ethernet links. "Volume Copy Timeout Time" may need to be increased when the task count is increased.

The following COPY settings are supported:

Keyword 3	Keyword 4	Description
IMMSNS	ALL/ UNEXP/ NONE	Immediate-Deferred State Reporting Method This is the control method to report immediate-deferred state in the CCW (RUN) ERA35 sense data. With Release 1.6 code, TS7700 reports all the immediate-deferred state in the CCW (RUN) ERA35 sense data. Since Release 1.7PGA5 (8.7.0.155) or 2.0 codes, the reporting method can be modified. If a keyword 4 of ALL is specified, all the immediate-deferred state is reported in the ERA35 sense data as same as Release 1.6. If a keyword 4 of UNEXP is specified, only the immediate-deferred state induced unexpectedly is reported in the ERA35 sense data. If a keyword 4 of NONE is specified, no immediate-deferred state is reported in the ERA35 sense data except the case where no valid source to copy is available. The default value is NONE. Reporting immediate-deferred state in the CCW (RUN) ERA35 sense data is discussed in the IBM TS7700 Series Best Practice Understanding, Monitoring and Tuning the TS7700 Performance.
TIMEOUT	Value	Volume Copy Timeout Time This is the timeout value, in minutes for logical volume copies between clusters to complete. A volume copy between TS7700s has a timeout value. If a volume copy starts and doesn't finish in the allotted time, it is canceled and returned to the copy queue for retry. Prior to Release 2.0 codes, the timeout value is fixed at 180 minutes. Starting with Release 2.0 codes, the timeout value can be modified. The default value is 180 minutes, however a larger timeout value may be needed when copy tasks increase or limited bandwidth is available. The allowed values for copy timeout are from 30 to 999 minutes. When a volume whose size is larger than 25GB is copied, the timeout value is automatically scaled. For the scaling, the timeout value is multiplied by $(volume\ size / 25GB)$.

Keyword 3	Keyword 4	Description
SCRATCH	ALWAYS/ NEVER/ NONTDLY	Control the replication of logical volumes in the scratch category This is the control method to change the logical volume replication behavior in the grid. Prior to Release 3.1 code, the logical volume is replicated in the grid based off of the copy mode assigned by the management class regardless of the volume category. This setting can be used to suppress the replication if the logical volumes in the grid is in the scratch category. If a keyword 4 of ALWAYS is specified, all logical volumes regardless of the category are replicated in the grid as same as the previous code level. If a keyword 4 of NEVER is specified, no logical volumes in a scratch category are replicated in the grid. If a keyword 4 of NONTDLY is specified, only the logical volumes in a scratch category with Time Delayed copy policy aren't replicated in the grid. The logical volumes with other copy modes are still replicated. The default value is ALWAYS. The request is supported only when all the clusters in the grid have the code level of 8.32.X.X or above. Note: If the logical volume is returned to a scratch volume while its copy job does not yet complete, the copy job will be deleted by the periodical cleanup process which runs once per 12 hours. Note: When the category of the logical volume is changed from a scratch to private while its copy job was deleted, the copy job is brought back automatically. Note: 8.51.1.x introduces the LWORM retention. The scratch volumes in a retained hold state are treated as private and always replicated regardless of this LI REQ setting.

3.1.12.7 Link Failover Settings

If a second keyword of LINK is specified, the cluster will modify how to react in case of a Link Failure during a Remote Mount.

The following LINK settings are supported:

Keyword 3	Keyword 4	Description
FAILOVER	ENABLE/ DISABLE	IP Link Failover for Remote Mounts If a keyword 4 of ENABLE is specified, then a cluster at code level 8.21.x.x or greater shall use the failover capability in case of a Link Failure during a Remote Mount. A keyword 4 of DISABLE shall remove the failover capability. Note that the default behavior is for this feature to be ENABLED.

3.1.12.8 Delexp Count Settings

In response to this request where a composite library is specified, the Delete-Expire setting will be modified as described in the table below

Keyword 3	Keyword 4	Description
COUNT	Value	Delete Expire Count This is the maximum number of the logical volumes that can be expired per hour. The Delete Expire Count can be set to any value from the default (1000) to the maximum (5000) value (the maximum count is changed from 2000 to 5000 at 8.50.x.x). If the Value given is smaller than the default value, it will be set to the default value. If the Value given is greater than the maximum value, it will be set to the maximum value.

3.1.12.9 Existdel Settings

In response to this request where a distributed library is specified, the cluster will modify how to handle the data of 'E' (Exist) copy mode volume.

A volume could have consistent data although it's assigned to 'N' (No Copy) copy mode. It occurs when the volume was previously assigned to a valid copy mode (RUN, Sync, Deferred or Time Delayed), but the copy mode was changed to 'N' copy mode and a private mount for read operation occurs. 'E' (Exist) copy mode is then assigned to the volume after the mount operation to indicate the data still exists although 'N' copy mode is assigned.

A possible scenario is 4-way Grid where C0 data is replicated to C2, and C1 data is replicated to C3. Copy mode is defined as 'DNDN' on C0 and 'NDND' on C1. Data is written from C0 then read mount occurs from C1 without Retain Copy Mode option. The copy mode of the volume is then changed from 'DNDN' to 'EDED' and all four clusters have consistent data.

In Release 8.31, EXISTDEL setting provides the option to delete such a consistent 'E' copy at mount/demount operation. But we found some use cases that 'E' copy isn't remove expectedly with this method. The example is the data move by COPYRFSH: CL1 is added to the existing cluster (CL0) and the data is moved from CL0 to CL1. When the copy mode is changed from "DN" to "ND" and COPYRFSH is used, CL0 copy mode is changed from 'D' to 'E'. But the data on CL0 will not be deleted by this setting until the volume has a mount/demount.

In Release 8.32, EXISTDEL is enhanced to cover such use cases and two options in the keyword 3 are added. The third keyword "CRITERIA" with the fourth keyword "STALE", "ALL" and "NONE" is used as same as the previous release 8.31 and it determines which status of 'E' copy mode volume can be deleted. A new third keyword "WHEN" is added and it determines when 'E' copy mode volume is deleted.

In Release 8.41.200.x, when EXISTDEL function deletes the logical volume, the logical volume is marked as "removed" and "REM" field in LVOL response should show 'Y'.

Keyword 3	Keyword 4	Description
CRITERIA	STALE/ ALL/ NONE	The status of 'E' copy mode volume which is feasible to delete This is the setting to determine which condition of 'E' copy mode volume can be deleted. When TS7700 finds 'E' copy mode volume, it deletes the volume based off of the setting: STALE: Only delete 'E' copy mode volume when it is inconsistent (down level). Prior to 8.31.x.x, TS7700 handles 'E' copy mode volume this way ALL: Always delete 'E' copy mode volume so long as all other non-'E' copy mode sites are consistent. "ALWAYS" was used in the previous release 8.31. NONE: Never delete consistent or inconsistent 'E' copy mode volumes. "NEVER" was used in the previous release 8.31. The request is supported only when all clusters in the domain have the code level of 8.31.x.x or above.

		The default is "STALE".
WHEN	ATCLOSE/ AUTO	When 'E' copy mode volume is deleted This is the setting to determine when 'E' copy mode volumes which satisfy the condition set by "CRITERIA" keyword can be deleted. 'E' copy mode volume is deleted at the timing based off of the setting: ATCLOSE: 'E' copy volume is deleted at the volume mount/demount. It is the same behavior with the previous release 8.31. AUTO: In addition to the volume mount/demount timing, TS7700 periodically checks 'E' copy mode volume, then deletes it if it satisfies the condition set by CRITERIA. The check runs once per 24 hours and it deletes up to 100 'E' copy mode volumes all at once. With 8.50.x.x or later, the interval of the check and the maximum number of 'E' copy volumes deleted at once can be changed by "AUTINTVL", "AUTCOUNT" keyword respectively. The request is supported only when all clusters in the domain have the code level of 8.32.x.x or above. The default is "ATCLOSE".
AUTINTVL	Value	The interval to check 'E' copy mode volume. This is the interval in unit of hours to check 'E' copy mode volume to be deleted in case "AUTO" is set by "WHEN" keyword. The value from 1 to 24 hours can be set. This keyword is supported only on clusters with 8.50.x.x or later. The default is 24.
AUTCOUNT	Value	The maximum number of 'E' copy volumes to be deleted at once. This is the maximum number of 'E' copy mode volumes to be deleted at once in case "AUTO" is set by "WHEN" keyword. The value from 1 to 10000 volumes can be set. This keyword is supported only on clusters with 8.50.x.x or later. The default is 100.

The response lines to a SETTING request are formatted as follows in this example. Just entering the first SETTING keyword with no second, third and fourth keywords provides this response:

SETTING V8 .0

ALERTS

COPYLOW = 0 COPYHIGH = 800
 PDRVLOW = 0 PDRVCRT = 0
 PSCRLOW = 0 PSCRCTRIT = 0
 RESDLOW = 0 RESDHIGH = 0
 RSDTLOW = 0 RSDTHIGH = 0
 PCPYLOW = 0 PCPYCRIT = 1800
 DEFDEG = ENABLED LINKDEG = DISABLED
 REMOVMSG = ENABLED

 CACHE CONTROLS

COPYFSC = DISABLED
 RECLPG0 = DISABLED
 PMPRIOR = 1600 PMTHLVL = 2000
 REMOVE = DISABLED
 REMVTHR = 2000
 CPYPRIOR = DISABLED
 CPYPRITH = 100
 RBPRIOR = 0
 RBTHLVL = 200

 THROTTLE CONTROLS

COPYFT = ENABLED
 ICOPYT = ENABLED
 DCOPYT = 125
 DCTAVGTD = 100 DCTCURTD = 100

 RECLAIM CONTROLS

RCLMMAX = 0 RCLMSMAX = 0

 DEVALLOC CONTROLS

SCRATCH = ENABLED
 PRIVATE = ENABLED (CLUSTER FAMILY: RESPECTED)

 COPY CONTROLS

CPYCNT RUN = 5
 CPYCNT DEF = 120
 IMMSNS = NONE
 TIMEOUT = 180
 SCRATCH = ALWAYS

 LINK CONTROLS

FAILOVER = ENABLED

 DELEXP CONTROLS

COUNT = 1800

 EXISTDEL CONTROLS

CRITERIA = STALE WHEN = AUTO (100/24H)

The following table details the format of the data reported.

Line	Bytes	Name	Description
1	0:8	Header Info	'SETTING V'
	9:10	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to: - 2 at 8.21.x.x - 3 at 8.30.x.x - 4 at 8.30.1.x - 5 at 8.31.x.x - 6 at 8.32.x.x - 7 at 8.40.x.x - 8 at 8.42.2.x/8.50.x.x
	11	Dot	'.'
	12:13	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	14:69	Blanks	
2	0	Blank	
	1:6	Header	'ALERTS'
	7:69	Blanks	
3	0	Blank	
	1:11	Header Info	'COPYLOW = '
	12:19	Uncopied Data Low Warning Threshold	The current COPYLOW setting value. The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	20	Blank	
	21:31	Header Info	'COPYHIGH = '
	32:39	Uncopied Data High Warning Threshold	The current COPYHIGH setting value. The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	40:69	Blanks	
4	0	Blank	
	1:11	Header Info	'PDRVLOW = '
	12:19	Available Physical Drive Low Warning Threshold	The current PDRVLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	'PDRVCRIT = '

Line	Bytes	Name	Description
	32:39	Available Physical Drive Critical Warning Threshold	The current PDRVCRIT setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
5	0	Blank	
	1:11	Header Info	`PSCRLOW = '
	12:19	Physical Volume Scratch Low Warning Threshold	The current PSCRLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`PSCRCRIT = '
	32:39	Physical Volume Scratch Critical Warning Threshold	The current PSCRCRIT setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
6	0	Blank	
	1:11	Header Info	`RESLOW = '
	12:19	Resident Data Low Warning Threshold	The current RESLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`RESHIGH = '
	32:39	Resident Data High Warning Threshold	The current RESHIGH setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
7	0	Blank	
	1:11	Header Info	`RSDTLOW = '
	12:19	Resident Data Low Warning Threshold for TS7700T tape attached partitions	The current RSDTLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`RSDTHIGH = '
	32:39	Resident Data High Warning Threshold for TS7700T tape attached partitions	The current RSDTHIGH setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
8	0	Blank	
	1:11	Header Info	`PCPYLOW = '
	12:19	Resident Data Low Warning Threshold	The current PCPYLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`PCPYCRIT = '

Line	Bytes	Name	Description
	32:39	Resident Data High Warning Threshold	The current PCPYCRIT setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
9	0	Blank	
	1:11	Header Info	`DEFDEG = '
	12:19	Degraded state reporting setting	The current DEFDEG setting value describing whether preventing a composite library from entering 'degraded' state when SyncDeferred or ImmediateDeferred condition occurs. Possible values are: `ENABLED` `DISABLED`
	20	Blank	
	21:31	Header Info	`LINKDEG = '
	32:39	Link degraded state reporting setting	The current LINKDEG setting value describing whether preventing a composite library from entering 'link degraded' state when Grid link degradation occurs. Possible values are: `ENABLED` `DISABLED`
	40:69	Blanks	
10	0	Blank	
	1:11	Header Info	`REMOVMSG = '
	12:19	Auto Removal state reporting setting	The current REMOVMSG setting value describing whether preventing Auto Removal start/stop notification from being surfaced to operator intervention event (operator message). Possible values are: `ENABLED` `DISABLED`
	20:69	Blanks	
11	0:69	Separator	All dash '-' characters
12	0	Blank	
	1:14	Header	`CACHE CONTROLS`
	15:69	Blanks	
13	0	Blank	
	1:11	Header Info	`COPYFSC = '
	12:19	Copies to Follow Storage Class	The current COPYFSC setting value (ENABLED or DISABLED). The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	20:69	Blanks	
14	0	Blank	
	1:11	Header Info	`RECLPG0 = '
	12:19	Recalls Preferred to be Removed from Cache	The current RECLPG0 setting value (ENABLED or DISABLED). The value is left justified and padded with blanks.

Line	Bytes	Name	Description
	20:69	Blanks	
15	0	Blank	
	1:11	Header Info	`PMPRIOR = '
	12:19	Premigration Priority Threshold	The current PMPRIOR setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`PMTHLVL = '
	32:39	Premigration Throttling Threshold	The current PMTHLVL setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
16	0	Blank	
	1:11	Header Info	`REMOVE = '
	12:19	RemovalSetting	The current REMOVE setting value describing whether Removal capability is enabled or disabled. Possible values are: `ENABLED ' `DISABLED'
	20:69	Blanks	
17	0	Blank	
	1:11	Header Info	`REMOVTHR = '
	12:19	Removal Threshold setting	Shows the current Removal Target, by default this value will be set to the lower limit of 2TB. The upper value is limited to The total of cache space but 3TB
	20:69	Blanks	
18	0	Blank	
	1:11	Header Info	`CPYPRIOR = '
	12:19	Premigration limitation setting	The current CPYPRIOR setting value (ENABLED or DISABLED), The value is left justified and padded with blanks.
	20:69	Blanks	
19	0	Blank	
	1:15	Header Info	`CPYPRITH = '
	16:19		The current CPYPRITH setting value (10 to 1000 (MB/s)). The value is left justified and padded with blanks.
	20:69	Blanks	
20	0	Blank	
	1:18	Header Info	`RBPRIOR = '
	19		The current RBPRIOR setting value (0 to 5).
	20:69	Blanks	
21	0	Blank	
	1:15	Header Info	`RBTHLVL = '
	16:19		The current RBTHLVL setting value (100 to 5000 (MB/s)). The value is left justified and padded with blanks.

Line	Bytes	Name	Description
	20:69	Blanks	
22	0:69	Separator	All dash '-' characters
23	0	Blank	
	1:17	Header	'THROTTLE CONTROLS'
	18:69	Blanks	
24	0	Blank	
	1:11	Header Info	'COPYFT = '
	12:19	Full Cache Copy Throttling	The current COPYFT setting value (ENABLED or DISABLED). The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	20:69	Blanks	
25	0	Blank	
	1:11	Header Info	'ICOPYT = '
	12:19	Immediate Copy Throttling	The current ICOPYT setting value (ENABLED or DISABLED). The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	20:69	Blanks	
26	0	Blank	
	1:11	Header Info	'DCOPYT = '
	12:19	Deferred Copy Throttle	The current DCOPYT setting value. The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	20:69	Blanks	
27	0	Blank	
	1:11	Header Info	'DCTAVGTD = '
	12:19	DCT Average Threshold	The current DCT average threshold setting value for 20 min. The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	20	Blanks	
	21:31	Header Info	'DCTCURTD = '
	32:39	DCT Current Threshold	The current DCT average threshold setting value for 30 sec. The value is left justified and padded with blanks. This value is only relevant on Grid configurations since non-grid configurations do not perform copies.
	40:69	Blanks	
28	0:69	Separator	All dash '-' characters
29	0	Blank	
	1:16	Header	'RECLAIM CONTROLS'
	27:69	Blanks	
30	0	Blank	

Line	Bytes	Name	Description
	1:11	Header Info	'RCLMMAX = '
	12:18	Reclaim Maximum Tasks	The current value of the maximum number of reclaims that the cluster is to limit to. The value is left justified and padded with blanks. Note that if limiting is disabled, this field is set to '0'.
	19	Blank	
	20:30	Header Info	'RCLMSMAX = '
	31:37	Reclaim Maximum Tasks for Sunset Media	The current value of the maximum number of concurrent reclamation tasks for sunset media. The value is left justified and padded with blanks. A value of zero states no limit is configured.
	38:69	Blanks	
31	0:69	Separator	All dash '-' characters
32	0	Blank	
	1:17	Header	'DEVALLOC CONTROLS'
	18:69	Blanks	
33	0	Blank	
	1:11	Header Info	'SCRATCH = '
	12:18	Device Allocation for Scratch Volumes	The current SCRATCH setting value (ENABLED/DISABLED). The value is left justified and padded with blanks. When ENABLED the cluster or Grid configuration will handle Device Allocation Assistance commands for Scratch volumes. Note that the default behavior is DISABLED.
	19:69	Blanks	
34	0	Blank	
	1:11	Header Info	'PRIVATE = '
	12:19	Device Allocation for Private Volumes	The current PRIVATE setting value (ENABLED/DISABLED). The value is left justified and padded with blanks. When DISABLED the cluster or Grid configuration will inhibit and not handle Device Allocation Assistance commands for Private Volumes. Note that the default behavior is ENABLED.
	20	Blank	
	21:37	Header Info	'(CLUSTER FAMILY: '
	38:46	Cluster Family consideration of Device Allocation for Private Volumes	The current FAMILY setting value (RESPECTED/IGNORED). When RESPECTED, the cluster family configuration is taken into account to determine the optimal cluster order for Device Allocation to private volumes. When IGNORED, the cluster family configuration is ignored. Note that the default behavior is RESPECTED.
	47	Header Info	') '
	48:69	Blanks	
35	0:69	Separator	All dash '-' characters
36	0	Blank	
	1:13	Header	'COPY CONTROLS'

Line	Bytes	Name	Description
	14:69	Blanks	
37	0	Blank	
	1:17	Header Info	`CPYCNT RUN =`
	18:20	Max RUN Copy Threads	Number of concurrent threads activated for processing RUN copies. The allowed values for this field are 5-128 copy threads.
	21:69	Blanks	
38	0	Blank	
	1:17	Header Info	`CPYCNT DEF =`
	18:20	Max Deferred Copy Threads	Number of concurrent threads activated for processing Deferred copies. The allowed values for this field are 5-128 copy threads.
	21:69	Blanks	
39	0	Blank	
	1:17	Header Info	`IMMSNS =`
	18:25	Immediate-Deferred State Reporting Method	The current IMMSNS setting value (ALL/UNEXP/NONE). The value is left justified and padded with blanks.
	26:69	Blanks	
40	0	Blanks	
	1:17	Header Info	`TIMEOUT =`
	18:25	Volume Copy Timeout Value	The current TIMEOUT setting value. The value is left justified and padded with blanks.
	26:69	Blanks	
41	0	Blanks	
	1:17	Header Info	`SCRATCH =`
	18:25	Scratch Volume Copy Control	The current SCRATCH setting value. The value is left justified and padded with blanks.
	26:69	Blanks	
42	0:69	Separator	All dash '-' characters
43	0	Blank	
	1:13	Header Info	`LINK CONTROLS`
	14:69	Blanks	
44	0	Blank	
	1:11	Header Info	`FAILOVER =`
	12:19	Failover setting	The current FAILOVER setting value describing whether Link Failover is enabled or disabled. Possible values are: `ENABLED` `DISABLED`
	20:69	Blanks	
45	0:69	Separator	All dash '-' characters

Line	Bytes	Name	Description
46	0	Blank	
	1:15	Header Info	'DELEXP CONTROLS'
	16:69	Blanks	
47	0	Blank	
	1:11	Header Info	'COUNT = '
	12:19	Delete Expire Count	The current Delete Expire Count value
	20:69	Blanks	
48	0:69	Separator	All dash '-' characters
49	0	Blank	
	1:17	Header Info	'EXISTDEL CONTROLS'
	18:69	Blanks	
50	0	Blank	
	1:11	Header Info	'CRITERIA = '
	12:19	Existdel CRITERIA setting	The current EXISTDEL CRITERIA setting value describing the conditions of 'E' copy mode volumes which can be deleted. Possible values are: ' STALE' ' ALL' ' NONE'
	20	Blank	
	21:31	Header Info	'WHEN = '
	32:39	Existdel WHEN setting	The current EXISTDEL WHEN setting value describing when 'E' copy mode volumes can be deleted. Possible values are: ' ATCLOSE' ' AUTO'
	40:41	Header Info	' (' if the current EXISTDEL WHEN setting value is AUTO, otherwise padded with blanks.
	42:46	Existdel AUTOCOUNT setting	The current EXISTDEL AUTOCOUNT setting value describing how many 'E' copy mode volumes can be deleted. Possible values are from 1 to 10000. The value is left justified and padded with blanks. If the current EXISTDEL WHEN setting value is not AUTO, this field is padded with blanks.
	47	Header Info	'/' if the current EXISTDEL WHEN setting value is AUTO, otherwise padded with blanks.
	48:49	Existdel AUTINTVL setting	The current EXISTDEL AUTINTVL setting value describing how often (interval) EXISTDEL function runs periodically. Possible values are from 1 to 24. The value is left justified and padded with blanks. If the current EXISTDEL WHEN setting value is not AUTO, this field is padded with blanks.
	50	Header Info	') ' if the current EXISTDEL WHEN setting value is AUTO, otherwise padded with blanks.

Line	Bytes	Name	Description
	51:69	Blanks	

3.1.13 SETTING2

In the response for the SETTING request, only up to 50 lines can be provided. In the release 8.32, all 50 lines are used up. The new first keyword SETTING2 is added to provide more SETTING type of requests in the release 8.32.

The following table lists the keywords supported for the SETTING2 request, whether the request is applicable to a composite and/or distributed library, and whether the request is applicable on a cluster with physical tape and/or a disk-only configuration.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
SETTING2	SCRATCH	PFRLOCO	ENABLE/ DISABLE	N/A	Y	Y	Y	Y	Y
SETTING2	CACHE	MAXLGMC	Value	N/A	Y	Y	Y	Y	Y
SETTING2	CACHE	OVRSPVL	Value	N/A	Y	N	N	Y	Y
SETTING2	ALERT	PDRVSLOW	Value	N/A	Y	Y	N	Y	N
SETTING2	ALERT	PDRVSCRT	Value	N/A	Y	Y	N	Y	N
SETTING2	ALERT	LMTDTHR	Value	N/A	Y	N	Y	Y	Y
SETTING2	ALERT	RSDOLOW	Value	N/A	Y	N	Y	Y	Y
SETTING2	ALERT	RSDOHIGH	Value	N/A	Y	N	Y	Y	Y
SETTING2	ALERT	CAGAHIGH	Value	N/A	Y	Y	Y	Y	Y
SETTING2	ALERT	CAGALOW	Value	N/A	Y	Y	Y	Y	Y
SETTING2	ALERT	CAGHIGH	Value	N/A	Y	Y	Y	Y	Y
SETTING2	ALERT	CAGLOW	Value	N/A	Y	Y	Y	Y	Y
SETTING2	THROTTLE	LNKSPEED	Value	N/A	Y	Y	Y	Y	Y
SETTING2	PHYSLIB	TVCWDEG	DISABLE/ EQUAL/ LOWER	N/A	Y	N	N	Y	N
SETTING2	PHYSLIB	CPYWDEG	ENABLE/ DISABLE	N/A	Y	N	N	Y	N
SETTING2	PHYSLIB	PRETHDEG	ENABLE/ DISABLE	N/A	Y	N	N	Y	N
SETTING2	PHYSLIB	SLDPMPRI	ENABLE/ DISABLE	N/A	Y	N/A (*1)	N	Y	Y
SETTING2	PHYSLIB	MAINT	ENABLE/ DISABLE	N/A	Y	N/A (*1)	N	Y	N
SETTING2	RECALL	BLKRCLSZ	Value	N/A	Y	Y	N	Y	N
SETTING2	BVIR	TIMEOUT	Value	N/A	Y	Y	Y	Y	Y
SETTING2	CENCRYPT	CIPHER	DISABLE/ AES128/ AES256	N/A	Y	N/A (*1)	Y	Y	Y

SETTING2	COPY	COPYRFSH/ DEF	PAUSE/ UNPAUSE	N/A	Y	N/A (*1)	Y	Y	Y
SETTING2	VOLINVT	LPAGRP	ENABLE/ DISABLE	N/A	Y	N/A (*1)	Y	Y	Y

(*1) This command is not supported on TS7740 since it is new to R5.0 and R5.0 is not supported on TS7740.

When the request has the first keyword “SETTING2” only, the response just provides the current settings on the target distributed library specified in the target library name. If the composite library is specified in the target library name, the returned setting values depend on the distributed library which receives the request then the unexpected setting values could be returned because it may vary which distributed library receives the request. The distributed (not composite) library name must be always specified when it is required to get the setting values from the specific distributed library.

3.1.13.1 Scratch Settings

If a second keyword of SCRATCH is specified, the cluster will follow the setting when a scratch mount is requested.

The following SCRATCH setting is supported

Keyword 3	Keyword 4	Description
PFRLOCO	ENABLE, DISABLE	Prefer locally owned logical volumes in the scratch mount TS7700 adopts the logical volume ownership mechanism and favors locally owned volume to the host scratch mount request. This is to reduce the ownership exchange overhead and contention on peer clusters. There may be possible that the ownerships are unevenly distributed among the clusters in the domain and it's not optimal to always favor locally owned volumes from delete expire function viewpoint. This setting allows to set to ignore favoring locally owned volumes and older volume is selected at the scratch mount regardless of the ownership status. When the ENABLE keyword is specified, TS7700 prefers the locally owned logical volumes to remotely owned logical volumes in the scratch mount. When the DISABLE keyword is specified, TS7700 doesn't prefers the locally owned logical volumes and older volume regardless of the ownership status is favored in the scratch mount. The default is enabled.

If this request is issued to a stand-alone cluster, the following error text is returned:

'CLUSTER NOT PART OF A GRID CONFIGURATION'

3.1.13.2 Alert Settings

As same as SETTING, ALERT keywords, SETTING2, ALERT sets the alert threshold specified by the third keyword.

The following ALERT thresholds are supported:

Keyword 3	Keyword 4	Description
PDRVSCRT	Value	Available Sunset Physical Drive Critical Warning Limit Similar to “SETTING, PDRVCRT”, this is the threshold in number of physical drives at which the TS7700 will generate a message indicating that the number of available physical drives has fallen below the critical warning limit. But, <u>this setting is only applicable to the number of the sunset drives</u>. If the cluster has both non-sunset and sunset drives, this setting is only applicable to the sunset drives. The SETTING, PDRVCRT setting must be independently set for non-sunset drives. For this message to be generated, the number of available sunset physical drives must be below the value specified in keyword 4 for 15 minutes. As long as the number of available sunset physical drives is below the threshold, the below the threshold message is repeated every 15 minutes. If the message had been generated and the number of available sunset physical drives is at or has risen above that threshold for 15 minutes, a message is generated indicating that the number of available sunset physical drives is above the critical warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than the number of installed sunset physical drives minus 1 cannot be set. If a value is specified that is greater than the number of installed sunset physical drives minus 1, the threshold is set to the number of installed sunset physical drives minus 1. A value less than 3 cannot be set. If a value is specified that is less than 3, the threshold is set to 3. If the available sunset physical drive low warning limit is not zero and a value is specified that is not 1 less than the available sunset physical drive low warning limit, the available sunset physical drive low warning limit will be changed so that it is 1 more than the value specified (but never more than the number of installed sunset physical drives). Message Text (xx is the number of available sunset drives, yy is the threshold): When fallen below the threshold:

Keyword 3	Keyword 4	Description
		“AL5017 Available sunset physical drives of xx is below critical limit of yy.” When risen above the threshold: “AL5018 Available sunset physical drives is no longer below critical limit of yy.”
PDRVSLOW	Value	Available Sunset Physical Sunset Drive Low Warning Limit Similar to “SETTING, PDRVLOW”, this is the threshold, in number of physical drives at which the TS7700 will generate a message indicating that the number of available physical drives has fallen below the low warning limit. But <u>this setting is only applicable to the number of the sunset drive</u>. If the cluster has both non-sunset and sunset drives, this setting is only applicable to the sunset drives. The SETTING, PDRVLOW setting must be independently set for non-sunset drives. For this message to be generated, the number of available sunset physical drives must be below the value specified in keyword 4 for 15 minutes. If the message had been generated and the number of available sunset physical drives is at or has risen above that threshold for 15 minutes, a message is generated indicating that the number of available sunset physical drives is above the low warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. If the available sunset physical drive critical warning limit is set to 0, then the available sunset physical drive low warning limit cannot be set to a value greater than the number of installed sunset physical drives or less than 3. If a value is specified that is greater than the number of installed sunset physical drives, the threshold is set to the number of installed sunset physical drives. If a value is specified that is less than 3, the threshold is set to 3. If the available sunset physical drive critical warning limit is not zero and a value is specified that is not 1 greater than the available sunset physical drive critical warning limit, the available sunset physical drive low warning limit will be changed so that it is 1 more than the available sunset physical drive critical warning limit (but never greater than the number of installed sunset physical drives). Message Text (xx is the number of available sunset drives, yy is the threshold): When fallen below the threshold:

Keyword 3	Keyword 4	Description
		“AL5019 Available sunset physical drives of xx is below low limit of yy.” When risen above the threshold: “AL5020 Available sunset physical drives is no longer below low limit of yy.”
LMTDTHR	Value	Limited Cache threshold This setting allows you to configure the threshold of when a TS7700D or TS7700T enters the limited cache free space warning state. This warning threshold provides an opportunity to take action before a cluster enters the out of disk cache read-only state. The cluster enters the limited cache warning state when the free space in the TS7720D disk cache or TS7720T/7700C CP0 partition drops below a value of: (TS7700D) LMTDTHR + 1,000GB (TS7700T/7700C) LMTDTHR + MIN(1000GB, 5% of CP0_size) The cluster exits the limited cache warning state when the free space in the TS7720D disk cache or TS7720T CP0 partition exceeds a value of: (TS7700D) LMTDTHR + 1,000GB + 500GB (TS7700T/7700C) LMTDTHR + MIN(1000GB, 5% of CP0 size) + 500GB The minimum LMTDTHR value allowed is, (TS7700D) 1,000GB (TS7700T/7700C) MIN(1,000GB, 5% of CP0 size) The maximum LMTDTHR value allowed is, (TS7700D) cache size – 2,072GB (TS7700T/7700C) CP0 size – MIN(1,072GB, 5% of CP0 size) – 1,000GB LMTDTHR = 0 is supported and informs the TS7700 to use internally defined default values. These values are: (TS7700D) 2,000GB (TS7700T/7700C) MIN(2000GB, 10% of CP0 size) In a TS7700T/7700C, the CP0 size can be changed at a later time. If the newly configured CP0 size conflicts with the currently defined LMTDTHR values minimum or maximum, the

Keyword 3	Keyword 4	Description
		LMTDTHR value will remain unchanged, but the TS7700 will internally adjust it to a legal value. The "SETTING2" LMTDTHR output will continue to display the value set prior to the CP0 size change. For example, let's assume CP0 currently has a size of 10TB and the LMTDTHR setting was set to 800 (where the allowed minimum was 500). Next, the size of CP0 is changed to 30TB. Since the new allowed minimum of LMTDTHR is 1000 and the currently assigned setting of 800 is below the new minimum, a minimum value of 1000 is used internally as a limited cache threshold while LMTDTHR continues to display 800 in the "SETTING2" output. If the LMTDTHR value is later adjusted, it will then enforce the new minimum and maximum values.
RSDOHIGH	Value	Resident Data High Warning Limit For Objects This is the threshold, in GBs of resident data for all objects cache partitions, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a high warning limit. This threshold is similar to RESDHIGH except that this value is applied to objects. Message Text (xxxxxxx is the amount of resident data, yyyyyyyy is the threshold): When above the threshold: "AL5021 Sum of resident data in objects partitions of xxxxxxxx GB above high warning limit of yyyyyyyy GB." When below the threshold: "AL5022 Sum of resident data in objects partitions no longer above resident data high warning limit of yyyyyyyy GB." Note: This setting has been moved to OBJSET1,ALERT as of R8.52.200.x microcode level. See the OBJSET1 section for any additional changes that may apply if at this microcode level or higher.
RSDOLOW	Value	Resident Data Low Warning Limit For Objects This is the threshold, in GBs of resident data for all objects cache partitions, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a low warning limit. This threshold is similar to RESDLOW except that this value is applied to objects. Message Text (xxxxxxx is the amount of resident data, yyyyyyyy is the threshold): When above the threshold: "AL0023 Sum of resident data in objects partitions of xxxxxxxx GB above low warning limit of yyyyyyyy GB."

Keyword 3	Keyword 4	Description
		When below the threshold: “AL0024 Sum of resident data in objects partitions no longer above low warning limit of yyyyyyyy GB.” Note: This setting has been moved to OBJSET1,ALERT as of R8.52.200.x microcode level. See the OBJSET1 section for any additional changes that may apply if at this microcode level or higher.
CAGAHIGH	Value	Inbound All Copy Backlog Age High Warning Limit This is the threshold, in hours that the longest inbound copy backlog age on all copy jobs, at which the TS7700 will generate a message indicating that the longest inbound copy backlog age has been at or exceeded a high warning limit. The above warning limit message is not repeated. If the message had been generated and the longest copy backlog age falls below that threshold, a message is generated indicating that the longest inbound copy backlog age is below the high warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than 8760 which is 365 days cannot be set. Message Text (Z is the distributed library name, YYYY is the longest inbound copy backlog age, XXXX is the threshold, VVVVVV is the volser which has the longest copy backlog age): When above the threshold: "AL0014 Distributed Library Z has an inbound all copy type backlog that exceeds the high warning limit of XXXX hours. Volser VVVVVV is the oldest copy job in this group at YYYY hours." When below the threshold: "AL0015 Distributed Library Z has falling below the inbound all copy type backlog high limit of XXXX hours." Note: Attempting the lower CAGAHIGH value than CAGALOW may fail with the following error message: THE PROVIDED HIGH VALUE MUST BE ABOVE THE CURRENT LOW VALUE
CAGALOW	Value	Inbound All Copy Backlog Age Low Warning Limit This is the threshold, in hours that the longest inbound copy backlog age on all copy jobs, at which the TS7700 will generate a message indicating that the longest inbound copy backlog age has been at or exceeded a low warning limit. The above warning limit

Keyword 3	Keyword 4	Description
		message is not repeated. If the message had been generated and the longest copy backlog age falls below that threshold, a message is generated indicating that the longest inbound copy backlog age is below the low warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than 8760 which is 365 days cannot be set. Message Text (Z is the distributed library name, YYYY is the longest inbound copy backlog age, XXXX is the threshold, VVVVVV is the volser which has the longest copy backlog age): When above the threshold: "AL0016 Distributed Library Z has an inbound all copy type backlog that exceeds the low warning limit of XXXX hours. Volser VVVVVV is the oldest copy job in this group at YYYY hours." When below the threshold: "AL0017 Distributed Library Z has falling below the inbound all copy type backlog low limit of XXXX hours." Note: Attempting the higher CAGALOW value than CAGAHIGH may fail with the following error message: THE PROVIDED LOW VALUE MUST BE BELOW THE CURRENT HIGH VALUE
CAGHIGH	Value	Inbound Non Time-delayed Copy Backlog Age High Warning Limit This is the threshold, in hours that the longest inbound copy backlog age on non time-delayed copy jobs, at which the TS7700 will generate a message indicating that the longest inbound copy backlog age has been at or exceeded a high warning limit. The above warning limit message is not repeated. If the message had been generated and the longest copy backlog age falls below that threshold, a message is generated indicating that the longest inbound copy backlog age is below the high warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than 8760 which is 365 days cannot be set. Note: A time-delayed copy job is a copy job with Time Delayed copy mode which is still delayed until the delay is expired. Once

Keyword 3	Keyword 4	Description
		the delay is expired, the copy job is changed to a non time-delayed copy job. Message Text (Z is the distributed library name, YYYY is the longest inbound copy backlog age, XXXX is the threshold, VVVVVV is the volser which has the longest copy backlog age): When above the threshold: "AL0018 Distributed Library Z has an inbound non time-delayed copy type backlog that exceeds the high warning limit of XXXX hours. Volser VVVVVV is the oldest copy job at YYYY hours." When below the threshold: "AL0019 Distributed Library Z has falling below the inbound non time-delayed copy type backlog high limit of XXXX hours." Note: Attempting the lower CAGHIGH value than CAGLOW may fail with the following error message: THE PROVIDED HIGH VALUE MUST BE ABOVE THE CURRENT LOW VALUE
CAGLOW	Value	Inbound Non Time-delayed Copy Backlog Age Lowrning Limit This is the threshold, in hours that the longest inbound copy backlog age on non time-delayed copy jobs, at which the TS7700 will generate a message indicating that the longest inbound copy backlog age has been at or exceeded a low warning limit. The above warning limit message is not repeated. If the message had been generated and the longest copy backlog age falls below that threshold, a message is generated indicating that the longest inbound copy backlog age is below the low warning limit. The default value is 0, which indicates no warning limit is set and no messages are to be generated. A value greater than 8760 which is 365 days cannot be set. Note: A time-delayed copy job is a copy job with Time Delayed copy mode which is still delayed until the delay is expired. Once the delay is expired, the copy job is changed to a non time-delayed copy job. Message Text (Z is the distributed library name, YYYY is the longest inbound copy backlog age, XXXX is the threshold, VVVVVV is the volser which has the longest copy backlog age): When above the threshold:

Keyword 3	Keyword 4	Description
		"AL0020 Distributed Library Z has an inbound non time-delayed copy type backlog that exceeds the low warning limit of XXXX hours. Volser VVVVVV is the oldest copy job at YYYY hours." When below the threshold: "AL0021 Distributed Library Z has falling below the inbound non time-delayed copy type backlog low limit of XXXX hours." Note: Attempting the higher CAGLOW value than CAGHIGH may fail with the following error message: THE PROVIDED LOW VALUE MUST BE BELOW THE CURRENT HIGH VALUE

3.1.13.3 CACHE Settings

If a second keyword of CACHE is specified, the cluster will modify how it controls the workflow and content of the tape volume cache.

The following CACHE setting is supported

Keyword 3	Keyword 4	Description
MAXLGMC	Value	Limit the concurrent large volume (25GB) mount count In Release 8.32, 25GB volume size which can be defined in the storage constructs (data class) is fully supported without RPQ. It may be possible if a TS7700 has a lot of this large size (25GB) volume mount concurrently then the cache usage could be “overrun” and maybe migration cannot act fast enough. The risk is even more realistic when up to 496 virtual devices are supported in Release 8.32 (496 devices feature is supported only via RPQ). This is the number of the concurrent mount of 25GB volume size that uses the local cache as Tape Volume Cache, at which TS7700 will start failing the mount. The mount fails with the specific reason code X'42' (Excessive Concurrent Large Volume Mounts Detected) and an the following operator message is surfaced the host console : <pre>'G0035 Distributed Library XXX mount failed due to excessive large volume mounts'</pre> The mount includes the local and remote mount. If it is a mount with Sync Copy mode and both TVC and Fork clusters are opened, the mount is counted and MAXLGMC setting is effective on both clusters. If a cluster exceeds the setting but the other cluster is still available for the volume with 25GB size, the mount should succeed by using the other cluster as Tape Volume Cache. The default is 128. The value from 0 to 65535 can be set.
OVRSPVLV	Value	Reserve CP0 free space to avoid automatic removal of CP0 data by limiting the overspill data from CPx This setting is applicable TS7700T/7700C only. If the size of the cache partition x (CPx) is smaller than the number of FC5274 is installed, the amount of data targeting CPx can exceed its configured capacity. This state is called the “overspill state” and the excess data targeting CPx is called “overspill data”. The “overspill data” content actually utilizes free space in CP0. If the pre-existing CP0 content and "overspill data" in CP0 causes the total usage size of CP0 to exceed the automatic removal threshold, the CP0 pre-existing content is exposed to auto removal. "Overspill data" cannot be auto removed given it has not yet been pre-migrated to back end tape, thus only pre-

Keyword 3	Keyword 4	Description
		existing CP0 resident only data is a candidate for auto removal. The OVRSPVLV value provides a method to limit how much total overspill can exist in CP0 for all partitions in order to prevent too much content overspilling into CP0 and causing automatic removal to remove CP0 targeted content. - When OVRSPVLV is set to 0 GB, CPx can overspill the data up to CP0 limited cache state. As a result, the auto-removal in CP0 may occur due to excessive “overspill data”. The figure below shows how OVRSPVLV (0 GB) works on CP0. - When OVRSPVLV is set to X GB, CPx can accept "overspill data", but (X GB + REMVTHR + 1000 GB) free space in CP0 must exist for this overspill acceptance to continue. Meaning, the actual limit is relative to the current REMVTHR value. Even if the auto-removal is disabled, the OVRSPVLV limit is still based on the current REMVTHR value. If the free space criteria cannot be met, the automatic overspill of CPx partitions into CP0 is disabled. If OVRSPVLV is set to the value larger than the CP0 allocation size, overspill will always remain disabled. The figure below shows how an OVRSPVLV values of 2000 GB works on CP0. The default is 2000 (GB). The value from 0 to 99999999 can be set. A good best practice would be to set this threshold to a value equal to the maximum amount of expected CP0 targeted content. For example, if CP0 is configured to 10TB and you don't expect CP0 resident only content to exceed

Keyword 3	Keyword 4	Description
		5TB, the OVR SPLVL should be set to 5000 GB allowing roughly 2TB of overspill (10TB - 2TB - 1TB - 5TB = 2TB). Note: Limited cache point illustrated above can be changed by LMTDTHR.

3.1.13.4 THROTTLE Settings

If a second keyword of THROTTLE is specified, the cluster will modify how it controls the different data flow rates into and out of the cluster.

The following THROTTLE settings are supported:

Keyword 3	Keyword 4	Description
LNKSPEED	Value	Grid Link speed used to calculate Immediate Copy Throttling The Immediate Copy Throttling (also known as RUN copy throttling) mechanism, which can be enabled or disabled by the ICOPYT setting, helps prevent clusters from entering the immediate-deferred state by throttling host activity when the TS7700 predicts the immediate copy backlog may not be able to complete within 40 minutes. This prediction is based on the total amount of immediate copy queued data and an estimate of how fast the copy queue can be serviced. This LNKSPEED value is used as the cumulative single direction grid links speed in the above prediction calculation. The cumulative single direction grid link speed may not be the same as the configured or negotiated link speeds given the end to end configuration may not support full links speeds. For example, ISL links, link latency, shared network infrastructure or devices such as quality of service (QoS) devices can affect to the grid link speeds between two clusters. Smaller LNKSPEED value may trigger the Immediate Copy Throttling more aggressively while values too high may result in a cluster entering the immediate-deferred state. If the host happens to issue multiple demounts for a large number of large volumes (e.g. 6GB volumes) simultaneously, a very large jump in immediate copy backlog can occur in which a larger LINKSPEED value can help prevent excessive immediate mode copy throttling. The unit of LNKSPEED is in MB/sec. The minimum and maximum value allowed is 1 and 4000, respectively. LNKSPEED=0 informs the TS7700 that it should utilize an internal default, such as 200MB/s. Though every configuration is unique, below is a suggested starting point for different configurations where network infrastructure isn't a bottleneck. 2 x 1Gb links – LNKSPEED=100 4 x 1Gb links – LNKSPEED=200 2 x 10Gb links – LNKSPEED=1000 4 x 10Gb links – LNKSPEED=2000 The value of LNKSPEED is used only in calculation of Immediate Copy Throttling. It does not affect other operation in TS7700.

Keyword 3	Keyword 4	Description
		Note: A value of 200 MB/s has been used internally by all pre-8.41.200.x machines regardless of Grid link configurations. 8.41.200.X introduces the ability to tune this value.

3.1.13.5 PHYSLIB Settings

In Release 8.32, the TS7700T was supported and inherited many of the behaviors of the TS7740. One behavior is that inbound copies and Tape Volume Cache (TVC) selection is disabled against a TS7700T tape managed CPx partition when the physical library is unavailable or heavily degraded. Workloads and copies to CP0 or resident only partitions are always allowed to continue. Given that a CPx partition can contain a great deal of content, the inability to premigrate data to tape or recall data from tape may not be a significant problem and therefore not a reason to prevent copies or mounts from continuing to use the TS7700T. With large CPx partitions, a large premigration backlog can easily be retained in disk cache providing a great deal of buffer when physical tape operations are unavailable. The PHYSLIB provides options in R3.3 that allow you to alter the inbound copy and TVC selection behavior when a physical library becomes non-operational.

Keyword 3	Keyword 4	Description
TVCWDEG (TVC when DEGraded)	DISABLE/ EQUAL/ LOWER	Allow to select TS7700T with the degraded library as a mount TVC This setting determines how any TS7700 mount cluster in a Grid views this TS7700T for CPx TVC selection when this TS7700T cluster's attached physical libraries is in a non-operational or heavily degraded state. Meaning, this setting is set on the TS7700T and viewed by all peers and itself during a mount operation. When it is set to EQUAL or LOWER, the TS7700 cluster receiving the mount request will continue to allow this TS7700T cluster's CPx partitions to remain as candidates for TVC selection when this TS7700T's library is degraded as long as a recall from physical tape is not required. When set to EQUAL, this TS7700T is viewed as an equal candidate when compared to other physical tape based TVC choices. When set to LOWER, this TS7700T remains a candidate for TVC, but it will be lowered in preference within the TVC candidate list. When it is set to DISABLE, this TS7700T CPx partitions are excluded from the mount TVC candidate lists when the physical library is not operational. The default is DISABLE. When a TS7700T CPx partition with a degraded library is selected due to one of these settings, it will continue to allow inbound host writes to CPx partitions using up any available free space. The free space is made available through overspill into CP0 and/or through the migration of previously premigrated data. Once the free space is exhausted within any of the TS7700T's CPx partitions, premigration throttling will kick in at its highest setting to prevent any meaningful inbound host writes. Therefore, a non default setting makes most sense when CPx partitions are of a significant size to accommodate any workload that may arrive while the library remains inaccessible. The setting can be changed while the library remains degraded if a new behavior is required, such as to go from EQUAL to LOWER or to DISABLE.

Keyword 3	Keyword 4	Description
		If more than one TS7700T exists in a grid, each TS7700T's unique setting will be honored. But, this setting cannot be made unique based on which cluster was chosen as the mount cluster. For synchronous mode copies, both instances are chosen as part of TVC selection and each 'S' location's setting will be honored at mount time. This setting does not alter the behavior of CP0 workloads which are independent of the physical library state. For example, this is 2-way Grid (C0:TS7700T, C1:TS7700D). The physical tape library of C0 is in a heavily degraded state. 1. C0's setting is set to EQUAL C0 can be selected a TVC for mounts issued to C0 or C1 when: • Scratch Mount • Private Mount with no recall on C0 C0 can't be selected as a mount TVC when: • Private Mount with recall on C0 2. C0's setting is set to LOWER C0 can be selected as TVC, with lower preference, for mounts issued to C0 or C1 when: • Scratch Mount • Private Mount with no recall on C0 C0 can't be selected as a mount TVC when: • Private Mount with recall on C0 3. C0's setting is set to DISABLE C0 cannot be selected as TVC independent of mount point. This setting is unique per TS7700T. For example, if three TS7700Ts (C0/C1/C2) exist in the Grid, each cluster can have a different setting such as: C0:EQUAL, C1:LOWER, C2:DISABLE Note: The copy mode and cluster family configuration have precedence over PHYSLIB setting.

Keyword 3	Keyword 4	Description
		For example, this is 4-way Grid (all TS7700T). The physical tape library of C0 is in a heavily degraded state. CL0/CL1 are in the same family (FAM1) and CL2/CL3 are in another family (FAM2). C0's setting is set to LOWER. A scratch mount issued to C0 or C1: 1. C0/C1/C2/C3 copy mode = DEF/DEFDEF/DEF (DEF: Deferred) C0 can be selected as TVC, with lower preference. 2. C0/C1/C2/C3 copy mode = DEF/No/DEF/No (No: No copy) C0 can be selected as TVC. This is because C0 is the only cluster with the valid copy mode in FAM1 although C2 in the other family FAM2 also has DEF copy mode. 3. C0/C1/C2/C3 copy mode = RUN/DEF/RUN/DEF (RUN: Rewind Unload) C0 can be selected as TVC. This is because C0 is the only cluster with RUN copy mode in FAM1 although C2 in the other family FAM2 also has RUN copy mode (C1 in the same family FAM1 has DEF copy mode which is lower preference than RUN). At 8.33.x.x, DAA (Device Allocation Assistance) doesn't take into account the TVCWDEG setting. Once 8.40.x.x or later is installed, TVCWDEG is taken into account for device allocation assistance.
CPYWDEG (CoPY when DEGraded)	DISABLE/ ENABLE	Allow Incoming Copy on a TS7700T with the degraded library This setting determines how this TS7700T handles inbound replication when its attached physical library is in a non-operational or heavily degraded state. When it is set to ENABLE for this TS7700T, this TS7700T will continue to allow inbound replication into CPx partitions using up any available free space. The free space is made available through overspill into CP0 and/or through migration of previously premigrated data. Once the free space is exhausted within any CPx partition, premigration throttling will kick in at its highest setting to prevent any meaningful inbound replication. The default is DISABLE which means no inbound replication into CPx is allowed to this TS7700T with a heavily degraded physical library. Therefore, a non default setting makes most sense when CPx partitions are of a significant size to accommodate any inbound replication workload that may arrive while the library remains inaccessible. The setting

Keyword 3	Keyword 4	Description
		can be changed while the library remains degraded if a new behavior is required, such as to go from ENABLE to DISABLE. This setting only applies to RUN, and any type of Deferred copies (including Immed-deferred and synchronous-deferred copies). Synchronous mode copies are not controlled by CPYWDEG, but through the TVCWDEG setting. This setting does not alter the behavior of CP0 workloads which are independent of the physical library state. This setting is relative to the TS7700T cluster receiving the inbound copies. Each TS7700T in a Grid can be configured in a unique way.
PRETHDEG (PREmigration THrottling when DEGraded)	DISABLE/ ENABLE	Determine if premigration throttling should occur in a TS7700T when it has a degraded library This setting determines whether a TS7700T enforces premigration throttling when its attached physical library is non-operational or heavily degraded. This setting is only applicable to TS7700T configurations where either TVCWDEG or CPYWDEG is set to a value other than DISABLE. If both TVCWDEG and CPYWDEG is set to DISABLE, then the TS7700T will become read-only for data targeting CPx partitions preventing further inbound host writes and inbound copies. If either TVCWDEG or CPYWDEG is set to a value other than DISABLE, then inbound data will continue to target this TS7700T's CPx partitions and the PRETHDEG setting determines if premigration throttling should continue during this degraded period. When PRETHDEG is set to ENABLE, premigration throttling will continue to be applied during this degraded period as soon as the premigration queue exceeds PMTHLVL. When PRETHDEG is set to DISABLE, premigration throttling will be avoided during this degraded period. This is accomplished by temporarily adjusting the PMTHLVL during the degraded period. The temporary PMTHLVL or PMTH_TMP is automatically adjusted to be equal to the amount of data currently queued for premigration plus the current PMTHLVL value. As more data arrives and the premigration queue grows, the PMTH_TMP will continue to adjust. Therefore, the content in the premigration queue prior and during the degraded period is not included in the throttling calculation. Once the library returns to an operational state, the

Keyword 3	Keyword 4	Description
		PMTH_TMP will remain fixed until the premigration queue drops below the configured PMPRIOR value. Only then will the original PMTHLVL become active again. This behavior allows the TS7700T to operate during the degraded period without throttling while also providing the entire PMTHLVL of queue space once the library returns to an operational state. This also assumes there is efficient space in disk cache to accommodate this additional queued data. The current PMTH_TMP value can be retrieved using BVIR POINT IN TIME STATISTICS command. The default value of PRETHDEG is DISABLE. This setting is unique to each TS7700T cluster in a grid. Note: The SLDPMPRI should also be considered as a way to automatically adjust the PMPRIOR value too.
SLDPMPRI (SLiDe PMPRIor after physical library becomes operational)	DISABLE/ ENABLE	Automatically adjust CLDPRIOR in a TS7700C or PMPRIOR in a TS7700T when the attached cloud or tape library becomes degraded Please review the PRETHDEG setting first. The PRETHDEG is specific to a TS7700T, but a hidden TS7700C or cloud equivalent also exists. The TS7700C version of the setting is internally fixed to a value of DISABLE allowing the TS7700C to always operate in a non-throttling state during a cloud object store outage. In addition, the TS7700C has TVCWDEG and CPYWDEG internal settings fixed to a value of EQUAL. Therefore, understanding the PRETHDEG setting will help you understand why the SLDPMPRI setting exists for both a TS7700T and a TS7700C. If PRETHDEG is set to DISABLE for a TS7700T or you have a TS7700C, premigration throttling is disabled while an attached tape library or cloud is degraded. This occurs by automatically adjusting the PMTHLVL to a higher level. The SLDPMPRI value determines if the PMPRIOR setting for TS7700T or CLDPRIOR setting for TS7700C should also adjust automatically. When the library or cloud returns to an operational state, the TS7700T and TS7700C provide a full PMTHVL worth of queue space. But, the amount of data queued during the degraded period likely exceeds the PMPRIOR or CLDPRIOR value. Therefore, premigration will ramp up immediately once the library or cloud returns to an operational state. The SLDPMPRI value allows the PMPRIOR and CLDPRIOR settings to shift too during the degraded period so that premigration doesn't ramp up until PMPRIOR or CLDPRIOR amount of data is

Keyword 3	Keyword 4	Description
		additionally queued once the library or cloud returns to an operational state. When SLDPMPRI is set to DISABLE, the PMPRIOR and CLDPRIOR values remain fixed and are not adjusted during a library or cloud degraded period. When SLDPMPRI is set to ENABLE, the PMPRIOR and CLDPRIOR values are adjusted similarly to PMTHLVL during a library or cloud degraded period. The adjusted or temporary PMPRIOR value under a TS7700T is referred to as PMPRIOR_TMP. The adjusted or temporary CLDPRIOR value under a TS7700C is referred to as CLDPRIOR_TMP. The current PMPRIOR_TMP and CLDPRIOR_TMP value can be retrieved using BVIR POINT IN TIME STATISTICS command. The default value of SLDPMPRI is ENABLE.
MAINT (MAINTenance mode of a backend tape library)	DISABLE/ ENABLE	Enable/Disable maintenance mode of a backend tape library This setting determines if maintenance mode of a backend tape library need to be enabled in TS7700T. When MAINT is set to ENABLE, TS7700T suspends all physical tape activities such as premigration, read only recovery, reclamation, secure data erase, copy export, and offsite reclamation. If any activity is already running, it will not be canceled. Recall will also be suspended as this cluster will not be selected as TVC if recall is required. The default value is DISABLE.

If TVCWDEG or CPYWDEG request is issued to a stand-alone cluster, the following error text is returned (PRETHDEG, SLDPMPRI, and MAINT are applicable to a stand-alone cluster):

‘CLUSTER NOT PART OF A GRID CONFIGURATION’

If this request is issued to a non-TS7700T cluster, the following error test is returned:

‘ONLY SUPPORTED IN A TS7700 TAPE ATTACHED VIRTUALIZATION ENGINE’

3.1.13.6 RECALL Settings

If a second keyword of RECALL is specified, the cluster will modify recall settings.

Keyword 3	Keyword 4	Description
BLKRCLSZ	Value	Bulk Recall Size Limit Setting for the maximum GB of logical volume content included in each bulk copy recall request. Bulk Recall is a recall task triggered by Grid replication activity. When one or more peer clusters require a copy of a volume that is only on back end physical tape within a copy source cluster, a recall is required before the replication can be started. A source cluster keeps track of which volumes may need to be copied by all peers. When a recall request for Grid replication occurs, the source cluster can look for other volumes on the same physical tape which also need to be replicated to one or more peers. The source cluster will then recall the requested volume as well as other volumes on the same physical tape using a bulk recall mechanism which helps accelerate the process of recalling many volumes into the source cluster's disk cache for Grid replication purposes. The Bulk Recall mechanism has a default limit of 13 GB per bulk request. For example, if 1 GB logical volumes are used, the bulk recall mechanism can only recall up to 13 of these volumes per bulk request. As logical volume sizes increase, the number of volumes recalled per bulk request is reduced. For example, when 6 GB volumes are used, only two volumes per bulk request can be recalled. When 25 GB volumes are used, the benefits of bulk recall are lost given only one volume per bulk request will exceed the default 13 GB limit. This setting allows you to change the 13 GB default to a larger value. The default value is 13 GB and a value from 13 GB to 256 GB can be set. Maximum number of logical volumes recalled per bulk request is 50. More than 50 volumes per bulk request can't not be recalled even if total size of volumes is less than Bulk Recall size Limit. A recall from a host initiated operation can be elongated if it requires the recall of a common logical volume. This is because the host initiated recall will rely on the bulk recall process given it has already started. If the bulk recall is recalling a large amount of data, it may take a while before it recalls the volume the host is attempting to access.

Taking the above condition into account, the value you choose can be determined based on:

- a) How long you are willing to wait for a host job recall if a logical volume collision does occur.
- b) The installed drive model.

The value for BLKRCLSZ can be calculated as follows:

$$\text{BLKRCLSZ} = (\text{maximum time willing to wait for a host recall}) * (\text{estimated data transfer rate of the installed drives})$$

If you can wait for 10 minutes for a host initiated recall due to a collision with a logical volume currently being bulk recalled, and the installed drive type is TS1160, BLKRCLSZ can be set to 120 (= 600 seconds * 200 MB/s / 1000).

The following table shows estimated data transfer rate for each drive type.

Installed drive type	Estimated data transfer rate (MB/s)
TS1160	200
TS1150	180
TS1140	125
TS1130	80
TS1120	52
3592 J1A	20

3.1.13.7 BVIR Settings

If a second keyword of BVIR is specified, the cluster will modify the BVIR volume open timeout value.

The following BVIR settings are supported:

Keyword 3	Keyword 4	Description
TIMEOUT	Value	BVIR timeout value This is the timeout value in minutes for BVIR to complete. Some requests, for example the volume status, may take several hours to complete when many logical volumes exist and may not finish in the allotted time. If a request does not finish in the allotted time, BVIR returns no response data. Prior to Release 4.1.2, the timeout value is fixed at 300 minutes. Starting with Release 4.1.2, the timeout value can be modified. The default value is 300 minutes and a value from 300 minutes to 3000 minutes can be set.

3.1.13.8 CENCRIPT Settings

If a second keyword of CENCRIPT is specified, the cluster will modify the cipher key setting of the Secure Data Transfer function.

The following CENCRIPT settings are supported:

Keyword 3	Keyword 4	Description
CIPHER	DISABLE/ AES128/ AES256	Secure Data Transfer cipher key setting This setting determines what type of the cipher key of the Secure Data Transfer function is used on the target distributed library. When it is set to DISABLE, the target distributed library does not encrypt the data on Grid links. When it is set to AES128 or AES256, the target distributed library uses AES-128 or AES-256 as the cipher key to encrypt the data on Grid links. CENCRIPT is supported in the code level of 8.50.x.x or above. If the distributed library does not have the Secure Data Transfer feature code (FC 5281) enabled, setting AES128/AES256 fails with the following error text: "NO SECURE DATA TRANSFER FEATURE IS INSTALLED"

Keyword 3	Keyword 4	Description
		If it fails to get the current cipher key setting from the vital product data or set the new requested value to the vital product data, the following error text is provided: " FAILED TO GET SECURE DATA TRANSFER SETTING FROM VPD" " FAILED TO UPDATE SECURE DATA TRANSFER SETTING TO VPD" If enabling the encryption function by changing the cipher setting from DISABLE to AES128/AES256 and exporting the encryption key fails, the following error text is provided: " FAILED TO EXPORT ENCRYPTION KEY" Note: When the cipher key setting is changed from the management interface, the corresponding task is created. But when it's changed from the host command line request, no task is created.

3.1.13.9 COPY Settings

If a second keyword of COPY is specified, the cluster will modify the Grid copy function behavior. By using this LI REQ, it's achieved to pause the incoming Deferred and/or COPYRFSH Grid copy jobs.

The following 3rd and 4th keywords are supported:

Keyword 3	Keyword 4	Description
COPYRFSH/ DEF	PAUSE/ UNPAUSE	Pause/Unpause incoming COPYRFSH/Deferred Grid copies TS7700 has supported LI REQ GRIDCNTL command which controls the Grid copy behavior. But it's the knob to stop all incoming as well as outgoing Grid copies. There was a request to implement finer knob to stop specific category of the incoming copy jobs only. This is the knob to control the incoming Grid copies which are categorized in the Deferred and/or COPYRFSH copy jobs. The outgoing copies from the target cluster are still allowed opposed to GRIDCNTL. - When the second keyword 'COPYRFSH' is used, all COPYRFSH incoming copy jobs are suspended when the fourth keyword 'PAUSE' is issued and they're resumed when the fourth keyword 'UNPAUSE' is issued. - When the second keyword 'DEF' is used, all Deferred incoming copy jobs (including Immediate-Deferred/Sync-Deferred/COPYRFSH/Family-Deferred (i.e. except RUN copies)) are suspended with 'PAUSE' fourth keyword and resumed with 'UNPAUSE' fourth keyword. Any copy job which is already in-flight when the LI REQ is issued is not suspended and continues to copy the data from the copy source cluster. No library operational state change occurs when COPYRFSH/Deferred copies are suspended/resumed opposed to GRIDCNTL (when the entire copy is disabled/enabled by GRIDCNTL, the library operational state (Copy Operation Disabled state/Host Copy Operations Disabled state) are changed). But when COPYRFSH/Deferred incoming copy jobs are suspended/resumed, the corresponding events are created and closed with the following text: <pre>"OP0216 Inbound copy operations for the COPYRFSH copy jobs on the local cluster LLLLL are disabled by the host."</pre> <pre>"OP0217 Inbound copy operations for the COPYRFSH copy jobs on the local cluster LLLLL are enabled by the host."</pre>

Keyword 3	Keyword 4	Description
		"OP0218 Inbound copy operations for the deferred copy jobs on the local cluster LLLLL are disabled by the host." "OP0219 Inbound copy operations for the deferred copy jobs on the local cluster LLLLL are enabled by the host." The following pause transition request are supported: Current state -> New state (request) => Result DEF PAUSED -> PAUSE COPYRFSH => COPYRFSH PAUSED (1) COPYRFSH PAUSED -> PAUSE DEF => DEF PAUSED (2) (1) No need to unpause DEF then pause COPYRFSH to pause COPYRFSH only from "DEF PAUSED" state. (2) No Need to unpause COPYRFSH then pause DEF to pause DEF from "COPYRFSH PAUSED" state. The following unpause transition requests are not supported: Current state -> New state (request) => Result DEF PAUSED -> UNPAUSE COPYRFSH => ERROR (3) COPYRFSH PAUSED -> UNPAUSE DEF => ERROR (4) (3) The error text "UNPAUSE COPYRFSH FAILED BECAUSE CURRENTLY DEF IS PAUSED" is returned. (4) The error text "UNPAUSE DEF FAILED BECAUSE CURRENTLY COPYRFSH IS PAUSED" is returned. If updating the new pause/unpause state into the vital product data, the following error text is returned: "FAILED TO SET SETTING2 COPY DATA IN VPD"

3.1.13.10 VOLINVT Settings

If a second keyword of VOLINVT is specified, the cluster will modify behavior of the volume inventory request. At the code level of 8.42.x.x or below, TS7700 doesn't take account of the Library Port Access Group setting when returning the logical volume inventory data. The Library Port Access Group setting has been only applicable to mount the logical volume or update the logical volume attributes (storage constructs/category). At the code level of 8.50.x.x, the host requests to inquiry the logical volume inventory also takes account of the Library Port Access Group setting when this setting is enabled (the default is disabled).

Note: The term "logical volume inventory data" doesn't mean BVIR VOLUME STATUS/LI REQ, LVOL output. The host sends the command to inquiry the volume inventory data (for example, "/d sms, lib(xxxxxx), detail" shows the volume count in the scratch category) and the returned count takes account of Library Port Access Group configuration when this option is enabled.

The target of this LI REQ is a distributed (not composite) library. The setting may need to be enabled on all clusters in the Grid by this LI REQ when this setting needs to be enabled in the Grid.

The following 3rd and 4th keywords are supported:

Keyword 3	Keyword 4	Description
LPAGRP	ENABLE/ DISABLE	Volume Inventory request takes account of Library Port Access Group setting This setting controls if the TS7700 which receives the logical volume inventory request should take account of the Library Port Access Group setting. When it is set to ENABLE, the target distributed library takes account of the Library Port Access Group setting when returning the logical volume inventory data. When it is set to DISABLE, the target distributed library does not take account of the Library Port Access Group setting when returning the logical volume inventory data as same as the code level of 8.42.x.x or below. The default value is DISABLE.

If the target distributed/composite library does not enable SDAC feature yet, the following error text is returned:

```
FAILED TO ENABLE LPAGRP BECAUSE SDAC FEATURE IS NOT ENABLED
```

The response lines to a SETTING2 request are formatted as follows in this example. Just entering the first SETTING2 keyword with no second, third and fourth keywords provides this response:

```
SETTING2 V4 .0
  SCRATCH CONTROLS
  PFRLOCO = ENABLED
```

```
ALERTS
PDRVSLOW =          0 PDRVSCRT =          0
LMTDTHR  =          0
CAGALOW  =          1 CAGAHIGH =          4
CAGLOW   =          0 CAGHIGH  =          0
RSDOLOW  =         100 RSDOHIGH =         300
```

```
CACHE CONTROLS
MAXLGMC  =         128
OVRSPVL  =         2000
```

```
THROTTLE CONTROLS
LNKSPEED =          0
```

```
PHYSLIB CONTROLS
PRETHDEG = DISABLED CPYWDEG = DISABLED TVCWDEG = -----E
SLDPMPRI =  ENABLED MAINT   = DISABLED
```

```
RECALL
BLKRCLSZ =          13
```

```
BVIR
TIMEOUT  =          300
```

```
CENCRYPT CONTROLS
CIPHER   = DISABLED
```

```
COPY CONTROLS
COPYRFSH = UNPAUSED DEF      = UNPAUSED
```

```
VOLINVT CONTROLS
LPAGRP   = DISABLED
```

The following table details the format of the data reported.

Line	Bytes	Name	Description
1	0:9	Header Info	'SETTING2 V'
	10:11	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.32.x.x. The version is incremented to 2 at 8.33.x.x. The version is incremented to 3 at 8.42.x.x The version is incremented to 4 at 8.50.x.x
	21	Dot	'.'
	13:14	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	15:69	Blanks	
2	0	Blank	
	1:16	Header	'SCRATCH CONTROLS'
	17:69	Blanks	
3	0	Blank	
	1:11	Header Info	'PFRLOCO = '
	12:19	Prefer locally owned logical volumes in the scratch mount setting	The current PFRLOCO setting value describing whether a scratch mount prefers locally owned logical volume. Possible values are: 'ENABLED' 'DISABLED'
	20:69	Blanks	
4	0:69	Separator	All dash '-' characters
5	0	Blank	
	1:6	Header	'ALERTS'
	7:69	Blanks	
6	0	Blank	
	1:11	Header Info	'PDRVSLOW = '
	12:19	Available Sunset Physical Drive Low Warning Threshold	The current PDRVSLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	'PDRVSCRT = '
	32:39	Available Sunset Physical Drive Critical Warning Threshold	The current PDRVSCRT setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
7	0	Blank	
	1:11	Header Info	'LMTDTHR = '
	12:19	Limited Cache Threshold	The current LMTDTHR setting value. The value is left justified and padded with blanks.

Line	Bytes	Name	Description
	20:69	Blanks	
8	0	Blank	
	1:11	Header Info	`CAGALOW = '
	12:19		The current CAGALOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`CAGAHIGH = '
	32:39		The current CAGAHIGH setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
9	0	Blank	
	1:11	Header Info	`CAGLOW = '
	12:19		The current CAGLOW setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	`CAGHIGH = '
	32:39		The current CAGHIGH setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
10	0	Blank	
	1:11	Header Info	`RSDOLOW = '
	12:19		The current RSDOLOW setting value. The value is left justified and padded with blanks. (Note: this setting is only displayed in R5.0 and R5.1 code. It was moved to OBJSET1 in R5.22)
	20	Blank	
	21:31	Header Info	`RSDOHIGH = '
	32:39		The current RSDOHIGH setting value. The value is left justified and padded with blanks. (Note: this setting is only displayed in R5.0 and R5.1 code. It was moved to OBJSET1 in R5.22)
	40:69	Blanks	
11	0:69	Separator	All dash '-' characters
12	0	Blank	
	1:14	Header Info	`CACHE CONTROLS'
	15:69	Blanks	
13	0	Blank	
	1:11	Header Info	`MAXLGMC = '
	12:19	Limit the concurrent large volume (25GB) mount count	The current MAXLGMC setting value. The value is left justified and padded with blanks.
	20:69	Blanks	
14	0	Blank	

Line	Bytes	Name	Description
	1:11	Header Info	`OVRSPVL = '
	12:19	Overspill level (GB)	The current OVRSPVL setting value. The value is left justified and padded with blanks.
	20:69	Blanks	
15	0:69	Separator	All dash '-' characters
16	0	Blank	
	1:17	Header Info	`THROTTLE CONTROLS'
	18:69	Blanks	
17	0	Blank	
	1:11	Header Info	`LNKSPEED = '
	12:19	Grid Link Speed to calculate Immediate Copy Throttling	The current LNKSPEED setting value. The value is left justified and padded with blanks.
	20:69	Blanks	
18	0:69	Separator	All dash '-' characters
19	0	Blank	
	1:16	Header Info	`PHYSLIB CONTROLS'
	17:69	Blanks	
20	0	Blank	
	1:11	Header Info	`PRETHDEG = '
	12:19	Premigration Throttling on a TS7700T with the degraded library setting	The current PRETHDEG setting value describing whether a premigration throttling on a TS7700T with the degraded library is disabled or enabled. Possible values are: `ENABLED' `DISABLED'
	20	Blank	
	21:31	Header Info	`CPYWDEG = '
	32:39	Incoming Copy on a TS7700T with the degraded library setting	The current CPYWDEG setting value describing whether incoming copy on a TS7700T with the degraded library is allowed or not. Possible values are: `ENABLED' `DISABLED'
	40	Blank	
	41:51	Header Info	`TVCWDEG = '
	52:59	Mount TVC on a TS7700T with the degraded library setting	The current TVCWDEG setting value describing whether it allows to select the TS7700T with the degraded library as a mount TVC. This is an array of the setting for the eight clusters (C0 – C7). Possible values are: `D' : DISABLE `E' : EQUAL `L' : LOWER

Line	Bytes	Name	Description
			'-': Not applicable. The cluster is not TS7700T or at 8.33.x.x. Example: "EL-----" C0=7700T(8.33.x.x,EQUAL), C1=7700T(8.33.x.x, LOWER), C2=7700T(8.32.x.x), C3=7740, C4=7700D
	60:69	Blanks	
21	0	Blank	
	1:11	Header Info	'SLDMPRI = '
	12:19		The current SLDMPRI setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	'MAINT = '
	32:39		The current MAINT setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
22	0:69	Separator	All dash '-' characters
23	0	Blank	
	1:6	Header Info	'RECALL'
	7:69	Blanks	
24	0	Blank	
	1:11	Header Info	'BLKRCLSZ = '
	12:19	Bulk Recall Size Limit	The current BLKRCLSZ setting value. The value is left justified and padded with blanks.
	20:69	Blanks	
25	0:69	Separator	All dash '-' characters
26	0	Blank	
	1:4	Header Info	'BVIR'
	5:69	Blanks	
27	0	Blank	
	1:11	Header Info	'TIMEOUT = '
	12:19	BVIR Timeout Value	The current TIMEOUT setting value. The value is left justified and padded with blanks.
	20:69	Blanks	
28	0:69	Separator	All dash '-' characters
29	0	Blank	
	1:17	Header Info	'CENCRYPT CONTROLS'
	18:69	Blanks	
30	0	Blank	
	1:11	Header Info	'CIPHER = '

Line	Bytes	Name	Description
	12:19	Encryption Cipher Key Value	The current CIPHER setting value. The value is left justified and padded with blanks.
	20:69	Blanks	
31	0:69	Separator	All dash '-' characters
32	0	Blank	
	1:13	Header Info	'COPY CONTROLS'
	14:69	Blanks	
33	0	Blank	
	1:11	Header Info	'COPYRFSH = '
	12:19	Copyrfs copy job control state	The current COPYRFSH copy job control setting value. The value is left justified and padded with blanks.
	20	Blank	
	21:31	Header Info	'DEF = '
	32:39	Deferred copy job control state	The current DEF copy job control setting value. The value is left justified and padded with blanks.
	40:69	Blanks	
34	0:69	Separator	All dash '-' characters
35	0	Blank	
	1:16	Header Info	'VOLINVT CONTROLS'
	17:69	Blanks	
36	0	Blank	
	1:11	Header Info	'LPAGRP = '
	12:19	Library Port Access Group control state	The current LPAGRP control setting state (DISABLED or ENABLED). The value is left justified and padded with blanks.
	20:69	Blanks	
37	0:69	Separator	All dash '-' characters

3.1.14 STATUS**3.1.14.1 GRID Response:**

In response to the STATUS request where GRID is specified, the cluster specified in the request will examine the current state of its Grid connections and operations. The lines are formatted as follows:

If the request is issued to a cluster configured in a Grid:

```
GRID STATUS V3 .0
```

```
COMPOSITE LIBRARY VIEW
```

		IMMED-DEFERRED		OWNERSHIP-T/O		RECONCILE	HCOPY
LIBRARY	STATE	NUM	MB	MODE	NUM	NUM	ENB
TS001A	ON	0	0	-	0	0	Y
TS001B	SVC	12	8713	SOT	24	12	N
TS001C	ON	0	0	-	0	0	Y

```
-----
--
```

```
COMPOSITE LIBRARY VIEW
```

```
SYNC-DEFERRED
```

LIBRARY	NUM	MB
TS001A	0	0
TS001B	1	1
TS001C	0	0

```
-----
--
```

```
DISTRIBUTED LIBRARY VIEW
```

		RUN-COPY-QUEUE		DEF-COPY-QUEUE		LSTATE	PT	FAM
LIBRARY	STATE	NUM	MB	NUM	MB			
TS001A	ON	5	1493	34	18463	A	Y	1
TS001B	UN	-	-	-	-	-	-	-
TS001C	ON	0	0	17	9518	A	N	1

```
-----
--
```

		ACTIVE-COPIES		CODE-LEVEL
LIBRARY	RUN	DEF		
TS001A	5	20	8.33.2.9	
TS001B	-	-	8.33.2.9	
TS001C	0	17	8.21.0.178	

If the request is issued to a cluster that is not configured in a Grid:

```
GRID STATUS
```

```
CLUSTER NOT PART OF A GRID CONFIGURATION
```

Line	Bytes	Name	Description
1	0:12	Header Info	'GRID STATUS V'

Line	Bytes	Name	Description
	13:14	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x and 3 at the code level 8.33.2.9.
	15	Dot	\ . '
	16:17	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	18:69	Blanks	
If the library the request is sent to is not in a Grid configuration, the following line is returned			
2	0	Blank	
	1:40		'CLUSTER NOT PART OF A GRID CONFIGURATION'
	41:69	Blanks	
If the library the request is sent to is in a Grid configuration, the following lines are returned			
2	0	Blank	
	1:21	Header Info	'COMPOSITE LIBRARY VIEW'
	22:69	Blanks	
3	0:16	Blanks	
	17:30	Header Info	'IMMED-DEFERRED'
	31:32	Blanks	
	33:45	Header Info	'OWNERSHIP-T/O'
	46:47	Blanks	
	48:56	Header Info	'RECONCILE'
	57	Blank	
	58:62	Header Info	'HCOPY'
	63:69	Blanks	
4	0	Blank	
	1:7	Header Info	'LIBRARY'
	8:9	Blanks	
	10:14	Header Info	'STATE'
	15:18	Blanks	
	19:21	Header Info	'NUM'
	22:28	Blanks	
	29:30	Header Info	'MB'
	31:32	Blanks	
	33:36	Header Info	'MODE'
	37:42	Blanks	
	43:45	Header Info	'NUM'
	46:53	Blanks	
	54:56	Header Info	'NUM'

Line	Bytes	Name	Description
	57:59	Blanks	
	60:62	Header Info	'ENB'
	57:69	Blanks	
For each distributed library in the grid configuration a line is formatted as follows:			
N	0:7	Library name (LIBRARY)	The library name defined for the cluster
	8:10	Blanks	
	11:14	Library State (STATE)	The composite view of the state of the distributed library. It is the highest state that any of the clusters in the grid report for any other cluster in the grid. The following are the values that can be indicated: 'ON' The library is in the online state. 'ONP' The library is in the online pending state. 'SVCP' The library is in the service preparation state. 'SVC' The library is in service state. 'OFFP' The library is in the offline-pending state. 'OF' The library is in the offline state. 'UN' The library's state cannot be determined. The value is right justified and padded with blanks. Note: Once the library becomes service state successfully, the state 'SVC' remains though the library is put into offline state. The state persists until the service state is cancelled.
	15:16	Blanks	
	17:21	Immediate Deferred count (IMMED-DEFERRED NUM)	This field indicates the aggregate count of the number of volumes that have a copy policy of RUN but have been deferred for a specific library. The number is right justified and padded with blanks.
	22	Blank	
	23:30	Immediate Deferred MBs (IMMED-DEFERRED MB)	This field indicates the aggregate count of the number of MBs on the volumes that have a copy policy of RUN but have been deferred for a specific library. The number is right justified and padded with blanks.
	31:33	Blanks	
	34:36	Ownership Takeover Mode (OWNERSHIP-T/O MODE)	This field indicates the highest takeover mode for the library. The following are the values that can be indicated: '-' The library is not in a takeover mode. 'SOT' The library is in the service ownership takeover mode. 'ROT' The library is in the read only ownership takeover mode. 'WOT' The library is in the write/read ownership takeover mode. The value is right justified and padded with blanks.
	37:38	Blanks	

Line	Bytes	Name	Description
	39:45	Ownership Takeover Count (OWNERSHIP-T/O MB)	This field indicates the aggregate view from all other libraries of the number of volumes that have had ownership taken from the library. The number is right justified and padded with blanks.
	46:49	Blanks	
	50:56	Reconcile Count (RECONCILE NUM)	This field indicates the aggregate number of volume tokens updates that all libraries have against the library. The number is right justified and padded with blanks.
	57:61	Blanks	
	62	Host Set Copy State (HCOPY ENB)	This field indicates the copy state for the distributed library set by the host using the GRIDCNTL request. 'Y' indicates that the distributed library is enabled as the source for copies and will initiate copies where it is the target. 'N' indicates that the distributed library is disabled for all copy activities.
	63:69	Blanks	
N+1	0:69	Separator	All dash '-' characters
N+2	0	Blank	
	1:21	Header Info	'COMPOSITE LIBRARY VIEW'
	22:69	Blanks	
N+3	0:11	Blanks	
	12:24	Header Info	'SYNC-DEFERRED'
N+4	0	Blank	
	1:7	Header Info	'LIBRARY'
	8:11	Blank	
	12:14	Header Info	'NUM'
	15:21	Blank	
	22:23	Header Info	'MB'
	24:69	Blank	
P	0:7	Library Name (LIBRARY)	The library name defined for the cluster.
	8:11	Blanks	
	12:14	Sync-Def Copy Count (SYNC-DEFERRED NUM)	This field indicates the number of copy jobs in the process of copying SYNC-DEFERRED copy mode data.
	15:18	Blanks	
	19:21	Sync-Deferred Copy Queue MBs (SYNC-DEFERRED MB)	This field indicates the number of MBs on the volumes that are queued to copy to the library with a copy consistency point of sync-deferred. If there are more than 99,999,999 volumes, the display stays at 99,999,999. The number is right justified and padded with blanks.
	22:69	Blanks	

Line	Bytes	Name	Description
P+1	0	Blank	
	1:23	Header Info	'DISTRIBUTED LIBRARY VIEW'
	24:69	Blanks	
P+2	0:16	Blanks	
	17:30	Header Info	'RUN-COPY-QUEUE'
	31:32	Blanks	
	33:46	Header Info	'DEFERRED-COPY-QUEUE'
	47:48	Blanks	
	49:54	Header Info	'LSTATE'
	55	Blank	
	56:57	Header Info	'PT'
	58:59	Blanks	
	60:62	Header Info	'FAM'
	63:69	Blanks	
P+3	0	Blank	
	1:7	Header Info	'LIBRARY'
	8:9	Blanks	
	10:14	Header Info	'STATE'
	15:18	Blanks	
	19:21	Header Info	'NUM'
	22:28	Blanks	
	29:30	Header Info	'MB'
	31:34	Blanks	
	34:37	Header Info	'NUM'
	38:44	Blanks	
	45:46	Header Info	'MB'
	47:69	Blanks	
For each distributed library in the grid configuration a line is formatted as follows:			
M	0:7	Library Name (LIBRARY)	The library name defined for the cluster.
	8:10	Blanks	
	11:14	Library State (STATE)	The state reported by each distributed library. The following are the values that can be indicated: 'ON' The library is in the online state. 'SVCP' The library is in the service preparation state. 'SVC' The library is in service state. 'OFFP' The library is in the offline-pending state. 'OF' The library is in the offline state. 'UN' The library's state cannot be determined.

Line	Bytes	Name	Description
			The value is right justified and padded with blanks.
	15:16	Blanks	
	17:21	Run Copy Queue Count (RUN-COPY-QUEUE NUM)	This field indicates the number of volumes that are queued to copy to the library with a copy consistency point of Rewind/Unload. If there are more than 99,999 volumes, the display stays at 99999. The number is right justified and padded with blanks.
	22	Blank	
	23:30	Run Copy Queue MBs (RUN-COPY-QUEUE MB)	This field indicates the number of MBs on the volumes that are queued to copy to the library with a consistency point of Rewind/Unload. If there are more than 99,999,999 volumes, the display stays at 99,999,999. The number is right justified and padded with blanks.
	31:32	Blanks	
	33:37	Deferred Copy Queue Count (DEFERRED-COPY-QUEUE NUM)	This field indicates the number of volumes that are queued to copy to the library with a copy consistency point of deferred. If there are more than 99,999 volumes, the display stays at 99999. The number is right justified and padded with blanks.
	38	Blank	
	39:46	Deferred Copy Queue MBs (DEFERRED-COPY-QUEUE MB)	This field indicates the number of MBs on the volumes that are queued to copy to the library with a copy consistency point of deferred. If there are more than 99,999,999 volumes, the display stays at 99,999,999. The number is right justified and padded with blanks.
	47:50	Blanks	
	54	Link State Library 0 (LSTATE)	This field indicates the distributed library's view of its link state. The following are the values that can be indicated: _ ' The link state is unknown. \A ' All links are Available. \D ' One or more of the links are Degraded. \U ' All links are Unavailable.
	55:56	Blanks	
	57	Physical Tape Attachment (PT)	This field indicates whether the distributed library has physical tape attached to it. \N ' The distributed library does not have physical tape attached to it. \Y ' The distributed library does have physical tape attached to it,
	58:61	Blanks	
	62	Cluster Family (FAM)	This field indicates if the distributed library has been configured as a member of a cluster family. _ ' The distributed library is not a member of a cluster family. \1 ' - \8 ' The distributed library is a member of cluster family \1 ' - \8 '.
	63:69	Blanks	
M+1	0:69	Separator	All dash '-' characters
M+2	0:8	Blanks	

Line	Bytes	Name	Description
	9:38	Header Info	'ACTIVE-COPIES CODE-LEVELS'
	39:69	Blanks	
M+3	0	Blank	
	1:7	Header Info	'LIBRARY'
	8:11	Blanks	
	12:14	Header Info	'RUN'
	15:18	Blanks	
	19:21	Header Info	'DEF'
	22:69	Blanks	
For each distributed library in the grid configuration a line is formatted as follows:			
L	0:7	Library Name (LIBRARY)	The library name defined for the cluster.
	8:11	Blanks	
	12:14	Active Run Copy Count (ACTIVE-COPIES RUN)	This field indicates the number of copy jobs in the process of copying RUN copy mode data.
	15:18	Blanks	
	19:21	Active Deferred Copy Count (ACTIVE-COPIES DEF)	This field indicates the number of copy jobs in the process of copying DEFERRED copy mode data.
	22	Blank	
	23:38	Distributed Library Code Level (CODE-LEVELS)	This field indicates the Latest Microcode Level installed on the Distributed Library. The level is right justified and padded with blanks.
	39:69	Blanks	

3.1.14.2 GRIDLINK Response:

In response to the STATUS request where GRIDLINK is specified, the cluster specified in the request will examine the current state of its Grid link connections to all of the available clusters in the Grid configuration. The information returned is from the perspective of the cluster receiving the request. A cluster can have up to four grid links. If a link is not installed, it will return 0's for the responses.

Three views of the status of the grid links as well as Secure Data Transfer settings are presented. The Secure Data Transfer settings are provided once the target distributed library is at 8.50.x.x or above.

Link View - A cluster may have up to 4 links included in its configuration for Grid interconnections. This view reports what the link reported that it moved in aggregate to other clusters. It provides information on whether or not the links capability (its speed) is being degraded due to a slower speed setting or capability in the entry point of the network the TS7700 connects to (this could happen if

the TS7700 connects to a switch and the switch only has a 100mb/s capability). The state of the link to the peer TS7700 is indicated as well.

Link Path Latency View - For each grid interconnection, there is a logical path across each link between the cluster reporting the data and its peer TS7700 clusters. For each logical path, there is latency inherent in that path through the customer network. That latency has a direct impact on the peak data transfer capability on the path. This view reports the latency that is captured every 5 minutes.

Cluster View - One important measurement about the quality of the links between the reporting cluster and its peers, is the number of packets that have to be retransmitted. Some amount of packet transmission is normal, but excessive retransmission can significantly degrade the overall copy performance of the cluster. The retransmission values are captured over a 5 minute interval.

Secure Data Transfer View – When the target distributed library is at 8.50.x.x or above, this view is present and provides the Secure Data Transfer function settings (the encryption protocol and the cipher key which are currently used).

Note: TS7700 supports up to four links per cluster since R3.0. If the cluster supports two links, the status for only two links is provided.

The lines are formatted as follows:

If the request is issued to a cluster configured in a Grid (this is an example of a cluster with two installed links configured in a three cluster Grid):

GRIDLINK STATUS V2 .2

CAPTURE TIMESTAMP: 2008-02-08 12:45:32

LINK VIEW

LINK NUM	CFG	NEG	READ MB/S	WRITE MB/S	TOTAL MB/S	ERR	LINK STATE
							01234567
0	1000	1000	87.2	102.4	189.6	0	-AA
1	1000	1000	74.9	104.6	179.5	0	-AA
2	0	0	0.0	0.0	0.0	0	
3	0	0	0.0	0.0	0.0	0	

LINK PATH LATENCY VIEW

LIBRARY	LINK 0	LINK 1	LINK 2	LINK 3
LATENCY IN MSEC				
TS001B	6	7	0	
TS001C	19	20	0	

CLUSTER VIEW

DATA PACKETS SENT:	103948956
DATA PACKETS RETRANSMITTED:	496782
PERCENT RETRANSMITTED:	0.4778

LOCAL LINK IP ADDRESS

LINK 0 IP ADDR:	9.11.200.60
LINK 1 IP ADDR:	9.11.200.61
LINK 2 IP ADDR:	
LINK 3 IP ADDR:	

GRIDLINK STATUS V2 .2

CAPTURE TIMESTAMP: 2019-11-25 05:48:50

LINK VIEW

LINK NUM	CFG	NEG	READ MB/S	WRITE MB/S	TOTAL MB/S	ERR	LINK STATE
							01234567
0	0	1000	0.1	0.1	0.2	0	-AA A
1	0	1000	0.0	0.0	0.0	0	-AA A
2	0	1000	0.0	0.0	0.0	0	-AA
3	0	1000	0.0	0.0	0.0	0	-AA

LINK PATH LATENCY VIEW

LIBRARY	LINK 0	LINK 1	LINK 2	LINK 3
LATENCY IN MSEC				

Lipizzan	0	0	0	0
Arabian	0	0	0	0
Oak	0	0	0	0
Palomino	0	0	0	0

CLUSTER VIEW

DATA PACKETS SENT: 53011
 DATA PACKETS RETRANSMITTED: 0
 PERCENT RETRANSMITTED: 0.0000

LOCAL LINK IP ADDRESS

LINK 0 IP ADDR: 9.11.219.31
 LINK 1 IP ADDR: 9.11.219.158
 LINK 2 IP ADDR: 9.11.219.162
 LINK 3 IP ADDR: 9.11.219.163

SECURE DATA TRANSFER

LIBRARY	PROTOCOL	CIPHER KEY
Lipizzan	TLS1.2	DISABLED
Arabian	TLS1.2	AES-128
Oak	TLS1.2	DISABLED
Palomino	TLS1.2	AES-128

If the request is issued to a cluster that is not configured in a Grid:

GRIDLINK STATUS

CLUSTER NOT PART OF A GRID CONFIGURATION

The following two error messages are supported in 8.33.x.x:

If the Grid link status hasn't been updated for more than 15 minutes because an IBM Service Representative is currently logged in the cluster for maintenance, the following response will be provided:

GRID LINK STATUS NOT AVAILABLE WHILE LIBRARY IS UNDER MAINTENANCE

If the Grid link status hasn't been updated for more than 15 minutes since the last update, the following response will be provided. It may require IBM Support be called to resolve this issue.

GRID LINK STATUS NOT AVAILABLE

The following table details the format of the data reported.

Line	Bytes	Name	Description
1	0:16	Header Info	'GRIDLINK STATUS V'
	17:18	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to 2 at the code level 8.21.x.x.
	19	Dot	'.'
	20:21	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0. The revision is incremented to 1 at 8.33.x.x. The revision is incremented to 2 at 8.50.x.x.
	22:69	Blanks	
If the cluster the request is sent to is not in a Grid configuration, the following line is returned			
2	0	Blank	
	1:40		'CLUSTER NOT PART OF A GRID CONFIGURATION'
	41:69	Blanks	
If the cluster the request is sent to is in a Grid configuration, the following lines are returned			
2	0	Blank	
	1:18		'CAPTURE TIMESTAMP:'
	19	Blank	
	20:38	Timestamp	The timestamp at the end of the last 5 minute sample interval. The timestamp is formatted as follows: YYYY-MM-DD HH:MM:SS for example: 2008-01-23 19:34:23
	39:69	Blanks	
3	0	Blank	
	1:9	Header Info	'LINK VIEW'
	10:69	Blanks	
4	0:1	Blanks	
	2:9	Header Info	'LINK NUM'
	10:13	Blanks	
	14:16	Header Info	'CFG'
	17:19	Blanks	
	20:22	Header Info	'NEG'
	23:26	Blanks	
	27:30	Header Info	'READ'
	31:32	Blanks	
	33:37	Header Info	'WRITE'
	38:39	Blanks	

Line	Bytes	Name	Description
	40:44	Header Info	'TOTAL'
	45	Blank	
	46:48	Header Info	'ERR'
	49:50	Blanks	
	51:60	Header Info	'LINK STATE'
	61:69	Blanks	
5	0:26	Blanks	
	27:30	Header Info	'MB/S'
	31:33	Blanks	
	34:37	Header Info	'MB/S'
	38:40	Blanks	
	41:44	Header Info	'MB/S'
	45:52	Blanks	
	53:60	Header Info	'01234567'
	61:69	Blanks	
For each of the four possible links of the cluster a line is formatted as follows:			
6:9	0:8	Blanks	
	9	Link Number (LINK NUM)	This field indicates the link number being reported. Link numbers start with 0.
	10:11	Blanks	
	12:16	Configured Link Speed (CFG)	This field indicates the configured transfer speed of the link in Mbits/sec. If the link is not included in the configuration, 0 is indicated. The value is right justified and padded with blanks.
	17	Blank	
	18:22	Negotiated Link Speed (NEG)	This field indicates the actual negotiated speed of the link in Mbits/sec. The negotiated speed is only the speed of the link between the TS7700 adapter and the first element of the customer's data network, other links in the network may not support this speed. If the link is not included in the configuration, 0 is indicated. The value is right justified and padded with blanks.
	23:24	Blanks	
	25:30	Read MB/Sec (READ)	This field indicates the data transfer rate of data into the cluster on the link over the last 5 minute sample interval. The data includes only the data where the transfer was initiated by the cluster. The rate is determined by the total number of bytes that the cluster caused to be read from other clusters in the Grid configuration over the last interval divided by 300 and rounded up to the nearest 1/10th of a MB/S. If another cluster initiated the transfer, for example a deferred copy operation or a remote mount, those bytes are not included in this value. The received data bytes include copy and remote mount data transfers initiated by this cluster as well as control and acknowledgement messages. If the link is not included in the configuration, 0.0 is indicated. The value is right justified and padded with blanks.
	31	Blank	

Line	Bytes	Name	Description
	32:37	Write MB/Sec (WRITE)	This field indicates the data transfer rate of data out of the cluster on the link over the last 5 minute sample interval. The data includes only the data where the transfer was initiated by the cluster. The rate is determined by the total number of bytes sent by the cluster to other clusters in the Grid configuration over the last interval divided by 300 and rounded up to the nearest 1/10th of a MB/S. The data written includes copy and remote mount data transfers initiated by this cluster as well as control messages. If the link is not included in the configuration, 0.0 is indicated. The value is right justified and padded with blanks.
	38	Blanks	
	39:44	Total MB/Sec (TOTAL)	This field indicated the combined read and write data transfer rate of the link. The value is right justified and padded with blanks.
	45	Blank	
	46:48	Link Errors (ERR)	This field indicates the number of link adapter detected errors for the link over the last 5 minute sample interval. If the link encountered more that 999 errors during the interval, 999 is indicated. If the link is not included in the configuration, 0 is indicated. The value is right justified and padded with blanks.
	49:52	Blanks	
	53	Link State to Cluster 0 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 0. The following are the values that can be indicated: \ - ' The link state of the reporting Cluster (it doesn't report on its own state). It is also indicated when the link state is unknown. \ A ' The link to the cluster is <u>A</u> ailable. \ D ' The link to the cluster is <u>D</u> egraded. \ U ' The link to the cluster is <u>U</u> navailable. \ ' The library is not part of the grid configuration.
	54	Link State to Cluster 1 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 1. See the definitions for link state Cluster 0.
	55	Link State to Cluster 2 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 2. See the definitions for link state Cluster 0.
	56	Link State to Cluster 3 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 3. See the definitions for link state Cluster 0.
	57	Link State to Cluster 4 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 4. See the definitions for link state Cluster 0.
	58	Link State to Cluster 5 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 5. See the definitions for link state Cluster 0.
	59	Link State to Cluster 6 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 6. See the definitions for link state Cluster 0.
	60	Link State to Cluster 7 (LINK STATE)	This field indicates the link state from the reporting Cluster to Cluster 7. See the definitions for link state Cluster 0.
	61:69	Blanks	
10	0:69	Separator	All dash '-' characters

Line	Bytes	Name	Description
11	0	Blank	
	1:22	Header Info	'LINK PATH LATENCY VIEW'
	23:69	Blanks	
12	0:1	Blanks	
	2:8	Header Info	'LIBRARY'
	9:17	Blanks	
	18:23	Header Info	'LINK 0'
	24:32	Blanks	
	33:38	Header Info	'LINK 1'
	39:47	Blanks	
	48:53	Header Info	'LINK 2'
	54:62	Blanks	
	63:68	Header Info	'LINK 3'
	69	Blank	
13	0:10	Blanks	
	11:25	Header Info	'LATENCY IN MSEC'
	26:69	Blanks	
For each of the clusters attached to this cluster a line is formatted as follows:			
J	0	Blank	
	1:8	Library Name (LIBRARY)	The library name defined for the cluster.
	9:18	Blanks	
	19:23	Link 0 Path Latency (LINK 0)	This field indicates the last test packet transfer latency, measured in millisecond, for the link. Every 5 minutes, the TS7700 sends a test data transfer on each of its active links to each of the available TS7700 peers in the Grid configuration to measure the link's transfer latency to that peer cluster. If a TS7700 peer is unavailable, this value will be 0. The value is right justified and padded with blanks.
	24:33	Blanks	
	34:38	Link 1 Path Latency (LINK 1)	This field indicates the last test packet transfer latency, measured in millisecond, for the link. Every 5 minutes, the TS7700 sends a test data transfer on each of its active links to each of the available TS7700 peers in the Grid configuration to measure the link's transfer latency to that peer cluster. If a TS7700 peer is unavailable, this value will be 0. The value is right justified and padded with blanks.
	39:48	Blanks	
	49:53	Link 2 Path Latency (LINK 2)	This field indicates the last test packet transfer latency, measured in milliseconds, for the link. Every 5 minutes, the TS7700 sends a test data transfer on each of its active links to each of the available TS7700 peers in the Grid configuration to measure the link's transfer latency to that peer cluster. If a TS7700 peer is unavailable, this value will be 0. The value is right justified and padded with blanks.
	54:63	Blanks	

Line	Bytes	Name	Description
	64:68	Link 3 Path Latency (LINK 3)	This field indicates the last test packet transfer latency, measured in millisecond, for the link. Every 5 minutes, the TS7700 sends a test data transfer on each of its active links to each of the available TS7700 peers in the Grid configuration to measure the link's transfer latency to that peer cluster. If a TS7700 peer is unavailable, this value will be 0. The value is right justified and padded with blanks.
	69	Blank	
J+14	0:69	Separator	All dash '-' characters
J+15	0	Blank	
	1:12	Header Info	'CLUSTER VIEW'
	13:69	Blanks	
J+16	0:1	Blanks	
	2:19	Header Info	'DATA PACKETS SENT:'
	20:34	Blanks	
	35:45	Data Packets Sent	This field indicates the number of packets that have been sent from this cluster in the last 5 minute sample interval. It includes all transmitted packets to other TS7700 peers in the Grid configuration as well as transmissions to the library manager, the enterprise key manager and the management interface. The value is left justified and padded with blanks.
	46:69	Blanks	
J+17	0:1	Blanks	
	2:28	Header Info	'DATA PACKETS RETRANSMITTED:'
	29:34	Blank	
	35:45	Data Packets Retransmitted	This field indicates the number of packets that had to be retransmitted from this cluster in the last 5 minute sample interval. It includes all retransmitted packets to other TS7700 peers in the Grid configuration as well as retransmissions to the library manager, the enterprise key manager and the management interface. The value is left justified and padded with blanks.
	46:69	Blanks	
J+18	0:1	Blanks	
	2:23	Header Info	'PERCENT RETRANSMITTED:'
	24:34	Blanks	
	35:41	Percent Retransmitted	This field indicates the percentage of the data packets sent that had to be retransmitted from this cluster in the last 5 minute sample interval, rounded up to the nearest 1/10000th of a percent. A packet retransmission rate of up to 1% may still provide acceptable performance. A high packet retransmission rate can be the result of network errors, quality of service limits or receiving address limitations. The value is left justified and padded with blanks.
	42:69	Blanks	
J+19	0:69	Separator	All dash '-' characters
J+20	0	Blank	
	1:21	Header Info	'LOCAL LINK IP ADDRESS'

Line	Bytes	Name	Description
	22:69	Blanks	
J+21	0:1	Blanks	
	2:16	Header Info	'LINK 0 IP ADDR:'
	17:18	Blanks	
	19:49	Link IP Address	This field indicates the IP address for link 0. The value is left justified and padded with blanks. If the link is not configured, the field will be all blanks.
	50:69	Blanks	
J+22	0:1	Blanks	
	2:16	Header Info	'LINK 1 IP ADDR:'
	17:18	Blanks	
	19:49	Link IP Address	This field indicates the IP address for link 1. The value is left justified and padded with blanks. If the link is not configured, the field will be all blanks.
	50:69	Blanks	
J+23	0:1	Blanks	
	2:16	Header Info	'LINK 2 IP ADDR:'
	17:18	Blanks	
	19:49	Link IP Address	This field indicates the IP address for link 2. The value is left justified and padded with blanks. If the link is not configured, the field will be all blanks.
	50:69	Blanks	
J+24	0:1	Blanks	
	2:16	Header Info	'LINK 3 IP ADDR:'
	17:18	Blanks	
	19:49	Link IP Address	This field indicates the IP address for link 3. The value is left justified and padded with blanks. If the link is not configured, the field will be all blanks.
	50:69	Blanks	
J+25	0:69	Separator	All dash '-' characters
J+26	0:20	Header Info	' SECURE DATA TRANSFER'
	21:69	Blanks	
J+27	0:34	Header Info	' LIBRARY PROTOCOL CIPHER KEY'
	35:69	Blanks	
For each of the clusters attached to this cluster a line is formatted as follows:			
K	0	Blank	
	1:8	Library Name (LIBRARY)	The library name defined for the cluster.
	9:13	Blanks	
	14:21	Encryption Protocol (PROTOCOL)	This field indicates the encryption protocol the library is set. In 8.50.x.x, only TLS1.2 is used.

Line	Bytes	Name	Description
			If the library is at 8.42.x.x or below and does not support Secure Data Transfer function, 'NA' is provided.
	22:24	Blanks	
	25:34	Encryption Cipher Key (CIPHER KEY)	This field indicates the encryption cipher key the library is set. In 8.50.x.x, the following value is provided: DISABLED - the library does not encrypt the data. AES-128 - the library uses AES-128 as the cipher key. AES-256 - the library uses AES-256 as the cipher key. NA - the library is at 8.42.x.x or below and does not support Secure Data Transfer function.
	35:69	Blanks	

3.1.14.3 RDP Response:

In response to a distributed library STATUS request where RDP is specified, Read Diagnostic Parameters (RDP) information for a given FICON port on the targeted cluster is returned. This includes information such as port speeds, port types, temperatures, power usage and other fields that can help diagnose FICON fabric health. Only one port's information can be returned at a time and therefore the particular FICON card and port within that card must be included as the third keyword. In addition to returning information about the local TS7700 ports, the STATUS RDP request can also return remote port information, but it is limited to only the first connected port outside of the TS7700. If not direct attached, any further ports beyond that can have RDP information returned by issuing the RDP command from z/OS host side instead.

Only 16Gb FICON adapters are supported.

A particular FICON card and port must be specified using the third keyword using a format FxPyz where x is a FICON adapter number on the cluster, y is port number of that adapter and z is L for local port or R for remote port. FICON adapter and port instances are zero index based.

For instance, F0P0L represents FICON 0, Port 0 on the cluster. F0P0R represents the first remote port linked to FICON 0, Port 0 on the cluster.

Each FICON adapter in a TS7700 cluster updates RDP information of each port and it's remote port periodically. Each local port is updated hourly while remote port information is updated every four hours. If you would like information to be more current, a fourth keyword REFRESH can be provided which will have the specified adapter and port be refreshed for both local and remote RDP information.

Note: The first RDP request occurs one hour after adapter online processing. Therefore, an RDP request during this first hour will not return valid information. A REFRESH should be issued if RDP information is needed during this first hour of operation.

The lines are formatted as follows:

If the request is issued to FICON 2, PORT 0 for the targeted cluster, an example output would like so:

RDP STATUS V1 .0

READ DIAGNOSTIC PARAMETERS OF LOCAL FICON2:PORT0:

```

LAST SFP READ TIME:          03/23/2018 07:44:13 CUT
ID:                          FICON2:0
WWPN:                        50050764050004CE
ATTACHED WWPN:              2096000533FDEA54
PHYSICAL TYPE:              FC-FS-3
LINK FAILURE ERROR:         0
LOSS OF SYNC ERROR:        0
LOSS OF SIGNAL ERROR:      0
PRIMITIVE SEQUENCE ERROR:  0
INVALID TRANSMISSION WORD ERROR: 0
CRC ERROR:                  0
FEC STATUS:                 ACTIVE
UNCORRECTED BLOCKS:        0
CORRECTED BLOCKS:          0
PORT SPEED CAPABILITIES:    4GB 8GB 16GB
PORT OPERATING SPEED:      16GB
ADVERTISED B-B CREDIT:     90
ATTACHED PORT B-B CREDIT:  8
NOMINAL RTT LINK LATENCY:  0
CONNECTOR TYPE:            SFP+
TX TYPE:                   LONG WAVE - LC

```

-SFP TX/RX FIELD--VALUE-----OPERATING RANGE-----

```

TEMPERATURE:      50.4 C          -5 TO 75 C
BIAS CURRENT:     42.8 MA         1 TO 70 MA
SUPPLY VOLTAGE:   3322.6 MV       3000 TO 3600 MV
RX POWER:        886.4 UW (-0.5 DBM) 10 TO 1995 UW
TX POWER:        1328.4 UW (1.2 DBM) 125 TO 2511 UW

```

-SFP PARAMETERS ALARM FLAGS-----

```

ELEMENT-----HIGH WARNING--LOW WARNING--HIGH ALARM--LOW
ALARM

```

```

TRANSCIEVER TEMPERATURE:  0          0          0          0
TX BIAS CURRENT:          0          0          0          0
TRANSCIEVER SUPPLY VOLTAGE: 0          0          0          0
RX POWER:                 0          0          0          0

```

C

If the request is issued to a composite library the following response will be provided:

REQUEST INVALID FOR COMPOSITE LIBRARY

If the invalid or no second keyword is specified, the following response will be provided:

AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 32768

If the specified FICON adapter does not support Extended Link Service, ELS, the following response will be provided. (8Gb FICON adapter and older ones do not support ELS.)

RDP STATUS V1 .0

READ DIAGNOSTIC PARAMETERS ARE NOT SUPPORTED ON THIS ADAPTER!

The following table details the format of the data reported.

Line	Bytes	Name	Description
1	0:11	Header Info	'RDP STATUS V'
	12:13	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1.
	14	Dot	'.'
	15:16	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	17:69	Blanks	
If the specified FICON adapter does not support Extended Link Service, the following line is returned			
2	0:60		'READ DIAGNOSTIC PARAMETERS ARE NOT SUPPORTED ON THIS ADAPTER!'
	61:69	Blanks	
If the specified FICON adapter supports Extended Link Service, the following line is returned			
2	0:29		'READ DIAGNOSTIC PARAMETERS OF '
	30:64	FICON port	'LOCAL FICONx:PORTy:' (If Keyword-3 is FxPyL) 'REMOTE PORT LINKED TO FICONx:PORTy:' (If Keyword-3 is FxPyR)
	65:69	Blanks	
3	0:1	Blanks	
	2:20	Field Name	'LAST SFP READ TIME:'
	21:34	Blanks	
	35:57	Date and Time	'MM/DD/YYYY hh:mm:ss ZZZ'
	58:69	Blanks	
4	0:1	Blanks	
	2:4	Field Name	'ID:'

Line	Bytes	Name	Description
	5:34	Blanks	
	35:42	Port ID	'FICONx:y' x: adapter number of the cluster y: port number of the adapter
	43:69	Blanks	
5	0:1	Blanks	
	2:6	Field Name	'WWPN:'
	7:34	Blanks	
	35:50	WWPN	'hhhhhhhhhhhhhhhhhh' WWPN of the port
	51:69	Blanks	
6	0:1	Blanks	
	2:15	Field Name	'ATTACHED WWPN:'
	16:34	Blanks	
	35:50	WWPN	'hhhhhhhhhhhhhhhhhh' WWPN of a port linked to the port
	51:69	Blanks	
7	0:1	Blanks	
	2:15	Field Name	'PHYSICAL TYPE:'
	16:34	Blanks	
	35:55	Port Physical Type	One of followings: 'No info' 'FC-FS-3' 'Lossless Ethernet MAC' 'Reserved'
	56:69	Blanks	
8	0:1	Blanks	
	2:20	Field Name	'LINK FAILURE ERROR:'
	21:34	Blanks	
	35:44	Number of Errors	The number of link failure error on the port
	45:69	Blanks	
9	0:1	Blanks	
	2:20	Field Name	'LOSS OF SYNC ERROR:'
	21:34	Blanks	
	35:44	Number of Errors	The number of loss of sync error on the port
	45:69	Blanks	
10	0:1	Blanks	
	2:22	Field Name	'LOSS OF SIGNAL ERROR:'

Line	Bytes	Name	Description
	23:34	Blanks	
	35:44	Number of Errors	The number of loss of signal error on the port
	45:69	Blanks	
11	0:1	Blanks	
	2:26	Field Name	'PRIMITIVE SEQUENCE ERROR:'
	27:34	Blanks	
	35:44	Number of Errors	The number of primitive sequence error on the port
	45:69	Blanks	
12	0:1	Blanks	
	2:33	Field Name	'INVALID TRANSMISSION WORD ERROR:'
	34	Blank	
	35:44	Number of Errors	The number of invalid transmission word error on the port
	45:69	Blanks	
13	0:1	Blanks	
	2:11	Field Name	'CRC ERROR:'
	12:34	Blanks	
	35:44	Number of Errors	The number of CRC error on the port
	45:69	Blanks	
14	0:1	Blanks	
	2:12	Field Name	'FEC STATUS:'
	13:34	Blanks	
	35:42	FEC Status	Status of "Forward Error Correction" setting on the port 'ACTIVE' or 'INACTIVE'
	43:69	Blanks	
15	0:1	Blanks	
	2:20	Field Name	'UNCORRECTED BLOCKS:'
	21:34	Blanks	
	35:44	Number of Blocks	The number of FEC uncorrected blocks Only when FEC STATUS is ACTIVE
	45:69	Blanks	
16	0:1	Blanks	
	2:18	Field Name	'CORRECTED BLOCKS:'
	19:34	Blanks	
	35:44	Number of Blocks	The number of FEC corrected blocks Only when FEC STATUS is ACTIVE
	45:69	Blanks	
17	0:1	Blanks	
	2:25	Field Name	'PORT SPEED CAPABILITIES:'
	26:34	Blanks	
	35:50	Capable Port Speeds	Row of capable port speeds of the port

Line	Bytes	Name	Description
	51:69	Blanks	
18	0:1	Blanks	
	2:22	Field Name	'PORT OPERATING SPEED:'
	23:34	Blanks	
	35:50	Operating Port Speed	Operating port speed
	51:69	Blanks	
19	0:1	Blanks	
	2:23	Field Name	'ADVERTISED B-B CREDIT:'
	24:34	Blanks	
	35:44	B-B Credit	Advertised BtoB credit of the port
	45:69	Blanks	
20	0:1	Blanks	
	2:26	Field Name	'ATTACHED PORT B-B CREDIT:'
	27:34	Blanks	
	35:44	B-B Credit	BtoB credit of a port linked to the specified port
	45:69	Blanks	
21	0:1	Blanks	
	2:26	Field Name	'NOMINAL RTT LINK LATENCY:'
	27:34	Blanks	
	35:44	Latency	Nominal RTT link latency
	45:69	Blanks	
22	0:1	Blanks	
	2:16	Field Name	'CONNECTOR TYPE:'
	17:34	Blanks	
	35:42	Connector Type	One of followings: 'Unknown' 'SFP+' 'Reserved'
	43:69	Blanks	
23	0:1	Blanks	
	2:9	Field Name	'TX TYPE:'
	10:34	Blanks	
	35:50	Transmitter Type	Transmitter type of SFP. One of followings: 'Short Wave Laser' 'Long Wave - LC' 'Long Wave - LL' 'Other'
	51:69	Blanks	
24	0	Blank	
	1:	Hyphen	'-'
	2:16	Column Title	'SFP TX/RX FIELD'

Line	Bytes	Name	Description
	17:18	Hyphen	'--'
	19:23	Column Title	'VALUE'
	24:42	Hyphen	'-----'
	43:57	Column Title	'OPERATING RANGE'
	58:66	Hyphen	'--'
	67:69	Blanks	
25	0:1	Blanks	
	2:13	Field Name	'TEMPERATURE: '
	14:18	Blanks	
	19:26	Temperature	Temperature of SFP [C]
	27:43	Blanks	
	44:59	Operating Range	Temperature range of operation Min TO Max C
	60:69	Blanks	
26	0:1	Blanks	
	2:14	Field Name	'BIAS CURRENT: '
	15:18	Blanks	
	19:27	Current	Bias current of transmitter [mA]
	28:43	Blanks	
	44:59	Operating Range	Current range of operation Min TO Max MA
	60:69	Blanks	
27	0:1	Blanks	
	2:16	Field Name	'SUPPLY VOLTAGE: '
	17:18	Blanks	
	19:27	Voltage	Voltage supplied on SFP [mV]
	28:43	Blanks	
	44:59	Operating Range	Voltage range of operation Min TO Max MV
	60:69	Blanks	
28	0:1	Blanks	
	2:10	Field Name	'RX POWER: '
	11:18	Blanks	
	19:27	Power	Receiver power on SFP [uW]
	28	Blank	
	29:40	Power	Receiver power on SFP [dBm] (N.N DBM)
	41:43	Blanks	
	44:59	Operating Range	Receiver power range of operation

Line	Bytes	Name	Description
			Min TO Max UW
	60:69	Blanks	
29	0:1	Blanks	
	2:10	Field Name	'TX POWER: '
	11:18	Blanks	
	19:27	Power	Transmitter power on SFP [uW]
	28	Blank	
	29:40	Power	Transmitter power on SFP [dBm] (N.N DBM)
	41:43	Blanks	
	44:59	Operating Range	Transmitter power range of operation Min TO Max UW
	60:69	Blanks	
30	0	Blank	
	1	Hyphen	'-'
	2:27	Table Title	'SFP PARAMETERS ALARM FLAGS'
	28:33	Hyphens	'-----'
	34:69	Blanks	
31	0:1	Blanks	
	2:8	Column Name	'ELEMENT'
	9:19	Hyphens	'-----'
	20:31	Column Name	'HIGH WARNING'
	32:33	Hyphens	'--'
	34:44:	Column Name	'LOW WARNING'
	45:46	Hyphens	'--'
	47:56	Column Name	'HIGH ALARM'
	57:58	Hyphens	'--'
	59:67	Column Name	'LOW ALARM'
	68:69	Blanks	
32	0:1	Blanks	
	2:25	Element	'TRANSCEIVER TEMPERATURE: '
	26:29	Blanks	
	30	High Warning	0 if temperature is under high warning threshold. 1 if temperature is over high warning threshold.
	31:40	Blanks	
	41	Low Warning	0 if temperature is over low warning threshold. 1 if temperature is under low warning threshold.
	42:51	Blanks	

Line	Bytes	Name	Description
	52	High Alarm	0 if temperature is under high alarm threshold. 1 if temperature is over high alarm threshold.
	53:62	Blanks	
	63	Low Alarm	0 if temperature is over low alarm threshold. 1 if temperature is under low alarm threshold.
	64:69	Blanks	
33	0:1	Blanks	
	2:17	Element	'TX BIAS CURRENT: '
	18:29	Blanks	
	30	High Warning	0 if current is under high warning threshold. 1 if current is over high warning threshold.
	31:40	Blanks	
	41	Low Warning	0 if current is over low warning threshold. 1 if current is under low warning threshold.
	42:51	Blanks	
	52	High Alarm	0 if current is under high alarm threshold. 1 if current is over high alarm threshold.
	53:62	Blanks	
	63	Low Alarm	0 if current is over low alarm threshold. 1 if current is under low alarm threshold.
	64:69	Blanks	
34	0:1	Blanks	
	2:28	Element	'TRANSCIVER SUPPLY VOLTAGE: '
	29	Blank	
	30	High Warning	0 if voltage is under high warning threshold. 1 if voltage is over high warning threshold.
	31:40	Blanks	
	41	Low Warning	0 if voltage is over low warning threshold. 1 if voltage is under low warning threshold.
	42:51	Blanks	
	52	High Alarm	0 if voltage is under high alarm threshold. 1 if voltage is over high alarm threshold.
	53:62	Blanks	
	63	Low Alarm	0 if voltage is over low alarm threshold. 1 if voltage is under low alarm threshold.
	64:69	Blanks	
35	0:1	Blanks	
	2:10	Element	'RX POWER: '

Line	Bytes	Name	Description
	11:29	Blanks	
	30	High Warning	0 if power is under high warning threshold. 1 if power is over high warning threshold.
	31:40	Blanks	
	41	Low Warning	0 if power is over low warning threshold. 1 if power is under low warning threshold.
	42:51	Blanks	
	52	High Alarm	0 if power is under high alarm threshold. 1 if power is over high alarm threshold.
	53:62	Blanks	
	63	Low Alarm	0 if power is over low alarm threshold. 1 if power is under low alarm threshold.
	64:69	Blanks	
36	0:1	Blanks	
	2:10	Element	`TX POWER: '`
	11:29	Blanks	
	30	High Warning	0 if power is under high warning threshold. 1 if power is over high warning threshold.
	31:40	Blanks	
	41	Low Warning	0 if power is over low warning threshold. 1 if power is under low warning threshold.
	42:51	Blanks	
	52	High Alarm	0 if power is under high alarm threshold. 1 if power is over high alarm threshold.
	53:62	Blanks	
	63	Low Alarm	0 if power is over low alarm threshold. 1 if power is under low alarm threshold.
	64:69	Blanks	

Note: If values could not be obtained, they are shown as '-'.

Best Practices: The following guidelines can be used to help determine if a fabric or port is running outside of a normal range. These are only basic guidelines, so exceptions may occur.

Cables should be cleaned up when the error count is larger than in the table below (growing larger over a 24 hour period) or the Rx received power level is lower than -3 dBm.

If after cleaning all cable segments and ports the received power level is still lower than -3 dBm, then the SFPs at the end of the paths should be replaced.

If after replacing the SFPs the received power is still below -3 dBm, then the card in the Hydra could be replaced , after that the card in the zSeries , or the blade in the switch could be replaced.

Errors value guidelines over 24 hour period:

Link Failures over 2 (monitor) over 6 take action

Loss of Sync over 4 (monitor) over 10 take action

Loss of Signal over 6 (monitor) over 12 take action

Primitive Seq Err over 1 (monitor) over 3 take action

Invalid Trans Word over 30 (monitor) over 100 take action

Invalid CRC over 1 (monitor) over 3 take action

FEC Uncorrected over 100 (monitor) over 500 take action

Rx Receive power if positive it is fine. If a small negative value of -1 or 2 that is acceptable. If it falls below -3dBm it should be corrected.

-6 dBm will definitely cause problems.

3.1.14.4 **GRLNKACT Response:**

In response to the STATUS request where GRLNKACT is specified, the cluster specified in the request will collect the point-in-time Grid link activity to all of the available clusters in the Grid configuration as well as cloud if the cluster is attached to cloud (even if the cluster is a standalone configuration, the request is accepted and provides the Grid link activity to cloud). The information returned is from the perspective of the cluster receiving the request.

This request internally refers to TS7700 Statistical Data (Hnode Grid Point in Time Record) to calculate the network activity to cloud. The record is summated and reset at each 15 seconds interval (i.e. at xx:yy:00, xx:yy:15, xx:yy:30, xx:yy:45). In order to be sync with it, Grid link activity is gathered and calculated for the next set of 15 seconds interval. For example, when LI REQ is issued at 00:10:02, Grid link activity from 00:10:15 to 00:10:30 are gathered and provided (the timestamp is internal in the cluster).

Four views of the Grid link activity are presented.

GRID LINK ESTABLISHED SOCKET CONNECTIONS - This view reports how many TCP sockets are currently established from each Grid interface on the local clusters to the peer cluster as well as cloud. It also provides the total established TCP connections which are used for MQ (Grid cluster-cluster Websphere message communication) and RFA (Grid cluster-cluster data (file) transfer). Starting with R5.1, the total established TCP connections used for GGM (Grid to Grid Migration) and DS8000 Transparent Cloud Tiering is shown as well.

- The number of the socket connections are provided per each Grid interface (L0:Primary, L1:Alternate, L2:Primary2, L3:Alternate2) with the actual resource name (enX) and IP

addresses. Cx columns show the connections used for Grid communications with the peer cluster Cx (Cx means cluster ID x).

- MQ column shows the total sum of the connections used for MQ.
- RFA column shows the total sum of the connections used for RFA.
- CLD column shows the total sum of the connections used for cloud. In the code level prior to R5.1, this sum might also include the ones for GGM activity with another copy source/target Grid if it's configured and used. On R5.1 and up, the sum for GGM are separated to the GGM column below.
- GGM column shows the total sum of the connections that are used for GGM.
- OBJ column shows the total sum of the connections that are used for DS8000 Transparent Cloud Tiering.

Note: The established socket may be left even after the actual communication completes for the packet reuse.

NET ACTIVITY - This view reports how much data (MB) has been sent or received on each Grid interface in the last 15 seconds interval. It also provides the current network request counts in Grid interface buffer. They are summated for MQ/RFA/Cloud/GGM/DS8000 Transparent Cloud Tiering usage. This view is also reported based off of each Grid interface as same as the view of established socket connections.

TxMBs shows the total transmitted network activity from the Grid interface in the last 15 seconds. The unit is "MB".

RxMBs shows the total received network activity into the Grid interface in the last 15 seconds. The unit is "MB".

MQ_REC/MQ_SND show the current network request (receiving/sending) counts in each Grid network interface buffer related for MQ. If the value is non-0, it must indicate there is a network activity for MQ activity.

GFA_REC/GFA_SND: The same counts related for GFA.

CLD_REC/CLD_SND: The same counts related for cloud. In the code level prior to R5.1, this may also include the ones for GGM activity as same as explained in the established socket connection view section. On R5.1 and up, the counts for GGM are separated to GGM_REC/GGM_SND columns below.

GGM_REC/GGM_SND: The same counts that are related for GGM.

OBJ_REC/OBJ_SND: The same counts that are related for DS8000 Transparent Cloud Tiering.

GRID LINK THROUGHPUT ESTIMATES - This view reports the estimated Grid link throughput estimation in the last 15 seconds interval. The values except with cloud are retrieved from TS7700 Statistical Data.

- Tx/Rx provide the transmitted network throughput to the remote target/received network throughput from the remote target. The unit is "MB/s".
- Cx columns mean the cluster ID.
- GRD_TOT column shows the total throughput with the remote clusters (Cx) in the Grid.

- CLD column shows the throughput against the cloud if configured and used.
- GGM column shows the throughput of GGM copy with the remote grid.
- OBJ column shows the throughput of DS8000 Transparent Cloud Tiering.
- TOT column shows the sum of GRD_TOT, CLD, GGM and OBJ.

(*) In the code level prior to R5.1, GGM activity can be added to CLD and/or TOT as same as explained in the established socket connection view section. On R5.1 and up, the GGM activity is separated to GGM column.

GRID CLOUD TIER EXPORT AND IMPORT ACTIVITY - This view reports the current cloud data premigration (export) and recall (import) status. The active premigration/recall volumes to/from cloud are provided.

- ACTIVE EXPORT VOLUME COUNT provides the total number of the active export (premigration) volumes. If it is non-0, the corresponding volsers are provided below.
- ACTIVE IMPORT VOLUME COUNT provides the total number of the active import (recall) volumes. If it is non-0, the corresponding volsers are provided below.

If it fails to retrieve the cluster ID internally, it returns:

```
INTERNAL ERROR DETECTING LOCAL CLUSTER INDEX
PLEASE TRY AGAIN.    RC=X
```

If it fails to retrieve the number of Grid interfaces internally, it returns:

```
INTERNAL ERROR DETECTING GRID IP PORT COUNT
PLEASE TRY AGAIN.    RC=X
```

If it detects the incorrect number of Grid interfaces on local internally, it returns:

```
INVALID LINK COUNT DETECTED:  COUNT=Z
PLEASE TRY AGAIN.    RC=X
```

If it detects the invalid cluster ID internally, it returns:

```
INVALID CLUSTER INDEX DETECTED:  COUNT=Z
PLEASE TRY AGAIN.    RC=X
```

If it fails to detect the primary/alternate/primary2/alternate2 Grid interface on local internally, it returns:

```
INTERNAL ERROR DETECTING <ADAPTER> CARD, <LOCATION> LINK
PLEASE TRY AGAIN.    RC=X
```

or

```
INTERNAL ERROR DETECTING <ADAPTER> CARD INTERFACE, FIRST LINK
PLEASE TRY AGAIN.    RC=X
```

or

INTERNAL ERROR DETECTING <ADAPTER> CARD LOCATION, FIRST LINK
PLEASE TRY AGAIN. RC=X

(*) <ADAPTER> is either "PRI" or "ALT". <LOCATION> is either "FIRST" or "SECOND".

(*) Primary Grid interface is "PRI" and "FIRST". Alternate Grid interface is "ALT" and "FIRST". Primary 2 Grid interface is "PRI" and "SECOND". Alternate 2 Grid interface is "ALT" and "SECOND".

If it fails to get TS7700 Statistical Data (Hnode Grid Point in Time Record) internally, it returns:

INTERNAL ERROR DETECTING INVALID STATS DATA
PLEASE TRY AGAIN.

If it detects the invalid cluster count internally, it returns:

INTERNAL ERROR DETECTING CLUSTER COUNT
PLEASE TRY AGAIN. RC=X

If it fails to get network activity internally, it returns:

INTERNAL ERROR RUNNING NETSTAT
PLEASE TRY AGAIN. RC=X

If it fails to get the established TCP socket connections internally, it returns:

INTERNAL ERROR GETTING ESTABLISHED CONNECTIONS
PLEASE TRY AGAIN. RC=X

If it fails to detect the clusters in the Grid internally, it returns:

INTERNAL ERROR DETECTING CLUSTER LIST
PLEASE TRY AGAIN. RC=X

If it detects invalid cluster ID internally, it returns:

INTERNAL ERROR, INVALID CLUSTER INDEX: X

If it detects invalid Grid interface link count on remote, it returns:

INTERNAL ERROR DETECTING LINK COUNT: X

If it detects the incorrect number of Grid interfaces on remote, it returns:

CLUSTER INDEX X HAS INVALID LINK COUNT: Z

If it fails to detect the primary/alternate/primary2/alternate2 Grid interface on remote, it returns:

INTERNAL ERROR DETECTING <ADAPTER> CARD, <LOCATION> LINK, CLZ
PLEASE TRY AGAIN. RC=X

(*) <ADAPTER> is either "PRI" or "ALT". <LOCATION> is either "FIRST" or "SECOND".

If any other unexpected internal error occurs, return code -1 is returned:

FAILED TO GET GRID LINK ACTIVITY RC=-1

If all result isn't contained due to the maximum lines of the output, the last line (50th line) has:

>>>> ADDITIONAL INFORMATION EXISTS

Starting with R5.1, the response lines are formatted as follows:

GRLNKACT STATUS V1 .0

CLUSTER INDEX: 1 LINK COUNT: 4 Time: Thu Nov 5 07:39:24 CUT 2020

GRID LINK ESTABLISHED SOCKET CONNECTIONS-----

LN	INTF	IP	C0	C1	C2	C3	C4	C5	C6	C7	MQ	RFA	CLD
L0	en0	10.30.1.21	108	0	0	0	0	0	0	0	108	0	0
L1	en8	10.31.1.21	0	0	0	0	0	0	0	0	0	0	0
L2	en1	10.32.1.21	0	0	0	0	0	0	0	0	0	0	0
L3	en9	10.33.1.21	0	0	0	0	0	0	0	0	0	0	0

LN	INTF	IP	GGM	OBJ
L0	en0	10.30.1.21	0	0
L1	en8	10.31.1.21	0	0
L2	en1	10.32.1.21	0	1
L3	en9	10.33.1.21	0	0

NET ACTIVITY -----TCP RECV/SEND ADAPTER BUFFER ACTIVITY BYTES-----

LN	TxMBs	RxMBs	MQ_REC	MQ_SND	GFA_REC	GFA_SND	CLD_REC	CLD_SND
L0	0	0	0	444	0	0	0	0
L1	0	0	0	0	0	0	0	0
L2	0	11	0	0	0	0	0	0
L3	0	0	0	0	0	0	0	0
TOT	1	11	-	-	-	-	-	-

LN	GGM_REC	GGM_SND	OBJ_REC	OBJ_SND
L0	0	0	0	0
L1	0	0	0	0
L2	0	0	11661	0
L3	0	0	0	0

GRID LINK THROUGHPUT ESTIMATES-MB/s-----

DIR	C0	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	11	11

GRID CLOUD TIER EXPORT AND IMPORT ACTIVITY-----

ACTIVE EXPORT VOLUME COUNT: 40

CL0651	CL0671	CL0677	CL0679	CL0681	CL0683	CL0685	CL0687
CL0689	CL0691	CL0693	CL0695	CL0697	CL0699	CL0701	CL0703
CL0705	CL0707	CL0709	CL0711	CL0713	CL0715	CL0717	CL0719
CL0721	CL0723	CL0725	CL0727	CL0729	CL0730	CL0731	CL0732
CL0733	CL0735	CL0737	CL0739	CL0741	CL0743	CL0745	CL0749
ACTIVE IMPORT VOLUME COUNT: 0							

The following table details the format of the data reported.

Line	Bytes	Name	Description
1	0:16	Header Info	'GRLNKACT STATUS V'
	17:18	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1.
	19	Dot	'.'
	20:21	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	22:69	Blanks	
2	0:15	Header Info	'CLUSTER INDEX: '
	16	Cluster ID	The cluster ID (0-7) which received the request
	17	Blank	
	18:30	Header Info	'LINK COUNT: '
	31	Grid interface link count	The number of configured Grid interfaces on the cluster received this request.
	32	Blank	
	33:38	Header Info	'Time: '
	39:66	Timestamp	The timestamp when the contents of the response is finalized. The timestamp is formatted as follows: WWW MMM DD HH:MM:SS CUT YYYY 'Mon Aug 27 06:17:54 CUT 2018'
	67:69	Blanks	
3	0:69	Header Info	'GRID LINK ESTABLISHED SOCKET CONNECTIONS-----'
4	0:69	Header Info	'LN INTF IP C0 C1 C2 C3 C4 C5 C6 C7 MQ RFA CLD '
5	0:2	Header Info	'L0 '
	3:6	Grid interface resource name (Primary)	Grid primary interface resource name on local cluster: "enxx".
	7	Blank	
	8:22	Grid interface IP address (Primary)	Grid primary interface IP address
	23:24	Blanks	

Line	Bytes	Name	Description
15	0:69	Header Info	`LN TxMBs RxMBs MQ_REC MQ_SND GFA_REC GFA_SND CLD_REC CLD_SND `
16	0:3	Header Info	`L0 `
	4:8	Amount of the transmitted data on primary interface	The amount of the transmitted data from the primary interface on the local cluster in the last 15 seconds.
	9	Blank	
	10:14	Amount of the received data on primary interface	The amount of the received data to the primary interface on the local cluster in the last 15 seconds.
	15	Blank	
	16:23	Network request counts to receive related to MQ on primary interface	The current network request counts to receive in the primary interface buffer related to MQ on the local cluster.
	24	Blank	
	25:32	Network request counts to send related to MQ on primary interface	The current network request counts to send in the primary interface buffer related to MQ on the local cluster.
	33	Blank	
	34:41	Network request counts to receive related to RFA on primary interface	The current network request counts to receive in the primary interface buffer related to RFA on the local cluster.
	42	Blank	
	43:50	Network request counts to send related to RFA on primary interface	The current network request counts to send in the primary interface buffer related to RFA on the local cluster.
	51	Blank	
	52:59	Network request counts to receive related to cloud on primary interface	The current network request counts to receive in the primary interface buffer related to cloud on the local cluster. In R4.2, this also includes the counts related to GGM.
	60	Blank	
	61:68	Network request counts to send related to cloud on primary interface	The current network request counts to send in the primary interface buffer related to cloud on the local cluster. In R4.2, this also includes the counts related to GGM.
	69	Blank	
17:19	Same network activity on alternate/primary2/alternamte2 interfaces with the 11 th line. The primary/alternate interfaces are only provided if only 2 Grid interfaces are configured.		
20	0:3	Header Info	`TOT `
	4:9	Total amount of the transmitted data on primary interface	The total amount of the transmitted data from the primary interface on the local cluster in the last 15 seconds.

Line	Bytes	Name	Description
	10	Blank	
	11:15	Total amount of the received data on primary interface	The total amount of the received data to the primary interface on the local cluster in the last 15 seconds.
	16:69	Header Info	`- - - - - - , - - - - -`
21	0:69	Header Info	`LN GGM_REC GGM_SND OBJ_REC OBJ_SND ,`
22	0:3	Header Info	`L0 `
	4:11	Network request counts to receive related to GGM copy on primary interface	The current network request counts to receive in the primary interface buffer related to GGM copy on the local cluster.
	12	Blank	
	13:20	Network request counts to send related to GGM copy on primary interface	The current network request counts to send in the primary interface buffer related to GGM copy on the local cluster.
	21	Blank	
	22:29	Network request counts to receive related to DS8000 Transparent Cloud Tiering on primary interface	The current network request counts to receive in the primary interface buffer related to DS8000 Transparent Cloud Tiering on the local cluster.
	30	Blank	
	31:38	Network request counts to send related to DS8000 Transparent Cloud Tiering on primary interface	The current network request counts to send in the primary interface buffer related to DS8000 Transparent Cloud Tiering on the local cluster.
	39:69	Blanks	
23:25	Same network activity on alternate/primary2/alternamte2 interfaces with the 22 th line. The primary/alternate interfaces are only provided if only 2 Grid interfaces are configured.		
26	0:69	Header Info	`GRID LINK THROUGHPUT ESTIMATES-MB/s----- -----`
27	0:69	Header Info	`DIR C0 C2 C3 C4 C5 C6 C7 GRD_TOT CLD GGM OBJ TOT `
			The "Cx" columns indicates the cluster ID from 0 to 7, where the local cluster is not included.
28	0:3	Header Info	`Tx `
	4:7	Grid link throughput estimation for the transmitted data to Cx	Estimated Grid link throughput for the transmitted data from the local cluster to Cx in the last 15 seconds. The "Cx" indicates the cluster ID from 0 to 7, where the local cluster is not included. The data is calculated from the TS7700 Statistical Data (Hnode Grid Point in Time Record).
	8	Blank	

Line	Bytes	Name	Description
	9:37	The estimated Grid link throughput for the transmitted data to remote clusters as same as above.	
	38	Blank	
	39:45	Total Grid link throughput estimation for the transmitted data to all peer clusters in the Grid	This is the sum of Grid link throughput for the transmitted data from the local cluster to all peer clusters in the last 15 seconds.
	46	Blank	
	47:50	Throughput estimation for the transmitted data to cloud	Estimated throughput for the transmitted data from the local cluster to cloud in the last 15 seconds.
	51	Blank	
	52:55	Throughput estimation for the transmitted data by GGM copy	Estimated throughput for the transmitted data from the local cluster to remote grid by GGM copy in the last 15 seconds.
	56	Blank	
	57:60	Throughput estimation for the transmitted data by DS8000 Transparent Cloud Tiering	Estimated throughput for the transmitted data from the local cluster to DS8000 by DS8000 Transparent Cloud Tiering in the last 15 seconds.
	61	Blank	
	62:68	Total Grid link throughput estimation for the transmitted data	This is the sum of all Grid link throughput for the transmitted data from the local cluster in the last 15 seconds.
	69	Blank	
29	0:3	Header Info	'Rx '
	4:7	Grid link throughput estimation for the received data to Cx	Estimated Grid link throughput for the received data to the local cluster from Cx in the last 15 seconds. The "Cx" indicates the cluster ID from 0 to 7, where the local cluster is not included. The data is calculated from the TS7700 Statistical Data (Hnode Grid Point in Time Record).
	8	Blank	
	9:37	The estimated Grid link throughput for the received data from remote clusters as same as above.	
	38	Blank	
	39:45	Total Grid link throughput estimation for the received data to all peer clusters in the Grid	This is the sum of Grid link throughput for the received data to the local cluster from all peer clusters in the last 15 seconds.
	46	Blank	

Line	Bytes	Name	Description
	47:50	Grid link throughput estimation for the received data to cloud	Estimated Grid link throughput for the received data to the local cluster from cloud in the last 15 seconds.
	51	Blank	
	52:55	Throughput estimation for the received data by GGM copy	Estimated throughput for the received data to the local cluster from remote grid by GGM copy in the last 15 seconds.
	56	Blank	
	57:60	Throughput estimation for the received data by DS8000 Transparent Cloud Tiering	Estimated throughput for the received data to the local cluster from DS8000 by DS8000 Transparent Cloud Tiering in the last 15 seconds.
	61	Blank	
	62:68	Total Grid link throughput estimation for the received data	This is the sum of all the throughputs for the received data to the local cluster in the last 15 seconds.
	69	Blank	
30	0:69	Header Info	'GRID CLOUD TIER EXPORT AND IMPORT ACTIVITY-- -----'
31	0:27	Header Info	'ACTIVE EXPORT VOLUME COUNT: '
	28:30	Active export volume count	The number of current active export (premig) volume count to cloud.
	31:69	Blanks	
32:X	0:54	Active export volumes	The current active export (premig) volumes to cloud are provided if exist. Each volume is separated by a space and up to 8 volumes are provided in a line.
33	0:27	Header Info	'ACTIVE IMPORT VOLUME COUNT: '
	28:30	Active import volume count	The number of current active import (recall) volume count from cloud.
	31:69	Blanks	
34:X	0:54	Active import volumes	The current active import (recall) volumes from cloud are provided if exist. Each volume is separated by a space and up to 8 volumes are provided in a line.

3.1.15 COPYRFSH

The COPYRFSH request provides a method to refresh copy policy without mounting/demounting a volume. The copy policy defined in the management class of the volume on the cluster associated with the distributed library in the request is assigned. The request queues a copy job on cluster(s) which need a copy of the volume and the copy job queued using this request will have a lower priority than newly created workload. No recalls will take place when the library processes the request. Instead, the efficient recalls for copies will happen as needed when the queued copy job is processed. The retain copy mode option will be ignored when the library processes the request.

Note1: The request is supported only when all the clusters in the grid have the code level of 8.21.0.118 (R2.1 PGA) or above.

Note2: A tool COPYRFSH which leverages the Library Request is available in IBMTTOOLS and it automates the process to refresh copy policy. PRESTAGE is no longer necessary to refresh copy policy by mounting/demounting volumes.

Note3: A third keyword CLDGHOST is supported at 8.50.x.x (R5.0).

The following table lists the keywords supported for the COPYRFSH request.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
COPYRFSH	zzzzzz	x (copy source cluster ID (0-7)) (optional) or CLDGHOST (optional)	NORECALL (optional) (*) When CLDGHOST is used in the 3rd KW, the 4th KW is unusable.	N/A	Y	N/A	Y	Y	Y

A second keyword is mandatory and the logical volume which will have a new copy policy.

A third and fourth keywords are optional and used to specify how a copy job will be attempted. A third keyword is the preferred copy source cluster ID (from 0 to 7). If it is specified, copy target cluster(s) will try to get the data from the cluster unless the cluster is unavailable. If no third keyword is specified, nor the specified copy source cluster isn't available, the best copy source cluster will be selected. A fourth keyword is only applicable when a third keyword is specified and there is only one valid option "NORECALL". When the logical volume is already migrated to a physical tape on the specified copy source, copy target cluster(s) will try to recall the volume on the specified copy source cluster if "NORECALL" option is not specified. But if "NORECALL" option is specified and cached cluster(s) exists, the copy target cluster(s) will try to get the data from other cached cluster.

Note1: If all the copy source cluster(s) already migrated the volume, a recall will be performed on the specified copy source cluster even if "NORECALL" option is specified.

Note2: Basically, a cached cluster is preferred over migrated cluster, but determining which copy source cluster is used internally would depends on COPYRFSH options (copy source cluster ID and NORECALL), cluster family setting and copy source volume state (i.e. migrated or cached):

[A] COPYRFSH copy target cluster does not belong to any cluster family

"NORECALL" option is specified with "copy source cluster ID = migrated cluster"	• If cached cluster exists, the cached cluster is preferred. • If no cached cluster exists, the specified copy source cluster is preferred.
No "NORECALL" option is specified with "copy source cluster ID = migrated cluster"	The specified copy source cluster is preferred.
No copy source cluster ID is specified	Cached cluster (if exists) is preferred over migrated cluster.

[B] COPYRFSH copy target cluster belongs to a cluster family

"NORECALL" option is specified with "copy source cluster ID = migrated cluster"	Specified copy source cluster is outside of local family: • If a consistent cluster exists in the local family, the local family member is preferred. • If no consistent cluster exists in the local family and a cached cluster exists outside of local family, the cached cluster is preferred.
	Specified copy source is in the same local family: • If a cached cluster exists in the same local family, the cached cluster is preferred. • If no cached cluster exists in the same local family, the specified copy source cluster is preferred.
No "NORECALL" option is specified with "copy source cluster ID = migrated cluster"	Specified copy source cluster is outside of local family: • If a consistent cluster exists in the local family, the local family member is preferred. • If no consistent cluster exists in the local family, the specified copy source cluster is preferred even if a cached cluster exists outside of local family.
	Specified copy source is in the same local family: Even if a cached cluster exists in the same local family, the specified copy source cluster is preferred.
No copy source cluster ID is specified	A consistent cluster in the same local family is preferred. If a cached and migrated clusters exist in the same local family, a cached cluster is preferred over migrated cluster.

If a copy target cluster belongs to a cluster family and cached copy source cluster only exists outside of the local cluster family, a recall will be performed on the specified copy source cluster even if "NORECALL" option is specified, but if cached copy source cluster exists in the local cluster family, it will be used as a copy source cluster instead of recalling a volume from the specified copy source cluster.

Here are some command samples and how it works:

```
LI REQ, XXXX, COPYRFSH, ZZZZZZ
```

The new copy policy defined in the logical volume (ZZZZZZ) management class on the distributed library XXXX will be assigned to the volume.

All the copy target cluster(s) will follow a normal copy source selection algorithm to determine the best copy source cluster.

```
LI REQ, XXXX, COPYRFSH, ZZZZZZ, 0
```

All the copy target cluster(s) will try to get the data from cluster 0. If cluster 0 has already migrated the volume to a physical tape when the copy is attempted, a recall will be performed on cluster 0.

```
LI REQ, XXXX, COPYRFSH, ZZZZZZ, 0, NORECALL
```

All the copy target cluster(s) will try to get the data from cluster 0. If cluster 0 has already migrated the volume to a physical tape when the copy is attempted, other cached cluster(s) (if exists) could be used as a copy source depending on the cluster family setting.

If the request is successfully processed, the following text is returned:

```
COPYRFSH V3 .0
```

```
LOGICAL VOLUME ZZZZZZ COPYRFSH REQUEST COMPLETED
```

If the request is issued to a cluster that is not configured in a Grid, the following error text is returned:

```
CLUSTER NOT PART OF A GRID CONFIGURATION
```

If the specified volume is already used (the volume is mounted or the volume's attribute is being modified), the following error text is returned (the distributed library nickname (sequence number) which is using the volume is shown with the logical device number):

```
VOLUME IN USE BY LIBRARY Lipizzan(#BA92A), DEV x4
```

If the specified volume is used internally in the library, the following error text is returned:

```
VOLUME IN USE INTERNALLY
```

If the optional third keyword is specified and the copy source cluster does not have a valid/consistent data, the following error text is returned (the specified distributed library nickname (sequence number) is shown):

```
DISTRIBUTED LIBRARY Arabian(#BA92B) HAS NO CONSISTENT DATA
```

If the optional third keyword is specified and the cluster is unavailable when the request is issued, the following error text is returned (the specified distributed library nickname (sequence number) is shown):

```
COPY SOURCE CLUSTER Arabian(#BA92B) UNAVAILABLE
```

If not all the clusters in the grid satisfies the code level 8.21.0.118 or above, the following error text is returned:

```
CAPABILITY NOT SUPPORTED FOR THIS COMPOSITE LIBRARY
```

If the volume specified is not in the library, the following error text is returned:

```
LOGICAL VOLUME                ZZZZZZ
VOLUME NOT IN LIBRARY SPECIFIED
```

If the current management class name assigned to the volume is missing, the following error text is returned:

```
ASSIGNED MANAGEMENT CLASS NAME IS NOT DEFINED
```

If any unexpected error occurs and the request is not completed, the following error text is returned:

```
AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = XX
```

If this error message is returned, please retry the request after a while. If the error message persists, please call IBM support with the return code in the error message to get the further assistance.

Line	Bytes	Name	Description
1	0:9	Header Info	'COPYRFSH V'
	10:11	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1. The version is incremented to: - 2 at the code level 8.30.x.x - 3 at the code level 8.31.x.x to reflect the error text change.
	12	Dot	','
	13:14	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	12:69	Blanks	
If the library request is successfully completed, the following lines are returned			

Line	Bytes	Name	Description
2	0	Blank	
	1:48		'LOGICAL VOLUME ZZZZZZ COPYRFSH REQUEST COMPLETED'
	49:69	Blanks	

The copy source selection algorithm of COPYRFSH copy job was changed at the code level 8.21.0.130 (R2.1 PGA2) since the original implementation at the code level 8.21.0.118 (R2.1 PGA1) when the command uses a valid copy source cluster (keyword 3) with the cluster family configuration. The following example explains what was changed at the code level 8.21.0.130 :

This is 3-way Grid configuration. C0 is the copy source cluster at the production site and C1 and C2 are the COPYRFSH copy target clusters which are in the same cluster family at DR site.

A COPYRFSH command is issued to C0 for the volume YYYYYY with copy source cluster ID C0 and its copy job is queued on C1 and C2 :

```
LI REQ, XXXX, COPYRFSH, YYYYYY, 0
```

Given that C1 finished a copy prior to C2, C2 sees 2 consistent copy sources (C0 and C1) when it attempts a copy. At 8.21.0.118 (R2.1 PGA1), C2 ignores the cluster family setting because the copy source cluster ID (C0) was specified in the request. As a result, C2 tries to pull the data from C0 as long as C0 is available. At 8.21.0.130 (R2.1 PGA2), C2 checks any consistent copy source in the local cluster family and tries to pull the data from C1 instead of C0 as long as C1 is available. In summary, COPYRFSH copy works as same as the normal copy job in regard to the copy source selection with the cluster family configuration by 8.21.0.130 algorithm change.

Note: If no copy source cluster ID is specified in the request, C2 always tries to select C1 as a copy source cluster regardless of the code level.

The copy job queuing method was changed at the code level 8.32.x.x.

Prior to 8.31.x.x, COPYRFSH puts a copy job based off of the assigned copy mode change only. There was an error case that a logical volume had an unexpected valid copy mode (like 'Deferred') on the cluster which just joined to an existing Grid although it should have had 'No copy' mode. Applying COPYRFSH to such a logical volume didn't put a copy job on the target cluster although the cluster had no data due to the unexpected copy mode. At 8.32.x.x, COPYRFSH checks the volume inconsistency as well as the copy mode change. Even if the volume had a valid copy mode and it's not changed (for example, from 'Deferred' to 'Deferred') by COPYRFSH, it will put a copy job on the copy target cluster if the cluster has inconsistent data. The removed or expired volume isn't recovered by COPYRFSH. COPY KICK (FORCE) needs to be used to replicate the data on such a logical volume.

There have been several requests from the field that it's desired to refresh (update) not only the copy modes but also other storage constructs when COPYRFSH is performed. At the code level 8.32.x.x, when a copy job generated by COPYRFSH finishes at the copy target cluster, the copy target cluster

updates the logical volume's attributed defined in the storage class, storage group and also the secondary pool defined in the management class as well.

At code level 8.42.x.x or above, TS7700 supports the cloud storage tier. The logical volume's attribute related to the cloud attach is bound to the cloud premigration rank defined in the storage group. As was mentioned in the previous paragraph, the logical volume's attributes are updated, but the logical volume's cloud attach attribute is not updated at the copy target cluster when a COPYRFSH copy job completes. A mount/demount of the logical volume is required to update the cloud attach attribute (cloud premigration rank).

For example, the existing Grid has the logical volume and it is not configured to premigrate to cloud. A new cluster is joined to the Grid and the storage group bound to the logical volume is updated to have a valid cloud premigration rank set to premigrate to cloud. When a COPYRFSH is issued to the logical volume, the cloud premigration rank is not updated on the target cluster so the volume does not premigrate to the cloud. A mount/demount of the volume is required to update the attribute so the volume will be premigrated to the cloud.

Please refer to Redpaper "IBM TS7760 R4.2 Cloud Storage Tier Guide" to get more details.

At the code level 8.50.x.x or above, the third keyword CLDGHOST is supported and it's usable only when all clusters in the Grid are at 8.50.x.x or above.

This option can be used when the copy source logical volume has been already premigrated to the cloud and the copy target cluster would like to see the logical volume on cloud without copying the actual data through the Grid links. CLDGHOST stands for "CLoud GHOST" which means it doesn't do the actual copy through the Grid link then makes the logical volume in the cloud accessible by changing the logical volume status on the copy target cluster. The logical volume status on C2 is changed to "migrated to cloud" instantly. This copy method is referred to a "cloud ghost" copy.

This is how the option is going to work:

1. The existing Grid is 2-way Grid (C0/C1). The logical volume is premigrated (or migrated) to cloud from C0 and/or C1.
2. Join C2 as a new cluster and the logical volume needs to be copied from C0/C1 to C2. C2 is configured to have the access to the same cloud which is already attached to C0/C1.
- 3-1. If the third keyword CLDGHOST is not specified and COPYRFSH is issued to copy the data from C0/C1 to C2, it may require recalling the logical volume from cloud if it's already migrated then the copy source volume is copied from C0/C1 to C2 through the Grid link.
- 3-2. If the third keyword CLDGHOST is specified and COPYRFSH is issued to copy the data to C2, C2 just updates the logical volume status so that C2 can see the logical volume in the cloud once the COPYRFSH copy is attempted and completed. No actual copy through the Grid link is done and C2 sees the logical volume as "migrated to cloud".
 - If C2 is not yet attached to the cloud, it attempts the normal Grid copy through the Grid links.
 - If C0/C1 has the logical volume on its cache when C2 attempts the copy, C2 attempts the normal Grid copy too.

- If C0/C1 is just premigrating the logical volume to the cloud and the data is not in the cloud yet when C2 attempts the copy, C2 just delays and retries the cloud ghost copy. When it's delayed, LI REQ, COPY, QUEUE output as well as Download Incoming Copy Queue page on the management interface GUI show the corresponding reason.

COPYRFSH operation has updated the copy modes defined in the management class on all clusters in the Grid and the storage group/storage class attributes on the copy target cluster too. When the cloud attach was introduced at 8.42.x.x, the cloud attribute (cloud premigration rank defined in the storage group) was not updated when COPYRFSH was issued. In 8.50.x.x., when CLDGHOST keyword is specified, the cloud attribute is also updated on all clusters in the Grid if the logical volume is not yet configured to attach to the cloud. For example,

1. The logical volume is in the cache on C0/C1. No cloud is attached.
2. C2 is joined and it's configured to attach to the cloud.
3. Assign the new storage group and storage class which are configured to premigrate the data to cloud to the logical volume. Or change the current storage group/storage class assigned to the logical volume to premigrate the data to cloud.
- 4-1. If the third keyword CLDGHOST is not specified and COPYRFSH is issued to copy the data from C0/C1 to C2, C2 copies the data through the Grid link and the copied data will stay in the cache on C2.
- 4-2. If the third keyword CLDGHOST is specified and COPYRFSH is issued to copy the data from C0/C1 to C2, C2 copies the data through the Grid link then the copied data will be premigrated to the cloud because the cloud attribute of the logical volume is updated on all clusters in the Grid.

Note: The cloud attribute update occurs by COPYRFSH with CLDGHOST operation only when the attribute is changed from "non-cloud attach" to "cloud attach". Although the logical volume already has the "cloud attach" attribute then the storage group setting is changed to "non-cloud attach" by COPYRFSH operation, the "cloud attach" attribute is not updated.

At the code level 8.50.x.x or above, the logical volume ownership movement behavior is changed. Prior to 8.50.x.x, the cluster which receives COPYRFSH request obtains the logical volume ownership then the ownership can move if the peer cluster has it prior to the COPYRFSH operation. At 8.50.x.x, the ownership is moved back to the original cluster if it's moved during the COPYRFSH operation so that the logical volume ownership is not moved by COPYRFSH operation.

When the copy source volume of COPYRFSH copy job is migrated (only on physical tape), the copy target cluster must request a remote recall on the source cluster. Each target cluster can issue up to four recalls in parallel among its peers. By alternating your COPYRFSH requests to rotate across four physical volumes, improved recall performance can be achieved through the four parallel recall tasks. If only LVOLs from one PVOL are queue sequentially, the four tasks can end up being queued against a single PVOL which creates a serialization point. Because up to 4 recalls can be issued from the copy target cluster at a time, changing the order of queuing the COPYRFSH copy jobs may contribute to the improving the performance.

For example, 100 logical volumes (LVOL00 – LVOL99) will be copied by COPYRFSH and all of them are already migrated on the copy source cluster. They are migrated to total 4 different physical volumes:

- LVOL00–LVOL24 on PVOL00
- LVOL25–LVOL49 on PVOL01
- LVOL50–LVOL74 on PVOL02
- LVOL75–LVOL99 on PVOL03

COPYRFSH copy jobs are basically attempted in FIFO order, then queuing COPYRFSH copy jobs with this order:

LVOL00->LVOL01->LVOL02->.....->LVOL99

may not use the 4 recall requests from the copy target cluster efficiently.

Instead of the above order, the following queuing order may use the 4 recall requests more efficiently:

(LVOL00->LVOL25->LVOL50->LVOL75)-> (LVOL01->LVOL26->LVOL51->LVOL76)->...->(LVOL24->LVOL49->LVOL74->LVOL99)

4 different physical volumes should be recalled concurrently so that it will have a good COPYRFSH recall (copy) performance. And, each PVOL recall will recall many LVOLs sequentially providing the most efficient operation.

To reach this logical volume sequence you have two possibilities:

1. Run one job with all the commands and then in the proposed order. To do so, use DFSORT to merge the four LVOL/PVOL mapping into one single input stream.
2. Submit four recall jobs (for each pvol one job) in the same second. Then the command interpreter of z/OS will interpret all four jobs concurrently – and this result in a similar distribution.

Both method result in nearly the same performance.

3.1.16 OTCNTL

OTCNTL command provides a method to transfer the ownership of the logical volumes in the background. One of the typical cases which this command is useful is FC 4016 (Remove Cluster from Grid). While FC 4016 is executed, all the logical volume ownership which resides on the removing cluster will be automatically transferred to one of the remaining cluster(s) to assure the access to the logical volumes after the cluster is removed. If a large number of logical volumes need to transfer the ownership, it may take a long time to finish the ownership transfer which may elongate the total elapsed time of the FC 4016 operation. By using OTCNTL command prior to the FC4016 operation, it can complete the necessary ownership transfer beforehand and help to shorten the required time of FC4016 operation.

Note: The request is supported only when all the clusters in the grid have the code level of 8.30.1.X or above.

A total 5 different second keywords are supported to control the ownership transfer operation:

1. START

The second keyword “START” initiates the ownership transfer. The distributed library which receives “START” request will start the ownership transfer in the background. The ownership of the logical volumes owned by the peer cluster which is specified in the third keyword will only be transferred to the local cluster. The third keyword (source cluster ID from which the local cluster gets the ownership) is mandatory. The fourth keyword can be used to limit the number of logical volumes to transfer the ownership by the request. This is optional and if no fourth keyword is specified, the command will transfer the ownership of all the logical volumes owned by the peer cluster specified in the third keyword.

Note 1: Only 1 ownership transfer operation instance can be running at a time in the grid. If “START” command is issued while another operation is already running, it must fail.

Note 2: When all the ownership transfer finishes, the operation stops automatically.

Note 3: Prior to R5.0PGA1 (8.50.1.25), the logical volumes to be transferred are ordered by volser. It's changed in R5.0 PGA1 that they are transfer in the order of their category time (the volume with the oldest category time is transferred first).

(START)

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OTCNTL	START	Source cluster ID (0-7)	Maximum number of logical volumes to transfer the ownership (option)	N/A	Y	Y	Y	Y	Y

The response lines are formatted as follows:

```
OWNERSHIP TRANSFER V3 .0
```

```
BACKGROUND OWNERSHIP TRANSFER HAS SUCCESSFULLY STARTED
```

The following table details the format of the data reported when “START” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:19	Header Info	'OWNERSHIP TRANSFER V'
	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.30.1.x). The version is incremented to 2 at 8.33.x.x and 3 at 8.50.1.25.
	22	Dot	'.'
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	25:69	Blanks	
2	0:53		If the operation starts successfully, the following text is returned: 'BACKGROUND OWNERSHIP TRANSFER HAS SUCCESSFULLY STARTED'
	54:69	Blanks	

Example of “START” command:

```
LI REQ, XXXX, OTCNTL, START, Y, 1000
```

The command initiates the ownership transfer of the logical volumes from cluster Y to the distributed library XXXX. It will transfer the ownership of up to 1000 logical volumes.

2. STOP

The second keyword “STOP” terminates the active ownership transfer operation. No third and fourth keywords are supported.

(STOP)

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OTCNTL	STOP			N/A	Y	Y	Y	Y	Y

The response lines are formatted as follows:

```
OWNERSHIP TRANSFER V3 .0
```

```
BACKGROUND OWNERSHIP TRANSFER HAS SUCCESSFULLY STOPPED
```

The following table details the format of the data reported when “STOP” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:19	Header Info	'OWNERSHIP TRANSFER V'

	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.30.1.x). The version is incremented to 2 at 8.33.x.x and 3 at 8.50.1.25.
	22	Dot	` . '
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	25:69	Blanks	
2	0:53	Ownership Transfer Status	If the active ownership transfer operation is halted successfully, the following text is returned: `BACKGROUND OWNERSHIP TRANSFER HAS SUCCESSFULLY STOPPED'
	54:69	Blanks	

3. MOVE1

The second keyword “MOVE1” initiates the ownership transfer for a single logical volume. The distributed library which receives “MOVE1” request will perform the ownership transfer of the specified volume. This operation is synchronous opposed to “START” command.

Note 1: The request is supported only when all the clusters in the grid have the code level of 8.30.1.X or above and the target cluster has the code level of R5.0 PGA1 or above.

Note 2: “MOVE1” and “START” operations work independently. While “START” operation is running in the Grid, “MOVE1” can be accepted even on the same cluster, and vice versa. But please make sure the same logical volume should not conflict between two operations. Otherwise, the unexpected ownership transfer result may occur.

(MOVE1)

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OTCNTL	MOVE1	ZZZZZZ		N/A	Y	Y	Y	Y	Y

where ZZZZZZ is the volser of the logical volume to be transferred.

The response lines are formatted as follows:

```
OWNERSHIP TRANSFER V3 .0
```

```
OWNERSHIP OF THE LOGICAL VOLUME ZZZZZZ HAS BEEN MOVED SUCCESSFULLY
```

The following table details the format of the data reported when “MOVE1” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:19	Header Info	`OWNERSHIP TRANSFER V'
	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Start with 3 (at the code level 8.50.1.25).
	22	Dot	` . '

	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. It starts with 0.
	25:69	Blanks	
2	0:53		If the operation completed successfully, the following text is returned: 'OWNERSHIP OF THE LOGICAL VOLUME ZZZZZZ HAS BEEN MOVED SUCCESSFULLY' where 'ZZZZZZ' is the volser of the logical volume transferred.
	54:69	Blanks	

4. DIST

The second keyword “DIST” provides the current ownership distribution status in the grid. The status shows the count of the logical volumes per the owner and each category.

At the code level of 8.33.x.x, the format of the output is changed and the third keyword (starting category) is supported.

Note 1: The list only shows the owner cluster and count of the category from 0x0001 to 0xFF00 and 0xFFFF.

Note 2: The output provides up to 45 entries. If more than 45 entries exist in the grid, they are not displayed in the output and “*ADDITIONAL CATEGORIES EXIST. NEXT CATEGORY IS CCCCX.” is provided at the bottom of the output. The user can specify the category “CCCC” in the third category so that another set of category information from the category can be obtained.

Note 3: Due to the field space limitation, up to 999999 count can be displayed in each count field. Even if more than 999999 volumes are in that category, the count must show 999999.

(DIST)

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OTCNTL	DIST	Category in hex (Starting category)		N/A	Y	Y	Y	Y	Y

The response lines are formatted as follows:

```

OWNERSHIP TRANSFER V3 .0
      TOTAL
CAT  COUNT      0      1      2      3      4      5      6      7
-----
0001      1      0      NC      0      NC      NC      1      0      NC
000F     21      1      NC      0      NC      NC     13      7      NC
001F    292      0      NC      0      NC      NC    291      1      NC
FF00   5389      0      NC      0      NC      NC   5389      0      NC
FFFF      1      0      NC      1      NC      NC      0      0      NC

```

The following table details the format of the data reported when “DIST” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:19	Header Info	'OWNERSHIP TRANSFER V'
	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.30.1.x). The version is incremented to 2 at 8.33.x.x and 3 at 8.50.1.25.
	22	Dot	'.'
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	25:69	Blanks	
2	0:42	Header Info	' TOTAL OWNER CLUSTER'
	43:69	Blanks	
3	0:66	Header Info	' CAT COUNT 0 1 2 3 4 5 6 7'
	67:69	Blanks	
4	0:66	Header Info	'----- -- -----'
	67:69	Blanks	
5	0:3	Category in hex	The category count status to provide on this line. The category is provided in hex.
	4	Blank	
	5:10	Total Logical Volume Count in the CAT category	The total number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster in OWNER column.
	11	Blank	
	12:17	Logical Volume Count owned by C0 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 0. If cluster 0 is not configured, 'NC' is provided.
	18	Blank	
	19:24	Logical Volume Count owned by C1 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 1. If cluster 1 is not configured, 'NC' is provided.
	25	Blank	
	26:31	Logical Volume Count owned by C2 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 2. If cluster 2 is not configured, 'NC' is provided.
	32	Blank	

	33:38	Logical Volume Count owned by C3 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 3. If cluster 3 is not configured, 'NC' is provided.
	39	Blank	
	40:45	Logical Volume Count owned by C4 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 4. If cluster 4 is not configured, 'NC' is provided.
	46	Blank	
	47:52	Logical Volume Count owned by C5 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 5. If cluster 5 is not configured, 'NC' is provided.
	53	Blank	
	54:59	Logical Volume Count owned by C6 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 6. If cluster 6 is not configured, 'NC' is provided.
	60	Blank	
	61:66	Logical Volume Count owned by C7 in the CAT category	The number of logical volumes which is assigned to the category in CAT column and is owned by the owner cluster by cluster 7. If cluster 7 is not configured, 'NC' is provided.
	67:69	Blanks	
6-49			The same information is provided until the line 49 (up to 50 entries can be displayed).
50	0:52		If more than 45 entries exist, '*ADDITIONAL CATEGORIES EXIST. NEXT CATEGORY IS CCCCX.' is provided.
	53:69	Blanks	

5. STAT

The second keyword "STAT" provides the current statistics of the active ownership transfer operation started by "START" on the target distributed library. If the operation is already stopped or no operation is running, the command returns the status and result of the last operation if any. No third and fourth keywords are supported.

Note 1: The count in "NUMBER OF CANDIDATE VOLS" and "TOTAL VOLUMES TO TRANSFER" is calculated when OTCNTL operation is initiated. There may be the case that the volume's category or owner is changed while the ownership transfer operation is in progress, then the count originally calculated when OTCNTL operation started is no longer correct. As a result, the ownership transfer status may not complete (i.e. the count provided in "ON GOING PROGRESS"

never reaches to the count of “NUMBER OF CANDIDATE VOLS and the “OWNERSHIP TRANSFER STATUS” stays at “ACTIVE status”). In such a case, it may be required to look at the current distribution status by “DIST” keyword to check if the required operation already completed or not. If the operation already completed, the operation needs to be terminated by “STOP” keyword. This case is taken into account to appropriately check the ownership transfer progress and complete the operation in the future code enhancement.

Note 2: “STAT” provides the current or previous ownership transfer operation initiated by “START” command only. Any operation issued by "MOVE1" command doesn't change “STAT” output (the result of “MOVE1” command is just reported synchronously).

(STAT)

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OTCNTL	STAT			N/A	Y	Y	Y	Y	Y

If a second keyword “STAT” is specified, the response lines are formatted as follows:

```

OWNERSHIP TRANSFER V3 .0
LIBRARY                  : elm
OWNERSHIP TRANSFER STATUS : ACTIVE
SOURCE CATEGORY          : ----
DELAY PER EXCHANGE       : 3500
NUMBER OF CANDIDATE VOLS  : 2718
REQUEST COUNT            : 1000
TOTAL VOLUMES TO TRANSFER : 1000
FAILED TRANSFER VOLS      : 1
ON GOING PROGRESS         : 19 (1.9 %)
TRANSFER STARTED AT      : 02/05 06:27:20
ELAPSED TIME              : 0 HOURS, 1 MINUTES

```

The following table details the format of the data reported when “STAT” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:19	Header Info	'OWNERSHIP TRANSFER V'
	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.30.1.x). The version is incremented to 2 at 8.33.x.x and 3 at 8.50.1.25.
	22	Dot	'.'
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	25:69	Blanks	
2	0:27	Library Info	'LIBRARY : '
	28:35	Distributed Library Info	The distributed library nickname which is returning the statistics information. The value is left justified and padded with blanks.

	36:69	Blanks	
3	0:27	Ownership Transfer Status	`OWNERSHIP TRANSFER STATUS : '
	28:37	Current Status	The current ownership transfer request status on the distributed library : `ACTIVE' : The ownership transfer request is in progress. `COMPLETED' : Previous ownership transfer request was completed. `STOPPED' : Previous ownership transfer request was halted. `FAILED' : Previous ownership transfer request was failed. `INACTIVE' : No ownership transfer is requested yet. The value is left justified and padded with blanks.
	38:69	Blanks	
4	0:27	Source Category Info	`SOURCE CATEGORY : '
	28:31	Source Category	The source category which is set by "CONFIG CAT" command is displayed. If no category is set, '----' is displayed.
	32:69	Blanks	
5	0:27	Interval Info	`DELAY PER EXCHANGE : '
	28:32	Interval value	The interval value in millisecond which is set by "CONFIG DELAY" command is displayed. If no interval is set, the default value 3500 is displayed. The value is left justified and padded with blanks.
	33:69	Blanks	
6	0:27	Candidate Volume Info	`NUMBER OF CANDIDATE VOLS : '
	28:34	Volume candidate value	Number of the candidate volumes when the current active transfer operation was requested. If no active request is running, the result of the previous attempt is provided.
	35:69	Blanks	
7	0:27	Request Count Info	`REQUEST COUNT : '
	28:34	Request Count Value	The maximum number of logical volumes to transfer the ownership which is set by the current active request. If no active request is running, the result of the previous attempt is provided.
	35:69	Blanks	
8	0:27	Total Volumes to Transfer Info	`TOTAL VOLUMES TO TRANSFER : '
	28:34	Total Count of the candidate logical volumes	The total count of the candidate logical volumes which is identified by the current active request. If no active request is running, the result of the previous attempt is provided.
	35:69	Blank	
9	0:27	Failed Transfer Volumes Info	`FAILED TRANSFER VOLS : '
	28:34	Total Count of Volumes Failed to Transfer	The total count of the logical volumes which could not transfer the ownership. If no active request is running, the result of the previous attempt is provided.

	35:69	Blank	
10	0:27	Ongoing progress Info	`ON GOING PROGRESS : `
	28:39	Progress information	If the ownership transfer operation is active, the following values are provided: `XX, (YY.Y %)` XX : total count of logical volume whose ownership was already transferred. YY.Y ; percentage of the completed transfer operation which is calculated by $(XX / \text{TOTAL VOLUMES TO TRANSFER} * 100)$. If the previous ownership transfer operation already completed or halted, its result is displayed. Otherwise, it is always `NA`.
	40:69	Blanks	
11	0:27	Transfer Starting Time Info	`TRANSFER STARTED AT : ` If there is no active, completed or stopped transfer operation, the line always shows: `STARTED AT : `
	28:41	Starting time	The data/time when it started is displayed in the following format: `MM/DD HH:mm:ss` MM/DD is Month(1-12)/Date(1-31) HH:mm:ss is Hour:Minute:Second. The time zone is GMT (TS7700 internal time). If the previous ownership transfer operation already completed or halted, the date/time when the last operation started is displayed. Otherwise, it is always `NA`.
	42:69	Blanks	
12	0:27	Elapsed Time Info	`ELAPSED TIME : `
	28:45	Elapsed time of the active operation	If the ownership transfer operation is active, its elapsed time is displayed in the following format: `XX HOURS, YY MINUTES` XX is hour and YY is minute. If the previous ownership transfer operation already completed or halted, the elapsed time of the last operation is displayed. Otherwise, it is always `NA`.
	46:69	Blank	

6. CONFIG

By default, “START” command transfers the ownership of all the logical volumes whose owner is the target cluster specified in the third keyword regardless of the category. It may need to transfer the ownership of the logical volumes which has a specific category only (for example, non scratch category). By using the second keyword “CONFIG” and third keyword “CAT”, it is allowed to

transfer the ownership of the logical volumes in the specific category only. The target category needs to be specified in the fourth keyword.

The ownership transfer function puts 3500 milliseconds delay between each volume ownership transfer operation by default in order to prevent the normal operation (host I/O, copy activity etc.) from getting the performance impact. But it may be better to shorten the delay and make the entire ownership transfer operation finish faster if the machine is relatively idle. By using the second keyword "CONFIG" and third keyword "DELAY", it is allowed to adjust the delay between each volume ownership transfer operation. The delay between 500 and 10000 milliseconds can be set in the fourth keyword. Since the code level of 8.33.x.x, the minimum delay 0 millisecond can be set.

The target category and delay set by "CONFIG" command can be checked in the output of "STAT" keyword (in the line of "SOURCE CATEGORY" and "DELAY PER EXCHANGE").

Note 1: Only one set of category and delay can be set per each distributed library. If another category or delay is set, the previous one is overwritten.

Note 2: The configured values are not persistent over a TS7700 reboot (also offline/online cycle). It's required to set them again if the cluster has a reboot or offline/online sequence after they're set.

Note 3: In order to delete the defined category, "FFFF" needs to be set by "CONFIG" and "CAT" keywords.

Note 4: The target category can't be changed while the ownership transfer request is running. On the other hand, the delay can be changed in the middle of the active ownership transfer and the new delay is applied to the current active operation.

At the code level of R3.3 PGA2 (8.33.2.9), the null category "0000" is used to delete the defined category (it has been "FFFF" in the prior levels). Then the category "FFFF" (VOLSER-Specific Category) can be set as a target category by CONFIG, CAT keywords.

(CONFIG)

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OTCNTL	CONFIG	CAT DELAY	yyyy (category) (mandatory for CAT) XX (milliseconds) (mandatory for DELAY)	N/A	Y	Y	Y	Y	Y

If the third keyword "CAT" is specified, the response lines are formatted as follows:

```
OWNERSHIP TRANSFER V3 .0
CATEGORY CHANGE SUCCESSFULLY APPLIED
```

If the third keyword "DELAY" is specified, the response lines are formatted as follows:

```
OWNERSHIP TRANSFER V1
```

DELAY PER EXCHANGE SUCCESSFULLY APPLIED

The following table details the format of the data reported when “CONFIG” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:19	Header Info	'OWNERSHIP TRANSFER V'
	20:21	Version	This is the version number for the response. The number is left justified and padded with blanks. Start with 1 (at the code level 8.30.1.x). The version is incremented to 2 at 8.33.x.x and 3 at 8.50.1.25.
	22	Dot	'.'
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.31.0.x and starts with 0.
	25:69	Blanks	
2	0:35		If the category is set successfully, the following text is returned: 'CATEGORY CHANGE SUCCESSFULLY APPLIED'
	36:69	Blanks	
2	0:38		If the delay is set successfully, the following text is returned: 'DELAY PER EXCHANGE SUCCESSFULLY APPLIED'
	39:69	Blanks	

Example of “CONFIG” command:

```
LI REQ, XXXX, OTCNTL, CONFIG, CAT, 0001
```

Set the target category to “0001” for the ownership transfer operation.

```
LI REQ, XXXX, OTCNTL, CONFIG, DELAY, 1000
```

Set 1000 milliseconds delay between each volume ownership transfer operation.

This is an example of the typical command sequences to initiate/check/stop the ownership transfer operation:

First, you should know from which cluster (c) you want to transfer the volume ownership.

- LI REQ, XXXX, OTCNTL, DIST (check the current ownership distribution status)

You may get the idea to which cluster the ownership should be transferred (also how many (nnn)) from the output. Determine the distributed library (LLLL) which needs to receive the request then the distributed library obtains the volume ownership from the cluster (c).

- LI REQ, LLLL, OTCNTL, CONFIG, CAT, xxxx (set the target category if needed)
- LI REQ, LLLL, OTCNTL, CONFIG, DELAY, yyyy (set the appropriate delay if needed)
- LI REQ, LLLL, OTCNTL, STAT (confirm the setting)
- LI REQ, LLLL, OTCNTL, START, c, nnn (initiate the operation to transfer the volume ownership of up to nnn volumes from cluster (c) to the distributed library (LLLL))
- LI REQ, LLLL, OTCNTL, STAT (check the progress)

- LI REQ, LLLL, OTCNTL, STOP (stop the operation if needed)
- LI REQ, XXXX, OTCNTL, DIST (check the ownership distribution status again)

The following error messages will be returned when the command is not accepted:

If the command is issued to a standalone cluster, the following error text is returned:

CLUSTER NOT PART OF A GRID CONFIGURATION

If all the clusters in the grid do not have the code level 8.30.1.x or above, the following error text is returned:

MINIMUM CODE LEVEL FOR DOMAIN IS NOT MET

If an unexpected error occurs during the process, the following error text is returned:

AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = XX

[“CONFIG” is the second keyword]

If the third keyword is “CAT” and the target category specified in the fourth keyword is invalid, the following error text is returned:

INVALID 4TH KEYWORD

(*) Applying no category, NULL (x0000) and xFF01 – xFFFE in the fourth keyword is considered as invalid. (xFFFF is used to reset the category setting).

If the target category change request is issued while the ownership transfer request is currently running, the following error text is returned:

REQUEST IS INVALID WHILE OWNERSHIP TRANSFER IS IN PROGRESS

If the third keyword is “DELAY” but no delay is specified in the fourth keyword, the following error text is returned:

INVALID 4TH KEYWORD

If the third keyword is “DELAY” and the specified delay value is out of range, the following error text is returned:

THE DELAY MUST BE BETWEEN 500 AND 10000

[“START” is the second keyword]

If no third keyword or invalid character is specified in the third keyword, the following error text is returned:

INVALID 3RD KEYWORD

If the source cluster ID specified in the third keyword does not exist, the following error text is returned:

DISTRIBUTED LIBRARY DOES NOT EXIST

If invalid character is specified in the fourth keyword, the following error text is returned:

INVALID 4TH KEYWORD

If the source cluster ID in the third keyword is the same cluster which the request is issued to, the following error text is returned:

SOURCE AND TARGET CLUSTER FOR OWNERSHIP TRANSFER ARE THE SAME

If the ownership transfer request is already running, the following error text is returned:

BACKGROUND OWNERSHIP TRANSFER IS ALREADY RUNNING

If the source cluster ID in the third keyword is not available, the following error text is returned:

SOURCE CLUSTER FOR BACKGROUND OWNERSHIP TRANSFER IS NOT ALIVE

If there are no candidates to do ownership transfer, the following error text is returned:

THERE ARE NO CANDIDATE VOLUMES TO PERFORM OWNERSHIP TRANSFER

[“STOP” is the second keyword]

If no active ownership transfer request is running in the grid, the following error text is returned:

BACKGROUND OWNERSHIP TRANSFER IS NOT ACTIVE

3.1.17 DRSETUP

At the code level 8.31.0.x, concurrent disaster recovery testing is improved with the Flash Copy for Disaster Recovery Testing function. This enables a Disaster Recovery host to perform testing against a point in time consistency snapshot while production continues. With Flash Copy, production data continues to replicate during the entire DR test and the same logical volume can be mounted at the same time by a DR host and a production host. Used in conjunction with Selective Write Protect 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 will be provided via a snapshot in time, or flash, of the logical volumes. In addition, a DR host will continue to have read access to production volumes and the volumes' original contents that have since been returned to scratch.

At the code level 8.33.x.x, a new keyword SELFLIVE is supported which provides a DR test host the ability to access its self-created content which has been moved into a write protected category when flash is enabled. This feature was previously only available through development/PFE enablement. The functionality is mainly designed for z/VSE environment, where no dedicated private categories can be established.

At the code level 8.41.x.x, a new keyword LIVEACC is supported which provides a DR test host the ability to access the live copy volume on the flash copy volume by changing its category to the one excluded from Write Protect Mode. This capability provides a DR test host to access the live copy volume created by the production host after Flash Copy is enabled.

At code level 8.41.200.113 or above, enabling/disabling Flash Copy and Write Protect is an asynchronous operation when all clusters in the Grid are at 8.41.200.113 or above. Due to various reasons, customers were experiencing a longer than expected time to disable Flash Copy which resulted in a timeout of the library command. The asynchronous operation returns immediately once the operation is issued. The progress of the operation can be monitored by issuing DRSETUP, SHOW. The result is reported through the operational message.

DRSETUP request provides the command to operate and monitor Flash Copy for DR testing function. Please refer to “IBM TS7700 Series Best Practices Flash Copy for Disaster Recovery Testing” on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.18 PARTRFSH

In Release 8.32, TS7700 Tape Attach (TS7700T) is supported. In addition to one resident only partition (partition 0), up to seven tape attached partitions (partition 1 – 7) can be defined in TS7700T. The cache partition to which the workload targets is controlled by outboard policy management (storage class definition). The target cache partition is determined when the volume has mount/demount.

At the code level 8.42.x.x, TS7700 Cloud Attach (TS7700C) is supported. TS7700C also has one resident only partition (partition 0) and up to seven cloud tier partitions (partition 1 – 7). LI REQ PARTRFSH is supported on TS7700C as well.

Similar to COPYRFSH, PARTRFSH can be used to refresh the cache partition by mimicking mount/demount. Prior to R3.3 PGA2, PARTRFSH works only when the actual cache partition assignment is changed. If the assignment is not changed, the command fails. In R3.3 PGA2 or later, PARTRFSH works regardless of the cache partition assignment change. This is useful for disabling delayed-pre-migration setting without changing the cache partition assignment prior to frame swap.

The PARTRFSH is not a real mount/demount process. The table below summarize how the data movement could occur by PARTRFSH cache partition change:

CP change operation by PARTRFSH	From CPx to CP0	From CP0 to CPx
TS7700T	When the volume exists only on physical tape, no recall process will be executed from physical tape. When the volume exists both on physical tape and cache, the data on tape will be deleted. Note: The data still remains on tape and it's recalled when the access to the data is required on the next mount and/or Grid copy operation.	Delay premigration is only honored if the grace period is not expired. If the grace period (volume creation / last accessed) is already expired, the volume will be put on the premigration queue immediately. Note: This data is counted against the installed number of FC 5274 – and can lead to premigration throttling
TS7700C	When the volume exists only on cloud, no recall process will be executed from cloud. When the volume exists both on cloud and cache, the data on cloud will not be deleted.	Cloud premigration will be executed only when a cloud premigration rank has been already configured and it's already assigned to the target volume at the last volume mount/demount. Note: If no cloud premigration rank has been assigned to the target

	Note: The data still remains on cloud and it's recalled when the access to the data is required on the next mount and/or Grid copy operation.	volume, the data will stay in cache under CPx.
--	---	--

Prior to R4.1.2, PARTRFSH only runs on TS7700T clusters. In R4.1.2 or later, PARTRFSH also runs on TS7700D clusters to refresh the Storage Class settings such as removal-policy.

Note: PARTRFSH only refreshes the cache partition assignment associated to the assigned storage class definition below the code level of R3.3 PGA2. But in R3.3 PGA2 or later PARTRFSH changes primary/secondary pool settings if the given volume is resident. If the target volume is already pre-migrated or migrated, PARTRFSH does not change primary/secondary pool settings.

At the code level 8.42.x.x, the cloud attach is supported. The cloud attributes defined in the storage group (cloud premig rank and object prefix) are bound to each logical volume and they're updated by a mount/demount of the logical volume. They are Grid scope and not updated by PARTRFSH request because PARTRFSH is a cluster scope request.

At the code level 8.50.x.x, the third keyword MMOUNT is supported. This option is usable only when all clusters in the Grid are at 8.50.x.x or above. It updates the cloud attributes by mimicking a mount/demount. This can be used when it's required to premigrate the logical volume to the cloud without a mount/demount.

For example:

1. 2-way Grid (C0/C1) has no cloud attach. The logical volumes are in the cache.
2. Configure the cloud attach and create a new storage group/storage class (or update the existing ones) for the cloud access.
3. Apply the storage group/storage class to the logical volumes so that C0/C1 are ready to premigrate the logical volume to the cloud.
4. Without a mount/demount, issue LI REQ, PARTRFSH, <volser>, MMOUNT to C0. C0 premigrates the logical volume to the cloud. Note: the cloud attributes are updated on the logical volume in the Grid but C1 doesn't recognize the data on cloud yet because its cache partition information is not yet updated.
5. Issue LI REQ, PARTRFSH, <volser> to C1, then C1 recognizes the data on the cloud which C0 has premigrated to the cloud.

Another example is:

1. C0 has no cloud attach. The logical volumes are in the cache.
2. C1 is joined to the Grid.
3. Configure the cloud attach and create a new storage group/storage class (or update the existing ones) for the cloud access.

3. Apply the storage group/storage class to the logical volumes so that C0/C1 are ready to premigrate the logical volume to the cloud.
4. Without a mount/demount, issue LI REQ, PARTREFSH, <volser>, MMOUNT to C0. C0 premigrates the logical volume to the cloud.
5. Issue LI REQ, PARTREFSH, <volser>, MMOUNT to C1 too.
5. Assign a management class to the logical volume so that C1 has a valid copy mode.
6. Issue LI REQ, COPYREFSH, <volser>, CLDGHOST to C0 or C1. The cloud ghost copy is attempted on C1 then it recognizes the data on the cloud without copying the actual data through the Grid links.

Note: MMOUNT option updates the cloud attributes of the logical volume on the target cluster only when the target logical volume is not configured to be premigrated to the cloud. If the logical volume already has the cloud attributes and the new storage group with no cloud attributes is assigned and MMOUNT option is used, the cloud attributes bound to the logical volume is not changed.

The following table lists the keywords supported for the PARTREFSH request.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7700D	7740	7700T	7700C
PARTREFSH	zzzzzz (volser)	MMOUNT (option)	N/A	N/A	Y	Y	N	Y	Y

Here is a command sample and how it works:

```
LI REQ, XXXX, PARTREFSH, ZZZZZZ
```

The new cache partition defined in the logical volume (ZZZZZZ) storage class on the distributed library XXXX will be assigned to the volume.

If the request is successfully processed, the following text is returned:

```
PARTITION REFRESH V1 .0
  UPDATED CACHE PARTITION OF LOGICAL VOLUME ZZZZZZ
```

If the request is issued to non TS7700T cluster, the following error text is returned:
ONLY SUPPORTED IN A TS7700 TAPE ATTACHED VIRTUALIZATION ENGINE

If the request is issued to a logical volume and its cache partition defined in the storage class is not changed, the following error text is returned:

```
CACHE PARTITION OF LOGICAL VOLUME Z00000 IS NOT CHANGED (CP1).
```

If the request is issued to a logical volume which is currently mounted, the following error text is returned:

```
PARTITION REFRESH V1 .0
```

LOGICAL VOLUME ZZZZZZ IS MOUNTED NOW

Line	Bytes	Name	Description
1	0:19	Header Info	'PARTITION REFRESH V'
	20:21	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.32.x.x.
	22	Dot	'.'
	23:24	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	25:69	Blanks	
If the library request is successfully completed, the following lines are returned			
2	0	Blank	
	1:48		'UPDATED CACHE PARTITION OF LOGICAL VOLUME ZZZZZZ'
	49:69	Blanks	

3.1.19 LOWRANK

This option provides the capability to override the TVC selection algorithm so that specific clusters are less preferred for mount TVC selection or copy source selection. For example, this can provide a way to have jobs move away from using a particular cluster for TVC when leading up to a planned service outage of that specific TVC. Or, for a method to have copy source selection prefer certain sources over others when all other factors appear equal.

LOWRANK TVC:

For example, a user has a 6-way Grid (C0-C5) and the host is connected to C0, C1 and C2. The user plans to have service-prep initiated against C1 later in the day. A large portion of data is only on physical tape within C0 but often in disk cache within C1 and C2. If mounts to C0 or C2 utilize C1 as a TVC, this can prolong the service-prep operation given it cannot complete until all mounts using C1's TVC completes. Mounts to C0 may frequently choose C1 for TVC given a recall is not required. By setting LOWRANK, TVC to 2 (00000010b), the TVC clusters are re-ordered internally so that C1 is always the last choice. As a result, only mounts where C1 is the only valid or accessible copy will utilize C1's TVC which should keep it TVC idle. This can help accelerate the future initiating of service-prep against C1.

LOWRANK COPY:

For example, a user has a 3-way Grid (C0-C2) and the host is connected to C0 and the data is created in C0's TVC. The copy mode for C0 is RRD so that C1 should receive a RUN copy too at volume close and C2 should receive a deferred copy sometime thereafter. With this setting, it is possible to have C2 favor C1 over C0 as its copy source by setting LOWRANK, COPY to 1 (00000001b) on C2. Only if C1 is unavailable or does not contain a copy will C2 utilize C0 as the source.

The first keyword is always "LOWRANK". The second keyword is either "TVC" or "COPY" to determine which behavior you are attempting to change. The third keyword can be "SHOW" or the cluster's mask value which represents the one or more clusters to be less preferred. For example, if it's a 6-way Grid and C0, C3 and C5 need to be set as "less preferred", the value "00101001b" = 41 (in decimal) should be set. Cluster0 is represented by the least significant bit or "00000001b." Cluster3 is the 4th significant bit or "00001000b." Cluster5 is the 6th least significant bit position or 00100000b. Bit-OR of the three masks creates "00101001b" in binary which is the value 41 in decimal. If the third keyword is "SHOW", it provides the current mask value. The default value is 00000000b or 0 decimal which states no preference change. The default value is 0. If 0 is set, LOWRANK setting is reset.

Additional consideration about LOWRANK setting:

1. TVC

- The mask setting takes precedence over the Sync Mode copy (if 'S' copy mode cluster is in LOWRANK, it may not be selected as one of the two 'S' TVC points. Then, the mount could fail or go synch-deferred based on the synch mode copy settings configured for the volume's management class within the TS7700 Management Interface.

- The mask setting takes precedence over the Immediate Mode copy, which means an update can result in the cluster entering the immediate-deferred state. The new immediate copy reason code (0x85: One or more RUN target clusters are in LOWRANK TVC setting) is returned in Byte 19 of ERA 35 sense data if LI REQ, SETTING, IMMSNS is set to ALL.
- “Force Local for private mounts” copy override settings take precedence over this mask. If the mount point cluster has "Force Local" configured and the mount-point cluster is also contained in the mask, the mask is ignored and the local cluster is still chosen as TVC. If you want it to honor the mask, you must uncheck "Force Local" settings in the copy override panel of the cluster's Management Interface
- “Prefer local cache for scratch (private) mount requests” copy override settings take precedence over the setting. If you want it to honor the mask, you must uncheck "Force Local" settings in the copy override panel of the cluster's Management Interface
- Device Allocation Assistance takes into account the setting, so host mounts for private volumes will favor devices assigned to clusters which are not included in the mask. Scratch Allocation Assistance will ignore the masks.
- The setting takes precedence over the cached/migrated TVC source status of private mounts, meaning a cached copy will be viewed less optimal than another cluster which requires a recall if the cluster with a cached copy is contained in the mask.

2. COPY

- The setting takes precedence over the cluster family setting (if the setting includes the cluster(s) in the same local family, the copy target cluster may select the copy source from outside of the family).
- The setting takes precedence over the cached/migrated copy source status, meaning a cached copy will be viewed less optimal than another cluster which requires a recall if the cluster with a cached copy is contained in the mask.
- The setting takes precedence over the copy source specified by LI REQ, COPYRFSH command (although LI REQ, COPYRFSH specifies the copy source cluster in the third keyword, the copy target cluster may ignore it if the setting includes the copy source cluster).

The following table lists the keywords supported for the LOWRANK request.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7700D	7740	7700T	7700C
LOWRANK	TVC/ COPY	SHOW/ xxx	N/A	N/A	Y	Y	Y	Y	Y

Note: Here is the cluster ID and mask conversion for the third keyword:

[Cluster ID => Mask value]

CL0 => 1

CL1 => 2

CL2 => 4

CL3 => 8

CL4 => 16

CL5 => 32

CL6 => 64

CL7 => 128

If CL2 needs to be set, the mask value to use is 4.

If CL0, CL3 and CL5 need to be set, the mask value to use is the sum of them: $1 + 8 + 32 = 41$.

If the request is successfully processed, the following text is returned:

(LOWRANK, TVC)

LOWRANK TVC INFORMATION V1 .1

LOWRANK 65 IS SET ON DISTRIBUTED LIBRARY cluster0(#BA92A)

(LOWRANK, COPY)

LOWRANK COPY INFORMATION V1 .1

LOWRANK 64 IS SET ON DISTRIBUTED LIBRARY cluster0(#BA92A)

The SHOW command provides the current setting:

(LOWRANK, TVC)

LOWRANK TVC INFORMATION V1 .1

DISTRIBUTED LIBRARY cluster0(#BA92A)

LOWRANK 65 (C0, C6)

(LOWRANK, COPY)

LOWRANK COPY INFORMATION V1 .1

DISTRIBUTED LIBRARY cluster0(#BA92A)

LOWRANK 64 (C6)

If the request (TVC/COPY) is issued to include all clusters in the Grid, the following text is returned (the example is the command for 4-way Grid (C0/C2/C5/C6)):

INVALID CLUSTER MASK 101. AT LEAST ONE CLUSTER MUST BE EXCLUDED.

If the request (COPY) is issued to include the local cluster, the following text is returned (the example is the command to C0):

INVALID CLUSTER MASK 1. LOCAL CLUSTER MUST BE EXCLUDED.

If the request (TVC/COPY) is issued to include the non-existing cluster, the following text is returned:

THE CLUSTER 10 IN THE MASK DOES NOT EXIST.

Line	Bytes	Name	Description
1	0:7	Header Info	'LOWRANK TVC INFORMATION V1 .1'
	8:11 or 8:12	Request type	'TVC ' : TVC request 'COPY ' : COPY request
	12:24 or 13:25	Header Info	'INFORMATION V'
	25:26 or 26:27	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.33.x.x.
	27 or 28	Dot	'.'
	28:29 or 29:30	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 at 8.33.x.x. The revision is incremented to 1 at 8.40.x.x.
	30:69 or 31:69	Blanks	
If the library request is successfully completed, the following lines are returned			
2	0	Blank	
	0:7	Header Info	'LOWRANK '
	8 or 8:9 or 8:10	Cluster mask value	The cluster mask value which is successfully set.
	9:39 or 10:40 or 11:41	Header Info	' IS SET ON DISTRIBUTED LIBRARY '
	40:69 or 41:69 or 42:69	Distributed library name	Distributed library name to set the value. The value is left justified and padded with blanks.

3.1.20 GGM

At the code level 8.33.0.x, Grid to Grid Migration (GGM) function is supported. GGM is used to replicate the logical volumes from one Grid to another Grid by using the existing Grid network interfaces. The control and monitor of all the GGM function can be done through LI REQ commands.

As part of the code release level 8.41.200.x (R4.1.2 GA), a new logical volume format was introduced. This new format must be supported by all clusters in a grid before it can be used. Therefore, any CTG must support this new format when a CSG has one or more volumes using the new format. It's prohibited to GGM copy any data from a 4.1.2 or later CSG to a CTG then the CTG has one or more clusters running a level prior to 4.1.2.

If the CSG copy source cluster is at 8.42.x.x, the CSG copy target cluster in CTG needs to be at 8.42.x.x or later because of the internal data structure change at 8.42.x.x. Also, a copy source volume may need to be recalled from cloud, then the new version of vtd_exec is required.

Please refer to “IBM TS7700 Series Grid To Grid Migration User's Guide” on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.21 DIAGDATA

At the code level of 8.33.2.9 (3.3 PGA2) (as of January, 2017), LI REQ option, DIAGDATA is supported. DIAGDATA option provides the command handshake elapsed time between two clusters in the Grid. The provided data may be used to diagnose and determine a problematic cluster in a Grid which could cause elongated mount times or slow running jobs.

At the code level of 8.41.200.x (R4.1.2 GA), DIAGDATA option provides the internal token handshake elapsed time instead of the miscellaneous handshake when all clusters in the Grid are at R4.1.2 or above.

Please refer to “IBM TS7700 Series z/OS Host Command Line Request DIAGDATA Guidance” on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.22 OTM

The OTM command provides a method for a user to enable an ownership takeover mode. If a TS7700 Cluster fails, ownership transfer of volumes owned by that cluster is not possible until an ownership takeover method is enabled. This can occur through the Autonomic Ownership Takeover Manager (AOTM) mechanism, through the web based management interface (MI) and through the LIBRARY REQUEST OTM command. Just like AOTM or the MI, one of the following two modes can be set by the OTM command. Just like the MI, a task is created when the mode is set by the OTM command.

Read Ownership Takeover:

When Read Ownership Takeover (ROT) is enabled for a failed cluster, volume ownership is allowed to be taken from a failed TS7700 Cluster by any other available TS7700 Cluster. Only read access to the volume is allowed through the remaining TS7700 Clusters in the event the volume is taken over. Once ownership for a volume has been taken in this mode, any operation attempting to modify data on that volume or change its attributes fails.

Write Ownership Takeover:

When Write Ownership Takeover (WOT) is enabled for a failed TS7700 Cluster, volume ownership is allowed to be taken from a failed TS7700 Cluster. Full access is allowed through the remaining available TS7700 Clusters in the Grid, meaning all volume attributes and the logical volume contents can be updated.

A cluster can have ROT or WOT enabled against it only when the cluster has failed. If the TS7700 Grid believes the cluster is still operational, an ROT or WOT request will be denied.

A cluster in ROT can be upgraded to WOT if needed. A cluster can also be moved from WOT to ROT, though all previously WOT taken volumes will remain in a full read-write status. Only future volumes taken will be read-only.

A cluster in ROT or WOT state can also be disabled so neither takeover mode is allowed. All volumes previously taken over remain in a taken state.

The three methods to enable ROT or WOT are all cross compatible. Meaning, AOTM, MI and LI REQ OTM can be used to achieve the same results as well as change or disable modes enabled by a different method. For example, if ROT is enabled to AOTM, WOT can later be enabled using the LI REQ OTM command.

The following table lists the keywords supported for the OTM request.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7700D	7740	7700T	7700C
OTM	SET	DISABLE/ ROT/ WOT	0-7 (cluster ID)	N/A	Y	Y	Y	Y	Y

	SHOW			N/A	Y	Y	Y	Y	Y
--	------	--	--	-----	---	---	---	---	---

1. SET, DISABLE

The second and third keywords “SET, DISABLE” disable ownership takeover for the cluster specified by the fourth keyword.

The response lines are formatted as follows:

```
OWNERSHIP TAKEOVER MODE CAPABILITIES V1 .0
```

```
DISABLING TAKEOVER MODE SUBMITTED
```

```
TO DISRIBUTED LIBRARY Cluster0(#BA92D)
```

If the cluster specified in the fourth keyword is currently available, the following text is returned:

```
DISABLING TAKEOVER MODE NOT SUBMITTED.
```

```
TARGET DISTRIBUTED LIBRARY Cluster0(#BA92D) IS CURRENTLY AVAILABLE
```

If the cluster specified in the fourth keyword is in service mode, the following text is returned:

```
DISABLING TAKEOVER MODE NOT SUBMITTED.
```

```
TARGET DISTRIBUTED LIBRARY Cluster0(#BA92D) IS IN SERVICE MODE
```

The following table details the format of the data reported when “SET,DISABLE” is specified in the second and third keywords:

Line	Bytes	Name	Description
1	0:37	Header Info	'OWNERSHIP TAKEOVER MODE CAPABILITIES V'
	38:39	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 (at the code level 8.41.x.x).
	40	Dot	'.'
	41:42	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.41.x.x)
	43:69	Blanks	
2	0:32		If the operation starts successfully, the following text is returned: 'DISABLING TAKEOVER MODE SUBMITTED'
	33:69	Blanks	
3	0:20		'TO DISRIBUTED LIBRARY'
	21	Blank	
	22:69	Distributed library name	Distributed library name to disable ownership takeover. The value is left justified and padded with blanks.

2. SET, ROT/WOT

The second and third keywords “SET, ROT” enable Read Ownership Takeover for the cluster specified by the fourth keyword.

The response lines are formatted as follows:

```
OWNERSHIP TAKEOVER MODE CAPABILITIES V1 .0
```

ROT TAKEOVER MODE CHANGE SUBMITTED
 TO DISRIBUTED LIBRARY Cluster0(#BA92D)

If the cluster specified in the fourth keyword is currently available, the following text is returned:

ROT NOT ENABLED.
 TARGET DISTRIBUTED LIBRARY Cluster0(#BA92D) IS CURRENTLY AVAILABLE

If the cluster specified in the fourth keyword is in service mode, the following text is returned:

ROT NOT ENABLED.
 TARGET DISTRIBUTED LIBRARY Cluster0(#BA92D) IS IN SERVICE MODE

The second and third keywords “SET, WOT” enable Write Ownership Takeover for the cluster specified by the fourth keyword.

The response lines are formatted as follows:

OWNERSHIP TAKEOVER MODE CAPABILITIES V1 .0
 WOT TAKEOVER MODE CHANGE SUBMITTED
 TO DISRIBUTED LIBRARY Cluster0(#BA92D)

If the cluster specified in the fourth keyword is currently available, the following text is returned:

WOT NOT ENABLED.
 TARGET DISTRIBUTED LIBRARY Cluster0(#BA92D) IS CURRENTLY AVAILABLE

If the cluster specified in the fourth keyword is in service mode, the following text is returned:

WOT NOT ENABLED.
 TARGET DISTRIBUTED LIBRARY Cluster0(#BA92D) IS IN SERVICE MODE

The following table details the format of the data reported when “SET,ROT” or “SET,WOT” is specified in the second and third keywords:

Line	Bytes	Name	Description
1	0:37	Header Info	'OWNERSHIP TAKEOVER MODE CAPABILITIES V'
	38:39	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 (at the code level 8.41.x.x).
	40	Dot	'.'
	41:42	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.41.x.x)
	43:69	Blanks	
2	0:2	Ownership takeover mode	'ROT' or 'WOT'
	3	Blank	
	4:33		'TAKEOVER MODE CHANGE SUBMITTED'
3	0:20		'TO DISRIBUTED LIBRARY'
	21	Blank	

	22:69	Distributed library name	Distributed library name to disable ownership takeover. The value is left justified and padded with blanks.
--	-------	--------------------------	---

3. SHOW

The second keyword “SHOW” provides the current ownership takeover status. No third and fourth keywords are supported.

The response lines are formatted as follows:

```
OWNERSHIP TAKEOVER MODE CAPABILITIES V1 .0
DISTRIBUTED LIBRARY Cluster0(#BA92A)
```

CLUSTER	0	1	2	3	4	5	6	7
ConState	N	NA	N	NA	NA	N	F	NA
TOMode	NA	NA	NA	NA	NA	NA	ROT	NA

Legend

ConState: Connection State

N: The cluster is in normal state.

S: The cluster is in service mode.

F: The cluster is in failed state.

NA: The cluster is not configured.

TOMode: Current Ownership Takeover Mode

NA: No Ownership Takeover

ROT: Read Ownership Takeover

WOT: Write Ownership Takeover

The following table details the format of the data reported when “SHOW” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:37	Header Info	'OWNERSHIP TAKEOVER MODE CAPABILITIES V'
	38:39	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 (at the code level 8.41.x.x).
	40	Dot	'.'
	41:42	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.41.x.x)
	43:69	Blanks	
2	0:18		'DISTRIBUTED LIBRARY'
	19	Blank	
	20:69	Distributed library name	Distributed library name to which the request is issued. The value is left justified and padded with blanks.

3	0:55	Separator	All dash '-' characters
	56:69	Blanks	
4	0:6	Header Info	'CLUSTER'
	7:12	Blanks	
	13	Header Info	'0'
	14:18	Blanks	
	19	Header Info	'1'
	20:24	Blanks	
	25	Header Info	'2'
	26:30	Blanks	
	31	Header Info	'3'
	32:36	Blanks	
	37	Header Info	'4'
	38:42	Blanks	
	43	Header Info	'5'
	44:48	Blanks	
	49	Header Info	'6'
	50:54	Blanks	
	55	Header Info	'7'
	56:69	Blanks	
5	0:55	Separator	All dash '-' characters
	56:69	Blanks	
6	0:7	Connection State Info	'ConState'
	8:11	Blanks	
	12:13	Connection State of cluster-0.	'N': The cluster is in normal state. 'S': The cluster is in service mode. 'F': The cluster is in failed state. 'NA': The cluster is not configured. The value is right justified and padded with blanks.
	14:17	Blanks	
	18:19	Connection State of cluster-1.	'N': The cluster is in normal state. 'S': The cluster is in service mode. 'F': The cluster is in failed state. 'NA': The cluster is not configured. The value is right justified and padded with blanks.
	20:23	Blanks	
	24:25	Connection State of cluster-2.	'N': The cluster is in normal state. 'S': The cluster is in service mode.

			'F' : The cluster is in failed state. 'NA' : The cluster is not configured. The value is right justified and padded with blanks.
	26:29	Blanks	
	30:31	Connection State of cluster-3.	'N' : The cluster is in normal state. 'S' : The cluster is in service mode. 'F' : The cluster is in failed state. 'NA' : The cluster is not configured. The value is right justified and padded with blanks.
	32:35	Blanks	
	36:37	Connection State of cluster-4.	'N' : The cluster is in normal state. 'S' : The cluster is in service mode. 'F' : The cluster is in failed state. 'NA' : The cluster is not configured. The value is right justified and padded with blanks.
	38:41	Blanks	
	42:43	Connection State of cluster-5.	'N' : The cluster is in normal state. 'S' : The cluster is in service mode. 'F' : The cluster is in failed state. 'NA' : The cluster is not configured. The value is right justified and padded with blanks.
	44:47	Blanks	
	48:49	Connection State of cluster-6.	'N' : The cluster is in normal state. 'S' : The cluster is in service mode. 'F' : The cluster is in failed state. 'NA' : The cluster is not configured. The value is right justified and padded with blanks.
	50:53	Blanks	
	54:55	Connection State of cluster-7.	'N' : The cluster is in normal state. 'S' : The cluster is in service mode. 'F' : The cluster is in failed state. 'NA' : The cluster is not configured. The value is right justified and padded with blanks.
	56:69	Blanks	
7	0:5	Ownership Takeover Mode Info	'TOMode'
	6:10	Blanks	
	11:13	Current Ownership	'NA' : No Ownership Takeover

		Takeover Mode of cluster-0.	'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	14:16	Blanks	
	17:19	Current Ownership Takeover Mode of cluster-1.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	20:22	Blanks	
	23:25	Current Ownership Takeover Mode of cluster-2.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	26:28	Blanks	
	29:31	Current Ownership Takeover Mode of cluster-3.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	32:34	Blanks	
	35:37	Current Ownership Takeover Mode of cluster-4.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	38:40	Blanks	
	41:43	Current Ownership Takeover Mode of cluster-5.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	44:46	Blanks	
	47:49	Current Ownership Takeover Mode of cluster-6.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	50:52	Blanks	
	53:55	Current Ownership Takeover Mode of cluster-7.	'NA' : No Ownership Takeover 'ROT' : Read Ownership Takeover 'WOT' : Write Ownership Takeover The value is right justified and padded with blanks.
	56:69	Blanks	

3.1.23 FENCE

As part of code release level 8.41.200.x (R4.1.2 GA), the LI REQ FENCE option is introduced. The FENCE option and all its underlying settings provide the capability to configure whether unhealthy remote clusters can be fenced in the event certain criteria are met. This enhancement is part of the Grid Resiliency improvements added in the 4.1.2 level of code.

Please refer to “IBM TS7700 Series Grid Resiliency Improvements User's Guide” on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.24 CUIR/LDRIVE

As part of code release level 8.41.200.x (R4.1.2 GA), the CUIR (Control Unit Initiated Reconfiguration) function was introduced. The LI REQ CUIR and LDRIVE options are introduced as part of this enhancement. Through the LI REQ (keyword “CUIR”) command, different CUIR functions can be enabled, disabled and configured. In addition, the LI REQ “LDRIVE” commands provide a series of new outputs summarizing host LPAR centric information about connected devices. For example, it can be used to determine which LPARs are connected to which devices and which devices are varied online to a given LPAR. This information was initially introduced to help track CUIR events, but it can be used in any system z configuration to simply view the current pathing and device states of your environment.

Please refer to “IBM TS7700 Series CUIR User's Guide” on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.25 CLDSET

As part of code release level 8.42.x.x (R4.2 GA), the cloud attachment function was introduced. The LI REQ CLDSET options provide the ability to manage data movement between the TS7700C and the cloud device(s). For example, it can be used to change the number of concurrent data premigration tasks to cloud. Also it can enable/disable the data premigration/recall to/from cloud.

There are also CLDSET options that provide information about many of the current cloud data workflow and management settings. The CLDSET option is only applicable to TS7700C. In the response of the CLDSET request, the settings on the cluster associated with the distributed library will be modified based on the keywords specified. If no additional keywords are specified, the request will just return the current settings.

A best practice of attaching cloud to TS7700 is available on Techdocs. Search for TS7700.
<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

Note: All settings are persistent across machine restarts, service actions or code updates. The settings are not carried forward as part of Disaster Recovery from Copy Exported tapes or the recovery of a system.

All requests are applicable to a TS7700C distributed library only. If the distributed library specified is not a TS7700C, the following error text is returned:

```
'ONLY SUPPORTED IN CLOUD ENABLED TS7700 VIRTUALIZATION ENGINE'
```

If the composite library is specified, the following error text is returned:

```
'REQUEST INVALID FOR COMPOSITE LIBRARY'
```

3.1.25.1 CPMCNTH/CPMCNTL

If a second keyword of CPMCNTH or CPMCNTL is specified, the cluster will set the number of premigrate processes which runs in parallel to premigrate the logical volumes from the cluster to cloud.

Keyword 2	Keyword 3	Description
CPMCNTH	Value	Cloud Premigration Count High This is the highest number of premigration processes that the TS7700C initiates in parallel at any given time when the premigration to cloud needs to be done in a high priority mode. The high/low priority mode of cloud premigration is described in the CLDPRIOR section.

Keyword 2	Keyword 3	Description
		The default value is 40. The maximum value is 128 and the minimum value is 1. A value lower than CPMCNTL cannot be set. If it is attempted, CPMCNTL will automatically be set to the same value as CPMCNTL. If the provided value is out of range (less than 1, or more than 128), the following error is returned: 'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'
CPMCNTL	Value	Cloud Premigration Count Low This is the lowest number of premigration processes that the TS7700C initiates in parallel at any given time when the premigration to cloud needs to be done in a low priority mode. The high/low priority mode of cloud premigration is described in the CLDPRIOR section. The default value is 0. The maximum value 128 and the minimum value is 0. A value higher than CPMCNTL cannot be set. If it is attempted, CPMCNTL will automatically be set to the same value as CPMCNTL. If the provided value is out of range (less than 0, or more than 128), the following error is returned: 'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'

3.1.25.2 CLDPRIOR

If a second keyword of CLDPRIOR is specified, the cluster will set the threshold to change TS7700C cloud data premigration behavior (high or low priority mode).

Keyword 2	Keyword 3	Description
CLDPRIOR	Value	Cloud Premigration Priority Threshold

Keyword 2	Keyword 3	Description
		This is the threshold, in GBs of unpremigrated data to cloud at which the TS7700C will begin increasing the number of cloud premigration tasks that will be allowed to compete with host I/O for cache and processor cycles. The amount of cloud unpremigrated data must be above the value specified in keyword 3 for 150 seconds before the additional cloud premigration tasks are added. As the amount of data to premigrate to cloud continues to grow above this threshold setting, so do the number of enabled cloud premigration tasks until the maximum is reached. If the amount of cloud unpremigrated data subsequently falls below this threshold for at least 150 seconds, the number of cloud premigration tasks may be reduced depending on host I/O demand. If I/O host demand is high, the number of premigration tasks will eventually be reduced to a minimum of one. The default value is 0. The maximum value can be set up to the total size of the active premigration queue. For example, if FC 5274 (1 TB Active Premigration Queue) x 10 plus FC 5279 (5 TB Active Premigration Queue) feature codes are installed, the total size of the active premigration queue is 15 TB, then up to $15 * 1000 = 15000$ in GBs can be set. A value higher than the total size of the active premigration queue (P) cannot be set. If it is attempted, CPMCNTL will automatically be set to "P".

3.1.25.3 CRCCNT

If a second keyword of CRCCNT is specified, the cluster will set the number of recall processes which runs in parallel to recall the object data from cloud to the cluster.

Keyword 2	Keyword 3	Description
CRCCNT	Value	Cloud Recall Count This is the number of recall processes that the TS7700C initiates in parallel at any given time when the recall from cloud needs to be done. The default value is 20. The maximum value is 32 and the minimum value is 1. If the provided value is out of range (less than 1, or more than 32), the following error is returned:

Keyword 2	Keyword 3	Description
		'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'

3.1.25.4 CDELCNT

If a second keyword of CDELCNT is specified, the cluster will set the number of stale (unnecessary) data delete processes on cloud which runs in parallel from the cluster.

Keyword 2	Keyword 3	Description
CDELCNT	Value	Cloud Delete Count This is the number of stale (unnecessary) data delete processes that the TS7700C initiates in parallel at any given time when the object data deletion on cloud needs to be done. The default value is 0, which no user setting is applied and the TS7700C uses an internal default value of 5 to delete object data on cloud. The maximum value is 16 and the minimum value is 1. If the provided value is out of range (more than 16), the following error is returned: 'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'

3.1.25.5 CPMTOUT

If a second keyword of CPMTOUT is specified, the cluster will set the cloud premigration timeout value which represents the maximum amount of time, in seconds, a premigration process will wait to move 1GiB of data from the cluster to cloud.

Keyword 2	Keyword 3	Description
CPMTOUT	Value	Cloud Premigration Timeout This is the timeout value, in seconds, to premigrate 1GiB of data from the cluster to cloud. For example, if a 6GiB volume is premigrated, the premigration process will wait up to 6 x this timeout value before it times out. The more tasks TS7700 runs in parallel, the longer the TS7700C can take to complete premigrations. Therefore, when adjusting this timeout value, the maximum number of premigration processes defined by CPMCNTL/ CPMCNTL should be taken into account.

Keyword 2	Keyword 3	Description
		When a cloud premigration timeout occurs, an event is posted to MI: <pre>'Cloud pre-migration for virtual volume <volser> to cloud pool <cloud_pool_nickname > timed out with <timeout> seconds, where CPMTOUT is <CPMTOUT> and size of the virtual volume is <lvol size> GiB'</pre> (*) Given multiple cloud premigration timeouts can occur at a time, only 1 event will be posted at a specific interval. New events can be posted every interval. The default value is 1800 (seconds) (i.e. 30 minutes per 1GiB of data). The maximum value is 99999999 and the minimum value is 1. If the provided value is out of range (less than 1), the following error is returned: <pre>'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'</pre> <pre>'Cloud recall for virtual volume <volser> from cloud pool <cloud_pool_nickname > timed out with <timeout> seconds, where CRCTOUT is <CPCTOUT> and size of the virtual volume is <lvol size> GiB'</pre>

3.1.25.6 CRCTOUT

If a second keyword of CRCTOUT is specified, the cluster will set the cloud recall timeout value which represents the maximum amount of time, in seconds, a recall process will wait to move 1GiB of data from cloud to the cluster.

Keyword 2	Keyword 3	Description
CRCTOUT	Value	Cloud Recall Timeout This is the timeout value, in seconds, to recall 1GiB of data from cloud to the cluster. For example, if a 6GiB volume is recalled, the recall process will wait up to 6 x this timeout value before it times out. The more tasks TS7700 runs in parallel, the longer the TS7700C can take to complete recalls. Therefore when adjusting this timeout

Keyword 2	Keyword 3	Description
		value, the maximum number of recall processes defined by CRCCNT should be taken into account. When a cloud recall timeout occurs, an event is posted to MI: 'Cloud recall for virtual volume <volser> from cloud pool <cloud_pool_nickname> timed out with <timeout> seconds, where CRCTOUT is <CPCTOUT> and size of the virtual volume is <lvol size> GiB' (*) Given multiple cloud recall timeouts can occur at a time, only 1 event will be posted at a specific interval. New events can be posted every interval. The default value is 1800 (seconds) (i.e. 30 minutes per 1GiB of data). The maximum value is 42900 and the minimum value is 1. If the provided value is out of range (less than 0 or more than 42900), the following error is returned: 'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'

3.1.25.7 CDELTOUT

If a second keyword of CDELTOUT is specified, the cluster will set the stale (unnecessary) cloud data delete timeout value which represents the maximum amount of time, in seconds, a delete process will wait to delete an object on cloud.

Keyword 2	Keyword 3	Description
CDELTOUT	Value	Cloud Delete Timeout This is the timeout value, in seconds, to delete an object on cloud which is no longer required because it's already stale data level. When a cloud data delete timeout occurs, an event is posted to MI: 'Deleting object <volser> from container <container_name> of cloud pool <cloud_pool_nickname> timed out with CDELTOUT (<CDELTOUT> seconds)'

Keyword 2	Keyword 3	Description
		The default value is 3600 (seconds) (1 hour). The maximum value is 99999999 and the minimum value is 1. If the provided value is out of range (less than 1), the following error is returned: <code>'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 2'</code>

3.1.25.8 CENABLMT

If a second keyword of CENABLMT is specified, the cluster will enable/disable the cloud data handling operations.

At the code level of 8.42.x.x, the request enables or disables all three of the cloud data operations (cloud data premigration/recall/delete). The ability to enable or disable each operation individually is a future item and is not available at this time. The third keyword “ALL” is only supported but if the third keyword of CLDPM, CLDRCALL or CLDDEL is specified, it will be accepted (with no error), however, all three operations will always be affected (either all enabled or all disabled).

Keyword 2	Keyword 3	Keyword 4	Description
CENABLMT	ALL	ENABLE/ DISABLE	Cloud Enablement This is to enable or disable the cloud data handling operations (cloud data premigration/recall/delete) to/from the cluster. When it's set to “DISABLED”, no cloud data premigration/recall/delete to/from the cluster will occur. The default is “ENABLED”.

If any invalid second keyword is used (or invalid third or fourth keyword with the second CENABLMNT), the following output is returned to provide all supported keywords:

```
INVALID/UNKNOWN KEYWORD
SUPPORTED KEYWORDS
CLDSET
CLDSET,CLDPRIOR,val
CLDSET,CPMCNTH,val
CLDSET,CPMCNTL,val
CLDSET,CRC CNT,val
CLDSET,CDELCNT,val
CLDSET,CPMTOUT,val
CLDSET,CRCTOUT,val
CLDSET,CDELTOUT,val
CLDSET,PFRCCPG0,[ENABLE|DISABLE]
CLDSET,PFRCCCTL,[ENABLE|DISABLE]
CLDSET,PFRCCTRY,[ENABLE|DISABLE]
CLDSET,CENABLMNT,ALL,[ENABLE|DISABLE]
CLDSET,USEHIGH,cloud_pool_name,val
CLDSET,USELOW,cloud_pool_name,val
CLDSET,USESHOW,val
```

The response lines to a CLDSET request are formatted as follows in this example. Just entering the first CLDSET keyword with no second, third and fourth keywords provides this response as well:

```
CLOUD SETTINGS V2 .0
CPMCNTH   =          20      CPMCNTL   =          5
CLDPRIOR  =           0
CRC CNT   =          20
CDELCNT   =           1
CPMTOUT   =         1800
CRCTOUT   =         1800
CDELTOUT  =         3600
-----
CENABLMNT Controls
CLDPM     =    ENABLED
CLDRCALL  =    ENABLED
CLDDEL    =    ENABLED
```

The following table details the format of the data reported.

Line	Bytes	Name	Description
1	0:15	Header Info	'CLOUD SETTINGS V'
	16:17	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.42.x.x. The version is incremented to 2 at the code level 8.51.x.x.
	18	Dot	'.'
	19:20	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	21:69	Blanks	
2	0	Blank	
	1:10	Header	'CPMCNTH ='
	11:20	Cloud premig count high value	The current CPMCNTH setting value describing the highest number of premigration processes that the TS7700C initiates in parallel at any given time. The number is right justified and padded with blanks.
	21:24	Blanks	
	25:34	Header	'CPMCNTL ='
	35:44	Cloud premig count low value	The current CPMCNTL setting value describing the lowest number of premigration processes that the TS7700C initiates in parallel at any given time. The number is right justified and padded with blanks.
	45:69	Blanks	
3	0	Blank	
	1:10	Header	'CLDPRIOR ='
	11:20	Cloud premig priority threshold value	The current CLDPRIOR setting value describing the threshold, in GBs of unpemigrated data to cloud at which the TS7700C will begin increasing the number of cloud premigration tasks. The number is right justified and padded with blanks.
	21:69	Blanks	
4	0	Blank	
	1:10	Header	'CRCCNT ='
	11:20	Cloud recall count value	The current CRCCNT setting value describing the number of recall processes that the TS7700C initiates in parallel at any given time. The number is right justified and padded with blanks.
	21:69	Blanks	
5	0	Blank	
	1:10	Header	'CDELCNT ='
	11:20	Cloud delete count value	The current CDELCNT setting value describing the number of stale (unnecessary) data delete processes that the TS7700C initiates in parallel at any given time. The number is right justified and padded with blanks.
	21:69	Blanks	
6	0	Blank	
	1:10	Header	'CPMTOUT ='

Line	Bytes	Name	Description
	11:20	Cloud premig timeout value	The current CPMTOUT setting value describing the timeout value, in seconds, to premigrate 1GiB of data from the cluster to cloud. The number is right justified and padded with blanks.
7	0	Blank	
	1:10	Header	`CRCTOUT ='
	11:20	Cloud recall timeout value	The current CRCTOUT setting value describing the timeout value, in seconds, to recall 1GiB of data from cloud to the cluster. The number is right justified and padded with blanks.
	21:69	Blanks	
8	0	Blank	
	1:10	Header	`CDELTOUT ='
	11:20	Cloud delete timeout value	The current CDELTOUT setting value describing the timeout value, in seconds, to delete an object on cloud. The number is right justified and padded with blanks.
	21:69	Blanks	
9	0	Blank	
	1:10	Header	`PFRCCPG0 ='
	11:20	Ghost copy state during PG0 replication	The current ghost copy state during PG0 replication on the local cluster. ENABLED or DISABLED.
	21:69	Blanks	
10	0	Blank	
	1:10	Header	`PFRCCCTL ='
	11:20	Ghost copy state during time delayed replication	The current ghost copy state during time delayed replication on the local cluster. ENABLED or DISABLED.
	21:69	Blanks	
11	0	Blank	
	1:10	Header	`PFRCCTRY ='
	11:20	Number of copy retries to be done before ghost copy	The number of normal copy retries before ghost copy is done. If the value is set to "0", no ghost copy is run.
	21:69	Blanks	
12	0:69	Separator	All dash '-' characters
13	0	Blank	
	1:17	Header Info	`CENABLMT Controls'
	18:69	Blanks	
14	0:3	Blanks	
	4:12	Header	`CLDPM ='
	13:22	Cloud data premig operational state	The current cloud data premig operation state on the local cluster. ENABLED or DISABLED.
	23:69	Blanks	

Line	Bytes	Name	Description
15	0:3	Blanks	
	4:12	Header	'CLDRCALL ='
	13:22	Cloud data recall operational state	The current cloud data recall operation state on the local cluster. ENABLED or DISABLED.
	23:69	Blanks	
16	0:3	Blank	
	4:12	Header Info	'CLDDEL ='
	13:22	Cloud data delete operational state	The current cloud data delete operation state on the local cluster. ENABLED or DISABLED.
	23:69	Blanks	

3.1.25.9 PFRCCPG0

If a second keyword of PFRCCPG0 is specified, the cluster will enable/disable ghost copy during replication of volumes to be assigned PG0 from a remote cluster.

Keyword 2	Keyword 3	Description
PFRCCPG0	ENABLE/ DISABLE	Ghost Copy Enablement during PG0 Replication This is to enable or disable the ghost copy when the local cluster is trying to replicate volumes to be assigned PG0 from a remote cluster. When it's set to "DISABLE", no ghost copy will occur unless another trigger condition is met. The default is "ENABLE".

Note: Ghost copy is a method to let the local cluster be aware of and accessible to a volume already premigrated on the cloud by another cluster. In this method, in contrast to the normal copy, the cluster don't need to recall, copy over the grid network and migrate the volume; just marks the volume as migrated in the local database.

Note: TS7700 treats the copied volumes as PG0 regardless of their storage constructs to save the amount of cache available to host operations. This behavior can be changed to follow the storage constructs by enabling COPYFSC option. If the option is enabled, no ghost copy will occur with the PFRCCPG0 setting.

If any invalid third keyword is used, all supported keyword list is returned (see the "[CENABLMT](#)" section for sample output).

The response lines to a CLDSET request are described in the "[CENABLMT](#)" section.

3.1.25.10 PFRCCCTL

If a second keyword of PFRCCCTL is specified, the cluster will enable/disable ghost copy during time delayed replication of volumes from a remote cluster.

Keyword 2	Keyword 3	Description
PFRCCCTL	ENABLE/ DISABLE	Ghost Copy Enablement during Time Delayed Replication This is to enable or disable the ghost copy when the local cluster is trying time delayed replication of a volume from a remote cluster. If this option is ENABLE, ghost copy will occur when the time delayed copy job expires. When it's set to "DISABLE", no ghost copy will occur unless another trigger condition is met. The default is "ENABLE".

If any invalid third keyword is used, all supported keyword list is returned (see the "[CENABLMT](#)" section for sample output).

The response lines to a CLDSET request are described in the "[CENABLMT](#)" section.

3.1.25.11 PFRCCTRY

If the local cluster got an error during normal copy, it continues to retry the copy until it will successfully complete. If a second keyword of PFRCCTRY is specified, the cluster will set the maximum number of normal copy retries to be run. If all the copy trials failed, then ghost copy is run.

Keyword 2	Keyword 3	Description
PFRCCTRY	Value	Number of Copy Attempts before Ghost Copy This is the maximum number of normal copy retries to be run. If the local cluster got an error during normal copy, it retries normal copy up to PFRCCTRY times until it successfully completes. If all the copy attempts failed, then ghost copy is applied. If this value is set to "0", the local cluster continues to run a normal copy.

Keyword 2	Keyword 3	Description
		The default value is 0. The maximum value is 65535 and the minimum value is 0.

If any invalid third keyword is used, all supported keyword list is returned (see the "[CENABLMT](#)" section for sample output).

The response lines to a CLDSET request are described in the "[CENABLMT](#)" section.

3.1.25.12 USEHIGH

If a second keyword of USEHIGH is specified, the cluster will set the high warning level of amount of data stored in a cloud that the cloud nickname points to.

Keyword 2	Keyword 3	Keyword 4	Description
USEHIGH	Cloud Nickname	Value	High Warning Level of Amount of Data Stored in Cloud Set the high warning level of amount of data stored in a cloud that the cloud nickname points to. The specified value is recognized in TB unit. When the cloud usage exceeds the USEHIGH value for more than 5 minutes, an event message "The amount of data stored in the cloud pool crossed USEHIGH value." is created. The message will repeat every 15 minutes. When the cloud usage falls below USEHIGH, another event message "The amount of data stored in the cloud pool is below USEHIGH value." is created. The default value is 0. If the value is set to 0, no operator intervention is surfaced.

On successful completion, the current USEHIGH/USELOW setting values are displayed in the following format.

```
CLOUD POOL SETTINGS V1 .0
```

```
NICKNAME =      MYPOOL      USEHIGH  =           0      USELOW   =           0
```

Line	Bytes	Name	Description
1	0:20	Header Info	'CLOUD POOL SETTINGS V'
	21:22	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.51.x.x.
	23	Dot	'.'
	24:25	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.

Line	Bytes	Name	Description
	26:69	Blanks	
2	0	Blank	
	1:10	Header	`NICKNAME` ='
	11:20	Cloud Nickname	The nickname of the cloud pool to set the warning level.
	21:24	Blanks	
	25:34	Header	`USEHIGH` ='
	35:44	High Warning Level	The high warning level of amount of data stored in a cloud that the cloud nickname points to.
	45:48	Blanks	
	49:58	Header	`USELOW` ='
	59:68	Low Warning Level	The low warning level of amount of data stored in a cloud that the cloud nickname points to.
	69	Blank	

If any invalid keyword is used, all supported keyword list is returned (see the "[CENABLMT](#)" section for sample output).

3.1.25.13 USELOW

If a second keyword of USELOW is specified, the cluster will set the low warning level of amount of data stored in a cloud that the cloud nickname points to.

Keyword 2	Keyword 3	Keyword 4	Description
USELOW	Cloud Nickname	Value	Low Warning Level of Amount of Data Stored in Cloud Set the low warning level of amount of data stored in a cloud that the cloud nickname points to. The specified value is recognized in TB unit. When the cloud usage exceeds the USELOW value for more than 5 minutes, an event message “The amount of data stored in the cloud pool crossed USELOW value.” is created. When the cloud usage falls below USELOW, another event message “The amount of data stored in the cloud pool is below USELOW value.” is created. The default value is 0. If the value is set to 0, no operator intervention is surfaced.

On successful completion, the current USEHIGH/USELOW setting values are displayed in the following format.

```
CLOUD POOL SETTINGS V1 .0
```

```
NICKNAME =      mypool      USEHIGH =          0      USELOW  =          0
```

Line	Bytes	Name	Description
1	0:20	Header Info	'CLOUD POOL SETTINGS V'
	21:22	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.51.x.x.
	23	Dot	'.'
	24:25	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	26:69	Blanks	
2	0	Blank	
	1:10	Header	'NICKNAME ='
	11:20	Cloud Nickname	The nickname of the cloud pool to set the warning level.
	21:24	Blanks	
	25:34	Header	'USEHIGH ='
	35:44	High Warning Level	The high warning level of amount of data stored in a cloud that the cloud nickname points to.
	45:48	Blanks	
	49:58	Header	'USELOW ='
	59:68	Low Warning Level	The low warning level of amount of data stored in a cloud that the cloud nickname points to.
	69	Blank	

If any invalid keyword is used, all supported keyword list is returned (see the "[CENABLM T](#)" section for sample output).

3.1.25.14 USESHOW

If USESHOW <page index> is specified, the cluster will display the list of USEHIGH and USELOW settings for all the cloud pools. The list is separated into pages made up of 141 records using 3 columns per each. Only one page can be displayed at a time, and the page index can be specified by the third keyword. The default page index 1 is applied when the third keyword is not set.

The list of USEHIGH/USELOW setting values are displayed in the following format.

```
CLDSET USESHOW V1 .0
```

```

POOLNAME  UHIGH    ULOW  POOLNAME  UHIGH    ULOW  POOLNAME  UHIGH    ULOW
-----
MYPOOL0      0      0  MYPOOL1      0      0  MYPOOL2     200     10
MYPOOL3      0      0  MYPOOL4     200     10  MYPOOL5     200     10
0 more cloud pool usage settings exist
```

Line	Bytes	Name	Description
1	0:15	Header Info	'CLDSET USESHOW V'

Line	Bytes	Name	Description
	16:17	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 at 8.51.x.x.
	18	Dot	`.'`
	19:20	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0.
	21:69	Blanks	
2	0:69	Header	`POOLNAME UHIGH ULOW POOLNAME UHIGH ULOW POOLNAME UHIGH ULOW`
3	0:69	Separator	`----- - -----`
The USEHIGH/USELOW setting values are displayed in the following format. Each line could provide up to 3 cloud pool setting values. Up to 47 lines (N is from 1 to 47) are displayed.			
3+N	0:7	Cloud Pool Nickname	The cloud pool nickname.
	8	Blank	
	9:14	USEHIGH Value	The current USEHIGH value.
	15	Blank	
	16:21	USELOW Value	The current USELOW value.
	22:23	Blanks	
	24:31	Cloud Pool Nickname	The cloud pool nickname.
	32	Blank	
	33:38	USEHIGH Value	The current USEHIGH value.
	39	Blank	
	40:45	USELOW Value	The current USELOW value.
	46:47	Blanks	
	48:55	Cloud Pool Nickname	The cloud pool nickname.
	56	Blank	
	57:62	USEHIGH Value	The current USEHIGH value.
	63	Blank	
	64:69	USELOW Value	The current USELOW value.
The following message is displayed in the end of the output.			
		Message	`X more cloud pool usage settings exist` where `X` is the number of remaining setting values to show, having 3 characters width, right justified and padded with blanks. The message string itself is left justified and padded with blanks.

If any invalid keyword is used, all supported keyword list is returned (see the "[CENABLMT](#)" section for sample output).

3.1.26 CLDVR

As part of code release level of 8.51.1.x (R5.1PGA1), the Logical Volume Version Restore function was introduced. This function provides the capability to restore older versions of volumes retained in the cloud. The LI REQ CLDVR commands are provided to configure and perform restore operations against the retained versions in the cloud.

Please refer to “IBM TS7700 Cloud Storage Tier Export Recovery and Testing Guide” on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.27 SERVICE

The SERVICE command provides a method for a user to initiate Service-Prep. When an operational TS7700 cluster must be taken offline for service, the TS7700 Grid first must be prepared for the loss of the resources that are involved to provide continued access to data. The controls to prepare a TS7700 for service (Service-Prep) are provided through the web based management interface (MI) and through the LIBRARY REQUEST SERVICE command. Just like the MI, a task is created when a Service-Prep is initiated by the SERVICE command. The task is created with user of “_LiReqService_”.

The following table lists the keywords supported for the SERVICE request.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7700D	7740	7700T	7700C
SERVICE	ENTER	FORCE		N/A	Y	Y	N/A	Y	Y
	CANCEL			N/A	Y	Y	N/A	Y	Y
	SHOW			N/A	Y	Y	N/A	Y	Y

1. ENTER

The second keyword ENTER initiates Service-Prep on a TS7700 cluster of the distributed library. If optional third keyword FORCE is specified, it initiates a Forced Service-Prep. Since it may cause data conflict between clusters, using FORCE option is not recommended, however, you can select the option if you believe an operation has stalled and is preventing the cluster from entering Service-Prep. No fourth keyword is supported.

The response lines are formatted as follows:

```
SERVICE ENTER V1 .0
Service Prep on cluster1 has been initiated.
```

If the cluster is not online, the following text is returned:

```
DISTRIBUTED LIBRARY UNAVAILABLE
```

If another cluster in the grid is in Service-Prep, the following text is returned:

```
ANOTHER CLUSTER IS GOING TO ENTER SERVICE
```

If no other cluster remains online, the following text is returned:

```
THE CLUSTER IS THE LAST ONLINE CLUSTER
```

If there is at least one cluster with paused or offline library, the following text is returned:

```
THERE IS AT LEAST ONE CLUSTER WITH PAUSED OR OFFLINE LIBRARY
```

The following table details the format of the data reported when ENTER is specified in the second and third keywords:

Line	Bytes	Name	Description
1	0:14	Header Info	'SERVICE ENTER V'

	15:16	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 (at the code level 8.50.x.x).
	17	Dot	`.'`
	18:19	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.50.x.x)
	20:69	Blanks	
2	0:22		If the operation starts successfully, the following text is returned: `Service Prep on cluster`
	23:23	Cluster ID	The cluster ID.
	24:43		` has been initiated.`
	44:69	Blanks	

2. CANCEL

The second keyword CANCEL is intended to cancel Service-Prep as long as Service-Prep is not completed. It cannot be used to cancel Service if Service-Prep is completed. No third and fourth keywords are supported.

The response lines are formatted as follows:

```
SERVICE CANCEL V1 .0
Service Cancel on cluster1 has been initiated.
```

If the cluster is not online, the following text is returned:

```
DISTRIBUTED LIBRARY UNAVAILABLE
```

If Service-Prep is not in progress on the cluster, the following text is returned:

```
THE CLUSTER IS NOT IN SERVICE
```

If no other cluster remains online, the following text is returned:

```
THE CLUSTER IS THE LAST ONLINE CLUSTER
```

If another cancellation task is in progress, the following text is returned:

```
ANOTHER CANCELLATION TASK IS IN PROGRESS
```

The following table details the format of the data reported when CANCEL is specified in the second and third keywords:

Line	Bytes	Name	Description
1	0:15	Header Info	`SERVICE CANCEL V`
	16:17	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 (at the code level 8.50.x.x).
	18	Dot	`.'`
	19:20	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.50.x.x)
	21:69	Blanks	
2	0:24		If the operation starts successfully, the following text is returned:

			'Service Cancel on cluster'
	25:25	Cluster ID	The cluster ID.
	26:45		' has been initiated.'
	46:69	Blanks	

3. SHOW

The second keyword SHOW provides the current Service-Prep status. No third and fourth keywords are supported.

If the cluster is Offline, Service or finalizing Service-Prep, the following text is returned:

```
DISTRIBUTED LIBRARY UNAVAILABLE
```

The response lines are formatted in three sections.

The first section is formatted as follows:

```
SERVICE STATUS V1 .0
```

```
-----
Current service state: SERVICE-PREP
Cluster1: SERVICE-PREP
Cluster2: NORMAL
-----
```

The definitions of the states are:

NORMAL: The cluster indicated is functional, and system should be online/operational. If the system was varied offline without properly initiating service mode, then this definition is not accurate.

SERVICE-PREP: The cluster indicated is currently transitioning to SERVICE mode. Do NOT perform any disruptive actions on the indicated cluster (such as offline) until the state is SERVICE.

SERVICE: The cluster indicated is in SERVICE mode. This does NOT mean the cluster is offline, however actions can be performed on the cluster without disrupting the Grid.

If the current service state is not SERVICE-PREP, the second and the third sections are omitted.

If it has just started Service Prep processing, the third section is omitted, the second section is formatted as follows:

```
Starting service preparation processing.
```

If it is about to start finalizing Service Prep, the third sections is omitted, the second section is formatted as follows:

```
Finalizing service preparation.
```

If service cancellation has been requested, the third sections is omitted, the second section is formatted as follows:

```
Beginning service cancellation.
```

If Service Prep is being processed, the second section is formatted as follows:

Service preparation will continue until
the following operations have completed:

If Service Prep is being processed, the third section is formatted as follows:

The cluster being prepared has:

```

1 devices mounted,
1 pending mounts, demounts or ejects,
0 LPARs which still wait for CUIR to vary-off of    0 devices,
1 LPARs which may require manual vary-off of    10 devices,
7 pending token updates against peer clusters,
8 pending token updates against unavailable cluster,
1 global resources it must surrender to peer clusters,
9 copies currently in progress and
1 copy thread change request in progress.

```

The peer sites have:

```

1 devices remotely mounted to this cluster,
0 commands in progress,
0 immediate copies or          0 MiBs of data pending,
16 deferred copies or          180 MiBs of data pending, and
4 copies currently in progress against this cluster.

```

The grid has:

```

0 logical volume deletes, imports, or exports in progress.

```

The following table details the format of the data reported when “SHOW” is specified in the second keyword:

Line	Bytes	Name	Description
1	0:15	Header Info	'SERVICE STATUS V'
	16:17	Version	The version number for the response. The number is left justified and padded with blanks. Starts with 1 (at the code level 8.50.x.x).
	18	Dot	'.'
	19:20	Revision	The revision number for the response. The number is left justified and padded with blanks. Starts with 0 (at the code level 8.50.x.x)
	21:69	Blanks	
The first section:			
2	0:69	Separator	All dash '-' characters
3	0	Blank	
	1:22	Header Info	'Current service state: '
	23:23	Blank	
	24:35	Current service state	Current service state of the cluster where the library request is directed to. One of the following. 'SERVICE-PREP' 'SERVICE ' 'NORMAL '

	36:69	Blanks	
4	0:1	Blanks	
	2:9	Cluster ID	'ClusterX'
	10	Colon	':'
	11	Blank	
	12:23	Cluster service state	Current service state of the cluster. One of the following. 'SERVICE-PREP' 'SERVICE' 'NORMAL'
	24:69	Blanks	
4+N	0:1	Blanks	
	2:9	Cluster ID	'ClusterY'
	10	Colon	':'
	11	Blank	
	12:23	Cluster service state	Current service state of the cluster. One of the following. 'SERVICE-PREP' 'SERVICE' 'NORMAL'
4+N +1	0:69	Separator	All dash '-' characters
The second section:			
4+N +2	0	Blank	
	1:69	Service state description	If it has just started Service Prep processing: 'Starting service preparation processing.' , If it is about to start finalizing Service Prep: 'Finalizing service preparation.' , If service cancellation has been requested: 'Beginning service cancellation.' , If Service Prep is being processed: 'Service preparation will continue until' ,
4+N +3	0	Blank	
	1:69	Service state description	If Service Prep in not being processed: All blanks If Service Prep is being processed: 'the following operations have completed:' ,
	1:69		
The third section:			
4+N +4	0:69	Separator	All dash '-' characters
4+N +5	0	Blank	

	1:31	Header	'The cluster being prepared has:'
	32:69	Blanks	
4+N +6	0	Blank	
	1:7	Right adjusted number	Number of mounted devices
	8	Blank	
	9:24	Description of operation to complete	'devices mounted,'
	25:69	Blanks	
4+N +7	0	Blank	
	1:7	Right adjusted number	Number of pending mounts, demounts or ejects
	8	Blank	
	9:43	Description of operation to complete	'pending mounts, demounts or ejects,'
	44:69	Blanks	
4+N +8	0	Blank	
	1:7	Right adjusted number	Number of LPARs which still wait for CUIR to vary-off
	8	Blank	
	9:54	Description of operation to complete	'LPARs which still wait for CUIR to vary-off of'
	55	Blank	
	56:59	Right adjusted number	Number of devices to vary-off
	60	Blank	
	61:68	Description of operation to complete	'devices,'
	69	Blank	
4+N +9	0	Blank	
	1:7	Right adjusted number	Number of LPARs which may require manual vary-off
	8	Blank	
	9:50	Description of operation to complete	'LPARs which may require manual vary-off of'
	51	Blank	

	52:55	Right adjusted number	Number of devices to vary-off
	56	Blank	
	57:64	Description of operation to complete	<code>`devices,`</code>
	65:69	Blanks	
4+N +10	0	Blank	
	1:7	Right adjusted number	Number of pending token updates against peer clusters
	8	Blank	
	9:52	Description of operation to complete	<code>`pending token updates against peer clusters,`</code>
	53:69	Blanks	
4+N +11	0	Blank	
	1:7	Right adjusted number	Number of pending token updates against unavailable cluster
	8	Blank	
	9:58	Description of operation to complete	<code>`pending token updates against unavailable cluster,`</code>
	44:69	Blanks	
4+N +12	0	Blank	
	1:7	Right adjusted number	Number of global resources it must surrender to peer clusters
	8	Blank	
	9:60	Description of operation to complete	<code>`global resources it must surrender to peer clusters,`</code>
	61:69	Blanks	
4+N +13	0	Blank	
	1:7	Right adjusted number	Number of copies currently in progress
	8	Blank	
	9:40	Description of operation to complete	<code>`copies currently in progress and`</code>
	41:69	Blanks	
4+N +14	0	Blank	

	1:7	Right adjusted number	Number of copy thread change request in progress
	8	Blank	
	9:47	Description of operation to complete	<code>`copy thread change request in progress.'</code>
	48:69	Blanks	
4+N +15	0:69	Separator	All dash '-' characters
4+N +16	0	Blank	
	1:20	Header	<code>`The peer sites have:'</code>
	21:69	Blanks	
4+N +17	0	Blank	
	1:7	Right adjusted number	Number of devices remotely mounted to this cluster
	8	Blank	
	9:43	Description of operation to complete	<code>`devices remotely mounted to this cluster,'</code>
	44:69	Blanks	
4+N +18	0	Blank	
	1:7	Right adjusted number	Number of commands in progress
	8	Blank	
	9:29	Description of operation to complete	<code>`commands in progress'</code>
	50:69	Blanks	
4+N +19	0	Blank	
	1:7	Right adjusted number	Number of immediate copies
	8	Blank	
	9:27	Description of operation to complete	<code>`immediate copies or'</code>
	28	Blank	
	29:36	Right adjusted number	MiBs of immediate copy data pending
	37	Blank	
	38:58	Description of operation to complete	<code>`MiBs of data pending,'</code>

	59:69	Blanks	
4+N +20	0	Blank	
	1:7	Right adjusted number	Number of deferred copies
	8	Blank	
	9:26	Description of operation to complete	<code>'deferred copies or'</code>
	27	Blank	
	28:35	Right adjusted number	MiBs of deferred copy data pending
	37	Blank	
	37:61	Description of operation to complete	<code>'MiBs of data pending, and'</code>
	62:69	Blanks	
4+N +21	0	Blank	
	1:7	Right adjusted number	Number of copies currently in progress against this cluster
	8	Blank	
	9:58	Description of operation to complete	<code>'copies currently in progress against this cluster.'</code>
	59:69	Blanks	
4+N +22	0:69	Separator	All dash '-' characters
4+N +23	0	Blank	
	1:13	Header	<code>'The grid has:'</code>
	14:69	Blanks	
4+N +24	0	Blank	
	1:7	Right adjusted number	Number of deletes, imports or exports currently in progress against this cluster
	8	Blank	
	9:64	Description of operation to complete	<code>'logical volume deletes, imports, or exports in progress.'</code>
	65:69	Blanks	
4+N +25	0:69	Separator	All dash '-' characters

3.1.28 LWORMR

As a part of code release level of 8.51.1.x (R5.1PGA1), the LI REQ, LWORMR keyword is introduced. When all clusters in the Grid are at R5.1PGA1 or above (including a stand-alone configuration), LWORM retention function can be enabled through iRPQ/SCORE. IBM support is required to enable and set up the LWORM retention special settings based off of the user's requirement which are applied to each data class name.

The LWORMR command provides a method to display the LWORM retention special settings configured by IBM support. It can be used to verify and audit the retention settings by the user.

Please refer to "IBM TS7700 LOGICAL WRITE-ONCE, READ-MANY (LWORM) RETENTION FUNCTION USER'S GUIDE" on Techdocs for the details. Search for TS7700.

<https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library>

3.1.29 OCOPY

As part of code release level 8.52.200.x (R5.22), the LI REQ, OCOPY keyword is introduced to support TS7700 Advanced Object Store (FC5283) functionality. The OCOPY keyword is only supported on clusters that completed FC5283 MES process, which enables the Advanced Object Store feature.

In response to the OCOPY request, the cluster associated with the distributed library in the request will collect object copy information based on the second keyword specified in the request. This request only includes information about copies for objects and does not include logical volume copies (see the COPY command for logical volume copies). Currently, only the SUMMARY second keyword is supported and returns a summary of all object copy jobs associated with the distributed library.

3.1.29.1 OCOPY, SUMMARY

No third and fourth keywords are applicable. The response lines are formatted as follows:

```
OCOPY SUMMARY V1.0
  DISTRIBUTED LIBRARY Lipizzan CLUSTER 0
-----
COPY COUNT BY ORIGINATING CLUSTER BY OBJECT COPY QUEUE
CLUSTER      0      1      2      3      4      5      6      7      *
-----
SyncDef      NA      0  91330      NA      NC      NC      NC      NC      0
  Def        NA  121M   11K      NA      NC      NC      NC      NC      0
CpyRfsh      NA      0      0      NA      NC      NC      NC      NC  5424M
-----
COPY SIZE GB BY ORIGINATING CLUSTER BY OBJECT COPY QUEUE
CLUSTER      0      1      2      3      4      5      6      7      *
-----
```

The first and second matrices show the total copy count and size grouped by the object copy queue (per each line) and originating cluster (per each column) on the distributed library in the request.

- SyncDef : A copy job with Synchronous copy mode. A copy job is posted when any event occurs which prevents the synchronous updates to either cluster assigned to synchronous copy mode.
- Def : A copy job with Deferred copy mode
- CpyRfsh : A copy job posted by “COPYRFSH” from a migration event

If no originating cluster is recorded when the copy job is initially created, it is displayed in the most right column ‘*’. This is the typical case for a Copy Refresh copy job.

The longest queue age line at the bottom shows the longest copy queue age for objects, in seconds, for all copy jobs.

No record matched the searching criteria

- The distributed library “Lipizzan” (Cluster ID = 0) returns the object copy job status.
- There are 91230 Sync-Deferred copies for objects whose originating cluster is 2, with an approximate size is 23GB.

- There are approximately 121 million objects in the Deferred object copy queue with a total size of 71GB, currently set to originate from Cluster 1.
- There are approximately 5424M Copy Refresh object copies in the queue with no originating cluster at this time. Their total size is approximately 21TB.
- The longest copy queue age for objects is 1499 seconds.
- Cluster 3 is configured in this grid but is not enabled with FC5283.
- Cluster IDs 4 to 7 are not configured in this Grid.

Both the count and size outputs have a maximum display size of 6-characters per cluster index. Due to the fact that there can be many small objects for this workload, it is possible to exceed the maximum value of 999999. Therefore, the values will be rounded up as the sixth digit is introduced. Examples of how the data may be display:

An object count of 91,230 objects will be displayed as “91230”.

An object count of 88,889,999,900 will be displayed as “88890M”.

A total size of 3,733 bytes will be displayed as “3733”.

A total size of 100,000,000,009 bytes will be displayed as “101G”.

A total size of 5,423,009,999 bytes will be displayed as “5424M”.

Copy counts may include the following identifiers at the end:

“K” = kilos (thousands)

“M” = millions

“B” = billions

“T” = Trillions

Copy total size may include the following identifies at the end:

“K” = kilobytes

“G” = gigabytes

“M” = megabytes

“T” = terabytes

“P” = petabytes

If any unexpected error occurs and the request is not completed, the following error text is returned:

AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = XX

Line	Bytes	Name	Description
1	0:13	Header Info	\OCOPY SUMMARY V'
	14:15	Version	The version number for the response. The number is left justified and padded with blanks. Start with 1 (Code level 8.52.2x.x).

	16	Dot	`.'`
	17:18	Revision	The revision number for the response. The number is left justified and padded with blanks. The revision is introduced at 8.52.2x.x and starts with 0.
	19:69	Blanks	
2	0	Blank	
	1:19	Distributed Library Info	`DISTRIBUTED LIBRARY`
	20	Blank	
	21:28	Distributed Library Name	The distributed library name which responds to the request. The name is left justified and padded with blanks.
	29	Blank	
	30:36	Cluster Info	`CLUSTER`
	37	Blank	
	38	Distributed Library ID	Distributed Library ID (1-8)
	39:69	Blanks	
3	0:69	Separator	All dash '-' characters
4	0:53	Header Info	`COPY COUNT BY ORIGINATING CLUSTER BY OBJECT COPY QUEUE`
	54:69	Blanks	
5	0:69	Header Info	<pre> `CLUSTER      0      1      2      3      4      5  6          7      *` </pre>
6	0:69	Separator	All dash '-' characters
7	0:6	"SyncDef"	`SyncDef`
	7:69	Sync Deferred copy count for objects	The line shows the Sync Deferred object copy count grouped by each originating cluster. Each field has "1 blank + 6-char copy count (including indicator if needed)". Copy count is right justified and padded with blanks. See the examples in the section above for information regarding the values that could be displayed. The line includes a total of 9 fields for each originating cluster.
8	0:3	Blanks	
	4:6	"Def"	`Def`
	7:69	Deferred copy count for objects	The line shows the Deferred object copy count grouped by each originating cluster. Each field has "1 blank + 6-char copy count (including indicator if needed)". Copy count is right justified and padded with blanks. See the examples in the section above for information regarding the values that could be displayed. The line includes a total of 9 fields for each originating cluster.
9	0:6	"CpyRfsh"	`CpyRfsh`
	7:69	Copy Refresh copy count	The line shows the Copy Refresh object copy count grouped by each originating cluster. Each field has "1 blank + 6-char copy count (including indicator if needed)". Copy count is right justified and padded with blanks. See the examples in the section above for information regarding the values that could be displayed. The line includes a total of 9 fields for each originating cluster.
10	0:69	Separator	All dash '-' characters

11	0:52	Header Info	'COPY SIZE BY ORIGINATING CLUSTER BY OBJECT COPY QUEUE'
	53:69	Blanks	
12	0:69	Header Info	'CLUSTER 0 1 2 3 4 5 6 7 *'
13	0:69	Separator	All dash '-' characters
14	0:6	"SyncDef"	'SyncDef'
	7:69	Sync Deferred copy size for objects	The line shows the total size of Sync Deferred object copy grouped by each originating cluster. Each field has "1 blank + 6-char copy size (including indicator if needed)". Copy size is right justified and padded with blanks. See the examples in the section above for information regarding the values that could be displayed. The line includes a total of 9 fields for each originating cluster.
15	0:3	Blanks	
	4:6	"Def"	'Def'
	7:69	Deferred copy size for objects	The line shows the total size of Deferred object copy grouped by each originating cluster. Each field has "1 blank + 6-char copy size (including indicator if needed)". Copy size is right justified and padded with blanks. See the examples in the section above for information regarding the values that could be displayed. The line includes a total of 9 fields for each originating cluster.
16	0:6	"CpyRfsh"	'CpyRfsh'
	7:69	Copy Refresh copy size for objects	The line shows the total size of Copy Refresh object copy grouped by each originating cluster. Each field has "1 blank + 6-char copy size (including indicator if needed)". Copy size is right justified and padded with blanks. See the examples in the section above for information regarding the values that could be displayed. The line includes a total of 9 fields for each originating cluster.
17	0:69	Separator	All dash '-' characters
18	0:30	Header Info	'LONGEST OBJECT COPY QUEUE AGE :'
	31:34	Blanks	
	35:42	Longest object copy queue age	The longest copy queue age for objects, in seconds, on the distributed library. The number is left justified and padded with blanks The age is calculated from all copy jobs.
	43:69	Blanks	

3.1.30 OBJSET1

As part of code release level 8.52.200.x (R5.22), the LI REQ, OBJSET1 keyword is introduced to support TS7700 Advanced Object Store (FC5283) functionality. The OBJSET1 keyword is only supported on clusters that completed FC5283 MES process, which enables the Advanced Object Store feature. This set of alerts is only for settings related to object management functionality.

The following table lists the keywords supported for the OBJSET1 request, whether the request is applicable to a composite and/or distributed library, and whether the request is applicable on a cluster with physical tape and/or a disk-only configuration.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
OBJSET1	ALERT	RSDOLOW	Value	N/A	Y	N	Y	Y	Y
OBJSET1	ALERT	RSDOHIGH	Value	N/A	Y	N	Y	Y	Y

When the request has the first keyword “OBJSET1” only, the response provides the current settings on the target distributed library. If the composite library is specified, an error message is displayed.

The following is an example of the LI REQ, <dist-library>, OBJSET1 response:

```
> SHOWING RESULTS FOR COMMANDS: OBJSET1
"OBJSET1 V1 .0"
" ALERTS"
" OCPYHIGH = 0 OCPYLOW = 0"
" RSDOHIGH = 2500 RSDOLOW = 2800"
> EXECUTING COMMANDS: OBJSET1
```

Note: OCPYHIGH/OCPYLOW shown in the example above are future alerts and will display with the OBJSET1 output but cannot be set at this time.

3.1.30.1 Alert Settings

A second keyword ALERT sets the alert threshold specified by the third keyword. The following ALERT settings are supported if the TS7700 Advanced Object Store feature is enabled:

Keyword 3	Keyword 4	Description
RSDOHIGH	Value	Resident Data High Warning Limit For Objects This is the threshold, in GBs of resident data for all objects cache partitions, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a high warning limit. This threshold is similar to RESDHIGH except that this value is applied to objects. Message Text (xxxxxxx is the amount of resident data, yyyyyyy is the threshold): When above the threshold: “AL5021 Sum of resident data in objects partitions of xxxxxxxx GB above high warning limit of yyyyyyy GB.” When below the threshold: “AL5022 Sum of resident data in objects partitions no longer above resident data high warning limit of yyyyyyy GB.”
RSDOLOW	Value	Resident Data Low Warning Limit For Objects This is the threshold, in GBs of resident data for all objects cache partitions, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a low warning limit. This threshold is similar to RESDLOW except that this value is applied to objects.

Keyword 3	Keyword 4	Description
		Message Text (xxxxxxx is the amount of resident data, yyyyyyyy is the threshold): When above the threshold: “AL0023 Sum of resident data in objects partitions of xxxxxxxx GB above low warning limit of yyyyyyyy GB.” When below the threshold: “AL0024 Sum of resident data in objects partitions no longer above low warning limit of yyyyyyyy GB.”

3.1.31 CRCSET

As part of code release level 8.53.0.x (R5.3), the LI REQ, CRCSET keyword is introduced to support CRC check functionality for logical volume. The CRCSET keyword is only supported on the grid all clusters are R5.3 or later, which is CRC check functionality supported environment.

The following table lists the keywords supported for the CRCSET request. TS7740 is not supported because R5.3 supports VED model only.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Comp	Dist	7740	7700D	7700T	7700C
CRCSET	GEN	LVOL	DISABLE/ CRC32	Y	N	N	Y	Y	Y
	VERIFY	LVOLWRR	DISABLE/ ENABLE	Y	N	N	Y	Y	Y
		LVOLWRL							
		LVOLCPYS							
		LVOLCPYT							
		LVOLTRCL							
		LVOLTPMG							
		LVOLRCLM							

When the request has the first keyword “CRCSET” only, the response provides the current settings on the target composite library. If the distributed library is specified, an error message is displayed.

The following is an example of the LI REQ, <comp-library>, CRCSET response:

```

> SHOWING RESULTS FOR COMMANDS: CRCSET
"CRC SETTINGS V1 .0
" GEN LVOL = CRC32
" LVOLWRR = ENABLED
" LVOLWRL = ENABLED
" LVOLCPYS = ENABLED
" LVOLCPYT = ENABLED
" LVOLTRCL = ENABLED
" LVOLTPMG = ENABLED
" LVOLRCLM = ENABLED
> EXECUTING COMMANDS: CRCSET

```

Note:

- The volume size does not change depending on CRC value generation is enabled or not. Because the calculated CRC values are saved in volume header and record header of existing volume format.
- When the CRC error occurs, the operation that CRC error occurred at fails and an MI event and host notification/attention message is just posted. The corrupted volume is not recovered automatically from another normal copy. The recovery from another normal copy needs IBM support. Currently, this is a function that just prevents the corrupted volumes from spreading unnoticed.
- The CRC calculation is fast because it uses POWER9 vector instructions. See the “IBM TS7700 Release 5.3 Performance White Paper” for the actual I/O performance with default settings (CRC generation and all the CRC checks are enabled).

3.1.31.1 GEN Settings

A second keyword GEN sets the enablement/disablement of CRC value generation for the data type specified by third keyword. Currently, only logical volume (keyword “LVOL”) is supported. The CRC value type is specified by the fourth keyword:

Keyword 3	Keyword 4	Description
LVOL	DISABLE/ CRC32	CRC value generation setting This setting determines what type of the CRC value is generated and added to logical volume. This value is used to verify the data during data movements over network or FICON, or between cache and tape. The operations for which data verification is performed is configured by LI REQ, CRCSET, VERIFY. When it is set to DISABLED, no CRC value is generated and added to logical volumes. As a result, no data verification is performed on the logical volumes. When it is set to CRC32, 32bit CRC value is generated and added to each record of the logical volume. The default value is “CRC32”. Note: The setting changes take effect when the logical volume is written from BOT. Until then, the old setting is retained.

3.1.31.2 VERIFY Settings

A second keyword VERIFY sets the enablement/disablement of the data verification using the CRC value. The operations for which data verification is performed is specified by third keyword:

Keyword 3	Keyword 4	Description
LVOLWRR	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed on remote cluster during remote host write. When it is set to DISABLED, no data verification is performed on the logical volumes on remote cluster during remote host write. When it is set to ENABLED, it verifies if the data is not corrupted before and after the data going through the network on remote cluster during remote host write. The default value is “ENABLED”. Note: The data verification is performed on only the logical volumes having CRC value generated when it was written.
LVOLWRL	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed on local cluster during local or remote host write. When it is set to DISABLED, no data verification is performed on the logical volumes on local cluster during local or remote host write. When it is set to ENABLED, it verifies if the data is not corrupted before and after the data going from host through the FICON on local cluster during local or remote host write. The default value is “ENABLED”. Note: The data verification is performed on only the logical volumes having CRC value generated when it was written. This verification is performed on uncompressed record data only.
LVOLCPYS	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed on copy source cluster during Grid Copy. When it is set to DISABLED, no data verification is performed on the logical volumes on copy source cluster during Grid Copy. When it is set to ENABLED, it verifies if the data reading from the cache on copy source cluster is not corrupted before sending it to network during Grid Copy.

Keyword 3	Keyword 4	Description
		The default value is “ENABLED”. Note: The data verification is performed on only the logical volumes having CRC value generated when it was written.
LVOLCPYT	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed on copy target cluster during Grid Copy. When it is set to DISABLED, no data verification is performed on the logical volumes on copy target cluster during Grid Copy. When it is set to ENABLED, it verifies if the data is not corrupted before and after the data going through the network on copy target cluster during Grid Copy. The default value is “ENABLED”. Note: The data verification is performed on only the logical volumes having CRC value generated when it was written.
LVOLTRCL	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed during Recall from tape. When it is set to DISABLED, no data verification is performed on the logical volumes during Recall from tape. When it is set to ENABLED, it verifies if the data reading from the tape is not corrupted before writing it to the cache during Recall from tape. The default value is “ENABLED”. Note: The data verification is performed on only the logical volumes having CRC value generated when it was written.
LVOLTPMG	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed during Premigration to tape.

Keyword 3	Keyword 4	Description
		When it is set to DISABLED, no data verification is performed on the logical volumes during Premigration to tape. When it is set to ENABLED, it verifies if the data reading from the cache is not corrupted before writing it to the tape during Premigration to tape. The default value is "ENABLED". Note: The data verification is performed on only the logical volumes having CRC value generated when it was written.
LVOLRCLM	DISABLE/ ENABLE	Data verification using CRC value setting This setting determines whether the data verification is performed during Reclaim. When it is set to DISABLED, no data verification is performed during Reclaim. When it is set to ENABLED, it verifies if the data reading from the tape is not corrupted before writing it to the another tape during Reclaim. The default value is "ENABLED". Note: The data verification is performed on only the logical volumes having CRC value generated when it was written.

3.1.32 zTape Air-Gap (FC5995)

zTape Air-Gap (Feature Code 5995) was introduced with Release 8.53. This feature provides a low cost alternative customer rack mount solution to replace C06/C07 IBM controllers. FC 5995 offers z Mainframe customers, a cost effective solution using LTO8 tape drives within a TS4300 an air-gap tape drive for exporting tapes for backup and restore.

The details for zTape Air-Gap may be found in the White Paper, "TS7700 zTape Air-Gap (FC5995)."

When FC5995 is installed, four new LI REQ commands are enabled. If FC5995 is not installed and these LI REQ command keywords are used, the TS7700 will report the standard error message for unknow request.

Error Responses

In processing a request, errors may be encountered due to problems with the request, the state of the subsystem receiving the request or the state of the Grid configuration. This section defines the responses generated based on the error condition encountered.

Invalid/Unknown Request

If the cluster receives a request that contains the first keyword that is not understood or valid, a response with the valid keywords is returned.

INVALID/UNKNOWN KEYWORD

SUPPORTED KEYWORDS

CACHE

STATUS, [GRID|GRIDLINK]

LVOL, volser

LVOL, volser, [PREFER|MIGRATE|REMOVE(, INFO|PROMOTE)|COPY(, KICK|FORCE)]

PDRIVE

POOLCNT(, pool number)

PVOL, volser(, DELETE)

RECALLQ(, volser(, PROMOTE))

COPYEXP, volser, [RECLAIM|DELETE]

GRIDCNTL, COPY, [ENABLE|DISABLE]

SETTING, ALERT, [COPYLOW|COPYHIGH], Val

SETTING, ALERT, [PDRVLOW|PDRVCRIT], Val

SETTING, ALERT, [PSCRLOW|PSCRCRIT], Val

SETTING, ALERT, [RESLOW|RESHIGH|RSDTLOW|RSDTHIGH], Val

SETTING, ALERT, [PCPYLOW|PCPYCRIT], Val

SETTING, ALERT, [DEFDEG|REMOVMSG|LINKDEG], [ENABLE|DISABLE]

SETTING, CACHE, [COPYFSC|RECLPG0|REMOVE|CPYPRIOR], [ENABLE|DISABLE]

SETTING, CACHE, [PMPRIOR|PMTHLVL|REMOVTHR|CPYPRITH|RBPRIOR|RBTHLVL], Val

SETTING, THROTTLE, [COPYFT|ICOPYT], [ENABLE|DISABLE]

SETTING, THROTTLE, [DCOPYT|DCTAVGTD|DCTCURTD], Val

SETTING, RECLAIM, [RCLMMAX|RCLMSMAX], Val

SETTING, DEVALLOC, [SCRATCH|PRIVATE], [ENABLE|DISABLE]

SETTING, DEVALLOC, FAMILY, [RESPECT|IGNORE]

SETTING, CPYCNT, [RUN|DEF], Val

SETTING, COPY, [TIMEOUT|IMMSNS|SCRATCH], Val

SETTING, LINK, FAILOVER, [ENABLE|DISABLE]

SETTING, DELEXP, COUNT, Val

SETTING, EXISTDEL, [CRITERIA, [STALE|ALL|NONE]|WHEN, [ATCLOSE|AUTO]]

RRCLSUN, [ENABLE|DISABLE|STATUS]

COPYRFSH, volser(, srcID(, NORECALL))

COPY, [ACTIVITY, [RUN|SDEF|IDEF|DEF]|SUMMARY]

OTCNTL, [CONFIG|DIST|START|STOP|STAT]

```

DRSETUP, [SHOW|ADD|REMOVE|WP|FLASH|LIVECOPY|DOALL|SELFLIVE |LIVEACC]
SETTING2,SCRATCH,PFRLOCO, [ENABLE|DISABLE]
SETTING2,CACHE, [MAXLGMC|OVRSPVLV],Val
SETTING2,THROTTLE,LNKSPEED,Val
SETTING2,ALERT, [PDRVSLOW|PDRVSCRT|LMTDTHR|RSDOLOW|RSDOHIGH],Val
SETTING2,PHYSLIB, [TVCWDEG|CPYWDEG|PRETHDEG],Val
SETTING2,PHYSLIB, [SLDMPRI|MAINT], [ENABLE|DISABLE]
SETTING2,RECALL,BLKCLSZ,Val
SETTING2,BVIR,TIMEOUT,Val
PARTRFSH,volser
GGM
LOWRANK, [TVC|COPY], [SHOW|Val]
CUIR, [SETTING|AONLINE] (, [SERVICE|FENCE|ALL], [ENABLE|DISABLE])
LDRIVE(,GROUP,Val)
SERVICE, [ENTER|SHOW|CANCEL] (,FORCE)
OCOPY,SUMMARY
OBJSET1,ALERT, [RSDOLOW|RSDOHIGH],VAL

```

Line	Bytes	Name	Description
1	0:22	Header Info	'INVALID/UNKNOWN KEYWORD'
	23:69	Blanks	
2	0:17	Blank	'SUPPORTED KEYWORDS'
	18:69	Blanks	
For each of the supported keywords and combinations, a line is included with the allowable format			
N	0	Blanks	
	1:69	Supported keywords	Text of the supported keywords. The text is left justified and padded with blanks.

Library Type Incompatible with the Request

If the request is issued to a library type that is not compatible with the request type, one of the following text lines is returned:

If the request is issued for a distributed library, but it is only valid for a composite library.

```
REQUEST INVALID FOR DISTRIBUTED LIBRARY
```

Line	Bytes	Name	Description
1	0:38	Header Info	'REQUEST INVALID FOR DISTRIBUTED LIBRARY'
	39:69	Blanks	

If the request is issued for a composite library, but it is only valid for a distributed library.

```
REQUEST INVALID FOR COMPOSITE LIBRARY
```

Line	Bytes	Name	Description
1	0:36	Header Info	'REQUEST INVALID FOR COMPOSITE LIBRARY'
	37:69	Blanks	

Library Specified in Request is Unavailable

If the request is issued to a distributed library that is unavailable, which could be because it is offline, in service, failed or the communication paths to it have failed, the following text is returned:

DISTRIBUTED LIBRARY UNAVAILABLE

Line	Bytes	Name	Description
1	0:30	Header Info	'DISTRIBUTED LIBRARY UNAVAILABLE'
	31:69	Blanks	

Library Specified Incompatible with Request

If the request is relevant only on a TS7700 with physical tape and is issued to a TS7700 without physical tape, the following text is returned:

NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE

Line	Bytes	Name	Description
1	0:57	Header Info	'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
	58:69	Blanks	

If the request is relevant only on a TS7700 without physical tape and is issued to a TS7700 with physical tape, the following text is returned:

NOT SUPPORTED IN A NON DISK-ONLY TS7700 VIRTUALIZATION ENGINE

Line	Bytes	Name	Description
1	0:57	Header Info	'NOT SUPPORTED IN A NON DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
	58:69	Blanks	

If the request is relevant only on a TS7700 TA and is issued to a non TS7700 TA, the following text is returned:

ONLY SUPPORTED IN A TS7700 TAPE ATTACHED VIRTUALIZATION ENGINE

Line	Bytes	Name	Description
1	0:61	Header Info	'ONLY SUPPORTED IN A TS7700 TAPE ATTACHED VIRTUALIZATION ENGINE'
	62:69	Blanks	

Unexpected Error Response

If the cluster the request was received on is operational, but an unexpected error, not covered in one of the previous cases is encountered, an error response is formatted as follows:

```
AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = 1
```

Line	Bytes	Name	Description
1	0:59	Header Info	'AN ERROR OCCURED PROCESSING THE HOST REQUEST. RETURN CODE = '
	60:61	Return Code	Indicates the area that detected the error. The return code is right justified and padded with blanks.
	62:69	Blanks	

4 Appendix – A. Summary table of SETTING/SETTING2/OBJSET1 keywords

This is the summary table of all the supported SETTING commands.

KW3	KW4	Cmp/Dist/7740/ 7700 D /7700T/7700C/7 7000	Code	Description	Default	Possible Values	Message/Note
KW2 : ALERT							
COPYHIGH	Value (GB)	N/Y/Y/Y/Y/Y/Y	R1.5	- High warning limit of uncopied data in cache. - A message (1) is generated when it has exceeded the threshold longer than 150 sec. - The message (1) is repeated every 15 min. - A message (2) is generated when it is at or falls below the threshold longer than 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size- 500GB) ~ 100GB	(1) AL5000 Uncopied data of xxxxxxxx GB above high warning limit of yyyyyyyyy GB. (2) AL5001 No longer above uncopied data high warning limit of yyyyyyyyy GB.
COPYLOW				- Low warning limit of uncopied data in cache. - A message (1) is generated when it exceeds the threshold longer than 150 sec. - A message (2) is generated when it is at or falls below the threshold longer than 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size- 500GB) ~ 100GB or COPYHIGH – 100GB	(1) AL0000 Uncopied data of xxxxxxxx GB above low warning limit of yyyyyyyyy GB. (2) AL0001 No longer above uncopied data low warning limit of yyyyyyyyyy GB.
(COPYHIGH) - Set COPYHIGH to (enabled cache size – 500GB) automatically when trying to set COPYHIGH to less than (enabled cache size – 500GB). - Set COPYHIGH to 100GB automatically when trying to set COPYHIGH to less than 100GB. [COPYLOW > 0] - Change COPYLOW to (COPYHIGH – 100GB) automatically when trying to set COPYHIGH to less than (COPYLOW + 100GB). (COPYLOW) [COPYHIGH = 0] - Set COPYLOW to (enabled cache size – 500GB) automatically when trying to set COPYLOW to less than (enabled cache size – 500GB). - Set COPYLOW to 100GB automatically when trying to set COPYLOW to less than 100GB. [COPYHIGH > 0] - Set COPYLOW to (COPYHIGH – 100GB) automatically when trying to set COPYLOW to more than (COPYHIGH – 100GB).							
PDRVCRT	Value (num of drives)	N/Y/Y/N/Y/N/Y	R1.5	- Critical warning limit of available physical drive. - A message (1) is generated when it has fallen below the threshold longer than 15 min. - The message (1) is repeated every 15 min. - A message (2) is generated when it is at or rises above the threshold longer than 15 min. - No message is generated when the default (0 drive) is set.	0	(installed drives -1) ~ 3	(1) AL5004 Available physical drives xx below critical limit of yy. (2) AL5005 Available physical drives no longer below critical limit of yy.
PDRVLOW				- Low warning limit of available physical drive. - A message (1) is generated when it has fallen below the threshold longer than 15 min.	0	(installed drives) ~ 3 or	(1) AL0004 Available physical drives xx below low limit of yy. (2) AL0005 Available physical drives no longer below low limit of yy.

				- A message (2) is generated when it is at or rises above the threshold longer than 15 min. - No message is generated when the default (0 drive) is set.		PDRVCRIT + 1	
(PDRVCRIT) - Set PDRVCRIT to (installed drives – 1) automatically when trying to set PDRVCRIT to the number of installed drives. - Set PDRVCRIT to 3 automatically when trying to set PDRVCRIT to less than 3. [PDRVLOW > 0] - Change PDRVLOW to (PDRVCRIT + 1) automatically when trying to set PDRVCRIT to more than (PDRVLOW - 1). (PDRVLOW) [PDRVCRIT = 0] - Set PDRVLOW to (the number of installed drives) automatically when trying to set PDRVLOW to more than (the number of installed drives). - Set PDRVLOW to 3 automatically when trying to set PDRVLOW to less than (the number of installed drives - 3). [PDRVCRIT > 0] - Set PDRVLOW to (PDRVCRIT + 1) automatically when trying to set PDRVLOW to less than (PDRVCRIT + 1).							
PSCRCRIT	Value (num of scratch pvol)	N/Y/Y/N/Y/N/Y	R1.5	- Critical warning limit of available physical scratch volume. - A message (1) is generated when it has fallen below the threshold longer than 16 min. - The message (1) is repeated every 16 min. - A message (2) is generated when it is at or rises above the threshold longer than 16 min. - No message is generated when the default (0 scratch pvol) is set.	0	5 ~ 190	(1) AL5006 Available physical scratch volumes of xxx below critical limit of yyy for pool zz. (2) AL5007 Available physical scratch volumes no longer below critical limit of yyy for pool zz.
PSCRLOW				- Low warning limit of available physical scratch volume. - A message (1) is generated when it has fallen below the threshold longer than 16 min. - A message (2) is generated when it is at or rises above the threshold longer than 16 min. - No message s generated when the default (0 scratch pvol) is set.	0	200 ~ 5 or PDRVCRIT + 10	(1) AL0006 Available physical scratch volumes of xxx below low limit of yyy for pool zz. (2) AL0007 Available physical scratch volumes no longer below low limit of yyy for pool zz.
All the active general pools (pool 1 – 32) are monitored. All media types allowed for borrowing are considered. (PSCRCRIT) - Set PSCRCRIT to 190 automatically when trying to set PSCRCRIT to more than 190. - Set PSCRCRIT to 3 automatically when trying to set PSCRCRIT to less than 3. [PSCRLOW > 0] - Change PSCRLOW to (PSCRCRIT + 10) automatically when trying to set PSCRLOW to less than (PSCRCRIT + 10). (PSCRLOW) [PSCRCRIT = 0] - Set PSCRLOW to 200 automatically when trying to set PSCRLOW to more than 200. - Set PSCRLOW to 5 automatically when trying to set PSCRLOW to less than 5. [PDRVCRIT > 0] - Set PSCRLOW to (PSCRCRIT + 10) automatically when trying to set PSCRLOW to less than (PSCRCRIT + 10).							
RESHIGH	Value (GB)	N/Y/Y/Y/Y/Y/Y	R1.5	- High warning limit of resident data. - A message (1) is generated when it has exceeded the threshold longer than at least 150 sec. - The message (1) is repeated every 15 min.	0	(enabled cache size-500GB) ~ 100GB	(1) AL5008 Resident data of xxxxxxxx GB above high warning limit of yyyyyyyy GB. (2) AL5009 No longer above resident data high warning limit of yyyyyyyy GB.

				- A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.			
RESLOW				- Low warning limit of resident data. - A message (1) is generated when it exceeds the threshold longer than at least 150 sec. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size-500GB) ~ 1GB or RESDHIGH – 100GB	(1) AL0008 Resident data of xxxxxxxx GB above low warning limit of yyyyyyyyy GB. (2) AL0009 No longer above resident data low warning limit of yyyyyyyyy GB.
On TS7740, the resident data is equal to the data that has not been copied to the physical volumes. On TS7700D, it is equal to the data in cache. On TS7700T, the resident data is equal to the data in the cache partition 0. (RESDHIGH) ● Set RESDHIGH to (enabled cache size – 500GB) automatically when trying to set RESDHIGH to less than (enabled cache size – 500GB). ● Set RESDHIGH to 100GB automatically when trying to set RESDHIGH to less than 100GB. [RESLOW > 0] ● Change RESLOW to (RESDHIGH – 100GB) automatically when trying to set RESDHIGH to less than (RESLOW + 100GB). (RESLOW) [RESDHIGH = 0] ● Set RESLOW to (enabled cache size – 500GB) automatically when trying to set RESLOW to less than (enabled cache size – 500GB). ● Set RESLOW to 100GB automatically when trying to set RESLOW to less than 100GB. [RESDHIGH > 0] Set RESLOW to (RESDHIGH – 100GB) automatically when trying to set RESLOW to more than (RESDHIGH – 100GB).							
RSDTHIGH	Value (GB)	N/Y/N/N/Y/Y/Y	R3.2	- High warning limit of resident data in the tape partitions on TS7700T. - A message (1) is generated when it has exceeded the threshold longer than at least 150 sec. - The message (1) is repeated every 15 min. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size-500GB) ~ 100GB	(1) AL5015 Sum of resident data in tape partitions of xxxxxxxx GB above high warning limit of yyyyyyyyy GB. (2) AL5016 Sum of resident data in tape partitions no longer above resident data high warning limit of yyyyyyyyy GB.
RSDTLOW				- Low warning limit of resident data. - A message (1) is generated when it exceeds the threshold longer than at least 150 sec. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size-500GB) ~ 1GB or RSDTHIGH – 100GB	(1) AL0012 Sum of resident data in volume partitions of xxxxxxxx GB above low warning limit of yyyyyyyyy GB. (2) AL0013 Sum of resident data in volume partitions no longer above low warning limit of yyyyyyyyy GB.
This is only applicable to TS7700T and the resident data is equal to the sum of data in the tape partitions from 1 to 7. (RSDTHIGH) ● Set RSDTHIGH to (enabled cache size – 500GB) automatically when trying to set RSDTHIGH to less than (enabled cache size – 500GB). ● Set RSDTHIGH to 100GB automatically when trying to set RSDTHIGH to less than 100GB. [RSDTLOW > 0] ● Change RSDTLOW to (RSDTHIGH – 100GB) automatically when trying to set RSDTHIGH to less than (RSDTLOW + 100GB). (RSDTLOW) [RSDTHIGH = 0]							

- Set RSDTLOW to (enabled cache size – 500GB) automatically when trying to set RSDTLOW to less than (enabled cache size – 500GB). - Set RSDTLOW to 100GB automatically when trying to set RSDTLOW to less than 100GB [RSDTHIGH > 0] - Set RSDTLOW to (RSDTHIGH – 100GB) automatically when trying to set RSDTLOW to more than (RSDTHIGH – 100GB).							
PCPYLOW	Value (GB)	N/Y/Y/Y/Y/Y/Y	R2.1	- Low warning limit of pending copy data. - A message (1) is generated when it has exceeded the threshold. - The message (1) is repeated every 15 min. - A message (2) is generated when it has fallen below the threshold. - No message is generated when the default (0 GB) is set.	0	500 ~ 99999999	(1) AL0010 Distributed Library xx has a total pending inbound copy backlog of yyyy GB which is above the low warning limit of zzzz GB. (2) AL0011 Distributed Library xx has successfully fallen below the inbound copy backlog low warning limit of zzzz GB.
PCPYCRIT				- Critical warning limit of pending copy data. - A message (1) is generated when it has exceeded the threshold. - A message (2) is generated when it has fallen below the threshold. - No message s generated when the default (0 GB) is set.	0	1000 ~ 99999999	(1) AL5010 Distributed Library xx has a total pending inbound copy backlog of yyyy GB which is above the critical warning limit of zzzz GB. (2) AL5011 Distributed Library xx has successfully fallen below the inbound copy backlog critical warning limit of zzzz GB.
The check is done every 15 min. When the check detects the condition, the message is generated immediately. (PCPYLOW) - Set PCPYLOW to 500GB automatically when trying to set PCPYLOW to less than 500GB. [PCPYCRIT > 0] - Set PCPYLOW to (PCPYCRIT – 500GB) automatically when trying to set PCPYLOW to more than (PCPYCRIT – 500GB). (PCPYCRIT) - Set PCPYCRIT to 1000GB automatically when trying to set PCPYCRIT to less than 1000GB. [PCPYLOW > 0] - Change PCPYLOW to (PCPYCRIT – 500GB) automatically when trying to set PCPYCRIT to less than (PCPYLOW + 500GB).							
DEFDEG	ENABLE/ DISABLE	Y/N/Y/Y/Y/Y/Y	R3.0	- Suppress “SyncDeferred” or “ImmediateDeferred” state - A composite library enters “degraded” state when the default (ENABLE) is set. - A composite library does not enter “degraded” state when DISABLE is set.	ENABLE	ENABLE/ DISABLE	(These operator messages are still reported even when DISABLE is set) - G0005 Distributed library xx has entered the immediate deferred state. - G0032 Distributed Library xx has entered the synchronous deferred state due to volser yyyyyy.
LINKDEG	ENABLE/ DISABLE	Y/N/Y/Y/Y/Y/Y	R3.1	- Suppress “Grid link degraded” state - A composite library enters “Grid link degraded” state when the default (ENABLE) is set. - A composite library does not enter “degraded” state when DISABLE is set.	ENABLE	ENABLE/ DISABLE	(These operator messages are still reported even when DISABLE is set) - G0030 Library XXXXXX,PPP,AAA Grid Link are degraded. - G0031 All grid links for this cluster have left the degraded state.
REMOVMSG	ENABLE/ DISABLE	N/Y/N/Y/Y/Y/Y	R3.0 R2.1 PGA4	- Suppress Auto Removal start/stop operator intervention (operator message) to MI (host). - A distributed library reports Auto Removal start/stop to MI/Host the default (ENABLE) is set. - A distributed library does not report Auto Removal start/stop to MI/Host when DISALBE is set.	ENABLE	ENABLE/ DISABLE	(These operator messages are suppressed when DISABLE is set) - OP0164 Auto removal of volumes has begun on this cluster. - AL5012 Auto removal of volumes has completed on this cluster.
KW2 : CACHE							
COPYFSH	ENABLE /DISABLE	N/Y/Y/N/Y/Y/Y	R1.5	- Copies to follow storage class preference - Copied volume follows the defined storage class constructs when ENABLE is set. - Copied volume is managed as PG0 when DISABLE is set.	DISABLE	ENABLE/ DISABLE	
PMPRIOR	Value (GB)	N/Y/Y/N/Y/N/Y	R1.5	- Premigration Priority Threshold	1600	Less than PMTHLVL	

				- Increase the number of premigration tasks when unpre migrated data has exceeded the threshold longer than 150 sec. - Decrease the number of premigration tasks when unpre migrated data falls below the threshold longer than at least 150 sec. - Set PMPRIOR to PMTHLVL automatically when PMPRIOR is higher than PMTHLVL.			
PMTHLVL	Value (GB)	N/Y/Y/N/Y/N/Y	R1.5	- PremigrationThrottling Threshold - Start premigration throttling (delay to host write) when unpre migrated data has exceeded the threshold longer than 150 sec. - Throttle (host write delay) starts from a few to up to 2000 milliseconds. - Stop premigration throttling when unpre migrated data falls below the threshold longer than at least 150 sec. - Set PMTHLVL to (enabled cache size – 500GB) automatically when trying to set it to more than (enabled cache size – 500GB). - Set PMTHLVL to (total amount of FC 5274) automatically when trying to set it to more than (total amount of FC 5274) for TS7700T.	2000	(enabled cache size- 500GB) ~ 0 GB (total amount of FC 5274) GB ~ 0GB for TS7700T	
RECLPG0	ENABLE /DISABLE	N/Y/Y/N/Y/Y/Y	R1.5	- Recalls Preferred to be Removed from Cache - Recalled volumes follows the defined storage class constructs when DISABLE is set. - Recalled volumes are managed as PG0 when ENABLE is set.	DISABLE	ENABLE/ DISABLE	
REMOVE	ENABLE /DISABLE	N/Y/N/Y/Y/Y/Y	R1.7	- Enable or disable automatic removal function. - Automatic removal is enabled/disabled when ENABLE/DISABLE is set.	ENABLE	ENABLE/ DISABLE	
REMVTHR	Value (GB)	N/Y/N/Y/Y/Y/Y	R2.1	- Automatic removal cache usage threshold - Automatic removal starts when the free cache space size is less than (REMVTHR+1000GB). - Set REMVTHR to 2000 automatically when trying to set it to less than 2000. - Set REMVTHR to 10000 automatically when trying to set it to more than 10000 (TS7700T). - Set REMVTHR to (cache size – 2000) automatically when trying to set it to more than (cache size – 2000).	3000	2000 ~ 10000	(TS7700D) The minimum (2000) and maximum (cache size – 2000) GB can be set. (TS7700T) The minimum (2000) and maximum (10000) GB can be set.
CPYPRIOR	ENABLE /DISABLE	N/Y/Y/N/Y/N/Y	R3.0P GA1	- Limit premigration resources under Grid copy activity. - Premigration resources are limited when ENABLE is set. - Default is DISABLE (pre-8.41.x.x) / ENABLE (8.41.x.x or later).	DISABLE/ ENABLE (depends on the code level)	ENABLE/ DISABLE	
CPYPRITH	Value (MB/s)	N/Y/Y/N/Y/N/Y	R3.0P GA1	- Premigration resources are limited when Grid copy activity crosses the threshold. - Premigration resources start being limited when the data transfer rate (MB/s) to read/write from/into the local cache is more than CPYPRITH. - An error is returned when the value less than 10 or more than 1000 is given and the current value is not changed.	100	10 ~ 1000	
RBPRIOR	Value	N/Y/Y/Y/Y/Y/Y	R3.1	- Cache rebuild priority. - Cache controller in TS7700 controls the cache rebuild priority based off of the setting. - An error is returned when the value less than 0 or more than 5 is given and the current value is not changed.	0	0 ~ 5	
RBTHLVL	Value (MB/s)	N/Y/Y/Y/Y/Y/Y	R3.1	- Cache rebuild priority threshold. - RBTHLVL is only applicable when RBPRIOR is set to 0 (dynamic cache rebuild prioritization). - When disk I/O rate crosses the threshold defined by "RBTHLVL", cache rebuild priority is dynamically changed.	3500	100 ~ 5000	

				- Set RBTHLVL to 2500 automatically when trying to set it to less than 100 or more than 5000.			
KW2 : THROTTLE							
COPYFT	ENABLE /DISABLE	N/Y/Y/Y/Y/Y/Y	R1.5	- Full Cache Copy Throttling - Apply cache full throttling when the cache is full of uncopied data and when ENABLE is set. - No cache full throttling occurs when DISABLE is set. - The throttle trigger is the uncopied data has exceeded 95% of the cache size and TS7700 has been up more than 24 hours.	ENABLE	ENABLE/ DISABLE	
DCOPYT	Value (millisecond s)	N/Y/Y/Y/Y/Y/Y	R1.5	- Deferred Copy Throttling Delay - Deferred copy throttling introduces the delay (threshold) for each 32KB copy transfer. - No deferred copy throttling occurs when 0 is set. - The throttle trigger is host throughput more than DCTAVGTD or idle CPU usage less than 15%.	125	0 ~ 250	
ICOPYT	ENABLE /DISABLE	N/Y/Y/Y/Y/Y/Y	R1.5	- Immediate Copy Throttling - Immediate copy throttling is enabled when ENABLE is set. - No immediate copy throttling occurs when DIABLE is set. - The throttle trigger is the load and age of the pending immediate copies.	ENABLE	ENABLE/ DISABLE	
DCTAVGTD	Value (MB/s)	N/Y/Y/Y/Y/Y/Y	R1.5	- Deferred Copy Throttling Average Threshold - This is the average host throughput over a 20 min period to determine the deferred copy throttling needs to be applied. - The default (100 MB/s) is set when 0 is set.	100	1 ~ 500	
DCTCURTD	Value (MB/s)	N/Y/Y/Y/Y/Y/Y	R4.1.2	- Deferred Copy Throttling Current Threshold - This is the average host throughput over a 30 sec period to determine the deferred copy throttling needs to be applied. - The default (100 MB/s) is set when 0 is set.	100	1 ~ 500	
KW2 : RECLAIM							
RCLMMAX	Value (MB/s)	N/Y/Y/N/Y/N/Y	R1.6	- Reclaim Maximum Tasks Limit - A cluster limits the maximum number of reclaims to the threshold. - The limitation is disabled when 0 is set and the maximum number in the table is used.	100	1 ~ (Max number in the table-1)	
RCLMSMAX	Value (MB/s)	N/Y/Y/N/Y/N/Y	R3.2	- Reclaim Maximum Tasks Limit for Sunset Media - A cluster limits the maximum number of reclaims for the sunset media to the threshold. - The limitation is disabled when 0 is set and the maximum number in the table is used.	0	0 ~ MIN(Availa ble non- sunset drives, sunset drives)-1	
KW2 : DEVALLOC							
SCRATCH	ENABLE /DISABLE	Y/N/Y/Y/Y/Y/Y	R2.0	- Devise Allocation Assist (DAA) for Scratch Volumes - DAA for Scratch Volumes is enabled/disabled when ENABLE/DISABLE is set.	DISABLE	ENABLE/ DISABLE	
PRIVATE				- Devise Allocation Assist (DAA) for Private Volumes - DAA for Private Volumes is enabled/disabled when ENABLE/DISABLE is set.	ENABLE		

FAMILY	RESPECT /IGNORE	N/Y/Y/Y/Y/Y/Y	R3.0P GA1	- Device Allocation Assist (DAA) for private volume takes into account the cluster family configuration. - Cluster family configuration is taken into account when RESPECT is set and ignored when IGNORE is set.	RESPECT	RESPECT/ IGNORE	
KW2 : CPYCNT							
RUN	Value	N/Y/Y/Y/Y/Y/Y	R2.0	- The number of concurrent copy threads for RUN copies	20 or 40	5 ~ 128	- Default is 20 for V06/VEA and V07/VEB with 2x1Gb links.
DEF	(num of copy threads)			- The number of concurrent copy threads for DEF copies	20 or 40		- Default is 40 for V07/VEB with 4 x 1Gb or 2 x 10Gb links. - The copy timeout value may need to be adjusted accordingly.
KW2 : COPY							
IMMSNS	ALL/ UNEXP/ NONE	N/Y/Y/Y/Y/Y/Y	R1.7	- Immediate-Deferred State Reporting Method - How to report immediate-deferred state in ERA 35 sense data for RUN CCW is defined. - ALL : All the immediate-deferred state is reported to host. - UNEXP : Only the unexpected immediate-deferred state is reported to host. - NONE : No immediate-deferred state except “no valid copy source” case is reported to host.	NONE	ALL/ UNEXP/ NONE	- In R1.6, all the immediate-deferred state is reported in ERA 35 sense data for RUN CCW. - Since R1.7, the default is set “NONE”.
TIMEOUT	Value (min)	N/Y/Y/Y/Y/Y/Y	R2.0	- Volume Copy Timeout Time - A timeout value for each logical volume copy between clusters is defined.	180	30 ~ 999	- Prior to R2.0, the timeout time is fixed (180 min). - If large number of CPYCNT is set, the longer TIMEOUT value may need to be set to prevent unexpected copy timeout error.
SCRATCH	ALWAYS /NEVER /NONTDLY	Y/N/Y/Y/Y/Y/Y	R3.2	- Control the replication of logical volumes in the scratch category - ALWAYS: Replicate all logical volumes regardless of the category. - NEVER: Do not replicate logical volumes in the scratch category. - NONTDLY: Do not replicate logical volumes in Time Delayed copy policy only. - All clusters in the domain must be at 8.32.x.x or above.	ALWAYS	ALWAYS /NEVE /NONTDLY	- Prior to R3.2, TS7700 handles the replication of logical volumes as same as “ALWAYS” behavior.
KW2 : LINK							
FAILOVER	ENABLE /DISABLE	N/Y/Y/Y/Y/Y/Y	R2.1	- Immediate Copy Throttling - IP Link Failover for Remote Mount is enabled/disabled when ENABLE/DISABLE is set.	ENABLE	ENABLE/ DISABLE	
KW2 : DELEXP							
COUNT	Value (num of delete expire count)	Y/N/Y/Y/Y/Y/Y	R2.1	- Delete Expire Count - The number of delete expired volumes at a time is defined. - Set COUNT to 1000 automatically when trying to set it to less than 1000. - Set COUNT to 5000 automatically when trying to set it to more than 5000.	1000	1000 ~ 5000	- The maximum COUNT value we can set is changed from 2000 to 5000 at 8.50.x.x.
KW2 : EXISTDEL							
CRITERIA	STALE /ALL /NONE	N/Y/Y/Y/Y/Y/Y	R3.2	- Determine which status of ‘E’ copy mode volume is deleted. - STALE: Delete ‘E’ copy mode volume if it’s inconsistent (down level). - ALL: Always delete ‘E’ copy mode volume regardless of its consistency. - NONE: Never delete ‘E’ copy mode volume. - All clusters in the domain must be at 8.31.x.x or above.	STALE	STALE/ ALL/ NONE	- Prior to R3.1, TS7700 handles ‘E’ copy mode volume as same as ‘STALE’ setting. - In R3.2, the second keyword is changed to CRITERIA. The third keywords STALE/ALL/NONE are the same meaning with the second keywords STALE/ALWAYS/NEVER in R3.1.

				- From R3.1, the second keywords "STALE/ALWAYS/NEVER" can be used and it's compatible with CRITERIA keyword.			
WHEN	ATCLOSE /AUTO		R3.2	- Determine when 'E' copy mode volume is deleted. - ATCLOSE: Same behavior with R3.1. It's deleted at volume mount/demount. - AUTO: In addition to the volume mount/demount, periodical check deletes up to 100 'E' copy mode volumes per 24 hours if it finds the volumes. - With both ATCLOSE and AUTO, only 'E' copy mode volumes which satisfy the conditions set by CRITERIA are deleted. - All clusters in the domain must be at 8.32.x.x or above.	ATCLOSE	ATCLOSE /AUTO	
AUTINTVL	Value (EXISTDEL interval)	N/Y/NA/Y/Y/Y/Y	R5.0/R 4.2PG A2	- Determine how often (interval) 'E' copy mode volume is checked for the auto-deletion. - This setting is effective when "CRITERIA" is not set to "NONE" and ""WHEN" is set to "AUTO". - Prior to 8.42.x.x (8.50.x.x), it's a fixed value, 24 hours.	24	1-24	
AUTCOUNT	Value (EXISTDEL count)		R5.0/R 4.2PG A2	- Determine how many 'E' copy mode volume can be deleted for the auto-deletion. - This setting is effective when "CRITERIA" is not set to "NONE" and ""WHEN" is set to "AUTO". - Prior to 8.42.x.x (8.50.x.x), it's a fixed value, 100 volumes.	100	1-10000	

This is the summary table of all the supported SETTING2 commands.

KW3	KW4	Cmp/Dist/7740/ 7700D/7700T/7 700C/7700O	Code	Description	Default	Possible Values	Message/Note
KW2 : SCRATCH							
PFRLOCO	ENABLE /DISABLE	N/Y/Y/Y/Y/Y/Y	R3.2	- Change favoring locally owned logical volume in the scratch mount. - ENABLE: The locally owned volume is favored in the scratch mount. - DISABLE: Do not favor locally owned volume and older volume regardless of the ownership status is favored in the scratch mount.	ENABLE	ENABLE /DISABLE	
KW2 : ALERT							
PDRVSCRT	Value (num of drives)	N/Y/Y/N/Y/N/Y	R3.3	- Critical warning limit of available sunset physical drive. - A message (1) is generated when it has fallen below the threshold longer than 15 min. - The message (1) is repeated every 15 min. - A message (2) is generated when it is at or rises above the threshold longer than 15 min. - No message is generated when the default (0 drive) is set.	0	(installed drives -1) ~ 3	(1) AL5017 Available sunset physical drives of xx is below critical limit of yy (2) AL5018 Available sunset physical drives is no longer below critical limit of yy
PDRVSLOW				- Low warning limit of available sunset physical drive. - A message (1) is generated when it has fallen below the threshold longer than 15 min. - A message (2) is generated when it is at or rises above the threshold longer than 15 min. - No message is generated when the default (0 drive) is set.	0	(installed drives) ~ 3 or PDRVCRIT + 1	(1) AL5019 Available sunset physical drives of xx below low limit of yy. (2) AL5020 Available sunset physical drives no longer below low limit of yy.

(PDRVSCRT)

- Set PDRVSCRT to (installed sunset drives – 1) automatically when trying to set PDRVSCRT to the number of installed sunset drives.
- Set PDRVSCRT to 3 automatically when trying to set PDRVSCRT to less than 3.

[PDRVSLow > 0]

- Change PDRVLOW to (PDRVSCRT + 1) automatically when trying to set PDRVSCRT to more than (PDRVSLow - 1).

(PDRVSLow)[PDRVSCRT = 0]

- Set PDRVSLow to (the number of installed sunset drives) automatically when trying to set PDRVSLow to more than (the number of installed sunset drives).
- Set PDRVSLow to 3 automatically when trying to set PDRVSLow to less than (the number of installed sunset drives - 3).

[PDRVSCRT > 0]

Set PDRVSLow to (PDRVSCRT + 1) automatically when trying to set PDRVSLow to less than (PDRVSCRT + 1).

LMTDTHR	Value	N/Y/N/Y/Y/Y/Y	R4.1.2	- Limited Cache Threshold	0	0 (apply default) or range depends on cache/CP0 size (see right)	(TS7700D) Minimum: 1,000GB Maximum: cache size – 2,072GB (TS7700T) Minimum: MIN(1,000GB, 5% of CP0 size) Maximum: CP0 size – MIN(1,072GB, 5% of CP0 size) – 1,000GB
RSDOHIGH	Value (GB)	N/Y/NA/N/N/N/Y	These settings are only applicable in R5.0 and R5.1 code	- High warning limit of resident data in the object partitions. - A message (1) is generated when it has exceeded the threshold longer than at least 150 sec. - The message (1) is repeated every 15 min. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size-500GB) ~ 100GB	(1) AL5021 Sum of resident data in objects partitions of xxxxxxxx GB above high warning limit of yyyyyyyy GB. (2) AL5022 Sum of resident data in objects partitions no longer above resident data high warning limit of yyyyyyyy GB.
RSDOLOW				- Low warning limit of resident data. - A message (1) is generated when it exceeds the threshold longer than at least 150 sec. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size-500GB) ~ 1GB or RSDOHIGH – 100GB	(1) AL0023 Sum of resident data in objects partitions of xxxxxxxx GB above low warning limit of yyyyyyyy GB. (2) AL0024 Sum of resident data in objects partitions no longer above low warning limit of yyyyyyyy GB.

This is only applicable to a TS7700 with objects offload enabled (FC5282), and the resident data is equal to the sum of data in the objects partitions from 1 to 7. These settings are only applicable in R5.0 and R5.1 code.

(RSDOHIGH)

- Set RSDOHIGH to (enabled cache size – 500GB) automatically when trying to set RSDOHIGH to less than (enabled cache size – 500GB).
- Set RSDOHIGH to 100GB automatically when trying to set RSDOHIGH to less than 100GB.

[RSDOLOW > 0]

- Change RSDOLOW to (RSDOHIGH – 100GB) automatically when trying to set RSDOHIGH to less than (RSDOLOW + 100GB).

(RSDOLOW)

[RSDOHIGH = 0]

- Set RSDOLOW to (enabled cache size – 500GB) automatically when trying to set RSDOLOW to less than (enabled cache size – 500GB).
- Set RSDOLOW to 100GB automatically when trying to set RSDOLOW to less than 100GB

[RSDOHIGH > 0]

- Set RSDOLOW to (RSDOHIGH – 100GB) automatically when trying to set RSDOLOW to more than (RSDOHIGH – 100GB).

CAGAHIGH	Value (hours)	N/Y/NA/Y/Y/Y/Y	R5.0	- Inbound All Copy Backlog Age High Warning Limit - A message (1) is generated when it exceeds the threshold. - A message (2) is generated when it is at or falls below the threshold. - No message is generated when the default (0 hour) is set.	0	0 ~ 8760 hours	(1) AL0014 Distributed Library Z has an inbound all copy type backlog that exceeds the high warning limit of XXXX hours. Volser VVVVVV is the oldest copy job in this group at YYYY hours. (2) AL0015 Distributed Library Z has falling below the inbound all copy type backlog high limit of XXXX hours.
CAGALOW				- Inbound All Copy Backlog Age Low Warning Limit - A message (1) is generated when it exceeds the threshold. - A message (2) is generated when it is at or falls below the threshold. - No message is generated when the default (0 hour) is set.	0	0 ~ 8760 hours	(1) AL0016 Distributed Library Z has an inbound all copy type backlog that exceeds the low warning limit of XXXX hours. Volser VVVVVV is the oldest copy job in this group at YYYY hours. (2) AL0017 Distributed Library Z has falling below the inbound all copy type backlog low limit of XXXX hours.
CAGHIGH				- Inbound Non Time-delayed Copy Backlog Age High Warning Limit - A message (1) is generated when it exceeds the threshold. - A message (2) is generated when it is at or falls below the threshold. - No message is generated when the default (0 hour) is set.	0	0 ~ 8760 hours	(1) AL0018 Distributed Library Z has an inbound non time-delayed copy type backlog that exceeds the high warning limit of XXXX hours. Volser VVVVVV is the oldest copy job at YYYY hours. (2) AL0019 Distributed Library Z has falling below the inbound non time-delayed copy type backlog high limit of XXXX hours.
CAGLOW				- Inbound Non Time-delayed Copy Backlog Age Lowrning Limit - A message (1) is generated when it exceeds the threshold. - A message (2) is generated when it is at or falls below the threshold. - No message is generated when the default (0 hour) is set.	0	0 ~ 8760 hours	(1) AL0020 Distributed Library Z has an inbound non time-delayed copy type backlog that exceeds the low warning limit of XXXX hours. Volser VVVVVV is the oldest copy job at YYYY hours. (2) AL0021 Distributed Library Z has falling below the inbound non time-delayed copy type backlog low limit of XXXX hours.

KW2 : CACHE

MAXLGMCMC	Value	N/Y/Y/Y/Y/Y/Y/Y	R3.2	- Limit the concurrent large volume (25GB) mount count. - If the number of the concurrent large volume size mount exceeds this setting, the mount fails with X'42' (Excessive Concurrent Large Volume Mounts Detected) reason code. Also, the operator message in the Note is surfaced to the host console.	128	0 ~ 65535	(These operator messages is reported when the mount fails) G0035 Distributed Library xxx mount failed due to excessive large volume mounts
OVRSPVL	Value	N/Y/N/N/Y/Y/Y/Y	R4.0	- Limit the overspill data size from CPx to CP0 to avoid automatic removal of CP0 data.	2000	0 ~ 99999999	

KW2 : THROTTLE

LNKSPEED	Value	N/Y/Y/Y/Y/Y/Y/Y	R4.1.2	- Grid Link speed used to calculate Immediate Copy Throttling	200	0 (apply default) or 1 ~ 4000	
----------	-------	-----------------	--------	---	-----	--	--

KW2 : PHYSLIB

TVCWDEG	DISABLE /EQUAL /LOWER	N/Y/N/N/Y/N/Y	R3.3	- Allow to select TS7700T with the degraded library as a mount TVC. - EQUAL/LOWER: The TS7700T can be a candidate for the mount TVC as long as no recall is required. - EQUAL: The TS7700T is selected as same as other healthy cluster. - LOWER: The TS7700T is lowered in preference as the mount TVC. - DISABLE: The TS7700T is not selected as the mount TVC (default).	DISABLE	DISABLE /EQUAL /LOWER	
CPYWDEG	ENABLE /DISABLE			- Allow Incoming Copy on TS7700T with the degraded library - ENABLE: The TS7700T continues to allow the inbound replication into CPx. - DISABLE: The TS7700T stops the inbound replication into CPx.	DISABLE	ENABLE /DISABLE	
PRETHDEG	ENABLE /DISABLE			Disable Premigration Throttling on a TS7700T with the degraded library.	DISABLE	ENABLE /DISABLE	

				DISABLE: The premigration throttling is disabled when the library is degraded and unable to operate attached to the TS7700T. ENABLE; The premigration throttling is enabled when the library is degraded and unable to operate attached to the TS7700T.			
SLDPMPRI	ENABLE /DISABLE	N/Y/NA/N/Y/N/Y	R5.0	Slide PMPRIOR in a TS7700T/TS7700C after a physical library or a cloud becomes operational. DISABLE: PMPRIOR will not slide after a physical library or a cloud becomes operational. ENABLE: PMPRIOR will slide after a physical library or a cloud becomes operational.	ENABLE	ENABLE /DISABLE	
MAINT	ENABLE /DISABLE			Enable maintenance mode of a physical library to suspend all physical tape activities. DISABLE: Disable maintenance mode and resume physical tape activities. ENABLE: Enable maintenance mode and suspend physical tape activities.	DISABLE	ENABLE /DISABLE	
KW2 : RECALL							
BLKRCLSZ	Value	N/Y/Y/N/Y/N/Y	R4.1	Tunes the size of data to be recalled for a Grid copy in GB.	13	13 - 256	
KW2 : BVIR							
TIMEOUT	Value	N/Y/Y/Y/Y/Y/Y	R4.1.2	Tunes the timeout value of BVIR in minutes.	300	300-3000	
KW2 : CENCRYPT							
CIPHER	DISABLE/ AES128/ AES256	N/Y/NA/Y/Y/Y/Y	R5.0	Change the Secure Data Transfer encryption cipher key setting	DISABLE	DISABLE/ AES128/ AES256	
KW2 : COPY							
COPYRFSH	PAUSE/ UNPAUSE	N/Y/NA/Y/Y/Y/Y	R5.0	Pause/Unpause incoming COPYRFSH Grid copies	UNPAUSE	PAUSE/ UNPAUSE	(This event is created/closed when it's paused/unpaused) G0035 Distributed Library xxx mount failed due to excessive large volume mounts
DEF	PAUSE/ UNPAUSE			Pause/Unpause incoming Deferred Grid copies			
KW2 : VOLINVT							
LPAGRP	ENABLE/ DISABLE	N/Y/NA/Y/Y/Y/Y	R5.0	Take account of Library Port Access Group setting to the volume inventory command.	DISABLE	ENABLE/ DISABLE	

This is the summary table of all the supported OBJSET1 commands.

KW3	KW4	Cmp/Dist/7740/ 7700D/7700T/7 700C/7700O	Code	Description	Default	Possible Values	Message/Note
KW2 : ALERT							
RSDOHIGH	Value (GB)	N/Y/NA/N/N/N/Y	R5.22	- High warning limit of resident data in the object partitions. - A message (1) is generated when it has exceeded the threshold longer than at least 150 sec.	0	(enabled cache size-	(1) AL5021 Sum of resident data in objects partitions of xxxxxxxx GB above high warning limit of yyyyyyyy GB.

				- The message (1) is repeated every 15 min. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.		500GB) ~ 100GB	(2) AL5022 Sum of resident data in objects partitions no longer above resident data high warning limit of yyyyyyyy GB.
RSDOLOW				- Low warning limit of resident data. - A message (1) is generated when it exceeds the threshold longer than at least 150 sec. - A message (2) is generated when it is at or falls below the threshold longer than at least 150 sec. - No message is generated when the default (0 GB) is set.	0	(enabled cache size- 500GB) ~ 1GB or RSDOHIGH – 100GB	(1) AL0023 Sum of resident data in objects partitions of xxxxxxxx GB above low warning limit of yyyyyyyy GB. (2) AL0024 Sum of resident data in objects partitions no longer above low warning limit of yyyyyyyy GB.
This is only applicable to a TS7700 with objects offload enabled with FC5283, and the resident data is equal to the sum of data in the objects partitions from 1 to 7. (RSDOHIGH) - Set RSDOHIGH to (enabled cache size – 500GB) automatically when trying to set RSDOHIGH to less than (enabled cache size – 500GB). - Set RSDOHIGH to 100GB automatically when trying to set RSDOHIGH to less than 100GB. [RSDOLOW > 0] - Change RSDOLOW to (RSDOHIGH – 100GB) automatically when trying to set RSDOHIGH to less than (RSDOLOW + 100GB). (RSDOLOW) [RSDOHIGH = 0] - Set RSDOLOW to (enabled cache size – 500GB) automatically when trying to set RSDOLOW to less than (enabled cache size – 500GB). - Set RSDOLOW to 100GB automatically when trying to set RSDOLOW to less than 100GB [RSDOHIGH > 0] - Set RSDOLOW to (RSDOHIGH – 100GB) automatically when trying to set RSDOLOW to more than (RSDOHIGH – 100GB).							

5 Disclaimers:

© Copyright 2007 - 2023 by International Business Machines Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

The information provided in this document is distributed "AS IS" without any warranty, either express or implied. IBM EXPRESSLY DISCLAIMS any warranties of merchantability, fitness for a particular purpose OR NON INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (*e.g.*, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interpretability of any non-IBM products discussed herein. The customer is responsible for the implementation of these techniques in its environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. Unless otherwise noted, IBM has not tested those products in connection with this publication 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.

The provision of the information contained herein is not intended to, and does not grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

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

Trademarks

The following are trademarks or registered trademarks of International Business Machines in the United States, other countries, or both.

IBM, TotalStorage, DFSMS/MVS, S/390, z/OS, and zSeries.

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