



December 2018

IBM® TS7700 Series
Bulk Volume Information Retrieval Function
User's Guide Version 4.2

Owner: Takeshi Nohta
IBM, Japan
Shinsuke Mitsuma
IBM, Japan

Contents

Introduction.....	3
Summary of Changes	4
1. Overview.....	7
2. Code Requirements.....	8
3. General Logical Volume Requirements.....	8
4. Request Data Format.....	9
5. Response Data Format	11
5.1. Volume Status Information.....	15
5.2. Cache Contents Information	32
5.3. Physical Volume to Logical Volume Mapping Information	35
5.4. Point In Time Statistics Information.....	37
5.5. Historical Statistics Information	38
5.6. Physical Media Pools Information.....	40
5.7. Physical Volume Status Information	45
5.8. Copy Audit Information.....	51
5.9. GGM Information	62
5.10. Cloud Volume to Logical Volume Mapping Information	63
5.11. Unknown or Invalid Request	68
6. Example JCL Statements	69
Disclaimers:	75

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. As the capability of tape virtualization has grown, so has the need to efficiently manage the large number of logical volumes the system supports. Internal to the TS7700, a large amount of information is captured and maintained about the state and operational aspects of the resources within the TS7700. The TS7700 provides a management interface based on open standards through which a storage management application can request specific information the TS7700 maintains. The open standards are not currently supported for applications running under z/OS, so an alternative method is needed to provide the information to mainframe applications. This white paper describes the use of a facility of the TS7700 through which a z/OS application can obtain that information.

At the code level 8.40.x.x, the new model TS7760 is introduced. **At the code level 8.42.x.x, the cloud enablement feature is available.** The following model definitions will be used throughout this document:

- TS7720/TS7760: TS7720 or TS7760 (with or without tape).
- TS7720D: TS7720 Disk Only (without tape)
- TS7760D: TS7760 Disk Only (without tape)
- TS7700D: Disk only models (TS7720D and TS7760D) are included.
- TS7720T: TS7720 Tape Attach
- TS7760T: TS7760 Tape Attach
- **TS7760C: TS7760 Cloud Attach**
- TS7700T: Tape attach models (TS7720T and TS7760T) are included. (TS7740 is NOT included.)
- **TS7700C: TS7700 with cloud enablement feature (only TS7760 is supported).**
- TS7740: TS7740 (V06 and V07) only.
- TS7700 Flash Capable: Flash capable models (TS7700D and TS7700T) are included.
- TS7700: All models (TS7740, TS7700D, and TS7700T) are included.

It is assumed throughout this white paper that the reader is familiar with using standard labeled tape volumes and MVS tape utilities.

Summary of Changes

This is the initial version of this document. It is similar to the users guide for the B10/B20 VTS (version 3 of that document is the starting point for this document).

Changes for version 1.1:

- Added the JB media type in the examples.

Changes for version 1.3:

- Added a new query for physical volume status information

Changes for version 1.4:

- Added a new query that audits a group of TS7700s in a Grid configuration to determine if they do not have a valid copy of logical volumes.
- Added that in a Grid configuration, the request volume must only be written to the cluster in the Grid where the data is required from.
- Added a JCL example where the response records have a different record length than the request records.

Changes for version 1.5

- A new volume state was added for copy exported volumes that are manually reclaimed.
- Added a new physical volume state & new recording format values for the TS1130 drive.
- Corrected the definition for the takeover_active field in the volume status information record.

Changes for version 1.6

- New fields have been added to the data returned for logical volume status to indicate whether the volume has Write Once, Read Many attributes or other access control attributes.
- Add error message for TS7720 attached configuration
- Many new fields have been added to the data returned for logical volume status.
- Record 3 of the response record will be increased to version 3, indicating an increase in the total record length of records 6-N in the volume status report from 320 to 640
- Added volume size to CACHE CONTENTS request
- Added COPY AUDIT columns for distributed removal states and composite removal states
- Added review changes

Changes for version 1.7

- Add removal_policy to the data returned for logical volume status
- Update last_device_mounted to indicate this is the device associated with the last error
- Updated removal_timestamp description
- Updated removal_state descriptions and added a new removal_state
- Added code level support table in the Version Number of the response record

Changes for version 3.0

- Added recall_count, storage constructs and file size in volume status information.
- Added additional information about mes_flag in volume status information.
- Removed the word “LSVC/LSVA” of “subcmd” explanation in volume status information.

- Added removal function change in volume status information.
- Added additional recording_format and media_type in physical volume status information.
- Supported PRIMARY or BACKUP request in VOLUME MAP request.
- Added 'S' copy mode for copy mode in volume status information.
- Added an example for COPY AUDIT output and how to interpret it.

Changes for version 3.0a

- Added removal_state explanation to Copy Audit section

Changes for version 3.0b

- Correct the version number of CACHE CONTENTS from 5 to 3.

Changes for version 3.0c

- Updated "Example JCL Statements".

Changes for version 3.1

- Include all the changes implemented in R3.1 (8.31.0.xx).
 - Support new "Time Delayed" copy mode.
 - Add Storage Class Assigned Preference Group to CACHE CONTENTS
 - Change 'M' to 'MiB' (denotation of unit only) in VOLUME MAP.
 - Add DB_BACKUP_NAME field to PHYSICAL VOLUME STATUS.

Changes for version 3.2

- Include all the changes implemented in R3.2 (8.32.0.xx).
 - Support cache partition number and last data creation time in VOLUME STATUS output.

Changes for version 3.3

- Added a new physical volume state & new recording format values for the TS1150 drive.
- Added missing explanation of 'F' (Fork) value in "VOLUME STATUS" consistent_type field.
- Added a new "VOLUME STATUS" mes_flag value 'G' of GGM copied volume.

Changes for version 4.0

- Changed description to include TS7760.
- Added cases which cause domain_lock_timestamp to be updated.
- Added note that LWORM volume cannot be used for BVIR.

Changes for version 4.1.2

- Added compression method, logical volume format id, wrap counter handling, channel length bytes and maximum volume size to "VOLUME STATUS" output
- Added a note to removal_state description that the value 'R' is set even when the volume is deleted by EXISTDEL function.

Changes for version 4.1.2a

December 2018

- Added a note to “General Logical Volume Requirements” about the case BVIR response data is too large to be stored in the BVIR volume.

Changes for version 4.2

- Added description to include TS7700C.
- Added GGM BVIR request which was already supported with 8.33.x.x code level.
- Added CLOUD VOLUME MAP request which is supported with 8.42.x.x code level.

1. Overview

With the potential to support hundreds of thousands of logical volumes in a TS7700 subsystem, providing a set of information for all of those volumes through normal channel control type commands is not very practical. Luckily, the functions of a TS7700 subsystem that allows it to virtualize a tape volume, also allows for a simple and effective method to transfer the information to a requesting application. The TS7700 converts the format and storage conventions of a tape volume into a standard file managed by a file system within the subsystem. The Bulk Volume Information Retrieval (BVIR) facility uses an IBM standard labeled tape volume to both initiate a request for information and return the results. By using a standard tape volume, no special interfaces or access methods are needed for an application to use this facility. In practice, no specific applications are required, as standard IBM utilities, such as IEBGENER, provide the function needed to request and obtain the information.

Note: The BVIR function was first introduced on the prior generation of Virtual Tape Servers. Its use has been expanded with the introduction of the TS7700.

There are two steps to obtain information using this facility. First, a single data set with the information request is written to a logical volume. The logical volume can be any logical volume in the subsystem the information is to be obtained from. Either a scratch or specific volume request can be used. The data set contains a minimum of two records and a maximum of three records that specifies the type of data being requested. The records are in human readable form, i.e. lines of character data. The data set can be cataloged or uncataloged (although cataloging the data set can make it easier for subsequent access to the data). On close of the volume, the TS7700 server will recognize it as a request volume and ‘prime’ the subsystem for the next step.

Note: Some of the information obtained through this function is specific to the cluster the logical volume is written on. In a TS7700 Grid configuration with multiple clusters, management class for the volume must be set up if data to be obtained for the cluster remote from the cluster the volume is mounted on. See the section on Request Data Format for Grid Consideration examples.

Second, the request volume is again mounted, this time as a specific mount. Seeing that the volume was ‘primed’ for a data request, the TS7700 appends the requested information to the data set. The process of obtaining the information and creating the records to append can take up to several minutes, depending on the request and, from a host’s viewpoint, is part of the mount processing time. Once the TS7700 has completed appending to the data set, the host is notified that the mount has completed. The requested data can then be accessed like any other tape data set.

Note: In a JES2 environment, the JCL to perform the two steps can be combined into a single job, however, in a JES3 environment, they must be run in separate jobs. This is because the volume will not be demounted and remounted between job steps in a JES3 environment.

Once the response data set has been written to the request logical volume, that logical volume functions identically to any other logical volume in the subsystem. Subsequent mount requests and read accesses to the logical volume should have no effect on its contents. Subsequent mount requests and write accesses to the logical volume will overwrite its contents. It can be returned to scratch status and reused by any application.

Note: Due to the two step approach, BVIR volumes cannot be written with LWORM specifications. You need to assign a Data Class without LWORM for BVIR volumes.

The building of the response information does require a small amount of the resources of the TS7700. It is recommended that the BVIR function not be used to ‘poll’ for a specific set of information and that only one request be issued at a time. Some requests, for example the volume map, may take several minutes to complete and to prevent ‘locking’ out another request during that time, the TS7700 is designed to handle two concurrent requests. If more than two concurrent requests are issued, they will be processed as prior requests are completed.

Whereas the request data is always in a human readable format, depending on the request, the data returned from the TS7700 can be in human readable or binary form. Refer to the response sections for the specifics of the returned data.

2. Code Requirements

The BVIR function was introduced with the TS7700. Although there are no library manager code changes to support the function, there are other functions of the TS7700 that require a compatible level of library manager code. To support the new BVIR function, the TS7700 code level required is 8.4.1.x or later with library manager code level 535.x or later.

For the new recording formats of the TS1130 drive, the TS7700 code level required is 8.5.x.xx or later. There are no host software updates required for this function.

For the new fields introduced with version 3 the TS7700 code level required is 8.6.x.xx or later.

For the new fields introduced with version 4 the TS7700 code level required is 8.7.x.xx or later.

For the new fields introduced with version 5 the TS7700 code level required is 8.30.x.xx or later.

For the new fields introduced with version 6 the TS7700 code level required is 8.31.x.xx or later.

For the new fields introduced with version 7 the TS7700 code level required is 8.32.x.xx or later.

For the new fields introduced with version 8 the TS7700 code level required is 8.41.200.xx or later.

For the new fields introduced with version 9 the TS7700 code level required is 8.42.x.xx or later.

3. General Logical Volume Requirements

Any logical volume defined to a TS7700 can be used as the request/response volume. Logical volumes in a TS7700 are formatted as IBM Standard Labeled volumes. Although a user can reformat a logical volume with an ANSI Standard Label or as an unlabeled tape volume, those formats are not supported for use as a request/response volume.

There are no restrictions regarding the prior use of a volume used as a request/response volume and no restrictions regarding its subsequent use for any other application. It is recommended that normal scratch allocation methods are used for each request (ie DISP=(NEW,CATLG)). In this way, any of the available scratch logical volumes in the TS7700 can be used. Likewise, it is recommended that when the response volume's data is no longer needed, the logical volume is returned to scratch status through the normal methods (typically by deletion of the data set on the volume and a return to scratch policy based on data set deletion).

Note: BVIR volumes cannot be written with LWORM specifications. You need to assign a Data Class without LWORM for BVIR volumes.

Note: BVIR response data to some requests such as "Volume Status Information" and "Physical Volume to Logical Volume Mapping Information" will become large according to the number of virtual volumes defined. Please make sure that the BVIR volume size has enough capacity to store the response data. From R4.1.2 and R3.3 PGA3, if BVIR response data did not fit into the specified volume, mount operation against the volume will fail with a reason code X'43' (BVIR volume is too small to store the result). Here are some guidelines for the volume size of BVIR logical volume. The logical volume size requirement can depend on the number of the logical volumes defined in the Grid which is currently up to 4 million:

Number of logical volumes defined in the Grid	Recommendation of the minimum BVIR logical volume size
1 – 1,400,000	1GB
1,400,001 – 2,800,000	2GB
2,800,001 – 4,000,000	4GB

4. Request Data Format

Several types of data can be requested. The type of data requested is indicated in the request data set. The request data set must be the only data set on the volume and must be written with a record format of FB and a logical record size of 80 bytes. Request information is in EBCDIC character form, beginning in the first character position of the record and padded with blank characters on the right to fill out the record.

Important note: the request fields must be as shown, not beginning in the first character position of the record or extra blanks between words will result in the request being failed.

Note: Although the request data format uses fixed block records, not all response records are fixed block. For the point in time and historical statistics responses, the data records are of variable length and the record format used to read them is the Undefined (U) format. Refer to the sample JCL section.

The format for the request data set records are:

Record 1

VTS BULK VOLUME DATA REQUEST

Bytes	Name	Description
1-28	Request Identifier	'VTS BULK VOLUME DATA REQUEST'
29-80	Blanks	Blank character padding

Record 2

VOLUME STATUS zzzzzz **or** CACHE CONTENTS **or** VOLUME MAP **or** POINT IN TIME STATISTICS **or** HISTORICAL STATISTICS FOR xxx **or** HISTORICAL STATISTICS FOR xxx-yyy **or** PHYSICAL MEDIA POOLS **or** PHYSICAL VOLUME STATUS VOLUME zzzzzz **or** PHYSICAL VOLUME STATUS POOL xx **or** COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids

Bytes	Name	Description
1-80	Request	'VOLUME STATUS zzzzzz' or 'CACHE CONTENTS' or 'VOLUME MAP' or 'VOLUME MAP PRIMARY' or 'VOLUME MAP BACKUP' or 'POINT IN TIME STATISTICS' or 'HISTORICAL STATISTICS FOR xxx-yyy' or 'PHYSICAL MEDIA POOLS' or 'PHYSICAL VOLUME STATUS VOLUME zzzzzz' or 'PHYSICAL VOLUME STATUS POOL xx' or 'COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids' 'GGM COPY RESULT' or 'GGM COPY STATUS' or 'CLOUD VOLUME MAP xx' left justified, padded with blanks on the right.

For the Volume Status and Physical Volume Status Volume requests, 'zzzzzz' specifies the volume serial number mask to be used. By using the mask, one to thousands of volume records can be retrieved for the request. The mask must be 6 characters in length, with the '_' character representing a positional wildcard mask. For example, assuming that volumes in the range of ABC000 through ABC999 have been defined to the cluster, a request of VOLUME STATUS ABC1_0 would return database records that exist for ABC100, ABC110, ABC120, ABC130, ABC140, ABC150, ABC160, ABC170, ABC180 and ABC190.

For the Historical Statistics request, 'xxx' specifies the Julian day being requested. Optionally, '-yyy' can also be specified and indicates that historical statistics from xxx through yyy are being requested. Valid days are 001 through 366 (to account for leap

year). For leap years, February 29th is Julian day 060 and December 31st is Julian day 366, for other years, Julian day 060 is March 1st and December 31st is Julian day 365. If historical statistics do not exist for the day(s) requested, that will be indicated in the response record (this would occur if a request is issued for a day prior to the day the system was installed, day(s) the system was powered off or after the current day before a rolling year has been accumulated). If a request spans the end of the year, for example a request that specified: HISTORICAL STATISTICS FOR 364-002, responses are provided for days 364, 365, 366, 001 and 002, regardless of whether the year was a leap year.

For Copy Audit, INCLUDE or EXCLUDE is specified to indicate which TS7700s clusters in a Grid configuration are to be included or excluded from the audit. COPYMODE is an option for taking a volume's copy mode for a cluster into consideration. If COPYMODE is specified, a single space must separate it from INCLUDE or EXCLUDE. The parameter libids specifies the library sequence numbers of the distributed libraries associated with each of the TS7700 clusters either to include or exclude in the audit. The libids are separated by a comma ','. At least one libid must be specified.

For the Physical Volume Status Pool request, 'xx' specifies the pool for which the data is to be returned. If there are no physical volumes currently assigned to the specified pool that will be indicated in the response record. Data may be requested for pools 00 through 32. The pool number must be a 2-digit number.

For point in time and historical statistics requests, any additional characters provided in the request record past the request itself are retained in the response data, but otherwise ignored.

For volume map requests, additional key words PRIMARY and BACKUP are supported starting with code level 8.30.x.xx. If PRIMARY is specified, the volume map will only contain the list of physical volumes, each having at least one primary logical volume mapped to the physical volume. If BACKUP is specified, the volume map will only contain the list of physical volumes, each having at least one secondary logical volume mapped to the physical volume. If no PRIMARY or BACKUP keyword is specified, the volume map of all the physical volumes is provided.

In a TS7700 Grid configuration, the request volume must only be valid on the specific cluster the data is to be obtained from. Use a specific Management Class that has a copy policy defined to indicate that only the desired cluster is to have a copy of the data. By ensuring that there is a sole copy of the request volume, any virtual device address on any of the clusters in the same Grid configuration can be used to request and access the data. You do not have to have host connectivity to the specific cluster. If a Management Class is used that indicates that more than one cluster is to have a valid copy of the request volume, unpredictable response data results can occur.

For example, assume that information about its physical media pools is needed from cluster 1 in a three-site Grid configuration. You would define a management class name and define a copy policy for it to only have a copy on cluster 1. You could define, for example:

MCPMP001

Copy Consistency Point for Cluster 0: No Copy

Copy Consistency Point for Cluster 1: Rew/Unload

Copy Consistency Point for Cluster 2: No Copy

When you create the request volume, you would assign it the management class, MCPMP001, and this will cause the TS7700 to create that volume in cluster 1 only. If necessary, this will cause a remote mount if the host creating the volume is attached to a cluster different than cluster 1.

Refer to the "IBM® TS7700 Series Best Practices – Copy Consistency Points" whitepaper (WP101230) on Techdocs for more details concerning Copy Consistency Points.

5. Response Data Format

When the request data set has been written to the volume and subsequently closed and demounted, when mounted again, the TS7700 will validate the contents of the request volume and append the requested data records to the data set. Human readable appended records are 80 bytes in length. Binary data appended records can be variable in length of up to 24000 bytes. The data set is now a response data set. The appropriate block counts in the end of file (EOF) records will be updated to reflect the total number of records written to the volume. After appending the records and updating the EOF records, the host that requested the mount is signaled that the mount is complete and can read the contents of the volume. If the contents of the request volume is not valid, either one or more error description records will be appended to the data set or the data set will be unmodified prior to signaling the host that the mount completed, depending on the problem encountered.

All human readable response records begin in the first character position of the record and are padded with blank characters on the right to fill out the record.

All binary records are variable in length and are not padded.

Note: In the response records, the date and times presented are all based on the internal clock of the TS7700 handling the request. The internal clock of a TS7700 is not synchronized to the host, but is synchronized with any other TS7700.

The general format for the response data set is:

Records 1-2

Contents of request records 1-2.

Record 3

This record contains the date and time the response data set was created as well as a format version number for the results as follows:

08/20/2009 12:27:00 VERSION 03

Bytes	Name	Description																				
1-10	Date	The date the response was generated in the format of MM/DD/YYYY																				
11	Field Delimiter	Blank character																				
12-19	Time	The time the response was generated in the format HH:MM:SS																				
20	Field Delimiter	Blank character																				
21-27		'VERSION'																				
28	Field Delimiter	Blank character																				
29-30	Version Number	The Version Number of this response right justified and zero filled. <table border="0"> <thead> <tr> <th>Version Number</th> <th>Code Level Supported</th> </tr> </thead> <tbody> <tr> <td>01</td> <td>Original VTS</td> </tr> <tr> <td>02</td> <td>Release 1.5, Release 1.5PGA1–PGA3</td> </tr> <tr> <td>03</td> <td>Release 1.5 levels PGA4 and later, Release 1.6</td> </tr> <tr> <td>04</td> <td>Release 1.7, Release 2.0, Release 2.1</td> </tr> <tr> <td>05</td> <td>Release 3.0</td> </tr> <tr> <td>06</td> <td>Release 3.1</td> </tr> <tr> <td>07</td> <td>Release 3.2, Release 3.3, Release 4.0, Release 4.1.1</td> </tr> <tr> <td>08</td> <td>Release 4.1.2</td> </tr> <tr> <td>09</td> <td>Release 4.2</td> </tr> </tbody> </table>	Version Number	Code Level Supported	01	Original VTS	02	Release 1.5, Release 1.5PGA1–PGA3	03	Release 1.5 levels PGA4 and later, Release 1.6	04	Release 1.7, Release 2.0, Release 2.1	05	Release 3.0	06	Release 3.1	07	Release 3.2, Release 3.3, Release 4.0, Release 4.1.1	08	Release 4.1.2	09	Release 4.2
Version Number	Code Level Supported																					
01	Original VTS																					
02	Release 1.5, Release 1.5PGA1–PGA3																					
03	Release 1.5 levels PGA4 and later, Release 1.6																					
04	Release 1.7, Release 2.0, Release 2.1																					
05	Release 3.0																					
06	Release 3.1																					
07	Release 3.2, Release 3.3, Release 4.0, Release 4.1.1																					
08	Release 4.1.2																					
09	Release 4.2																					
31-80	Blanks	Blank character padding																				

NOTE: The Version Number indicates in which TS7700 microcode level the response data format for each request was last revised. For Example, the Version Number '02' of 'VOLUME MAP' means that its response format was revised in Release 1.5, 1.5PGA1-1.5PGA3, and has not revised in later microcode levels. In the release 3.1 level, each response has the following version number:

Request record	Version Number
VOLUME STATUS	08
CACHE CONTENTS	06
VOLUME MAP (PRIMARY or BACKUP)	06
POINT IN TIME STATISTICS	02
HISTORICAL STATISTICS FOR xxx-yyy	02
PHYSICAL MEDIA POOLS	04
PHYSICAL VOLUME STATUS VOLUME zzzzzz	06
PHYSICAL VOLUME STATUS POOL xx	06
COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids	02
GGM COPY RESULT/STATUS	01
CLOUD VOLUME MAP xx	09

Record 4

This record contains both the 5 character hardware serial number of the TS7700 and the 5 character Distributed Library sequence number of the cluster that generated the response.

S/N: 0F16F LIB ID: AB123

Bytes	Name	Description
1-4		'S/N:'
5	Field Delimiter	Blank character
6-10	Serial Number	Hardware serial number of the TS7700
11-12	Blanks	Blank character padding
13-19		'LIB ID:'
20	Field Delimiter	Blank character
21-25	Distributed Library ID	Distributed library sequence number of the TS7700
26-80	Blanks	Blank character padding

The 5 character TS7700 hardware serial number is configured during manufacturing of the TS7700 server node. The serial number is fixed in the subsystem firmware and reported as part of the BVIR response in record 4.

The Distributed Library ID field is a unique 5 character identifier assigned to each TS7700 cluster during pre-installation planning. The IBM CE uses the distributed library sequence number during installation configuration of the TS7700. It is called **Library ID** on the DFSMS Tape Library Define panel or as shown on the Tape Library Display panel.

Record 5

This record contains all character blanks.

Record 6-N

These records contain the specific response records based on the request. If the request could not be understood or was invalid, that will be indicated. Refer to the following sections for the response record definitions. Record length of each response data is listed in the table below.

Request	Record Length of Record 6-N
VOLUME STATUS	640 bytes
CACHE CONTENTS	80 bytes
VOLUME MAP (PRIMARY or BACKUP)	80 bytes
POINT IN TIME STATISTICS	24000 bytes
HISTORICAL STATISTICS FOR xxx-yyy	24000 bytes
PHYSICAL MEDIA POOLS	80 bytes
PHYSICAL VOLUME STATUS VOLUME zzzzzz	400 bytes
PHYSICAL VOLUME STATUS POOL xx	400 bytes
COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids	80 bytes

GGM COPY STATUS/RESULT	400 bytes
CLOUD VOLUME MAP xx	900 bytes

5.1. Volume Status Information

A database is maintained on each individual TS7700 Cluster that contains information related to the management of the logical volumes on the cluster and copy and resynchronization processes when the TS7700 are in a Grid configuration. Several of the database fields returned can be useful in handling operational exceptions at one or more of the clusters in a Grid configuration.

Data Inconsistent

This field indicates whether the cluster has a valid version of the data or not. If it indicates that the data on the logical volume is not valid, this means that the same volume on another TS7700 in the Grid has been modified and it has not yet been copied. For customers who use the deferred copy consistency point (which is typically when there is some significant distance between the TS7700 in the Grid configuration), there will be some number of volumes that are not consistent between the TS7700 at any point in time. If a situation occurs that renders inoperable the site where the source data resides, by issuing the Volume Status request to an operable TS7700, this field can be used to identify the volumes which were not copied prior to the situation so that appropriate recovery steps can be performed for them.

MES Volume

This field indicates that the logical volume was created in the TS7700 Cluster prior to it being merged into a Grid configuration. Volumes that existed in a TS7740 Cluster prior to being included in a Grid configuration are not automatically copied to the other TS7700 Clusters in the configuration until they have been accessed and closed. This field could be used to determine which volumes in each TS7700 Cluster that have not been copied and used to build a set of jobs to access them and force the copy.

Copy Required for Cluster n

This field indicates that a copy to another TS7700 Cluster in a Grid configuration is required. In cases where deferred mode copy is used, this field can be used to determine if a critical set of volumes has completed their copy operations to specific clusters.

Volume Ownership and Volume Ownership Taken

At any point in time a logical volume is owned by a specific cluster. Ownership is transferred as part of mount processing. Ownership can transfer in one of two ways, either through communication with the current owning cluster or through a recovery process called ownership takeover. Normally, the cluster receiving a mount command request that the current owning cluster transfer ownership (assuming that the cluster receiving the mount request for the volume does not already have ownership of the volume) and ownership is transferred. However, if the cluster receiving the mount request cannot communicate with the owning cluster, that method does not work. In this case the requesting clusters cannot determine whether the owning cluster has failed or just the communication paths to it have failed. Operator intervention is required to indicate that the owning cluster has failed and that ownership takeover by the other clusters is allowed. There are two types of ownership takeover, Write and Read Only. With write ownership takeover (WOT), the cluster taking over ownership of the volume has complete freedom to modify the contents of the volume or modify any of the properties associated with the volume. With read only takeover (ROT), the cluster taking over ownership of the volume is restricted to reading the volume's data only.

Current and Pending Category

One of the key properties associated with a volume is the category it is assigned. The primary usage for category is to group together scratch volumes. A volume's category assignment changes as the volume is used. The current category field indicates the category the volume is assigned to in the library manager associated with the cluster. The pending category field indicates that a new category assignment is in progress for the category. These fields can be used to determine whether the category assignments are in synch between the clusters and the host databases.

Data Deleted

As part of normal processing in a TS7700, a customer can specify that after a certain period of time after being returned to scratch, the contents of a volume can be deleted. This field indicates whether or not the data associated with the volume has been deleted on the cluster.

Removal State

As part of normal processing in a TS7700 Grid configuration where a mixture of both TS7740 and TS7720/TS7760 clusters exist, a data removal or migration process occurs where data is removed from TS7720/TS7760 clusters in order to prevent

TS7720/TS7760 clusters from overrunning their tape volume cache. This field, as well as the removal timestamp can be used to determine whether or not the data associated with the volume has been removed. A data removal can take place on TS7720/TS7760 where Grid consists of TS7720/TS7760 clusters only since R1.7 (8.7.0.x).

- **Hot**
This field represents the cluster's view of which clusters have down level token or volume meta-data information as a result of a cluster outage. When clusters are unavailable due to expected or unexpected outages, the remaining clusters will mark the unavailable cluster for pending reconciliation by updating this hot mask. The field represents both Insert/Eject pending updates and/or regular pending updates. Insert/Eject updates are related to volumes being inserted and/or ejected during the outage. Regular pending updates are for updates that occur to the volume during an outage as a result of normal operations such as host I/O. Each bit within the mask represents which clusters are viewed as needing reconciliation.

The volume status information returned represents the status of the volume on the cluster the request volume is written. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the volume status information for the individual clusters. A response record is written for each logical volume, selected based on the volume serial number mask specified in the request, that exists in the cluster.

A response record consists of the database fields defined in the following table. Fields are presented in the order defined in the table and are comma (,) separated. The overall length of each record is 640 bytes with blank padding after the last field as needed. For example, the first few fields of the record returned for volser ABC123 would be:

```
ABC123,0,2006-04-22-11.56.45.871263,0,0,-1,0,N,2548,N,8719,N...
```

Note: The generation of the response may take several minutes to complete depending on the number of volumes requested and how busy the TS7700 cluster is at the time of the request.

Record 6-N

If the request is for a specific volume and it doesn't exist in the cluster's database or the request is for a range of volumes and none of the volumes exist in the cluster's database, the following record is returned.

```
NO VOLUME RECORD(S) EXIST FOR THE REQUESTED VOLUME(S)
```

Bytes	Name	Description
1-53		'NO VOLUME RECORD(S) EXIST FOR REQUESTED VOLUME(S)'
54-80	Blanks	Blank character padding

For the requested volumes that do exist in the cluster's database, each of these records provide information for one logical volume.

Note: The size of each field does not include the byte needed for the comma that separates each field

Field Name	Description
volser	6 character volume serial number
domain_lock_cluster	The cluster identifier of the cluster which currently has ownership of the volume. 0 is the identifier for the first cluster in a Grid. domain_lock_cluster is updated when a new ownership information is assigned to the volume, such as the following cases. <ul style="list-style-type: none"> - Volume is mounted. - Volume removal is attempted. - Deletion of obsolete copy of the volume is attempted. - Library request COPYRFSH is requested. - Library request OTCNTL is requested. domain_lock_cluster and domain_lock_timestamp are updated simultaneously.

domain_lock_timestamp	<p>Timestamp of the last volume ownership information assignment to the volume. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>When a volume is inserted in to the library, this field is set to 1970-01-01-00.00.00.000000.</p> <p>domain_lock_timestamp and domain_lock_cluster are updated simultaneously. Refer to the description of domain_lock_cluster for the cases they may be updated.</p>
mounted_cluster	<p>If the volume is mounted, the cluster identifier of the cluster the mount request for the volume was issued to. 0 is the identifier for the first cluster in a Grid. If the volume is not mounted, -1 is indicated.</p>
mounted_vnode	<p>If the volume is mounted, the VNode identifier in the cluster that the mount request for the volume was issued to. 0 is the identifier for the first VNode in a cluster. If the volume is not mounted, -1 is indicated.</p>
mounted_device	<p>If the volume is mounted, the virtual device number in the VNode of the cluster that the mount request for the volume was issued on. 0 is the identifier for the first virtual device in a VNode. If the volume is not mounted, -1 is indicated.</p>
tvcluster	<p>The cluster identifier associated with the tape volume cache being used for the mounted volume. With remote mounting support, this cluster may not be the same cluster as the mounted_cluster. 0 is the identifier for the first cluster in a Grid.</p>
data_inconsistent	<p>This field indicates whether the cluster has an inconsistent copy of the volume. 'Y' indicates it does not have a valid copy, 'N' indicates that it has a valid copy.</p>
data_level	<p>The current data level value for the volume on the cluster. Every time a volume's data is modified, its data level is incremented.</p> <p>If the data level matches the corresponding data level current, this cluster's version of the volume is viewed as consistent.</p>
properties_inconsistent	<p>This field indicates whether the cluster has inconsistent properties for the volume. 'Y' indicates it has one or more properties that are down level, 'N' indicates that it does not. The properties consist of the category and storage constructs assigned to the volume.</p>
properties_level	<p>The current properties level value for the volume on the cluster. Every time a volume's properties are modified, its data level is incremented.</p>
cluster0_copy_required	<p>This field indicates whether cluster 0 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p>
cluster1_copy_required	<p>This field indicates whether cluster 1 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p>
cluster2_copy_required	<p>This field indicates whether cluster 2 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p>

cluster3_copy_required	<p>This field indicates whether cluster 3 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p> <p>This field is new to VOLUME STATUS version 3.</p>
cluster4_copy_required	<p>This field indicates whether cluster 4 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p> <p>This field is new to VOLUME STATUS version 3.</p>
cluster5_copy_required	<p>This field indicates whether cluster 5 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p> <p>This field is new to VOLUME STATUS version 3.</p>
cluster6_copy_required	<p>This field indicates whether cluster 6 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p> <p>This field is new to VOLUME STATUS version 3.</p>
cluster7_copy_required	<p>This field indicates whether cluster 7 needs to make a copy of the volume from this cluster. 'Y' indicates it does, 'N' indicates that either a copy is not required or one has already been made.</p> <p>This field is new to VOLUME STATUS version 3.</p>
cluster0_copy_mode	<p>This field indicates whether cluster 0 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' – Synchronous copy consistency point.</p> <p>'I' - Rewind unload (RUN) copy consistency point.</p> <p>'D' - Deferred copy consistency point.</p> <p>'T' – Time Delayed copy consistency point.</p> <p>'N' - No copy</p> <p>'X' – Same as an 'N'. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to 'N' since this 'E' copy is no longer valid.</p>

<p>cluster1_copy_mode</p>	<p>This field indicates whether cluster 1 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ – Synchronous copy consistency point.</p> <p>‘I’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘X’ – Same as an ‘N’. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to ‘N’ since this ‘E’ copy is no longer valid.</p>
<p>cluster2_copy_mode</p>	<p>This field indicates whether cluster 2 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ – Synchronous copy consistency point.</p> <p>‘I’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘X’ – Same as an ‘N’. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to ‘N’ since this ‘E’ copy is no longer valid.</p>
<p>cluster3_copy_mode</p>	<p>This field indicates whether cluster 3 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ – Synchronous copy consistency point.</p> <p>‘I’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘X’ – Same as an ‘N’. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read</p>

	<p>operation occurred against the volume. A private mount for write append will change the mode to 'N' since this 'E' copy is no longer valid.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>cluster4_copy_mode</p>	<p>This field indicates whether cluster 4 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' – Synchronous copy consistency point.</p> <p>'I' - Rewind unload (RUN) copy consistency point.</p> <p>'D' - Deferred copy consistency point.</p> <p>'T' – Time Delayed copy consistency point.</p> <p>'N' - No copy</p> <p>'X' – Same as an 'N'. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to 'N' since this 'E' copy is no longer valid.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>cluster5_copy_mode</p>	<p>This field indicates whether cluster 5 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' – Synchronous copy consistency point.</p> <p>'I' - Rewind unload (RUN) copy consistency point.</p> <p>'D' - Deferred copy consistency point.</p> <p>'T' – Time Delayed copy consistency point.</p> <p>'N' - No copy</p> <p>'X' – Same as an 'N'. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to 'N' since this 'E' copy is no longer valid.</p> <p>This field is new to VOLUME STATUS version 3.</p>

<p>cluster6_copy_mode</p>	<p>This field indicates whether cluster 6 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ – Synchronous copy consistency point.</p> <p>‘I’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘X’ – Same as an ‘N’. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to ‘N’ since this ‘E’ copy is no longer valid.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>cluster7_copy_mode</p>	<p>This field indicates whether cluster 7 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ – Synchronous copy consistency point.</p> <p>‘I’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘X’ – Same as an ‘N’. Only set for a logical volume that was migrated from B10/20 P2P to TS7700, and its copy had existed on only one side of B10/20 P2P.</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy and a private mount for read operation occurred against the volume. A private mount for write append will change the mode to ‘N’ since this ‘E’ copy is no longer valid.</p> <p>This field is new to VOLUME STATUS version 3.</p>

mes_flag	<p>This field indicates whether the volume was part of a MES merge operation.</p> <p>‘Y’ indicates the volume existed prior to merging the cluster into a Grid configuration. After code level 8.3 this value is no longer used and is replaced with ‘W’ and ‘M’.</p> <p>‘W’ indicates the volume existed prior to merging the cluster into a Grid configuration and has not been accessed since the MES merge operation. This field is only set within the cluster that was previously not aware of the volume prior to the merge.</p> <p>‘M’ indicates the volume existed prior to merging the cluster into a Grid configuration and has been mounted/demounted without being modified. The volume will be copied to the clusters specified in the copy_mode fields. This field is only set within the cluster that was previously not aware of the volume prior to the merge.</p> <p>‘N’ indicates that the volume was not part of the MES merge operation or if it was, has since been modified and successfully copied to the clusters specified in the copy_mode fields.</p> <p>‘R’ indicates the volume received z/OS Host Command Line Request COPYRFSH command with NORECALL option as 4th keyword. The volume will be copied to the local cluster through the copy refresh process. The field is only set within the cluster that needs to have a copy through the copy refresh process.</p> <p>‘D’ indicates that the volume received z/OS Host Command Line Request COPYRFSH command without NORECALL option as 4th keyword. The volume will be copied to the local cluster through the copy refresh process. The field is only set within the cluster that needs to have a copy through the copy refresh process.</p> <p>Refer to the “IBM® TS7700 Series z/OS Host Comand Line Request User’s Guide” whitepaper (WP101091) on Techdocs for more details concerning COPYRFSH request.</p> <p>‘G’ indicates that the volume received z/OS Host Command Line request GGM, QUEUE command and it was inserted into this Grid by GGM (Grid to Grid Migration) operation.</p> <p>Refer to the "IBM TS7700 Series Grid to Grid Migration User's Guide" on Techdocs for more details concerning GGM operation.</p>
recall_error	<p>This field indicated whether or not an error occurred while recalling this volume from a physical tape associated with the tvc_cluster. ‘Y’ indicates that an error occurred the last time a recall was attempted with the volume. ‘N’ indicates that the last recall of the volume was successful.</p>
read_error	<p>This field indicates whether the data associated with the volume is corrupted. ‘Y’ indicates that the volume’s data is corrupted. ‘N’ indicates that the data is not corrupted. A volume’s data is determined to be corrupted if an error is detected when reading the data.</p>
disaster_rec	<p>This field indicates the disaster recovery state for the volume. ‘Y’ indicates that the volume has been through the disaster recovery process. ‘N’ indicates that no disaster recovery processing has been done or is required for the volume.</p>
data_deleted	<p>This field indicated whether the data associated with the volume has been deleted by the delete expired volume data function. ‘Y’ indicates that the data has been deleted. ‘N’ indicates that the data has not been deleted.</p>

volume_damaged	The volume/token is determined to be damaged or unrecoverable without human intervention. ‘Y’ indicates that the volume/token has been damaged. ‘N’ indicates that the volume/token has not been damaged.
mount_operation_time	Timestamp of when the volume was last mounted or unloaded. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129 If the volume has not yet been mounted, this field is set to 1970-01-01-00.00.00.000000.
media_type	The media type defined when the volume was inserted. The values are: “0” - Cartridge System Tape (400MB) “1” - Enhanced Capacity Cartridge System Tape (800MB)
takeover_active	This field indicates whether the volume’s ownership has been taken over. 0 indicates that the volume’s ownership has transferred normally. A non-zero value indicates that the volume’s ownership was transferred by one of the ownership takeover methods. bit 0 - Reserved bits 1:7 -- Reserved bits 8:15 -- Steal ID bits 16:19 -- Transfer Version bits 20:21 -- Reserved bits 22:23 – WOT/ROT/SOT 0 – WOT or Write Ownership Takeover 1 – ROT or Read Ownership Takeover 2 -- SOT or Service Ownership Takeover bits 24:27 -- Cluster ID which was the original owner. bits 28:31 -- Cluster ID which took control of this volume
pending_category	The category the volume is to be set to within the library manager associated with the cluster.
current_category	The category the volume is currently assigned to within the library manager associated with the cluster. When this value and the pending_category are equal, then the category is consistent.

category_time	<p>Prior to code release 8.5, this timestamp represents when the pending_category field was last changed. For release 8.5 and later, this timestamp represents when the volume's category was last changed within the composite library. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>Prior to the first change in category, this field is set to 1970-01-01-00.00.00.000000.</p>
takeover_timestamp	<p>Timestamp of when the volume's ownership was last changed due to write or read ownership takeover. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>If the volume has not had its ownership taken, this field is set to 1970-01-01-00.00.00.000000.</p>
subcmd_cluster	<p>If a category or storage construct change request is currently active against the volume, this field represents the cluster identifier of the cluster the request for the volume was issued to. 0 is the identifier for the first cluster in a Grid. If the volume has no active category or storage construct change request against it, -1 is indicated.</p> <p>This field is new to VOLUME STATUS version 3.</p>
subcmd_vnode	<p>If a category or storage construct change request is currently active against the volume, this field represents the vnode identifier of the cluster the request for the volume was issued to. 0 is the identifier for the first vnode in a cluster. If the volume has no active category or storage construct change request against it, -1 is indicated.</p> <p>This field is new to VOLUME STATUS version 3.</p>
subcmd_device	<p>If a category or storage construct change request is currently active against the volume, this field represents the virtual device in the Vnode of the cluster the request for the volume was issued to. 0 is the identifier for the first virtual device in a Vnode. If the volume has no active category or storage construct change request against it, -1 is indicated.</p> <p>This field is new to VOLUME STATUS version 3.</p>
data_level_current	<p>The highest known level in the grid.</p> <p>This field is new to VOLUME STATUS version 3.</p>
data_level_last_modify	<p>Define mask to show what cluster, vnode and device last wrote to the volume.</p> <p>Cluster/Vnode/Hnode to last update properties version</p> <p>bit 0 – Reserved</p> <p>bits 1:3 – Cluster</p> <p>bits 4:7 Hnode</p> <p>bits 8:15 – Vnode</p> <p>bits 16:31 Logical Device</p> <p>This field is new to VOLUME STATUS version 3.</p>

<p>properties_level_current</p>	<p>The highest known level in the grid.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>properties_level_last_modify</p>	<p>Define mask to show what cluster, vnode and device last issued an LSVCLLSVA</p> <p>Cluster/Vnode/Hnode to last update properties version</p> <p>bit 0 – Reserved</p> <p>bits 1:3 – Cluster</p> <p>bits 4:7 - Hnode</p> <p>bits 8:15 – Vnode</p> <p>bits 16:31 – Logical Device</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>subcmd_operation_time</p>	<p>This timestamp represents when an LSVCLLSVA was issued to this volume. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>Prior to the first LSVCLLSVA for the volume, this field is set to 1970-01-01-00.00.00.000000.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>hot</p>	<p>Hot mask showing definition of regular hot and insert hot.</p> <p>bits 0:15 – Reserved</p> <p>bits 16:23 – Token Insert/Eject Hot</p> <p>From right to left, each bit position represents a cluster index starting with cluster0.</p> <p>If a cluster index is set to 1, that cluster is viewed as token hot due to an insert and/or eject operation.</p> <p>bits 24:31 – Token Standard Hot</p> <p>From right to left, each bit position represents a cluster index starting with cluster0.</p> <p>If a cluster index is set to 1, that cluster is viewed as token hot due a standard operation.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>hot_creation_time</p>	<p>Timestamp of when the hot mask became non-zero. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>Prior to the first change, this field is set to 1970-01-01-00.00.00.000000.</p> <p>This field is new to VOLUME STATUS version 3.</p>

<p>hot_update_time</p>	<p>Timestamp of when the hot mask had at least one bit set. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>Prior to the first change, this field is set to 1970-01-01-00.00.00.000000.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>steal_conflict</p>	<p>When two clusters communicate for the first time after an outage, an initial volume token update occurs where volumes that may have been taken over during the outage are flagged. This field determine the result of this initial token update that must occur prior to a cluster moving to the available state.</p> <p>‘Y’ – This volume may have been taken over by a peer cluster. Token reconciliation must resolve the condition before host access is allowed.</p> <p>‘N’ – ownership takeover is suspected for this volume.</p> <p>‘X’ – An illegal or double takeover was detected. The volume will be put into the volume damaged state.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>close_inconsistent</p>	<p>The volume is currently mounted by a device within the composite library. If the mounted_device value is not -1, then this cluster has the volume mounted.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>last_properties_modify_time</p>	<p>This field represents a timestamp of when this cluster’s volume last had its local properties updated either through a host operation or through a reconciliation process.</p> <p>The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>last_data_modify_time</p>	<p>This field represents a timestamp of when this cluster’s volume became consistent. The consistent_type field determines how it became consistent.</p> <p>The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>This field is new to VOLUME STATUS version 3.</p>

<p>expire_time</p>	<p>If delete-expired, this time represents when it was deleted. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>Prior to the first change, this field is set to 1970-01-01-00.00.00.000000.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>consistent_type</p>	<p>If the volume is consistent, this field states what method was used to bring this volume up to consistency.</p> <p>‘U’ – Unknown ‘T’ – This volume was the actual version written to directly by the last host write operation as a primary TVC cluster. ‘C’ – Grid replication was used to bring the volume up to consistency. ‘M’ – An MES process was used to bring the volume up to consistency. ‘R’ – The volume was recovered as part of a disaster recovery event. ‘F’ – This volume was the actual version written to directly by the last host write operation as a secondary TVC (Fork) cluster with Synchronous copy mode.</p> <p>This field is new to VOLUME STATUS version 3.</p>
<p>removal_timestamp</p>	<p>In a TS7700 Grid configuration, TS7720/TS7760 clusters may remove volumes from tape volume cache after replicating to peer clusters. If the removal state shows that this volume was removed, this timestamp represents the time of when it was removed. If not already removed, and removal is enabled, and this volume is not pinned, this time represents the earliest time of which it can be removed (last access time plus the configured minimum retention time).</p> <p>The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>Prior to the first change, this field is set to 1970-01-01-00.00.00.000000.</p> <p>This field is new to VOLUME STATUS version 3.</p>

removal_state	<p>TS7720 clusters may remove volumes from tape volume cache after replicating to peer clusters in a hybrid configuration if the code level is 8.6.x.x. With the code level 8.7.x.x or later, automatic removal can take place on TS7720/TS7760 clusters in a tapeless only configuration (not hybrid) as well. This field represents the current state of the volume within this cluster. TS7740 clusters will always state 'N'.</p> <p>'N' - The local TS7720/TS7760 cluster has not attempted to remove this volume. If removal is enabled, the removal_timestamp represents the earliest time of which it can be removed (last access time plus the configured minimum retention time).</p> <p>'R' - The local TS7720/TS7760 cluster removed this volume or deleted this volume by EXISTDEL¹ function. The removal_timestamp represents when it was removed.</p> <p>'X' - The local TS7720/TS7760 cluster determined that removal is not possible for this volume. This volume has been retained.</p> <p>'D' - The local TS7720/TS7760 attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>'P' - The local TS7720/TS7760 will never remove this volume because it was configured as a "Pinned" volume under removal_policy.</p> <p>This field is new to VOLUME STATUS version 3.</p>
world_id	<p>This field represents an ASCII version of a 12 byte unique world wide identifier associated with the volume. Only volumes bound as logical worm volumes utilize this field.</p> <p>This field is new to VOLUME STATUS version 3.</p>
write_mount_count	<p>This field represents how many times this volume, under the current world id instance, was mounted and at least one write operation occurred.</p> <p>This field is new to VOLUME STATUS version 3.</p>
removal_policy	<p>This field represents the value assigned to the logical volume via the constructs that were introduced with code level 8.7.x.x. The allowed values are:</p> <p>'0' - "Prefer remove" - If removal is enabled, volumes assigned to this removal policy will be removed first after the minimum retention time has passed. The order of which the volumes are removed is in least recently used (LRU) order when the TS7720/TS7760 is reaching full capacity.</p> <p>'1' - "Prefer keep" Group 1 (LRU) - (Default value and only value allowed in code level 8.6.x.x) - If removal is enabled, volumes assigned to this removal policy will be removed second after all "Prefer Remove" candidates have been removed. The same minimum retention time must have elapsed and the same least recently used access is used to determine order.</p> <p>'4' - Pinned (Disabled) - If removal is enabled, volumes assigned to this removal policy are never removed. A pinned volume is removed from cache.</p>

¹ A function which can be set up by library request. Refer to IBM TS7700 Series z/OS Host Command Line Request User's Guide for more details.

	<p>Note: Independent of being Pinned or having a Minimum Retention Timed defined, all volumes which have been returned to scratch become preferred candidates for removal over volumes which remain private.</p> <p>This field is new to VOLUME STATUS version 4.</p>
recall count	<p>This field represents how many times this volume has been recalled since the last scratch mount on the local. The recall count is reset to 0 every time the volume is used as a scratch volume.</p> <p>This field is new to VOLUME STATUS version 5.</p>
storage group	<p>This field represents the storage group name assigned to this logical volume. If a default storage group is assigned to this logical volume, the field is NULL.</p> <p>This field is new to VOLUME STATUS version 5.</p>
management class	<p>This field represents the management class name assigned to this logical volume. If a default management class is assigned to this logical volume, the field is NULL.</p> <p>This field is new to VOLUME STATUS version 5.</p>
storage class	<p>This field represents the storage class name assigned to this logical volume. If a default storage class is assigned to this logical volume, the field is NULL.</p> <p>This field is new to VOLUME STATUS version 5.</p>
data class	<p>This field represents the data class name assigned to this logical volume. If a default data class is assigned to this logical volume, the field is NULL.</p> <p>This field is new to VOLUME STATUS version 5.</p>
file size	<p>This field represents the compressed file size of this logical volume expressed in bytes as a decimal value.</p> <p>This field is new to VOLUME STATUS version 5.</p>
cache partition number	<p>This field represents the logical cache partition number of this logical volume residing on a TS7700T. A value of -1 indicates the logical volume has no active data. For TS7700 Disk Only and TS7740 models, the logical cache partition is always 0.</p> <p>This field is new to VOLUME STATUS version 7.</p>
last data creation time	<p>This field represents a timestamp of when the volume was created in the domain. The timestamp is updated when the volume is mounted as a scratch and the data is written. It's also updated when the host specifies writeFromBOT bit at the volume mount and the data is written.</p> <p>The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129</p> <p>This field is new to VOLUME STATUS version 7.</p>
compression method	<p>This field represents the "Compression Method" assigned to the logical volume via the Data Class constructs when the logical volume was initially created:</p>

	<p>‘0’ - UNKNOWN - This value is shown only when the logical volume is in an expired/deleted or newly inserted state.</p> <p>‘1’ - FICON - This volume was written from beginning of tape using the FICON compression method. This is the conventional compression method which has been supported since the initial VTS product released in 1997. The compression occurred within the FICON (previously ESCON) adapter. Logical volumes in a deleted or newly inserted state may always report this value until they are used.</p> <p>‘2’ - LZ4 - This volume was written from beginning of tape using the LZ4 compression method. This compression method was introduced in R4.1.2.</p> <p>‘3’ - ZSTD - This volume was written from beginning of tape using the ZSTD compression method. This compression method was introduced in R4.1.2.</p> <p>Note: If compression is disabled for this volume using host JCL or policy management, a compression type is still assigned even though no compression is actually taking place.</p> <p>This field is new to VOLUME STATUS version 8.</p>
<p>logical volume format id</p>	<p>This field represents the volume format ID which is used internally by the TS7700 to store the emulated tape data for this logical volume.</p> <p>‘5’ - This format is applied to all logical volumes that were initially written in any configuration where one or more clusters were running a level prior to 4.1.2. Older volumes will retain this version until they are returned to scratch and reused in a 4.1.2 or later configuration.</p> <p>‘6’ - This format is applied to all logical volumes that were initially written in any configuration where all clusters are running 4.1.2 or later.</p> <p>‘-2’ - This format is applied to all logical volumes on which data is not written yet, or which was already removed or deleted.</p> <p>‘-1’ - This format is applied to all logical volumes that were written with an unknown format due to error or something.</p> <p>Some fields within the STATUS payload may only be populated for logical volume format id version 6 or later volumes. Those fields will explicitly state if they are only populated for a given logical volume format id version.</p> <p>This field is new to VOLUME STATUS version 8.</p>
<p>wrap counter handling</p>	<p>This field represents the “Counters Handling” assigned to the logical volume via the Data Class constructs when it was initially created. It describes whether the TS7700 will allow internal TS7700 and Host counter values to wrap.</p> <p>These options determine what behavior will occur if the total size of channel bytes written (in any sequence during the same mount) to the logical volume exceeds 68GiB bytes or the total amount of channel bytes stored on a tape exceeds 68GiB.</p> <p>‘0’ - Unknown.</p> <p>‘1’ - Wrap Supported - The wrapping of internal host exchanged counters will be supported. Only applications that can properly handle such wrapping events should utilize this option.</p> <p>‘2’ - Wrap Not Supported - The wrapping of internal host exchanged counters</p>

	<p>will not be supported. An early Logical End of Tape (LEOT) will be surfaced prior to such counters wrapping. This option should be used for applications that cannot handle wrapping of counters which can lead to invalid TMS or SMF statistics.</p> <p>This field is new to VOLUME STATUS version 8.</p> <p>This field is only populated for logical volume format id=6 or later logical volumes.</p>
channel length bytes	<p>This field represents the total number of host written bytes stored within the entire logical volume prior to any enabled compression (decimal).</p> <p>This field is new to VOLUME STATUS version 8.</p> <p>This field is only populated for logical volume format ID=6 or later logical volumes.</p>
maximum volume size	<p>This field represents the “Maximum Virtual Volume Size” assigned to the logical volume via the Data Class constructs:</p> <p>‘0’ - 400 MiB or 800 MiB, based on media type (Default).</p> <p>‘10’ - 1,000 MiB.</p> <p>‘20’ - 2,000 MiB.</p> <p>‘40’ - 4,000 MiB.</p> <p>‘60’ - 6,000 MiB.</p> <p>‘250’ - 25,000 MiB.</p> <p>This field is new to VOLUME STATUS version 8.</p>
Blanks	As needed to pad record to 640 bytes

5.2. Cache Contents Information

Volumes accessed by a host are maintained in the tape volume cache managed by each cluster. The cache may be partitioned into up to 8 partitions. The TS7700 controls the movement of logical volumes out of a cache partition as space is needed for newly created or recalled volumes for that partition. The primary goal of the cache management algorithms in the TS7700 is to maximize the utilization of its cache for volumes that have some likelihood to be accessed again. The cache management function of the TS7700 arranges the volumes in a cache partition in the anticipated order they are to be removed when space is needed. In order to remove a volume from cache it must first have been premigrated (which means copied to a physical tape or cloud tier). For this reason, it is possible that volumes with a higher order number are removed from cache first. As part of the Advanced Policy Management functions of the TS7700, the Storage Class construct provides for customer control of the partition for a volume's data and cache preferencing policies for the management of the volume in cache. Two preferencing policies are supported:

- Preference Group 0 (PG0)

When space is needed in the cache, premigrated volumes assigned to preference group 0 are removed from cache before volumes assigned to preference group 1. Within preference group 0, the volumes are ordered for removal from cache by largest volumes first.

Note: Volumes assigned to preference group 0 may also be removed from the cache, independent of the need for cache space, as a background task within the TS7700.

- Preference Group 1 (PG1)

When space is needed in the cache and there are no premigrated preference group 0 volumes to remove, premigrated volumes assigned to preference group 1 are removed. Within preference group 1, the volumes are ordered for removal from cache based on time since last access (LRU).

Note: The order of removal of a volume from cache may also be influenced by other storage constructs settings for a volume, so the order presented in the response data should not be relied on to be exact.

The contents of the cache associated with the specific cluster the request volume is written to are returned in the response records. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the cache contents of all of the clusters.

The response records are written in 80 byte fixed block format.

Note: The generation of the response may take several minutes to complete depending on the number of volumes in the cache and how busy the TS7700 cluster is at the time of the request.

Record 6

This record provides a heading for the data records to follow:

ORDER VOLSER DATE/TIME IN CACHE PG PART SIZE PG-SG

Bytes	Name	Description
1-3		Blank characters
4-8		'ORDER' Heading for the volume order column
9	Field Delimiter	Blank character
10-15		'VOLSER' Heading for the volume serial number column
16	Field Delimiter	Blank character
17-34		'DATE/TIME IN CACHE' Heading for the data/time stamp columns
35-38	Field Delimiter	Blank characters
39-40		'PG' Heading for the currently assigned preference group column
41-42	Field Delimiter	Blank characters
43-46		'PART' Heading for the partition column
47-48	Field Delimiter	Blank characters This field is new to CACHE CONTENTS version 3.

49-52		'SIZE' Heading for the LVOL size column This field is new to CACHE CONTENTS version 3
53-61	Field Delimiter	Blank characters This field is new to CACHE CONTENTS version 6.
62-66		'PG-SG' Heading for the storage class assigned preference group column This field is new to CACHE CONTENTS version 6.
67-80	Blanks	Blank character padding

Record 7-N

Each of these records provides information for one logical volume. The cache contents information is provided by partition, then in the order in which the volumes are anticipated to be removed from the cache partition.

```

1 VOL020 11/30/2005 11:57:00 0 0 799997952 2
2 VOL019 11/29/2005 03:00:00 0 0 799997952 2
3 VOL023 11/20/2005 09:57:00 1 0 799997952 2
4 VOL016 11/20/2005 10:01:00 1 0 799997952 2
5 ABC309 11/20/2005 17:31:00 0 1 799997952 2
6 ABC333 11/20/2005 11:44:00 1 1 799997952 2
7 ABC789 11/20/2005 10:51:00 1 1 799997952 2
8 ABC234 11/20/2005 04:27:00 1 1 983040 2
9 ABC045 11/19/2005 21:45:00 1 1 9994240 2
    
```

Bytes	Name	Description
1-8	Order in Cache Partition	The order in which volumes are to be removed from the cache partition, right justified and blank filled.
9	Field Delimiter	Blank character
10-15	Logical Volser	Six character volume serial number
16	Field Delimiter	Blank character
17-26	Date in Cache	The date the volume was created or recalled into the cache in the format of MM/DD/YYYY.
27	Field Delimiter	Blank character
28-35	Time in Cache	The time the volume was created or recalled into the cache in the format of HH:MM:SS.
36-39	Field Delimiter	Blank characters
40	Currently Assigned Preference Group	Contains the current preference group the volume is assigned, which may or may not be the same as the Storage Class Assigned Preference Group. The Currently Assigned Preference Group for a logical volume may be different than the Storage Class Assigned Preference Group as a result of a recall by a background task such as a copy export reclamation initiated by a library request, or an automatic read-only recovery.
41-45	Field Delimiter	Blank characters
46	Partition	Contains the partition the volume is resident in
47-48	Field Delimiter	Blank characters This field is new to CACHE CONTENTS version 3
49-60	LVOL size	Size of the volume in cache after compression, right justified and blank filled. This data is specified in bytes as a decimal value. This field is new to CACHE CONTENTS version 3
61	Field Delimiter	Blank character
62	Storage Class Assigned Preference Group	Contains the preference group the volume is assigned as defined by the volume's storage class.
63-80	Blanks	Blank character padding

Note: The contents of the cache typically are all private volumes; however, it is possible that some may have been returned to scratch status soon after being written. The VTS does not filter the cache contents based on the private or scratch status of a volume.

5.3. Physical Volume to Logical Volume Mapping Information

The TS7700 maintains the mapping between logical and physical volumes in a database on each cluster. It is possible that there are inconsistencies in the mapping information provided with this function. This results when a logical volume is being moved from one physical volume to another. For a period of time, the volume is shown on more than one physical volume. This can result in a small number of logical volumes reported as being on physical volumes which they were located on in the past, but are not presently located on.

Even with some inconsistencies, the mapping data is useful to customers that want to design jobs that recall data efficiently off of physical volumes. If the logical volumes reported on a physical volume are recalled together, the efficiency of the recalls will be increased. If a logical volume with an inconsistent mapping relationship is recalled, it will recall correctly, but an additional amount of a different physical volume may be required.

The physical volume to logical volume mapping associated with the physical volumes managed by the specific cluster the request volume is written to are returned in the response records. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the mapping for all physical volumes.

Starting with code level 8.30.x.xx, the TS7700 supports up to 4 million logical volumes. If the total record of the volume mapping information exceeds 4 million lines, an operator informational message is surfaced to the attached hosts. The message should state “BVIR VOLUME MAP EXCEEDS 4000000 RECORDS”.

Refer to the “IBM® TS7700 Series Operator Informational Messages” white paper (WP101689) on Techdocs for more details.

The response records are written in 80 byte fixed block format.

Note: The generation of the response may take several minutes to complete depending on the number of active logical volumes in the library and how busy the TS7700 cluster is at the time of the request.

Record 6

If this is a TS7700 Disk Only configuration (**with or without cloud enablement feature**), the following record is returned.

Bytes	Name	Description
1-57		‘NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE’
58-80	Blanks	Blank character padding

This record provides a heading for the data records to follow:

PHYSICAL LOGICAL P/B ORDER PART SIZE

Bytes	Name	Description
1-8		‘PHYSICAL’ Heading for the physical volume serial number column
9-10	Field Delimiter	Blank characters
11-17		‘LOGICAL’ Heading for the logical volume serial number column
18	Field Delimiter	Blank character
19-21		‘P/B’ Heading for the primary or backup pool indicator column
22-23	Field Delimiter	Blank characters
24-28		‘ORDER’ Heading for the logical volume order column
29-31	Field Delimiter	Blank characters
32-35		‘PART’ Heading for the logical volume spanning indicator column
36-44	Blanks	Blank characters
45-48		‘SIZE’ Heading for the logical volume compresses size
49-80	Blanks	Blank character padding

The primary and backup pool indicator column indicates whether the volume being reported resides on a primary pool volume or the secondary pool volume as defined using the Selective Dual Copy function that is part of the TS7700's advanced policy management function.

The size field reports the number of MiBs, rounded to two places after the decimal point, a logical volume occupies on the physical volume. This includes the effect of the compression performed on the data by the TS7700, but does not include any effect of the compression performed by the physical drive. Any volume with a size of less than 5KB will report a size of 0.00. When a volume indicates that it spans, the size of the entire volume is indicated for each of the physical volumes.

Note: Only legacy data created on prior B10/B18/B20 VTSs and migrated under the control of a TS7700 will indicate that they span from one physical volume to another. Volumes created or recopied by the TS7700 do not span physical volumes.

Record 7-N

Each of these records provides information for a logical to physical volume association. The records are ordered alphanumerically (0-9, A-Z) by physical volser, then by logical sequence of the active logical volumes on the physical volume.

```
P00024      GK0000      P  000001  1  OF  1           23.45 MiB
P00024      GK0020      P  000002  1  OF  1           76.50 MiB
P00024      GK0010      P  000003  1  OF  1          145.70 MiB
P00024      GK0030      P  000004  1  OF  1          670.32 MiB
P00024      GK0040      P  000005  1  OF  1         1934.12 MiB
P00024      GK0060      P  000006  1  OF  1            0.00 MiB
P00024      GK0050      P  000007  1  OF  2          540.12 MiB
P00467      GK0050      P  000001  2  OF  2          540.12 MiB
```

Bytes	Name	Description
1-6	Physical Volser	Physical volser the logical volser is located on, left justified and padded with blanks.
7-10	Field Delimiter	Blank characters
11-16	Logical Volser	Logical volser, left justified and padded with blanks.
17-19	Field Delimiter	Blank characters
20	Pool Indicator	'P' indicates the logical volume is the primary copy, 'B' indicates that the volume is the backup copy.
21-22	Field Delimiter	Blank characters
23-28	Order	The relative order of the logical volume on the physical volume.
29	Field Delimiter	Blank character
30-35	Spanning	Indicates whether the logical volume spans to another physical volume or not and if so, which part.
36-41	Blanks	Blank characters
42-45	Size	Integer part of the size, right justified and padded with leading blanks
46	Decimal Point	'.'
47-48	Hundredths	Decimal part of the size
49	Field Delimiter	Blank character
50-52	Units Indicator	'MiB' indicates the size is in megabytes, which is 1024x1024 bytes
53-80	Blanks	Blank character padding

5.4. Point In Time Statistics Information

A TS7700 is continually logging information regarding the activities within it. The logged information is referred to as statistical information and is recorded in two forms, Point In Time and Historical. Point In Time statistics indicate the state and operation aspects of the TS7700 over a short interval of time. The time interval is currently approximately 15 seconds. A request for Point In Time statistics will respond with the data accumulated in the interval completed just prior to the request being processed. Because of this, the state information reported may lag the actual state of the TS7700 by an interval.

Other than an information header, Point In Time statistics are provided in a mixture of character and binary format fields. The record sizes and format of the statistical records are defined in the *IBM® TS7700 Series Statistical Data Format White Paper*.

The Point In Time statistics for all clusters are returned in the response records. In a TS7700 Grid configuration, this means that the request volume can be written to any cluster to obtain the information for the entire configuration.

Note: If a cluster or node is not available at the time the point in time statistics are recorded, except for the headers, all the data fields for that cluster or node will be zeroes.

Note: The request records are written in FB format. To read the response records, use the Undefined (U) format with a maximum blocksize of 24000. The response records are variable in length.

Record 6

This record provides a human readable header for the response record.

POINT IN TIME STATISTICS

Bytes	Name	Description
1-24		'POINT IN TIME STATISTICS'
25-80	Blanks	Blank character padding

Record 7-N

Each of these records provides point in time information for the nodes of the TS7740 Cluster. Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version and data type of the record. Since the point in time information is provided for all nodes in the configuration, including nodes in all clusters in a Grid configuration, the number of records returned depends on the configuration. There are two records for each VNode and two records for each HNode. The application processing the response volume records should not assume that the records for a V or H node are presented in any specific order.

Bytes	Name	Description
1-2	Length	Total length of the response record in binary
3	Version	Binary 1
4	Data Type	Identifies the type of response record data
5-length	Data	Response data

Data Type

The statistics for a node in the subsystem are subdivided into different data types. Refer to the *IBM® TS7700 Series Statistical Data Format White Paper* for the description of the data types.

Binary Response Data

Refer to the *IBM® TS7700 Series Statistical Data Format White Paper* for the size and format of the response data for each data type.

5.5. Historical Statistics Information

A TS7700 is continually logging information regarding the activities within it. The logged information is referred to as statistical information and is recorded in two forms, Point In Time and Historical. Historical statistics indicate the operational aspects of the TS7700 accumulated over a 15 minute interval of time. The data from each 15 minute interval is maintained and logged within the TS7700. A request for Historical statistics will result in a response file that contains all of the data logged up to that point for the requested julian day.

Other than an information header, Historical statistics are provided in character and binary format fields. The sizes and format of the statistical records are defined in the *IBM® TS7700 Series Statistical Data Format White Paper*.

The Historical statistics for all clusters are returned in the response records. In a TS7700 Grid configuration, this means that the request volume can be written to any cluster to obtain the information for the entire configuration.

Note: If a cluster or node is not available at the time the historical statistics are recorded, except for the headers, all the data fields for that cluster or node will be zeroes.

Note: The TS7700 retains 90 days worth of historical statistics. If you wish to keep statistics for a longer period of time, it is recommended that you retain the logical volumes used to obtain the statistics.

Note: The request records are written in FB format. To read the response records, use the Undefined (U) format with a maximum blocksize of 24000. The response records are variable in length.

Record 6-N

The historical statistical response records for each day requested are preceded by the following human readable header. If historical statistics are requested for a day which the TS7700 does not have historical data, the No Historical Statistics for xxx header is the only response record provided for the day.

HISTORICAL STATISTICS FOR xxx or NO HISTORICAL STATISTICS FOR xxx

Bytes	Name	Description
1-46		'HISTORICAL STATISTICS for xxx' or 'NO HISTORICAL STATISTICS FOR xxx', left justified, padded with blanks on the right.
47-80	Blanks	Blank character padding

Data Response Records

Each of these records provides historical information for the nodes of the TS7740 Cluster for the day indicated in the header record. Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version and data type of the record. Since the historical information is provided for all nodes in the configuration, including nodes in all clusters in a Grid configuration, the number of records returned depends on the configuration. There are two records for each VNode and from 4 to 7 records for each HNode (depends on the number of physical libraries attached to the cluster). The application processing the response volume records should not assume that the records for a V or H node are presented in any specific order.

Bytes	Name	Description
1-2	Length	Total length of the response record in binary
3	Version	Binary 1
4	Data Type	Identifies the type of response record data
5-length	Data	Response data

Data Type

The statistics for a node in the subsystem are subdivided into different data types. Refer to the *IBM® TS7700 Series Statistical Data Format White Paper* for the description of the data types.

Binary Response Data

Refer to the *IBM® TS7700 Series Statistical Data Format White Paper* for the size and format of the response data for each data type.

5.6. Physical Media Pools Information

The TS7700 supports separating the physical volumes it manages into pools. The supported pools include a pool that contains scratch (empty) volumes that are common and up to 32 pools that may contain scratch (empty) and data (filling/full) volumes. Pools can borrow and return volumes from the common scratch pool. Each pool can contain several types of media.

For pool 00 (common pool), since it only contains empty volume, only the empty count is returned. Volumes that have been borrowed from the common pool are not included.

For pools 1-32, a count of the physical volumes that are empty, are empty and waiting for erasure, are in the process of being filled or have been marked as full is returned. The count for empty includes physical volumes that have been specifically assigned to the pool as well as volumes that were borrowed from the common scratch pool but have not yet been returned. The count of volumes that are marked as Read Only or Unavailable (including destroyed volumes) are returned. Also, the full data volumes contain a mixture of valid and invalid data. Response records are provided for the distribution of active data on the data volumes marked as full for a pool.

Information is returned for the common pool and all other pool that are defined and have physical volumes associated with them.

The physical media pool information managed by the specific cluster the request volume is written to are returned in the response records. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the physical media pool information for all clusters.

The response records are written in 80 byte fixed block format.

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

Record 6

If this is a TS7700 Disk Only configuration (with or without cloud enablement feature), the following record is returned.

Bytes	Name	Description
1-57		'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
58-80	Blanks	Blank character padding

This record provides a heading for the data records to follow:

```
R POOL MEDIA EMPTY FILLING FULL ERASE ROR UNAVAIL
```

Bytes	Name	Description
1		'R' Heading for the record type column
2	Field Delimiter	Blank character
3-6		'POOL' Heading for the pool number column
7	Field Delimiter	Blank character
8-12		'MEDIA' Heading for the media type column
13-14	Field Delimiter	Blank characters
15-19		'EMPTY' Heading for the empty count column
20-21	Field Delimiter	Blank characters
22-28		'FILLING' heading for the filling count column
29-31	Field Delimiter	Blank characters
32-35		'FULL' Heading for the full count column
36-37	Field Delimiter	Blank characters
38-42		'ERASE' Heading for the waiting to erase count column

43-44	Field Delimiter	Blank characters
45-47		'ROR' Heading for the read only recovery count column
48-49	Field Delimiter	Blank characters
50-56		'UNAVAIL' Heading for the unavailable count column
57-80	Blanks	Blank character padding

Record 7-N

Each of these records provides the counts for a specific media type for each pool. Up to 8 media types may be reported for a pool. The count of ERASE, ROR, UNAVAIL physical volumes are also included in the count of EMPTY, FILLING, FULL physical volumes.

In the example below, the pool 01 has 746 empty, filling, full JA physical volumes (746=134+3+609), and 123 physical volumes of the 746 are in the read only recovery state and 4 physical volumes of the 746 are in the unavailable or destroyed state.

The pool 01 also has 4 empty, filling, full JB physical volumes, and 1 physical volume of the 4 needs to be erased.

```

C 00 J 00340
C 00 K 00523
C 00 JA 00083
C 00 JB 00051
C 00 JJ 00002
C 01 JA 00134 00003 00609 00000 00123 00004
C 01 JB 00001 00002 00001 00001 00000 00000
C 02 K 00023 00002 00325 00000 00000 00045
C 03 J 00103 00001 00336 00004 00000 00000
C 03 K 00034 00000 00185 00002 00001 00000
    
```

Bytes	Name	Description
1	Record Type	Record type identifier. 'C' for media count.
2-4	Field Delimiter	Blank character padding
5-6	Pool Number	The pool number. Pool number is right justified and padded with a leading zero.
7-10	Blanks	Blank character padding
11-12	Media Type	The media type defined for the pool. Media types with a single character are padded on the left with a blank.
13-14	Field Delimiter	Blank characters
15-19	Empty Count	The count of the physical volumes that are empty for the media type. The count is right justified and padded with leading zeros.
20-23	Field Delimiter	Blank characters
24-28	Filling Count	The count of the physical volumes that are in the filling state for the media type. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
29-30	Field Delimiter	Blank characters
31-35	Full Count	The count of the physical volumes that have been marked full for the media type. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
36-37	Field Delimiter	Blank characters
38-42	Erase Count	The count of the physical volumes that have been reclaimed, but need to be erased before they become empty. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
43-44	Field Delimiter	Blank characters
45-49	ROR Count	The count of the physical volumes that are in the read only recovery state.
50-51	Field Delimiter	Blank characters

52-56	Unavailable Count	The count of the physical volumes that are in the unavailable or destroyed state.
57-80	Blanks	Blank character padding

Record N+1

For pools 1-32, this record provides a heading for the data records to follow:

R POOL MEDIA 0+ 10+ 20+ 30+ 40+ 50+ 60+ 70+ 80+ 90+

Bytes	Name	Description
1		'R' Heading for the record type column
2	Field Delimiter	Blank character
3-6		'POOL' Heading for the pool number column
7	Field Delimiter	Blank character
8-12		'MEDIA' Heading for the media type column
13-16	Field Delimiter	Blank characters
17-18		'0+' Heading for the 0-10% count column
19-21	Field Delimiter	Blank characters
22-24		'10+' Heading for the 10-20% count column
25-27	Field Delimiter	Blank characters
28-30		'20+' Heading for the 20-30% count column
31-33	Field Delimiter	Blank characters
34-36		'30+' Heading for the 30-40% count column
37-39	Field Delimiter	Blank characters
40-42		'40+' Heading for the 40-50% count column
43-45	Field Delimiter	Blank characters
46-48		'50+' Heading for the 50-60% count column
49-51	Field Delimiter	Blank characters
52-54		'60+' Heading for the 60-70% count column
55-57	Field Delimiter	Blank characters
58-60		'70+' Heading for the 70-80% count column
61-63	Field Delimiter	Blank characters
64-66		'80+' Heading for the 80-90% count column
67-69	Field Delimiter	Blank characters
70-72		'90+' Heading for the 90+% count column
73-80	Blanks	Blank character padding

Record N+2-M

Each of these records provides the active data distribution counts for a specific media type. Up to 8 media types may be reported for a pool.

```
D 01 JA 00000 00002 00034 00056 00092 00078 00084 00065 00195 00183
D 01 JB 00000 00000 00000 00000 00000 00000 00000 00002 00003 00001
D 02 K 00000 00000 00000 00000 00003 00023 00033 00087 00076 00103
D 03 J 00000 00000 00000 00002 00023 00046 00003 00056 00072 00134
D 03 K 00000 00000 00003 00007 00015 00012 00041 00018 00024 00065
```

Bytes	Name	Description
1	Record Type	Record type identifier. 'D' for active media data distribution count.
2-4	Field Delimiter	Blank character padding
5-6	Pool Number	The pool number. Pool number is right justified and padded with a leading zero.
7-10	Blanks	
11-12	Media Type	The media type defined for the pool. Media types with a single character are padded on the left with a blank.
13	Field Delimiter	
14-18	0-9 Count	The count of the physical volumes that are marked full and have an active data percentage greater than 0 and less than 10. The count is right justified and padded with leading zeroes.
19	Field Delimiter	
20-24	10s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 10 and less than 20. The count is right justified and padded with leading zeroes.
25	Field Delimiter	
26-30	20s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 20 and less than 30. The count is right justified and padded with leading zeroes.
31	Field Delimiter	
32-36	30s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 30 and less than 40. The count is right justified and padded with leading zeroes.
37	Field Delimiter	
38-42	40s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 40 and less than 50. The count is right justified and padded with leading zeroes.
43	Field Delimiter	
44-48	50s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 50 and less than 60. The count is right justified and padded with leading zeroes.
49	Field Delimiter	
50-54	60s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 60 and less than 70. The count is right justified and padded with leading zeroes.
55	Field Delimiter	
56-60	70s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 70 and less than 80. The count is right justified and padded with leading zeroes.
61	Field Delimiter	
62-66	80s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 80 and less than 90. The count is right justified and padded with leading zeroes.
67	Field Delimiter	

68-72	90s Count	The count of the physical volumes that are marked full and have an active data percentage equal to or greater than 90. The count is right justified and padded with leading zeroes.
73-80	Blanks	Blank character padding

5.7. Physical Volume Status Information

A database is maintained on each individual TS7740 Cluster that contains information related to the management of the physical volumes on the cluster.

The physical volume status information returned represents the status of the volume(s) on the cluster the request volume is written. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the physical volume status information for the individual clusters. A response record is written for each physical volume, selected based on the volume serial number mask or pool number specified in the request, that exists in the cluster.

A response record consists of the database fields defined in the following table. Fields are presented in the order defined in the table and are comma (,) separated. The overall length of each record is 400 bytes with blank padding after the last field as needed. For example, the first few fields of the record returned for volser A03599 would be:

A03599, 2, FULL, READ-WRITE, 2007-05-05-06.40.08.030061, 2007-05-04-13.45.15.918473, . . .

Note: The generation of the response may take several minutes to complete depending on the number of volumes requested and how busy the TS7700 cluster is at the time of the request.

Record 6-N

If this is a TS7700 Disk Only configuration (with or without cloud enablement feature), the following record is returned.

Bytes	Name	Description
1-57		'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'
58-80	Blanks	Blank character padding

If the request is for a specific physical volume and it doesn't exist in the cluster's database or the request is for a range of physical volumes and none of the volumes exist in the cluster's database, or the request is for a pool and there are no physical volumes currently resident in the pool, the following record is returned.

NO VOLUME RECORD(S) EXIST FOR THE REQUESTED VOLUME(S)

Bytes	Name	Description
1-53		'NO VOLUME RECORD(S) EXIST FOR REQUESTED VOLUME(S)'
54-80	Blanks	Blank character padding

For the requested volumes that do exist in the cluster's database, each of these records provide information for each physical volume.

Field Name	Description
VOLSER	6 character volume serial number
CURRENT_POOL	This field indicates the pool that the volume is currently assigned to. Pool 0 is the common scratch pool. Pools 1-32 are the specific data pools.

<p>VOLUME_STATUS</p>	<p>This field indicates the volume's current capacity state. The following are the values that can be indicated:</p> <p>'EMPTY' The volume contains no data and is available for use as a physical scratch volume.</p> <p>'FILLING' The volume contains valid data, but is not yet full. It is available for additional data to be added to it.</p> <p>'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.</p> <p>'UNKNOWN' The volume's capacity state is unknown.</p>
<p>VOLUME_ACCESS</p>	<p>This field indicates the volume's current accessibility state. The following are the values that can be indicated:</p> <p>'READ-WRITE' The volume can be read from or written to.</p> <p>'READ-ONLY' The volume contains valid data, but writes to it are not allowed. This may because of an error that has been detected during a prior write operation to the volume, or a management operation, such as reclaim, is going to be performed against the volume.</p> <p>'UNAVAILABLE' The volume is unavailable to the TS7740.</p> <p>'DAMAGED' The volume has been damaged physically or logically such that it cannot be mounted and or read from.</p> <p>'COPY_EXPORTED' The volume has been exported using the copy export function and is not currently resident in the library.</p> <p>'CE_RECLAIM' The volume had a state of COPY_EXPORTED and a host console request was received requesting that the volume be reclaimed. It will remain in the CE_RECLAIM state during the reclaim process.</p>
<p>STARTED_EXPIRING_TIMESTAMP</p>	<p>Timestamp of when the data first started expiring on the physical volume. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>This field is set to 1970-01-01-00.00.00.000000 when the volume becomes empty (returned to scratch) or is inserted into the library.</p>

<p>BECAME_EMPTY_TIMESTAMP</p>	<p>Timestamp of when all of the data on the volume last become invalid and the volume is empty. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.</p>
<p>BECAME_FULL_TIMESTAMP</p>	<p>Timestamp of when the volume is last marked as full. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.</p>
<p>RECORDING_FORMAT</p>	<p>This field indicates the format that was used in recording the data on the volume. The following are the values that can be reported:</p> <p>‘4’ The volume is written in the 3592 J1A format.</p> <p>‘5’ The volume is written in the TS1120 E05 format.</p> <p>‘6’ The volume is written in the TS1120 E05E format (encrypted).</p> <p>‘7’ The volume is written in the TS1130 E06 format.</p> <p>‘8’ The volume is written in the TS1130 E06E format (encrypted).</p> <p>‘9’ The volume is written in the TS1140 E07 format.</p> <p>‘10’ The volume is written in the TS1140 E07E format (encrypted).</p> <p>‘11’ The volume is written in the TS1150 E08 format.</p> <p>‘12’ The volume is written in the TS1150 E08E format (encrypted).</p> <p>Values not defined are reserved.</p>
<p>PCT_UTILIZED</p>	<p>This field indicates the percent of the volume that contains active data. The percentage is determined using the active_size and the total_bytes_written fields. The number is reported in 1/10th of a percent, rounded up. This field is updated hourly (if needed). It is reset to 0 when the volume becomes empty (scratch).</p>
<p>TOTAL_BYTES_WRITTEN</p>	<p>This field indicates the number of bytes that were written to the volume when it was filled. It is reset to 0 when the volume becomes empty (scratch).</p>
<p>ACTIVE_SIZE</p>	<p>This field indicates the number of bytes of the active data on the volume. It is reset to 0 when the volume becomes empty (scratch).</p>
<p>ERROR_TYPE_FLAG</p>	<p>This field indicates whether the TS7700 has detected an error with the use of the volume. The error values are used by IBM service.</p>

SCRATCH_COUNT	This field indicates the number of times the volume has been returned to scratch status since it was inserted into the library.
MOUNT_COUNT	This field indicates the number of times the volume has been mounted since it was inserted into the library.
MEDIA_TYPE	<p>This field indicates the media type for the volume. The following are the values that can be reported:</p> <p>'5' IBM JA</p> <p>'6' IBM JJ</p> <p>'7' IBM JB</p> <p>'8' IBM JC</p> <p>'9' IBM JK</p> <p>'10' IBM JD</p> <p>'11' IBM JL</p> <p>Values not defined are reserved.</p>
BATE_ROR_REASON	If the volume has a volume access value of 'Read-Only', this field provides additional information on the reason for the access value. A volume may be marked as read-only for reasons other than a media error. These values may be used for IBM Service.
DVM_ROR_REASON	If the volume has a volume access value of 'Read-Only', this field provides additional information on the reason for the access value. A volume may be marked as read-only for reasons other than a media error. These values may be used for IBM Service.
ERASE_FLAG	This field indicates whether or not the volume needs to be secure data erased.
MOUNT_FAILURE_COUNT	This field indicates the number of time a mount failure occurred with this volume since it was inserted into the library.
LAST_DEVICE_MOUNTED	This field contains the device number the volume was last mounted on when an error occurred. This field will be null if there has been no error mounting this volume.
LAST_WRITE_TIMESTAMP	<p>Timestamp of when the volume was last written to. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p> <p>When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.</p>
LAST_READ_TIMESTAMP	<p>Timestamp of when the volume was last read from. The format of the timestamp is:</p> <p>Year-Month-Day-Hour.Minute.Second.Microsecond for example:</p> <p>2006-05-23-19.34.23.876129</p>

	When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.
KEKLABEL1	If the volume has been encrypted, this field contains the first key label. A key label can be up to 64 characters in length. If the volume is not encrypted, this field is null. It is set to null when the volume becomes empty (scratch).
KEKLABEL2	If the volume has been encrypted, this field contains the second key label. A key label can be up to 64 characters in length. If the volume is not encrypted, this field is null. It is set to null when the volume becomes empty (scratch).
INSERT_TIMESTAMP	Timestamp of when the volume was inserted into the library. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129
ADSM_FORMAT_FLAG	This field indicates whether the physical volume is written using the pre-TS7700 (B10/B20 VTS) or TS7700 format. The following are the values that can be indicated: '0' The volume is written in the TS7700 format. '1' The volume is written in the pre-TS7700 format.
DB_BACKUP_NAME	This field indicates whether the physical volume contains a TS7700 database backup or not. Only a copy exported volume contains the backup and its name is included. Otherwise, "NULL" is included to indicate the physical volume does not contain a database backup. This field is new to PHYSICAL VOLUME STATUS version 6.
Blanks	As needed to pad record to 400 bytes

Example of Physical Volume Status Volume output

```
VTS BULK VOLUME DATA REQUEST
PHYSICAL VOLUME STATUS VOLUME Q00794
0/03/2013 04:19:23 VERSION 06
S/N: 00C4A LIB ID: BA54A
```

```
Q00794, 1, FILLING, READ-WRITE, 2013-09-28-11.50.00.000000, 1970-01-01-00.00.00.000000,
1970-01-01-00.00.00.000000, 4, 988, 20778009783, 20541898411, 0, 0, 11, 5, 0, 0, N, 0, NULL,
2013 -03-04.06.39.000000, 2013-10-03-04.06.38.000000, NULL, NULL, 2013-09-26-06.46.27.064539, 0, NULL
```

Example of Physical Volume Status Pool output

```
VTS BULK VOLUME DATA REQUEST
PHYSICAL VOLUME STATUS POOL 0
```

10/03/2013 05:29:24 VERSION 06
S/N: 00C4A LIB ID: BA54A

Q00797,0,EMPTY,READ-WRITE,1970-01-01-00.00.00.000000,2013-10-02-08.50.00.000000,
1970-01-01-00.00.00.000000,0,0,0,0,0,0,1,0,5,0,0,N,0,NULL,1970-01-01-00.00.00.000000,
1970-01-01-00.00.00.000000,NULL,NULL,2013-09-26-06.46.27.783401,0,NULL
Q00798,0,EMPTY,READ-WRITE,1970-01-01-00.00.00.000000,2013-10-02-08.50.00.000000,
1970-01-01-00.00.00.000000,0,0,0,0,0,0,1,0,5,0,0,N,0,NULL,1970-01-01-00.00.00.000000,
1970-01-01-00.00.00.000000,NULL,NULL,2013-09-26-06.46.27.984190,0,NULL

5.8. Copy Audit Information

A database is maintained on each individual TS7700 Cluster that contains status information about the logical volumes defined to the grid. Two key pieces of information are whether or not the cluster contains a valid copy of a logical volume and whether the copy policy for the volume indicates that it should have a valid copy.

This request performs an audit of the databases on a set of specified TS7700 distributed libraries to determine if there are any volumes that do not have a valid copy on at least one of them. If the COPYMODE option is specified, whether or not the volume is supposed to have a copy on the distributed library is taken into account in determining whether that distributed library has a valid copy. If COPYMODE is specified and the copy policy for a volume on a specific cluster is 'S', 'R' or 'D', then that cluster is considered during the audit. If COPYMODE is specified and the copy policy for a volume on a specific cluster is 'N', then the volume's validity state is ignored because that cluster does not need to have a valid copy. The request then returns a list of any volumes that do not have a valid copy, subject to the copy mode if the COPYMODE option is specified, on the TS7700s clusters specified as part of the request. The specified clusters may not have a copy for several reasons:

- The copy policy associated with the volume did not specify that any of the clusters specified in the request were to have a copy and the COPYMODE option was not specified. This may be because of a mistake in defining the copy policy or really intended. For example, volumes used in a disaster recovery test only need to reside on the disaster recovery TS7700 and not on the production TS7700s. If the request specified only the production TS7700s, all of the volumes used in the test would be returned in the list.
- The copies have not yet been made from a source TS7700 to one or more of the specified clusters. This could be because the source TS7700 or the links to it are unavailable or because a copy policy of deferred was specified and a copy had not been completed when the audit was performed. In addition, one or more of the specified clusters may have completed their copy and then had their copy automatically removed as part of the TS7700 automated removal policy function. Automatic removal can only take place on TS7720/TS7760 clusters.
- Each of the specified clusters contained a valid copy at one time but has since removed them as part of the TS7700 automated removal policy function. Automatic removal can only take place on TS7720/TS7760 clusters in a hybrid configuration if the code level is 8.6.x.x or higher. With the code level 8.7.x.x or later, automatic removal can take place on TS7720/TS7760 clusters in a tapeless only configuration (not hybrid) as well.

The Copy Audit request is intended to be used for the following:

- A TS7700 is to be removed from a Grid configuration. Prior to its removal you want to ensure that the TS7700s that are to remain in the Grid configuration have a copy of all important volumes that were created on the TS7700 that is to be removed.
- A condition has occurred (could be a site disaster or as part of a test procedure) where one of the TS7700s in a Grid configuration is no longer available and you want to determine which, if any, volumes on the remaining TS7700s do not have a valid copy.

In the Copy Audit request, you need to specify which TS7700 clusters are to be audited. The clusters are specified by using their associated distributed library ID (this is the 5 character library sequence number defined when the TS7700 cluster was installed). If more than one distributed library ID is specified, they are separated by a comma. The following are the rules for determining which TS7700 clusters are to be included in the audit:

- When the INCLUDE parameter is specified, all specified distributed library IDs will be included in the audit. All clusters associated with these IDs must be available or the audit will be failed.
- When the EXCLUDE parameter is specified, all specified distributed library IDs will be excluded from the audit. All other clusters in the Grid configuration must be available or the audit will be failed.
- Distributed library IDs specified are checked for being valid in the Grid configuration. If one or more of the specified distributed library IDs are invalid, the Copy Audit is failed and the response will indicate the IDs that are considered invalid.
- Distributed library IDs must be specified or the Copy Audit is failed.

Here are some examples of valid requests (for these assume a three cluster Grid configuration with distributed library IDs of BA45A, BA45B and BA45C):

COPY AUDIT INCLUDE BA45A - Audits the copy status of all volumes on only the cluster associated with distributed library ID BA45A.

COPY AUDIT COPYMODE INCLUDE BA45A - Audits the copy status of volumes that also have a valid copy policy on only the cluster associated with distributed library ID BA45A.

COPY AUDIT INCLUDE BA45B,BA45C - Audits the copy status of volumes on the clusters associated with distributed library IDs BA45B and BA45C.

COPY AUDIT EXCLUDE BA45C - Audits the copy status of volumes on the clusters in the Grid configuration associated with distributed library IDs BA45A and BA45B.

The following example shows the expected COPY AUDIT result on a specific configuration to understand how COPY AUDIT request works

The configuration is a 6 cluster grid (distributed library IDs are D0000, D0001, D0002, D0003, D0004 and D0005 (cluster 0 – 5)) and there are 3 sites (1 production and 2 DR sites). The cluster 0 and 1 are in the production site, cluster 2 and 3 are in the DR site1 and cluster 4 and 5 are in the DR site2.

A total of 10 volumes exist and their copy mode, removal status and volume consistency are listed in the following table:

VOLSER	Copy Mode	Removal Status	Consistent	Note
VOL000	RRDDDD	NNNNNN	YYYYYY	All requested copies made and still exist
VOL001	RRDDDD	NNRRRN	YYNNYY	Automatic removal occurred on cluster 2 and 3.
VOL002	RRDDDD	NNRNRN	YYNYNY	Automatic removal occurred on cluster 2 and 4.
VOL003	RRDDDD	NNNNNN	YYNYNY	Deferred copy to cluster 3 and 5 are not yet completed.
VOL004	RRDDDD	NNNNNN	YYNNNN	Deferred copy to cluster 2, 3, 4, and 5 are not yet complete.
VOL005	RRNNNN	NNNNNN	YYNNNN	Only production site (clusters 0 and 1) has a valid copy mode.
VOL006	RNDNDN	NNNNNN	YNYNYN	All requested copies made and still exist
VOL007	RNDNDN	RNNNRN	NNYNNN	Automatic removal occurred on cluster 0 and 4.
VOL008	NNRNNN	NNNNNN	NNYNNN	Test volume on cluster 2 only
VOL009	NNRNRN	NNNNNN	NNYNYN	Test volume on cluster 2 and 4 only

□ COPY AUDIT for DR site information

This request can check the copy status of the DR sites and the result can be used in case that the production site experiences a disaster and no volumes can be accessed from the production site.

Request (No COPYMODE):

COPY AUDIT INCLUDE D0002, D0003, D0004, D0005 (or COPY AUDIT EXCLUDE D0000, D0001)

Result:

VOL004,R,R,D,D,D,N,N,N,N,N,N,N,N,N,N,N

VOL005,R,R,N,N,N,N,N,N,N,N,N,N,N,N,N,N,N

VOL004 is reported because a copy to the 4 clusters in DR sites are not yet completed. VOL005 is also reported because the 4 clusters in DR sites have no copy due to 'N' copy mode but COPYMODE is NOT specified in the request.

Request (with COPYMODE):

COPY AUDIT COPYMODE INCLUDE D0002 D0003 D0004, D0005 (or COPY AUDIT COPYMODE EXCLUDE D0000, D0001)

Result:

VOL004,R,R,D,D,D,N,N,N,N,N,N,N,N,N,N

VOL005 is NOT reported because the 4 clusters in DR sites are supposed to have no copy with 'N' copy mode setting and COPYMODE is specified in the request. COPY AUDIT expects the 4 clusters in DR sites to have no valid copy because of the copy mode. Only VOL004 is reported.

- COPY AUDIT request to get the information from all the clusters except one cluster.

This request can check the copy status of all the clusters except one cluster in the domain. The result can be used in the case where one cluster will be removed from the grid as part of a cluster removal MES operation. In this example, cluster 2 is to be removed from the domain.

Request (No COPYMODE) :

COPY AUDIT INCLUDE D0000, D0001, D0003, D0004, D0005 (or COPY AUDIT EXCLUDE D0002)

Result :

VOL007,R,N,D,N,D,N,N,R,N,N,R,N,N,N,R
VOL008,N,N,R,N,N,N,N,N,N,N,N,N,N,N,N,N

VOL007 is reported because only cluster 2 has a valid copy. It is because the cluster 0 and 4 had a valid copy but their copies were already removed. VOL008 is also reported because all the clusters except cluster 2 have no copy due to 'N' copy mode but COPYMODE is NOT specified in the request.

Request (with COPYMODE) :

COPY AUDIT COPYMODE INCLUDE D0000, D0001 D0003 D0004, D0005 (or COPY AUDIT COPYMODE EXCLUDE D0002)

Result:

VOL007,R,N,D,N,D,N,N,R,N,N,R,N,N,N,R

VOL008 is NOT reported because all the clusters except cluster 2 are supposed to have no copy with 'N' copy mode setting and COPYMODE is specified in the request. COPY AUDIT expects for 5 remaining clusters (0, 1, 3, 4, 5) 2 to have no valid copy because of the copy mode. Only VOL007 is reported.

On completion of the audit, a response record is written for each logical volume that did not have a valid copy on any of the specified clusters. Volumes that have never been used, have had their associated data deleted or have been returned to scratch are not included in the response records. The record includes the volume serial number, copy policy definitions, and removal state definitions for the volume. The volser, copy policy definitions, and removal state definitions are comma separated. For example:

L00001, R, D, D, N, N, N, N, N, **N, R, R, R, N, N, N, N, R**

The response records are written in 80 byte fixed block format.

Note: The output for copy audit includes copy consistency points and removal states for up to 8 TS7700 clusters. This is to provide for future expansion of the number of clusters supported in a TS7700 Grid to the architected maximum.

Note: Copy Audit may take more than an hour to complete depending on the number of logical volumes have been defined, how many clusters are configured in the Grid configuration and how busy the TS7700s are at the time of the request.

Record 6

This record lists the distributed libraries IDs by their 5 character library sequence number separated by a comma. The distributed library IDs listed are either the ones used for a successful Copy Audit operation or ones that the TS7700 had a problem with. For example:

BA45A,BA45B

Record 7-N

If the audit does not find any logical volumes that are not valid on the specified TS7700 clusters, the following record is returned.

NO INVALID VOLUMES ON SPECIFIED CLUSTERS

Bytes	Name	Description
1-40		'NO INVALID VOLUMES ON SPECIFIED CLUSTERS'
41-80	Blanks	Blank character padding

If errors or operational conditions are encountered by the TS7700 while processing the copy audit request, the audit cannot be performed and one of the following records is returned, depending on the condition (also see the error responses listed in the section Unknown or Invalid Request:

- One or more of the clusters associated with the specified distributed library IDs are not available because they are in service or offline. Record 6 lists the distributed library IDs specified that are associated with the unavailable cluster(s).

CLUSTERS NOT IN VALID STATE

Bytes	Name	Description
1-27		'CLUSTERS NOT IN VALID STATE'
28-80	Blanks	Blank character padding

- No distributed library IDs are included in the request. At least one is required. Record 6 will be all blanks.

NO DISTRIBUTED IDS SPECIFIED

Bytes	Name	Description
1-28		'NO DISTRIBUTED IDS SPECIFIED'
29-80	Blanks	Blank character padding

- One or more of the specified distributed library IDs are not defined to the Grid configuration. Record 6 lists the distributed library IDs specified that are not defined.

DISTRIBUTED ID NOT VALID

Bytes	Name	Description
1-22		'DISTRIBUTED ID NOT VALID'
23-80	Blanks	Blank character padding

- The EXCLUDE parameter was specified and the distributed library IDs specified includes all defined in the Grid configuration. Record 6 lists the distributed library IDs specified.

ALL CLUSTERS HAVE BEEN EXCLUDED

Bytes	Name	Description
1-31		'ALL CLUSTERS HAVE BEEN EXCLUDED'
32-80	Blanks	Blank character padding

- Neither the INCLUDE or the EXCLUDE parameter was specified.

INCLUDE OR EXCLUDE MUST BE SPECIFIED

Bytes	Name	Description
1-36		'INCLUDE OR EXCLUDE MUST BE SPECIFIED'
37-80	Blanks	Blank character padding

- For copy audit to be performed, it requires that no other major database update functions are being performed. If the TS7700 Grid is in the process of inserting logical volumes, performing a copy export operation or is already executing a copy export operation.

GLOBAL RESOURCE CURRENTLY BUSY - TRY AGAIN LATER

Bytes	Name	Description
1-48		'GLOBAL RESOURCE CURRENTLY BUSY - TRY AGAIN LATER'
49-80	Blanks	Blank character padding

- An internal error in the Grid resulted in no response being generated.

INTERNAL ERROR CAUSED NO RESPONSE

Bytes	Name	Description
1-33		'INTERNAL ERROR CAUSED NO RESPONSE'
34-80	Blanks	Blank character padding

When there are logical volumes that do not have a copy on any of the specified TS7700 clusters, each of these records provide information for each volume.

Field Name	Description
VOLSER	6 character volume serial number

<p>cluster0_copy_mode</p>	<p>This field indicates whether cluster 0 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ - Synchronous copy consistency point.</p> <p>‘R’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
<p>cluster1_copy_mode</p>	<p>This field indicates whether cluster 1 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ - Synchronous copy consistency point.</p> <p>‘R’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
<p>cluster2_copy_mode</p>	<p>This field indicates whether cluster 2 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ - Synchronous copy consistency point.</p> <p>‘R’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
<p>cluster3_copy_mode</p>	<p>This field indicates whether cluster 3 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>‘S’ - Synchronous copy consistency point.</p> <p>‘R’ - Rewind unload (RUN) copy consistency point.</p> <p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p>

	<p>'N' - No copy</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
cluster4_copy_mode	<p>This field indicates whether cluster 4 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' - Synchronous copy consistency point.</p> <p>'R' - Rewind unload (RUN) copy consistency point.</p> <p>'D' - Deferred copy consistency point.</p> <p>'T' - Time Delayed copy consistency point.</p> <p>'N' - No copy</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
cluster5_copy_mode	<p>This field indicates whether cluster 5 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' - Synchronous copy consistency point.</p> <p>'R' - Rewind unload (RUN) copy consistency point.</p> <p>'D' - Deferred copy consistency point.</p> <p>'T' - Time Delayed copy consistency point.</p> <p>'N' - No copy</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
cluster6_copy_mode	<p>This field indicates whether cluster 6 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' - Synchronous copy consistency point.</p> <p>'R' - Rewind unload (RUN) copy consistency point.</p> <p>'D' - Deferred copy consistency point.</p> <p>'T' - Time Delayed copy consistency point.</p> <p>'N' - No copy</p> <p>'E' - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
cluster7_copy_mode	<p>This field indicates whether cluster 7 is to have a copy of the volume and the copy consistency point defined for the volume. The values are:</p> <p>'S' - Synchronous copy consistency point.</p> <p>'R' - Rewind unload (RUN) copy consistency point.</p>

	<p>‘D’ - Deferred copy consistency point.</p> <p>‘T’ – Time Delayed copy consistency point.</p> <p>‘N’ - No copy</p> <p>‘E’ - The volume was previously assigned a copy consistency point of synchronous, rewind unload or deferred, but was changed to no copy.</p>
cluster0_removal_state	<p>This field indicates the removal state of the volume on cluster 0.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>
cluster1_removal_state	<p>This field indicates the removal state of the volume on cluster 1.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>

<p>cluster2_removal_state</p>	<p>This field indicates the removal state of the volume on cluster 2.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>
<p>cluster3_removal_state</p>	<p>This field indicates the removal state of the volume cluster 3.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>
<p>cluster4_removal_state</p>	<p>This field indicates the removal state of the volume on cluster 4.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>

<p>cluster5_removal_state</p>	<p>This field indicates the removal state of the volume on cluster 5.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>
<p>cluster6_removal_state</p>	<p>This field indicates the removal state of the volume on cluster 6.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>
<p>cluster7_removal_state</p>	<p>This field indicates the removal state of the volume on cluster 7.</p> <p>R - This TS7720/TS7760 cluster has automatically removed its consistent copy due to the volumes configured removal policy.</p> <p>N - This cluster has not removed its copy. A TS7740 configuration will always contain this value.</p> <p>X - This TS7720/TS7760 cluster has determined that removal is not possible for this volume. This volume has been retained.</p> <p>D – This TS7720/TS7760 cluster has attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.</p> <p>P - This TS7720/TS7760 cluster will never remove this volume because it was configured as a “Pinned” volume under removal policy.</p> <p>Note: If this cluster is not included in the copy audit, the value will default to ‘N’ even though its actual value may be ‘R’.</p> <p>This is new for version 3.</p>

composite_removal_state	<p>This field indicates the overall removal state of the volume on the Grid with respect to those clusters specified as part of the copy audit consistency check.</p> <p>R - All clusters included in the consistency check have a copy mode of 'S', 'R' or 'D', are TS7720/TS7760s, and are not consistent due to the removal of data. Only TS7720/TS7760 clusters can remove volumes.</p> <p>N - At least one cluster in the consistency check is not consistent due to an invalid copy mode policy or it has not yet completed its copy.</p> <p>This is new for version 3.</p>
-------------------------	---

5.9. GGM Information

At the code level 8.33.x., Grid to Grid Migration (GGM) is supported. GGM introduces a mechanism for two independent Grid configurations to intercommunicate for the purpose of replicating content from one Grid to another and addresses many of the migration issues below:

- Older levels of hardware may have limitations supported code levels.
- More than three code levels are not supported in the same Grid.
- Grid configurations are limited to eight clusters total.
- Grid merge requires all clusters be at the same code level.

Two BVIR requests are supported to retrieve GGM copy activity records:

- 'GGM COPY RESULT' provides the response of all GGM copy activity records which already completed the GGM copies.
- 'GGM COPY STATUS' provides the response of all GGM copy activity records whose GGM copy is still in progress.

The detail of the response data format is described in IBM TS7700 Series Grid To Grid Migration User's Guide which can be found on:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

5.10. Cloud Volume to Logical Volume Mapping Information

At the code level 8.42.x.x or above, a cloud enablement feature is supported and can be applied to TS7760D (Disk Only model) (the cluster with cloud enablement feature is referred to TS7760C). This new feature enables access to a public or private cloud object store. Logical volumes can be premigrated and/or recalled to/from the cloud. The data on cloud tier are policy managed through the storage constructs assigned to the logical volume.

A logical volume is premigrated to cloud based off of the assigned storage constructs. When the logical volume is premigrated to cloud, it has a specific and unique object name. The TS7700 maintains the mapping between logical volumes and data objects on cloud. The cloud tier mapping information is synchronized among all clusters in the Grid.

Please refer to Redpaper “IBM TS7760 R4.2 Cloud Storage Tier Guide” to get more details..

TS7760C provides the database entries with logical volumes which already have consistent data on cloud and can be recalled from cloud when the access is required. The logical volumes are ordered by volser. TS7760 can provide up to 4 million database entries to a BVIR request. If the total number of cloud volume map entries exceeds 4 million, using 4th keyword can provide the next 4 million entries. For example, when 4th keyword ‘0’ is specified (this is equivalent to no 4th keyword), the first 4 million entries (ordered by volser) are provided in the response. Then, if 4th keyword ‘1’ is specified, the next 4 million entries are provided in the response and so on.

Note: The generation of the response may take several minutes to complete depending on the number of active logical volumes on cloud and how busy the TS7700 cluster is at the time of the request.

Record 6

If BVIR response is retrieved from TS7700 which the cloud enablement feature is disabled (i.e. not TS7760C), the following record is returned.

Bytes	Name	Description
1-66		‘NOT SUPPORTED IN A NON-CLOUD-ATTACHED TS7700 VIRTUALIZATION ENGINE’
67-80	Blanks	Blank character padding

If BVIR response is retrieved from TS7760C, this record provides a heading for the data records to follow:

```
VOLSER INSERT_VERSION          DATA_LEVEL          SYSPLEX_NAME SYSTEM_NAME PROGRAM_NAME
FILE_SIZE          CLOUD_ACCOUNT_ID    CLOUD_ACCOUNT_NICKNAME ENCRYPTION_FLAG
ENCRYPTION_KEY_LABEL1          ENCRYPTION_KEY_LABEL2
CLOUD_DATA_FORMAT DELETE_REASON DELETED_TIME          PREMIGRATED_TIME
INITIATOR CLOUD_POOL          CLOUD_POOL_NAME CLOUD_PREMIG_RANK COMPOSITE_STATE
CLUSTER0_STATE COMP_CLU_MASK_REQS_UPD OBJECT_NAME
CONTAINER_NAME
```

Bytes	Name	Description
1-6		‘VOLSER’ Heading for the logical volume serial number
7	Field Delimiter	Blank character
8-21		‘INSERT_VERSION’ Heading for the insert version of the logical volume when it’s inserted into the Grid then premigrated to cloud
22-28	Field Delimiter	Blank characters
29-38		‘DATA_LEVEL’ Heading for the data level of the logical volume when it’s premigrated to cloud

39-49	Field Delimiter	Blank characters
50-61		'SYSPLEX_NAME' Heading for the customer defined host sysplex name for the host partition which issued the mount to the volume when it's premigrated to cloud
62	Field Delimiter	Blank character
63-73		'SYSTEM_NAME' Heading for the customer defined host system name for the host partition which issued the mount to the volume when it's premigrated to cloud
74	Field Delimiter	Blank character
75-86		'PROGRAM_NAME' Heading for the program name associated with the application that issued the mount when it's premigrated to cloud
87	Field Delimiter	Blank character
88-96		'FILE_SIZE' Heading for the size of the logical volume in bytes premigrated to cloud
97-108	Field Delimiter	Blank characters
109-124		'CLOUD_ACCOUNT_ID' Heading for the cloud account ID generated by TS7700 internally corresponding to the cloud account defined by the customer
125-128	Field Delimiter	Blank characters
129-150		'CLOUD_ACCOUNT_NICKNAME' Heading for the cloud account ID nickname defined by the customer
151	Field Delimiter	Blank character
152-166		'ENCRYPTION_FLAG' Heading for the cloud data encryption type. This is not used in 8.42.x.x and always 0.
167	Field Delimiter	Blank character
168-188		'ENCRYPTION_KEY_LABEL1' Heading for the encryption key label 1 when data on cloud is encrypted. This is not used in 8.42.x.x and always 'NULL'.
189-232	Field Delimiter	Blank characters
233-253		'ENCRYPTION_KEY_LABEL2' Heading for the encryption key label 2 when data on cloud is encrypted. This is not used in 8.42.x.x and always 'NULL'.
254-297	Field Delimiter	Blank characters
298-314		'CLOUD_DATA_FORMAT' Heading for the data format on cloud. This is not used in 8.42.x.x and always '1'.
315	Field Delimiter	Blank character
316-328		'DELETE_REASON' Heading for the delete reason when data on cloud is deleted. This is always '0'.
329	Field Delimiter	Blank character
330-341		'DELETED_TIME' Heading for the delete timestamp when data on cloud is deleted. This is always '1970-01-01-00.00.00.000000'.
342-356	Field Delimiter	Blank characters
357-372		'PREMIGRATED_TIME' Heading for the premigration timestamp when the logical volume was premigrated to cloud.
373-383	Field Delimiter	Blank characters
384-392		'INITIATOR' Heading for the cluster ID which premigrated the logical volume to cloud.
393	Field Delimiter	Blank character
394-403		'CLOUD_POOL' Heading for the cloud pool ID generated by TS7700 internally corresponding to the cloud pool defined by the customer where the logical volume is premigrated to
404-413	Field Delimiter	Blank characters
414-428		'CLOUD_POOL_NAME' Heading for the cloud pool nickname defined by the customer where the logical volume is premigrated to
429	Field Delimiter	Blank character
430-446		'CLOUD_PREMIG_RANK' Heading for the rank of the cloud pool defined in the storage construct (storage group) bound for the logical volume. This is always '1' at

		8.42.x.x.
447	Field Delimiter	Blank character
448-462		'COMPOSITE_STATE' Heading for the composite library view of data state on cloud for the logical volume. This is always '5376' (premigrated) at 8.42.x.x.
463	Field Delimiter	Blank character
464-477		'CLUSTERX_STATE' Heading for the distributed library (local cluster) view of data state on cloud for the logical volume. X is a cluster ID which returns this BVIR response.
478	Field Delimiter	Blank character
479-500		'COMP_CLU_MASK_REQS_UPD' Heading for the cluster bit mask which indicates which clusters have the down level entry of this cloud volume map entry.
501	Field Delimiter	Blank character
502-512		'OBJECT_NAME' Heading for the object name of this logical volume on cloud.
513-630	Field Delimiters	Blank characters
631-644		'CONTAINER_NAME' Heading for the container name defined by the customer where the logical volume is contained on cloud
645-900	Blanks	Blank characters

Record 7-N

If no cloud volume map entries exist, the following record is returned.

Bytes	Name	Description
1-32		'NO LVOL_TO_CLOUD RECORD(S) EXIST'
33-900	Blanks	Blank character padding

If any cloud map entries exist, a response record consists of the database fields defined in the following table. Each of these records provides information for a logical volume to cloud object association and the records are ordered alphanumerically (0-9, A-Z) by logical volser. Fields are space () separated and the overall length of each record is 900 bytes with blank padding after the last field as needed. For example, the first few fields of the record returned for volser Z99999 would be:

```

99999 1488                102
830751027                EFBBV20180821014456 lipcosa          0          NULL
NULL                    1          0
1970-01-01-00.00.000000 2018-09-04-09.42.00.000000 0          EFBBV20180821014421
lipcos                   1          5376          5376          0
lipizzan/5e22/BA092/BA92A/EFBBV/Z99999/1488/102/20180904094145
lipcsc
...
    
```

Bytes	Name	Description
1-6	VOLSER	Logical volume serial number with 6 bytes characters
7	Field Delimiter	Blank character
8-27	INSERT_VERSION	The insert version of the logical volume when it's inserted into the Grid then premigrated to cloud, left justified and padded with blanks.
28	Field Delimiter	Blank character
29-48	DATA_LEVEL	The data level of the logical volume when it's premigrated to cloud, left justified and padded with blanks.
49	Field Delimiter	Blank character
50-61	SYSPLEX_NAME	The customer defined host sysplex name for the host partition which issued the mount to the volume when it's premigrated to cloud, left justified and padded with blanks.

62	Field Delimiter	Blank character
63-73	SYSTEM_NAME	The customer defined host system name for the host partition which issued the mount to the volume when it's premigrated to cloud, left justified and padded with blanks.
74	Field Delimiter	Blank character
75-86	PROGRAM_NAME	The program name associated with the application that issued the mount when it's premigrated to cloud, left justified and padded with blanks.
87	Field Delimiter	Blank character
88-107	FILE_SIZE	The size of the logical volume in bytes premigrated to cloud, left justified and padded with blanks.
108	Field Delimiter	Blank character
109-127	CLOUD_ACCOUNT_ID	The cloud account ID generated by TS7700 internally corresponding to the cloud account defined by the customer with 19 bytes characters.
128	Field Delimiter	Blank character
129-150	CLOUD_ACCOUNT_NICKNAME	The cloud account ID nickname defined by the customer, left justified and padded with blanks.
151	Field Delimiter	Blank character
152-166	ENCRYPTION_FLAG	The cloud data encryption type. This is not used in 8.42.x.x and always 0.
167	Field Delimiter	Blank character
168-231	ENCRYPTION_KEY_LABEL1	The encryption key label 1 when data on cloud is encrypted. This is not used in 8.42.x.x and always 'NULL', left justified and padded with blanks.
232	Field Delimiter	Blank character
233-296	ENCRYPTION_KEY_LABEL2	The encryption key label 2 when data on cloud is encrypted. This is not used in 8.42.x.x and always 'NULL', left justified and padded with blanks.
297	Field Delimiter	Blank character
298-314	CLOUD_DATA_FORMAT	The data format on cloud. This is not used in 8.42.x.x and always '1', left justified and padded with blanks.
315	Field Delimiter	Blank character
316-328	DELETE_REASON	The delete reason when data on cloud is deleted. This is always '0', left justified and padded with blanks.
329	Field Delimiter	Blank character
330-355	DELETED_TIME	The delete timestamp when data on cloud is deleted. This is always '1970-01-01-00.00.00.000000'.
356	Field Delimiter	Blank character
357-382	PREMIGRATED_TIME	The premigration timestamp when the logical volume was premigrated to cloud.
383	Field Delimiter	Blank character
384-392	INITIATOR	The cluster ID which premigrated the logical volume to cloud, left justified and padded with blanks.
393	Field Delimiter	Blank character
394-412	CLOUD_POOL	The cloud pool ID generated by TS7700 internally corresponding to the cloud pool defined by the customer where the logical volume is premigrated to, left justified and padded with blanks.
413	Field Delimiter	Blank characters
414-428	CLOUD_POOL_NAME	The cloud pool nickname defined by the customer where the logical volume is premigrated to, left justified and padded with blanks.
429	Field Delimiter	Blank character
430-446	CLOUD_PREMIG_RANK	The rank of the cloud pool defined in the storage construct (storage group) bound for the logical volume, left justified and padded with blanks. This is always '1' at 8.42.x.x.
447	Field Delimiter	Blank character
448-462	COMPOSITE_STATE	The composite library view of data state on cloud for the logical volume, left justified and padded with blanks. This is always '5376' (this means "data is premigrated to cloud") at 8.42.x.x.

463	Field Delimiter	Blank character
464-477	CLUSTERX_STATE	<p>The distributed library (local cluster) view of data state on cloud for the logical volume, left justified and padded with blanks. X is a cluster ID which returns this BVIR response. The following value could be provided:</p> <p>0 – The local cluster does not premigrated the logical volume to cloud yet, or it’s not configured to premigrate the logical volume to cloud.</p> <p>4352 – The local cluster failed to premigrate the logical volume to cloud. Retry will be attempted automatically.</p> <p>4864 – The local cluster is premigrating the logical volume to cloud.</p> <p>5367 – The local cluster has premigrated the logical volume to cloud, or the local cluster can access the data on cloud when a peer cluster in the Grid has premigrated the data to cloud already.</p> <p>12544 – The data on cloud needs to be deleted.</p> <p>13056 – The local cluster failed to delete the data on cloud. Retry will be attempted automatically.</p> <p>13568 – The data on cloud is being deleted by the local cluster.</p> <p>14080 – The deletion of the data on cloud has completed.</p> <p>20736 – The deletion of the data on cloud has completed and this cloud volume map entry is ready to be deleted completely.</p>
478	Field Delimiter	Blank character
479-500	COMP_CLU_MASK_RE QS_UPD	<p>The cluster bit mask which indicates which clusters have the down level entry of this cloud volume map entry, left justified and padded with blanks.</p> <p>If it is 0, the synchronization of this cloud volume map entry is done in the Grid. If it is not 0, the corresponding cluster still has down level entry and the synchronization will be done automatically.</p>
501	Field Delimiter	Blank character
502-629	OBJECT_NAME	<p>The object name of this logical volume on cloud, left justified and padded with blanks.</p> <p>The format of the object name is: "Customer Provided Object Pre-Fix Name (defined in storage group bound for this logical volume)"/"Composite Library Sequence Number"/"Distributed Library Sequence Number"/"Logical Volume Serial Number"/Insert Version/Data Level/"Timestamp when this entry is created"</p>
630	Field Delimiter	Blank character
631-886	CONTAINER_NAME	The container name defined by the customer where the logical volume is contained on cloud, left justified and padded with blanks.
887	Field Delimiter	Blank character
888-900	Blanks	Blank characters

5.11. Unknown or Invalid Request

If the request file does not contain the correct number of records or the first record is incorrect, the request file on the volume is unchanged and no error is indicated.

If the request file contains the correct number of records and the first record is correct but the second is not, the response file will indicate that the request is unknown as follows:

Record 6

This record indicates that the request is unknown.

UNKNOWN REQUEST TYPE

Bytes	Name	Description
1-20		'UNKNOWN REQUEST TYPE'
21-80	Blanks	Blank character padding

If the request file contains the correct number of records, the first record is correct, the second is recognized but includes a variable that is not within the range supported for the request, the response file will indicate that the request is invalid as follows:

Record 6

This record indicates that the request is invalid.

INVALID VARIABLE SPECIFIED

Bytes	Name	Description
1-26		'INVALID VARIABLE SPECIFIED'
27-80	Blanks	Blank character padding

6. Example JCL Statements

The following are fragments of JCL that show how to use a standard IBM utility, IEBGENER, to request and process the requested data for the BVIR function.

The following sample JCL obtains a scratch volume to perform the request:

```
//VTSQUERY JOB ...
//*****
//* DO NOT USE COMPACTION WHEN WRITING THE REQUEST FILE
//*****
//* SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//* DATA SET IS CATALOGED
//*****
//* USING A LOGICAL SCRATCH VOLUME, CREATE THE REQUEST FILE WITH
//* THE 2 REQUIRED RECORDS. IN ORDER TO ENSURE THAT A SCRATCH VOLUME
//* IS ALLOCATED IN THE TARGET LIBRARY FOR THE QUERY OPERATION,
//* THE ACS ROUTINES NEED TO HAVE LOGIC TO ALLOCATE A TAPE
//* DRIVE IN THE TARGET LIBRARY. ONE WAY TO ACCOMPLISH THIS IS TO
//* HAVE A STORAGE GROUP UNIQUE TO EACH VTS LIBRARY PROVIDING A
//* 1 TO 1 RELATIONSHIP BETWEEN STORAGE GROUP AND LIBRARY.
//* THE ACS ROUTINES WOULD THEN NEED TO KEY OFF OF SOMETHING
//* UNIQUE IN THE DD STATEMENT (DATA SET NAME, DATA CLASS SPECIFICATION,
//* UNIT SPECIFICATION, ETC ...) TO GET THE CORRECT STORAGE GROUP
//* AND THE RIGHT TARGET LIBRARY SELECTED.
//*****
//* FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//* RECORDS MUST BE SPECIFIED AS ILLUSTRATED BELOW, STARTING IN
//* THE FIRST COLUMN:
//* SPECIFY THE SEQUENCE NUMBER OF THE TS7700 THE REQUEST IS TO GO TO AS A
//* CHECK THAT IT IS GOING TO THE CORRECT TS7700
//* //*****
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSIN DD DUMMY
//SYSUT2 DD DSN=SYSBADM.CQUERY,
// UNIT=3490,LABEL=(,SL),
// DISP=(NEW,CATLG),
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
CACHE CONTENTS
/*
...
```

The following sample JCL uses a specific volume to perform the request:

```
//VTSQUERY JOB ...
//*****
//* DO NOT USE COMPACTION WHEN WRITING THE REQUEST FILE
//*****
//* SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//* DATA SET IS CATALOGED
//*****
//* USING A LOGICAL SPECIFIC VOLUME, CREATE THE REQUEST FILE WITH
//* THE 2 REQUIRED RECORDS.
//*****
```

```

/** FILE SEQUENCE 1: REQUEST/RESPONSE FILE
/** RECORDS MUST BE SPECIFIED AS ILLUSTRATED BELOW, STARTING IN
/** THE FIRST COLUMN:
/** SPECIFY THE SEQUENCE NUMBER OF THE TS7700 THE REQUEST IS TO GO TO AS A
/** CHECK THAT IT IS GOING TO THE CORRECT TS7700

```

```

/** //*****
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=SYSBADM.CQUERY,
// UNIT=3490,LABEL=(,SL),
// DISP=(NEW,CATLG),VOL=SER=VTA001,
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
CACHE CONTENTS
/**
...

```

For the request types that return 80 byte fixed block records (cache contents, volume map, physical media pools and copy audit), the following sample JCL reads the response data and sends it to a printer:

```

//VTSRESP JOB ...
//*****
/** THE RESPONSE DATA IS NOT COMPACTED
//*****
/** SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//*****
/** USING THE DATA SET CATALOGED IN THE REQUEST JOB
//*****
/** FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//*****
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSIN DD DUMMY
//SYSUT1 DD DSN=SYSBADM.CQUERY,DISP=OLD

//SYSUT2 DD SYSOUT=A,
// DCB=(DSORG=PS,RECFM=FB,LRECL=80,BLKSIZE=0)
...

```

For the request types that return variable length unformatted records (point in time statistics and historical statistics), the following JCL reads the response data to a DASD data set:

```

//VTSRESP JOB ...
//*****
/** THE RESPONSE DATA IS NOT COMPACTED
//*****
/** SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//*****
/** USING THE DATA SET CATALOGED IN THE REQUEST JOB
//*****
/** FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//*****
//STEP1 EXEC PGM=IEBGENER

```

```
//SYSPRINT DD SYSOUT=A
//SYSIN DD DUMMY
//SYSUT1 DD DSN=SYSBADM.PITSTAT,DISP=OLD
// DCB=(RECFM=U,BLKSIZE=24000)

//SYSUT2 DD DSN=SYSBADM.PITSTAT.CAPTURE,
// UNIT=SYSDA,SPACE=(CYL,(2,1)),
// DISP=(NEW,CATLG),
// DCB=(RECFM=U,BLKSIZE=24000)
...

```

The following sample JCL for JES2 (this will not work for JES3 because it will not demount/mount the volume between steps) combines the request and read steps into a single job with REF=* and includes the response data in the job output:

```
//BVIRINFO JOB ...
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT2 DD DSN=TAPE.BVIR.LIB5MPLS,
// UNIT=B63M2N36,LABEL=(,SL),
// DISP=(NEW,KEEP),
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
PHYSICAL MEDIA POOLS
/*
//SYSIN DD DUMMY
//*
//STEP2 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TAPE.BVIR.LIB5MPLS,
// VOLUME=(,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2,
// DISP=(OLD),LABEL=(1,SL)
//SYSUT2 DD SYSOUT=A,
// DCB=(DSORG=PS,RECFM=FB,LRECL=80,BLKSIZE=0)
//SYSIN DD DUMMY
/*

```

The following sample JCL for JES2 (this will not work for JES3 because it will not demount/mount the volume between steps) combines the request and read steps into a single job where the request tape is cataloged and the response is written to an output file on DASD:

```
//BVHIS29 JOB . . .
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT2 DD DSN=TAPE.BVIR.BAR29HIS,
// UNIT=B29M2C36,LABEL=(,SL),
// DISP=(NEW,KEEP),
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
HISTORICAL STATISTICS FOR 250-260
/*
//SYSIN DD DUMMY
//*
//STEP2 EXEC PGM=IEBGENER

```

```
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TAPE.BVIR.BAR29HIS,
//      VOLUME=(, , REF=* .STEP2.SYSUT2) ,
//      DCB=(RECFM=U, BLKSIZE=24000) ,
//      DISP=(OLD) , LABEL=(1, SL)
//SYSUT2 DD DSN=TAPE.BVIR.BAR29HIS.OUT,
//      DISP=(NEW, CATLG) , SPACE=(CYL, (5, 5)) ,
//      UNIT=3390,
//      DCB=(RECFM=U, BLKSIZE=24000)
//SYSIN DD DUMMY
/*
```

Note that in the above sample JCL, the blocksize for the read step is different than the request records blocksize. This is because for the data records returned for statistical data are variable in length. In the following sample JCL, the blocksizes are different, but are still fixed in length:

```
//BVHIS29 JOB . . .
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT2 DD DSN=TAPE.BVIR.BAR29PV,
//      UNIT=B29M2C36, LABEL=(, SL) ,
//      DISP=(NEW, KEEP) ,
//      DCB=(RECFM=F, BLKSIZE=80, LRECL=80, TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
PHYSICAL VOLUME STATUS VOLUME P1_____
/*
//SYSIN DD DUMMY
/*
//STEP2 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TAPE.BVIR.BAR29PV,
//      VOLUME=(, , REF=* .STEP2.SYSUT2) ,
//      DCB=(RECFM=F, BLKSIZE=400, LRECL=400) ,
//      DISP=(OLD) , LABEL=(1, SL)
//SYSUT2 DD DSN=TAPE.BVIR.BAR29PV.OUT,
//      DISP=(NEW, CATLG) , SPACE=(CYL, (5, 5)) ,
//      UNIT=3390,
//      DCB=(RECFM=FB, BLKSIZE=0, LRECL=400)
//SYSIN DD DUMMY
/*
```

The following sample JCL for JES3 separates the create and read steps into two separate jobs:

```
JOB1:
//JS010 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
VOLUME MAP
/*
//SYSUT2 DD DSN=OUTPUT.DATASET.NAME,
//      DISP=(NEW, CATLG, DELETE) ,
//      UNIT=CTAPE,
```



```
//          RETPD=14,
//          DCB= (LRECL=80, BLKSIZE=80, TRTCH=NOCOMP)
//*
JOB2:
//JS020   EXEC PGM=IEBGENER
//SYSPRINT DD  SYSOUT=*
//SYSIN   DD  DUMMY
//SYSUT1  DD  DSN=OUTPUT.DATASET.NAME,
//          DISP=OLD
//SYSUT2  DD  SYSOUT=U, LRECL=80, RECFM=FB, BLKSIZE=0
```

The following sample JCL reads the bvirmes.cntl.txt records and writes them to a virtual volume. That volume is kept. Opened again and copied to disk. When this last open is issued, the TS7740 pulls the status of the requested lvol(s) and makes those records appear as if they had come from the tape so they can be copied to the disk file.:

```
//*JOB1
//*JOB2
//*JOB3
//*JOB4
/*JOBPARM SYSAFF=*
/*
/* TO INSURE THAT BVIR REQUESTS ARE SATISFIED FROM THE PROPER
/* CLUSTER, YOU SHOULD HAVE MANAGEMENT CLASSES FOR RNN(CL0), NRN(CL1)
/* AND NNR(CL2) SO THAT ONLY THE TARGET CLUSTER WILL HAVE A COPY OF
/* THE BVIR VOLUME. YOU CAN'T CONTROL WHERE THE INITIAL SCRATCH
/* MOUNT IS SATISFIED, BUT YOU CAN CONTROL WHICH TVC THE VOLUME WILL
/* BE IN WHEN THE SUBSEQUENT SPECIFIC MOUNT IS DONE. THE SPECIFIC
/* MOUNT COLLECTS THE BVIR INFORMATION, NOT THE INTIIAL SCRATCH MOUNT.
/*
/* IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/* STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO ALLOCATE ON THE
/* DESIRED GRID. USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/*
/* IF YOU ARE RUNNING JES3, YOU MUST RUN STEP3 AS A SEPARATE JOB
/* IN ORDER TO FORCE THE DISMOUNT OF THE TAPE IN STEP2. BVIR DATA
/* WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL DISMOUNT AND
/* RE-MOUNT FOR READ.
/*
/* THIS JOB ISSUES THE BULK VOLUME INFORMATION REQUEST (BVIR) FOR
/* ALL VIRTUAL VOLUMES BELONGING TO THE VTS ASSOCIATED WITH THE
/* VIRTUAL DRIVE ADDRESS USED. THE BVIR FEATURE MUST BE ACTIVATED
/* ON THE VTS RECEIVING THE REQUEST. THIS IS FOR TS7740 ONLY.
/* IF YOU ARE RUNNING AGAINST A PTP AND GETTING DATA FOR THE PTPSYNC
/* JOB, YOU NEED TO RUN THIS JOB TWICE, ONCE FOR EACH VTS.
/*
//BVIRMES PROC USERHLQ=USERID,          HI-LEVEL FOR USER DATA FILES
//          TOOLHLQ=TOOLID,           HLQ FOR LOAD AND CNTL
//          SITE=SITENAME,           2ND LEVEL QUALIFIER
//          MC=,                      DIRECT TO SPECIFIC VTS OR CLUSTER
//          VTSID=,                   VTS SERIAL NUMBER TO BE PART OF DSN
//          UNIT=VTAPE                UNITNAME ON THIS VTS
/*
//STEP1   EXEC PGM=IEFBR14
//DEL1    DD  UNIT=(&UNIT, , DEFER) , DISP=(MOD, DELETE) ,
//          DSN=&USERHLQ..&SITE..#&VTSID..BVIRMES
```

```
//DEL2      DD  UNIT=SYSDA, DISP=(MOD,DELETE), SPACE=(TRK,1),
//          DSN=&USERHLQ..&SITE..#&VTSID..BVIRDATA.MESFILE
//*
//STEP2     EXEC PGM=IEBGENER
//SYSPRINT  DD  SYSOUT=*
//SYSUT1    DD  DISP=SHR, DSN=&TOOLHLQ..IBMTOOLS.CNTL(BVIRMES)
//SYSUT2    DD  DSN=&USERHLQ..&SITE..#&VTSID..BVIRMES, MGMTCLAS=&MC,
//          UNIT=&UNIT, LABEL=(,SL), DISP=(NEW,CATLG),
//          DCB=(RECFM=FB, BLKSIZE=80, LRECL=80, TRTCH=NOCOMP)
//SYSIN     DD  DUMMY
//*
//STEP3     EXEC PGM=IEBGENER
//SYSPRINT  DD  SYSOUT=*
//SYSUT1    DD  DSN=&USERHLQ..&SITE..#&VTSID..BVIRMES, DISP=(OLD,DELETE),
//          DCB=(DSORG=PS, RECFM=F, BLKSIZE=640, LRECL=640)
//SYSUT2    DD  DSN=&USERHLQ..&SITE..#&VTSID..BVIRDATA.MESFILE,
//          UNIT=SYSDA, SPACE=(640, (500000, 200000), RLSE),
//          DISP=(,CATLG), DCB=(DSORG=PS, RECFM=FB, BLKSIZE=0, LRECL=640)
//SYSIN     DD  DUMMY
//      PEND
//*
//GETVOLS   EXEC BVIRMES, VTSID=VTS0#, MC=
//*
```

The following is the data in the text file read from the above sample:

```
VTS BULK VOLUME DATA REQUEST
VOLUME STATUS _____
```

References

DFSMS: Using Magnetic Tapes, SC26-7412-01 (SC23-6858-nn in z/OS 2.1 or higher)

DFSMSdfp Utilities, SC26-7414-02 (SC23-6864-nn in z/OS 2.1 or higher)

Disclaimers:

© Copyright 2006, 2018 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.