September 04, 2019

IBM® TS7700 Series
DS8000 Object Store User's Guide Version 1.0

Lourie Goodall
lourie.bryan@ibm.com
IBM TS7700 Architect

Table of Contents

September 4, 2019

Change History

Version 1.0 – Original Version (TS7700 R5.0.x)

.

1. Introduction

Today, a z/OS host works with many different applications to move data out of primary storage to free up space. One predominant application is DFSMShsm. The DS8000 disk storage system is one well-known primary storage system.

After writing to its primary disk cache system, the host may choose to migrate or back up the data to a long-term storage system. This storage system could be a physical tape library or a virtual tape library such as the TS7700. Backing up to tape or virtual tape has cost benefits and is extremely efficient as a long term storage strategy.

Since its origin, the TS7700 has always presented itself as an image of one or more virtual tape libraries, containing a range of virtual tape volumes and virtual drives, to a z/OS host. Virtual volumes are files contained in a VTS disk file system. A virtual volume consists of customer data, header data and metadata, which represents the entire image of a real tape volume. The customer data, header data and metadata in this volume consist of hundreds to hundreds of thousands of objects.

There is a cost for z/OS to coordinate and manage backend storage as it can require a large amount of the mainframe's processing power to package, move and then manage the data continuously. This resource utilization is measured in MIPS (Millions of Instructions Per Second).

In 2017, z/OS DFSMS and DS8000 introduced transparent cloud tiering. By having DFSMS and DS8000 transparent cloud tiering manage a portion of the backend storage workload, z/OS can reduce the MIPS it uses for backend storage and reserve it for more important tasks. The targeted data for the transparent cloud tiering offering is DFSMShsm data.

There is also a large amount of complexity in the processing of back end tape storage. A cloud solution allows for a simpler approach (Figure 1):

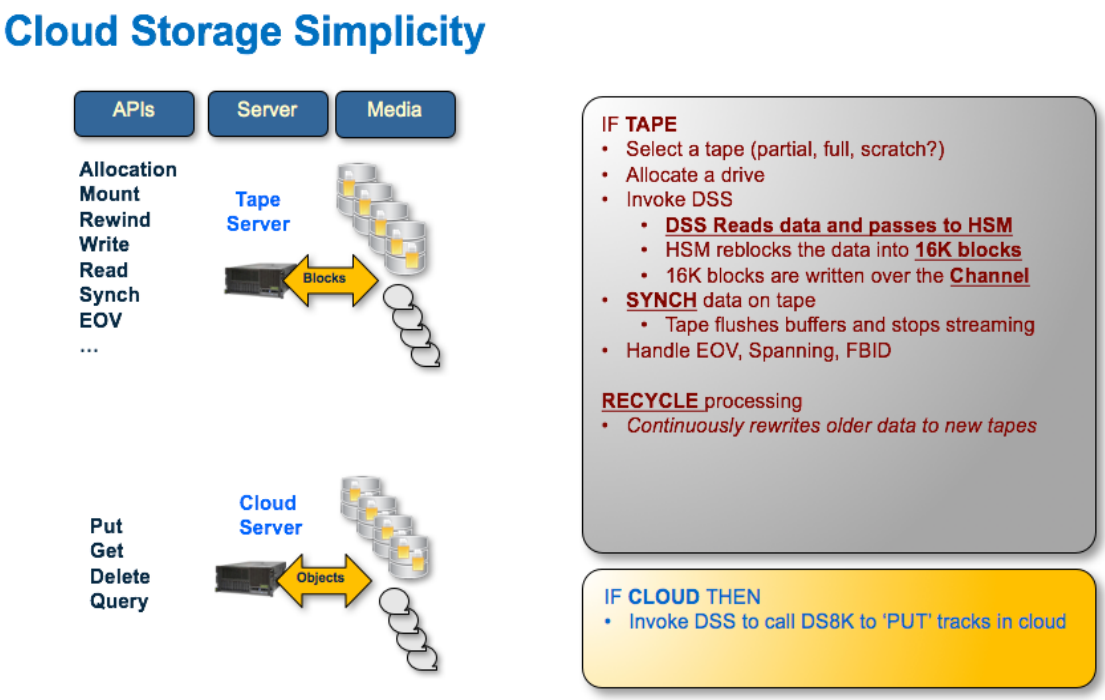

Figure 1 Cloud Storage Simplicity

Recalling that data using tape services is equally as complex and requires the data to move through multiple steps(Figure 2).

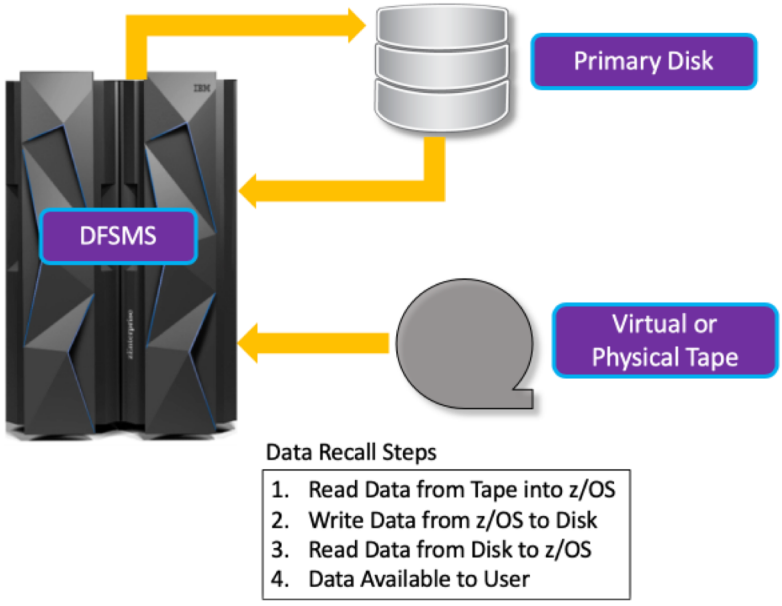

Figure 2 Tape Recall Steps

The process of migrating and recalling data through transparent cloud tiering has been significantly simplified. For instance, recalling data is reduced to a single step (Figure 3).

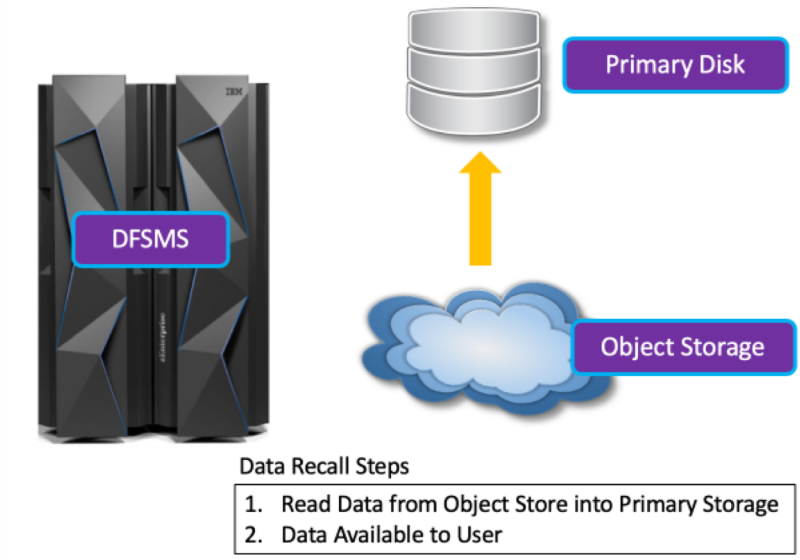

Figure 3 Cloud Recall Steps

The following lists some of the main HSM constraints or resource drains that can be eliminated by using an object store off the DS8000:

- Consumes a large amount of CPU resources
- Data passes through multiple layers of software
- Data is not able to be co-located
- Does not allow parallel access to data
- Continuous recycles
- 16K block size limit
- Affects <control data set> REORGs, Audits and backups

z/OS DFSMS and DS8000 transparent cloud tiering features enable the ability to move HSM data through an IBM System Storage DS8000 to a cloud storage device. That cloud storage device can be an IBM TS7700 Virtualization Engine through a TS7700 offering called DS8000 Object Store.

The concept of the TS7700 as a virtual tape server has not changed. However, beginning with microcode level 8.50.x.x, the TS7700 will also be able to store customer data, header data and meta data, in the form of objects, within the same cache as virtual tape volumes.

This solution not only resolves the constraints listed above but also allows TS7700 users to employ existing hardware to store the same HSM data as objects. This document will describe these DS8000 Object Store functions in more detail:

- Co-existence of logical volume data and object data within the same TS7700 cache system (no additional hardware needed)
- Logical cache partitions for object data

- Replication of object data via forking mechanism to up to two TS7700 clusters in the same grid.
- Error recovery mechanism for forked setups.
- Display cache capacity and cache utilization on the TS7700 Management Interface for object data
- Display Historical Statistics for object partitions on the TS7700 Management Interface
- Library Request Command to show object cache partition usage.
- Library Request Command to show throughput data for object workload
- Library Request Command to set object cache partition alerts
- DS8000 Object Store enablement and setup information.
- Restrictions and limitations

It is assumed throughout this white paper that the reader is familiar with the TS7700 and IBM Z host systems.

2. TS7700 Configuration Supported

The DS8000 Object Store function is supported in a stand-alone or Grid configuration up to an 8-way. DS8000 Object Store can target one or two TS7700s in the same grid for a DS8000 system. See the Grid Replication section for more information.

3. TS7700 Code Requirements

The DS8000 Object Store function was introduced with the TS7700 microcode level release 8.50.x.x. Target clusters must be at this microcode level or higher to enable DS8000 Object Store. Other clusters can reside in the grid at a lower microcode level but cannot be a DS8000 object target.

Note that microcode level release 8.50.x.x is only supported on TS7760(VEC), TS7770(VED) and future hardware models.

4. Host Considerations

In the z/OS transparent cloud tiering solution, DFSMS writes object data to the cloud object store via an IBM System Storage DS8000. To begin, z/OS DFSMS must be configured. This document will provide general information about object flow from z/OS to the TS7700 and how to configure z/OS for transparent cloud tiering.

For more detailed information about z/OS transparent cloud tiering, refer the following documents:

http://publibz.boulder.ibm.com/zoslib/pdf/OA51622.pdf

https://www.redbooks.ibm.com/redbooks/pdfs/sg248381.pdf

https://www.ibm.com/support/knowledgecenter/en/SSLTBW_2.3.0/com.ibm.zos.v2r3.idak100/cloud23.htm

The z/OS host must be at V2R2 with a minimum OA51622 PTF and REQs or V2R3 base level and higher.

5. DS8000 Considerations

The IBM DS8000 requires microcode level V8.5SP4 or higher with model DS8880 to support the TS7700 DS8000 Object Store feature. A zSynergy License Bundle is required to be enabled on the DS8880.

The DS8000 also has configuration steps. Details about DS8000 transparent cloud tiering and configuration requirements can be found at the following links:

www.redbooks.ibm.com › Cloud

https://www.ibm.com/support/knowledgecenter/ST5GLJ_8.5.0/com.ibm.storage.ssic.help.doc/f2c_configuring_trans_cloud_tiering.html

DS8000 currently only supports a single cloud account to a cloud storage device. However, multiple DS8000s can target a single TS7700 cloud device.

6. General Configuration Information

This section provides a brief summary of configuration information needed to migrate data from z/OS to a TS7700 object store via a DS8000 using transparent cloud tiering functions. It is a prelude to understanding the data flow for this feature and is not meant to include all requirements, configuration steps and limitations. Please review the z/OS and DS8000 publications linked above for detailed information about configuration and other transparent cloud tiering concepts.

- **z/OS**

As part of the z/OS transparent cloud tiering offering, a new cloud network connection construct was introduced. Cloud network connection construct can be defined in the ISMF Cloud panel (option 'S') in the SMS source configuration data set. Additionally, a new migration level has been added, Migration Level Cloud (MLC).

Unlike ML1 and ML2, data does not go through those tiers before it lands on MLC. Data goes directly from L0 to MLC.

The SMS Cloud network connection construct contains a list of cloud attributes and values that are used by DFSMShsm and DFSMSdss commands to manage data set migration. DFSMShsm's MIGRATE and HMIGRATE commands now include the keyword, CLOUD (*cloud_network_connection_name*). When this keyword is specified, the data sets are migrated to the requested cloud. The *cloud_network_connection_name* must match an existing SMS Cloud definition.

In Management Class, a new "Level 2 Days Non-usage" attribute is used for direct migration to the cloud storage for a *cloud_network_connection_name*. It currently supports two values, 0 and NOLIMIT. When NOLIMIT is specified, the data set will not migrate to cloud storage and will migrate to Level 2 (ML2) based on the value of "Level 1 Days Non-usage". When 0 is specified, the dataset will migrate to cloud storage (MLC) when the "Primary Days Non-usage" value is met as long as the data still resides on Level 0. "Level 2 Days Non-usage" takes priority over "Level 1 Days Non-usage". DFSMShsm does not support the movement of migration copies from ML1 or ML2 to MLC.

New migration actions can also be made based on the size of the data set (in tracks) as long as the "Primary Days Non-usage" criteria is met. "Size Less Than or Equal To" and "Size Greater Than" are optional fields that specify a size threshold. If the current size of the data set meets the threshold for "Size Less Than or Equal To" or "Size Greater Than" fields, then the specified action is taken. For example, if the "Primary Days Non-usage" criteria is met and the "Size Greater Than" field is set to 200 tracks, a data set of 250 tracks will trigger the ACTION specified. Actions include the ability to migrate to ML1, ML2 and now MLC.

Studies have shown that the larger the data, the more MIPS savings. The new migration options are provided as an additional way to tailor migrations to fit the needs of each user. Figure 4 shows the number of MIPS required for data movement by size during the z/OS DFSMS backup and archival process:

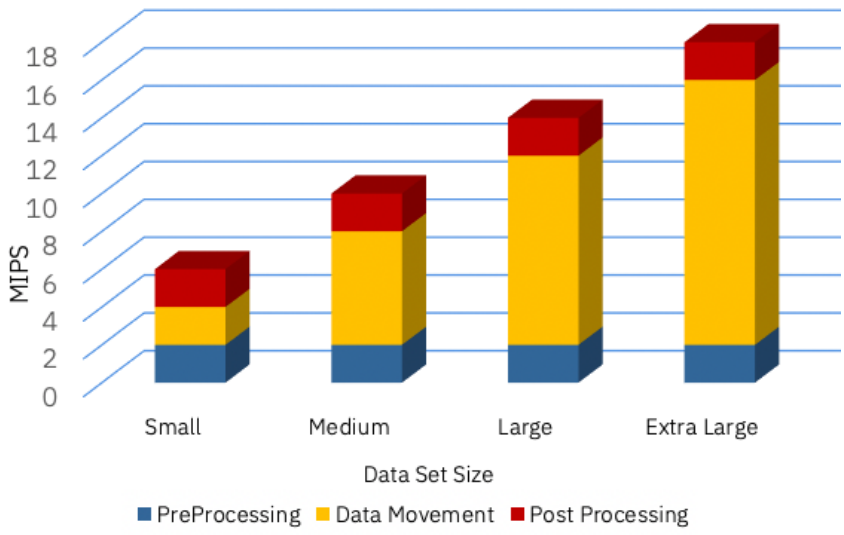

Figure 4 Number of MIPS by Data Set Size Using ML2

When using the z/OS with DS8000 transparent cloud tiering solution, the number of MIPS in the z/OS remains low and constant (Figure 5).

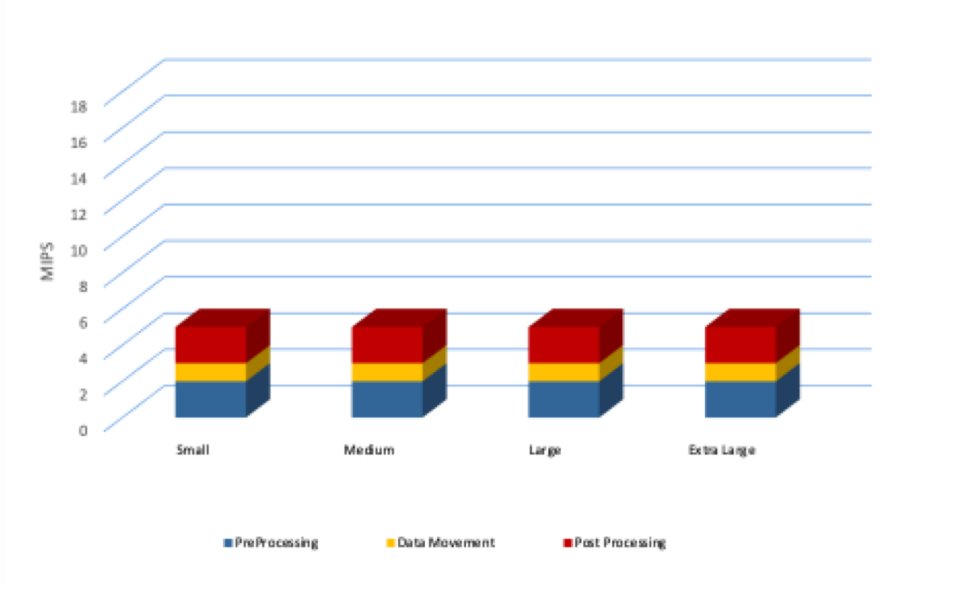

Figure 5 Number of MIPS by Data Set Size using MLC

One may choose to move all data sets larger than a certain size through the CLOUD option to gain the most MIPS benefit while leaving smaller data sets to follow another path (e.g. ML2 or remain on primary storage).

Once the Management Class migration criteria has been met (Days non-usage and/or size requirement), and the CLOUD action was specified, DFSMShsm request the data movement through the DS8000.

For more information on z/OS transparent cloud tiering settings, follow the document links provided in the Host Considerations section of this document.

- **DS8000**

Configuration for the DS8880 transparent cloud tiering uses the DS command-line interface (DS CLI) or the DS8000 Storage Management GUI to create a local user account to support the DS8880 as a proxy for metadata and administrative commands from the z/OS.

To configure with a TS7700 as a cloud target, use the DS CLI to enable the Ethernet ports and enter the TS7700 Grid link connection information using the mkcloudserver command with the -primary7700IPs parameter . This parameter will take a list of up to four comma-separated Grid IP Addresses from one TS7700. If two TS7700's will be targeted for synchronous copies, use the optional -secondary7700IPs parameter to list the IP Addresses for the second TS7700 clusters.

Example:

(single)

```
mkcloudserver -type TS7700 -primary7700IPs 1.1.1.1 -nossl
cloud_network_connection_name
```

or

(forking)
```
mkcloudserver -type TS7700 -primary7700IPs 1.1.1.1 -secondary7700IPs 2.2.2.2
-nossl cloud_network_connection_name
```

For more information on DS8000 transparent cloud tiering configuration, follow the document links provided in the DS8000 Considerations section of this document.

- **TS7700**

The TS7700 requires the activation of the DS8000 Object Store feature code 5282. This feature is activated on the TS7700 Management Interface (MI) Feature License panel once it is at microcode level of 8.50.x.x or later.

Once the feature code is activated, the TS7700 must be taken offline and an SSR must run an enablement procedure from the TS7700 service panel. The enablement process will fail if FC 5282 is not activated. The process is quick and once it completes the TS7700 can be brought back online. During the enablement process a single logical cache partition is created for objects. If all cache partitions are in use by tape or cloud functionality, the process will fail. Cache Partitions will be explained in more detail further in this document.

Next, an SSR must enter DS8000 primary and secondary (if exists) IP Addresses for each CEC into another TS7700 service panel. This can be done while the TS7700 is offline for the enablement process or once the TS7700 is fully online. For each DS8000, the following information will be requested:

- DS8K Machine Serial Number
- DS8K Primary IP Address for CEC1
- DS8K Primary IP Address for CEC2
- DS8K Alternate IP Address for CEC1 (optional)
- DS8K Alternate IP Address for CEC2 (optional)

DS8000 IP Addresses can be added, removed or deleted through the TS7700 service panels.

Once these configuration settings are established the TS7700 is ready to accept object data from the DS8000.

A TS7700 user should also allocate the desired cache partition size for object storage by accessing the Cache Partition panel on the TS7700 MI. See the TS7700 Cache Partitioning section in this document for more information on this topic.

7. DS8000 Object Store Flow

Once z/OS, DS8000 and the TS7700 are configured, DFSMShsm data is ready to be migrated to the TS7700.

Data sets in level 0 assigned to a Management Class that targets cloud as a migration tier will begin migrating MLC data using transparent cloud tiering if "Level 2 Days Non-usage" is set to 0 and the "Primary Days Non-usage" value is met. If the optional fields "Size Less Than or Equal To" and "Size Greater Than" are used, those criteria must also be met.

At this point DFSMShsm will initiate the migration of the data to be processed by DS8000. DS8000 will send the data, in the form of objects, to the TS7700 along with the *cloud_network_connection_name* using a PUT command.

When a recall is needed, DFSMShsm will initiate a restore of the data set via the DS8000 using a GET command to the TS7700.

If the Management Class "Level 2 Days Non-usage" is set to NOLIMIT, the data set will be migrated based on the values of "Level 1 Days Non-usage", "Size Less Than or Equal To" and "Size Greater Than". Data sets migrated to ML2 storage are otherwise handled the same as they are today (as logical volumes).

Figure 6 ML2 and MLC AutoMigation PathsFigure 6 shows these migration paths:

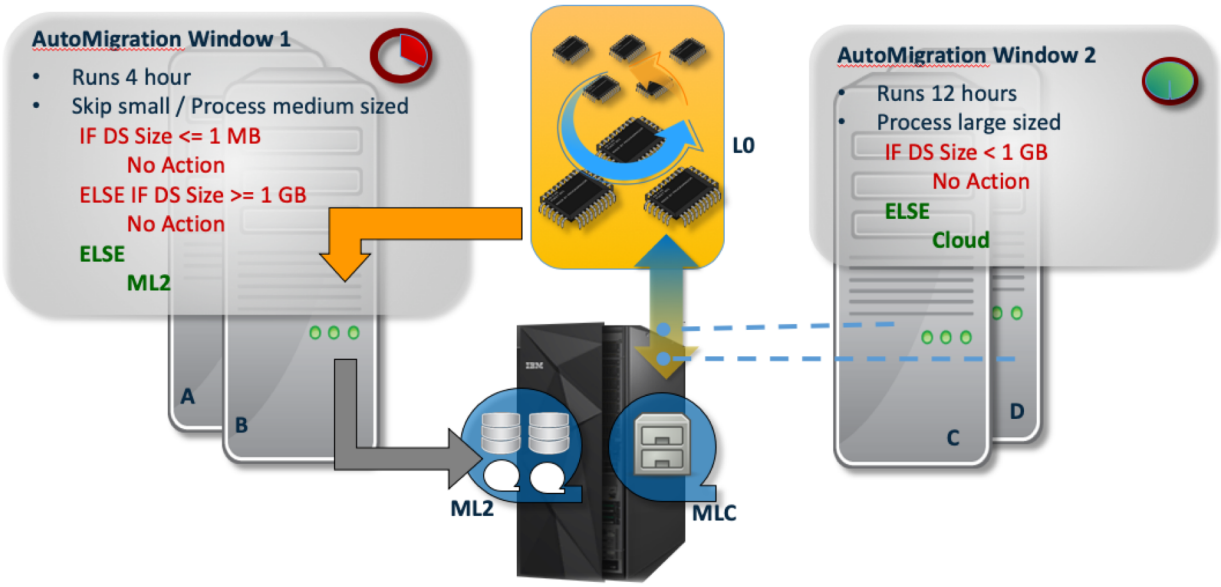

Figure 6 ML2 and MLC AutoMigation Paths

The DS8000 will read and write data to the TS7700 using existing Ethernet interfaces on both systems. On the TS7700, this would be the grid links and can be 1Gbit or 10Gbit Ethernet adapter cards. The DS8000 supports 1Gbit Ethernet with an optional upgrade to 10Gbit adapters.

8. TS7700 Cache Partitioning

- **Origin of Logical Cache Partitions**

TS7700 Logical Cache Partitioning is a concept that has been around since microcode release level 8.32.x.x. Cache partitioning has two main purposes. The first is to allow the ability to separate data into logical "buckets". What those buckets symbolize depends on business models and needs. For instance, separating data for a Sales division vs a Research division of a company or, if working in a multi-tenancy structures, the data may be separated by clients.

The second purpose of cache partitioning is to allow data to be migrated to a lower tier of storage for long-term archiving or backup purposes. At this time, the lower tier of storage is a physical tape library.

In the TS7700, Storage Class constructs are created to set Tape Volume Cache Preferences for a particular cache partition. These preference levels determine how soon volumes are removed from cache following their copy to tape. The Storage Class then targets a cache partition. This, along with other construct policies, allows for customization of data handling depending on the need. For instance, the Sales division of a company may want their data to remain in cache as long as it can for faster access while Research may be able to store the majority of their data into a more cost-efficient storage medium for archival purposes.

There are eight Logical Cache Partitions, cache partition 0 (CP0) through cache partition 7 (CP7). CP0 is reserved for resident only logical volume data - that is, data that comes over the FICON channel as logical volumes and is expected to remain in disk cache indefinitely and is not copied to physical tape.

Prior to microcode level 8.42.x.x, CP1 through CP7 were only activated with FC 5274 (Tape attachment). These cache partitions are managed in the Cache Partition MI page and assigned to logical volumes through the STORCLASS construct. The page was only viewable to a tape attached TS7700.

- **Logical Cache Partitions for Cloud Migration**

Beginning in TS7700 microcode level 8.42.x.x, partitions CP1-CP7 can also be used to export logical volume data to a cloud system through the TS7700 Cloud Storage Tier offering. Therefore, the Cache Partition page will be visible on a tapeless TS7700 although only when the Cloud Storage Tier feature code 5278 is activated. CP1-CP7 work the same for cloud exports as it did for migration to physical tape. For more information about Cloud Storage Tier see the *IBM TS7760 R4.2 Cloud Storage Tier Guide* Redpaper:

www.redbooks.ibm.com/redpapers/pdfs/redp5514.pdf

- **Logical Cache Partitions for Objects**

With the introduction of the DS8000 Object Store feature in the TS7700 8.50.x.x microcode release, cache partitions will again be used to separate data, although with some minor differences. In this

initial release, only one cache partition will be available for object data usage. This cache partition will be created during the DS8000 enablement process (see section 0. TS7700 above) with a default size of 3TB.

Object partition will have the following rules and restrictions:

- There must always be one object partition if DS8000 Object Store is enabled and it cannot be deleted.
- An object partition can be set as "primary object partition" and there must always be one primary object partition if DS8000 Object Store is enabled. The primary object partition is denoted by a green star to the left of the partition number on the Cache Partition page (see Figure 14). The concept of having a primary object partition is similar to the primary tape/cloud partition used today. However, in this initial microcode level 8.50.x.x, primary object partitions are purely symbolic since there is only one object partition.
- There must be one available cache partition for the DS8000 Object Store feature to be enabled. If all cache partitions are being used for logical volume workload, the enablement will fail.
- Only one object partition can exist in this release. Selecting the "Create Partition" button on a tape or cloud attached TS7700 will allow the creation of a virtual volume partition, as it does today, but not an object partition. If the TS7700 model is tapeless/cloudless, the "Create Partition" button is disabled.
- An object partition can be resized. Like all other tiered partitions (CP1-7), increasing the object partition size will decrease CP0's available space and decreasing the object partition will return the space back to CP0. There must be enough available space in CP0, minus any space taken due to flash copy operations, to meet the new allocation size of a tiered partition. Note: CP0 available space cannot be less than 2TB.
- An object partition can be renamed.

For additional information about the TS7700 MI Cache partition page, see the 0. **Cache Partition** section under Management Interface.

There is one new major concept introduced with object partitions. Object partitions are divided into even smaller sections. These sections are separated by the *cloud_network_connection_name* that was originally sent by DFSMShsm. In TS7700 microcode level 8.50.x.x, all *cloud_network_connection_names* will go into the one object partition that exists.

Since the DS8000 currently only supports one *cloud_network_connection_name*, each *cloud_network_connection_name* would be indicative of one DS8000 system. More than one DS8000 system can write data to a TS7700 but the *cloud_network_connection_name* must be unique.

The following image shows CP1 (an object partition) has two cloud names, Cloudname1 and Cloudname2:

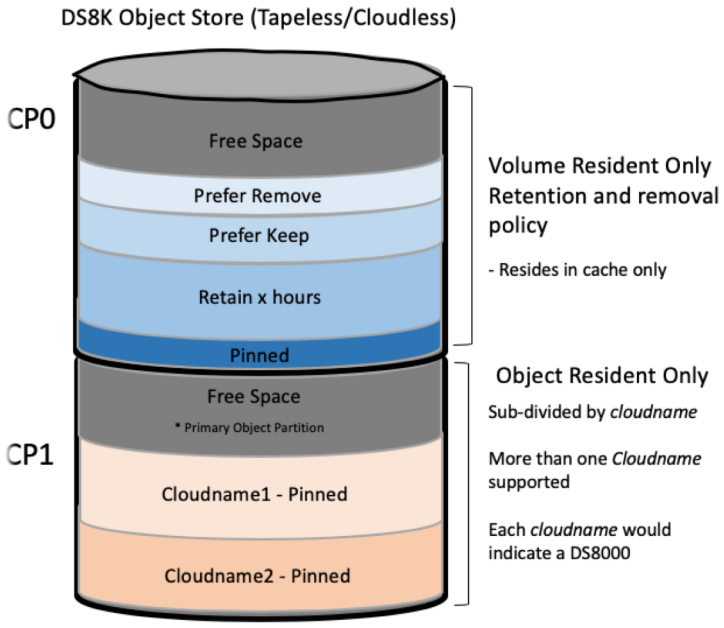

Figure 7 TS7700 Object Store Partition

Dividing a cache partition using *cloud_network_connection_name* will allow more flexibility for object data management in the future.

9. TS7700 Object Management

In TS7700 microcode level 8.50.x.x, TS7700 Advanced Policy Management is not available for object data coming from the DS8000. Data Class, Management Class, Storage Class and Storage Group cannot be assigned to objects. This section discusses how this affects managing object data.

- **Migrating Objects**

In this initial release of the DS8000 Object Store feature, object data will be resident only. In other words, objects will remain in cache and cannot be copied or moved to physical tape or to a cloud device attached to the TS7700 as part of the Cloud Storage Tier offering. Due to this reason, Storage Group and Storage Class options for data migration are not available for object data.

- **Data Redundancy**

The DS8000 can target one or two TS7700 VTSs to migrate object data to. If two TS7700s are targeted, a synchronous copy will be performed to both clusters using a forking mechanism. Both clusters must be in the same TS7700 Grid.

The forking process for synchronous copy is performed at the DS8000 level and does not use the TS7700 grid replication process (Management Class). TS7700 level Management Class policy for object data is not supported in this release.

A maximum of two TS7700s can be targeted by a single DS8000. Those TS7700s must be at microcode level 8.50.x.x or higher and can be a part of a larger grid.

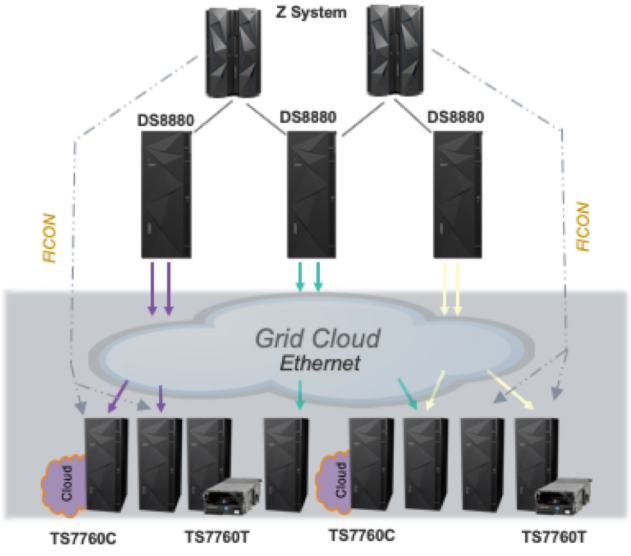

Figure 8 IBM Cloud Synergy - DS8000 Object Store

If a DS8000 is configured to migrate to one TS7700 and object data is written, it cannot be changed to forked copies later. All data would have to be removed and the DS8000 reconfigured to set up forking.

- **Compression**

Since object data comes through the Ethernet adapter and not through the FICON path, FICON compression cannot be used. Additionally, TS7700 software compression through Data Class is not available for DS8000 object data.

- **Encryption**

The TS7700 supports encryption for data at rest by either feature code 5272 (Disk Encryption with Local Key Management) or feature code 5276 (Disk Encryption with External Key Management). DS8000 object data written to these systems will benefit from these disk encryption features if they are enabled.

Secure data transfer of objects between a DS8000 and TS7700 is not supported in this release.

- **Removal Policies**

The TS7700 removal policy, known as Volume Copy Retention Group, provides a way to remove data from a disk only TS7700 as active data reaches full capacity. Volumes become candidate for removal as long as an adequate number of copies exist on peer clusters within a Grid and a specified retention period has elapsed.

Object data can only be removed from TS7700 cache by the DS8000 via DFSMShsm. In microcode level 8.50.x.x Volume Copy Retention Group (Storage Class) is not available for objects.

10. Disaster Recovery

If a DS8000 is targeting two TS7700s using forking, and one of the TS7700s becomes unavailable, DS8000 will continue to write to the available TS7700. DS8000 will send a system health message to the z/OS SYSLOGs about the out-of-sync condition and start cataloging any transactions that were not able to be completed. Once the TS7700 becomes available, DS8000 will begin to synchronize the object data. Once both TS7700s are back in sync, DS8000 will send another system health message to the z/OS SYSLOGs to indicate the out-of-sync condition is cleared. If the DS8000 cannot communicate with either TS7700, the job will fail.

The TS7700 has procedures to join new clusters to a grid, merge clusters that contain data into a grid or remove (unjoin) a cluster from a grid. A TS7700 could also go offline for a code upgrade or other scheduled outages. It is recommended to wait until the out-of-sync condition is cleared prior to any of these activities being performed. Activities such as unjoining a cluster from a grid that contains object data may result in unexpected data loss if data is out of sync. Note: The BVIR copy audit function does not detect data inconsistencies for DS8000 objects.

11. Management Interface

This section provides an overview of the changes to the TS7700 MI for the DS8000 Object Store feature.

- **Cluster Summary**

The cache tube widget displays a summary of how cache is being utilized within the TS7700. In clusters with no tape or cloud attachment, the cache tube provides a breakdown of cache by preference groups (prefer keep, prefer removal and pinned). In clusters with tape or cloud attached, the cache tube provides a breakdown of cache by cache partition.

Beginning with microcode level 8.50.x.x, if DS8000 Object Store is enabled, the cache tube will always show a breakdown by cache partitions even on a tapeless/cloudless TS7700.

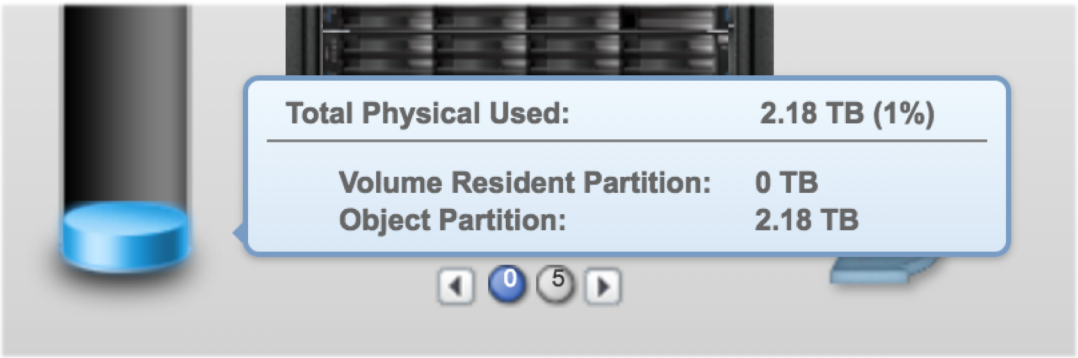

Figure 9 TS7700 MI Cluster Summary Cache Capacity

- **Historical Summary**

Historical Summary is located under **Monitor→Performance** on the TS7700 MI. The "Primary Used" and "Data in Cache" by partition fields under the **Storage** section of **Select Metrics** will include object data cache usage if DS8000 Object Store is enabled.

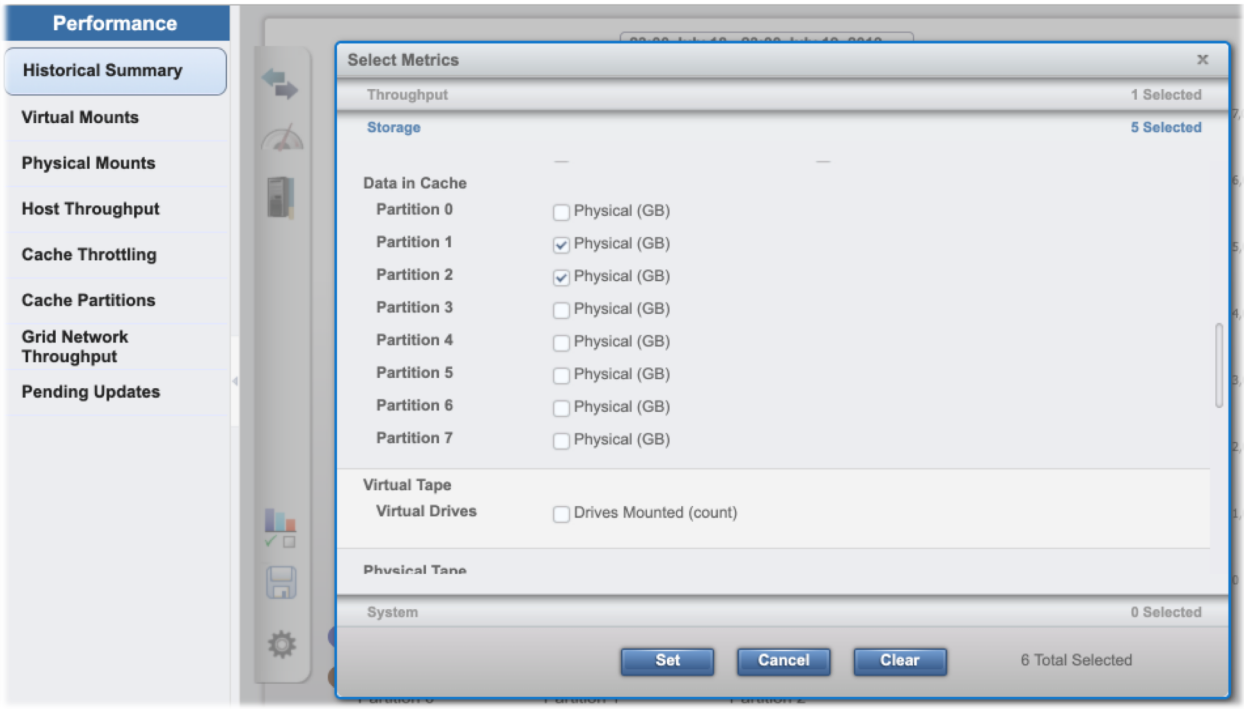

Figure 10 TS7700 MI Historical Summary Storage

- **Cache Utilization**

This Cache Utilization page is only displayed on a tapeless/cloudless TS7700. It will include a selection to show the total number of objects in cache.

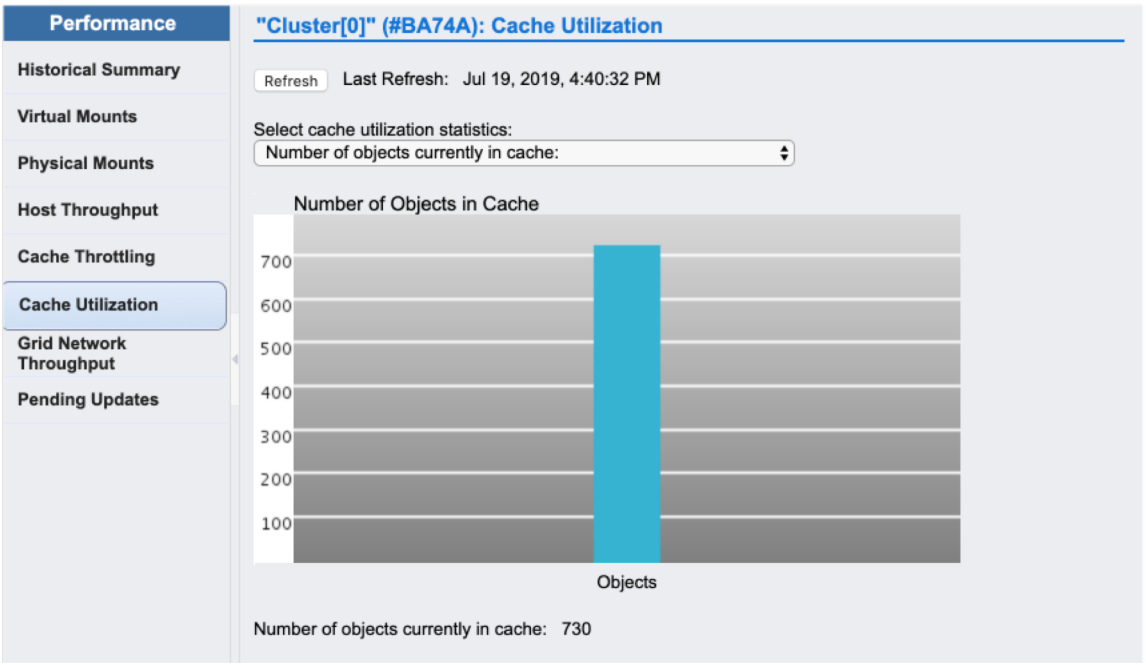

Figure 11 TS7700 MI Cache Utilization with Object Counts

It will also include the total amount of objects in cache as shown in Figure 12.

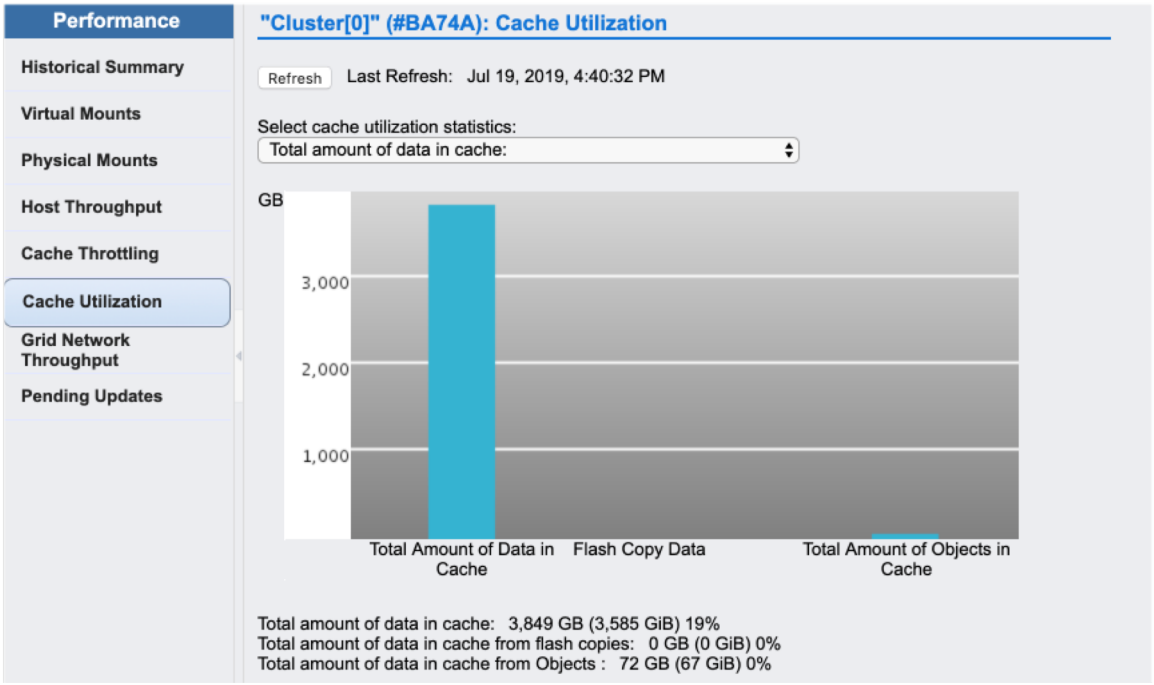

Figure 12 TS7700 MI Cache Utilization Cache Capacity

- **Cache Partition**

The Cache Partition page is located under **Virtual** in the MI navigation menu. When DS8000 Object Store is enabled, one object partition is created. This object partition can be resized by using the "Resize" action. Since objects cannot be premigrated to tape or cloud, there is no option to choose a delayed premigration limit as there is for a logical volume partition.

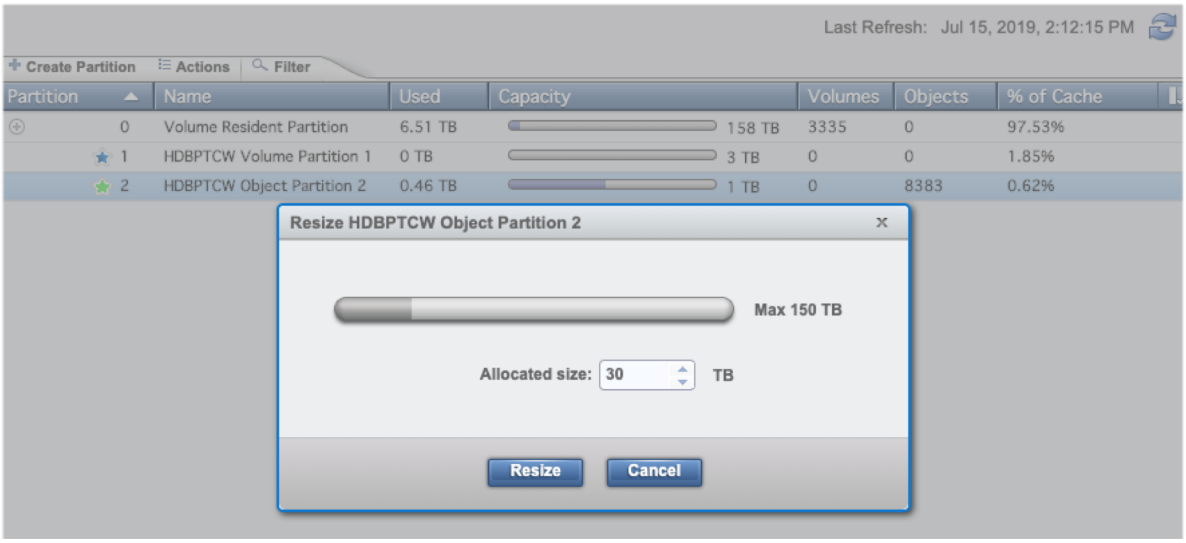

Figure 13 TS7700 MI Object Cache Partition Resize

There is also a new column to display the object count. This is the total number of objects that reside in an object cache partition. Figure 14 below shows the Cache Partition page on a cloud attached TS7700 with DS8000 Object Store enabled. Partition 2 is the object partition and contains 3351 objects while Partition 1 is a logical volume partition and contains no volume count.

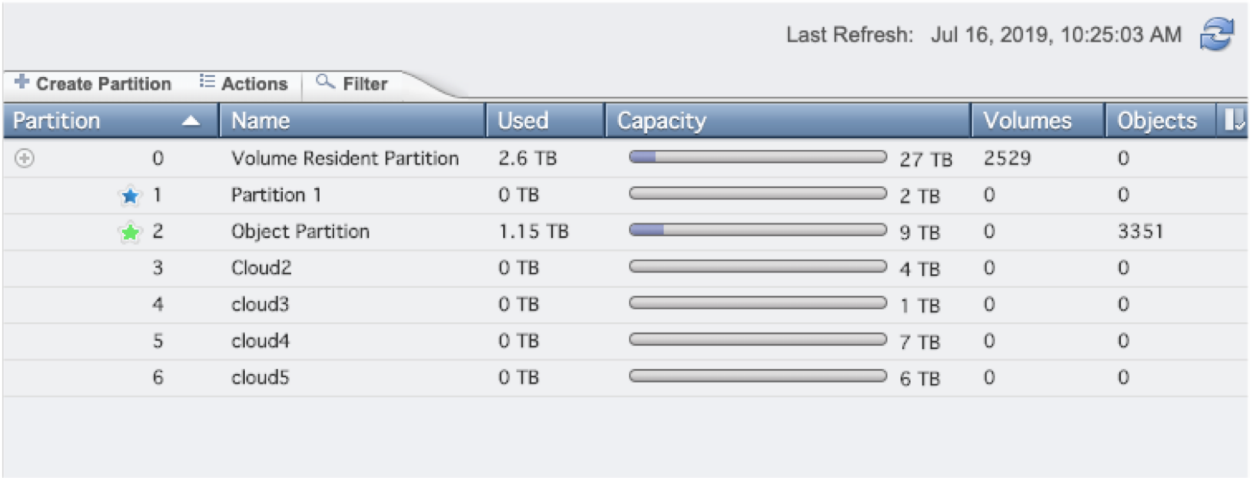

Figure 14 TS7700 MI Cache Partition with Objects

It should also be noted that resizing an object partition takes longer than resizing a logical volume partition. This is due to internal setup requirements and is expected. Progress can be monitored in the Task page found under **Monitor→Task** on the MI.

- **Storage Class**

Since object data cannot be migrated/premigrated to tape or cloud, Storage Class options are not available. When creating or modifying a Storage Class from the Storage Class page, the object partition will not be made available to select from the Partition dropdown. Only logical volume partitions will be displayed.

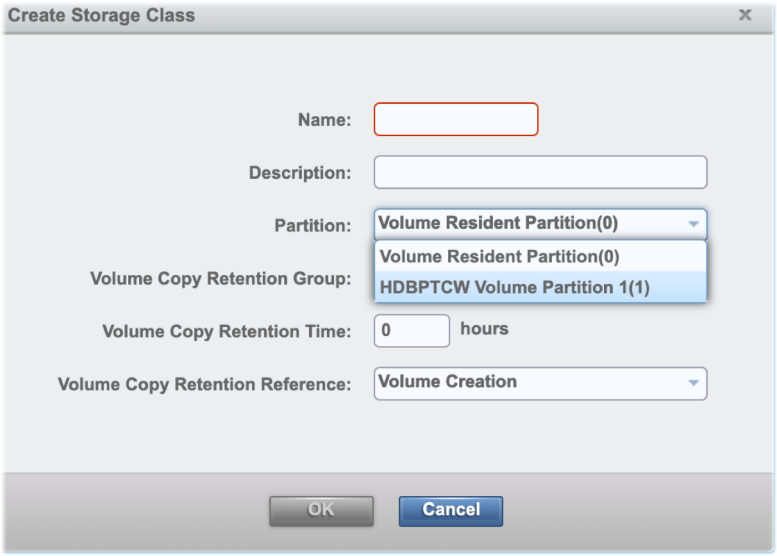

Figure 15 TS7700 MI Storage Class Page

- **Backup and Restore Settings**

Backup Settings and Restore Settings can be found under **Settings→Cluster Settings**. These pages provide an easy way to backup settings that can be restored on another TS7700. On a TS7700 tape or cloud attached, these settings include Cache Partitions. Object partitions have special configuration properties that are cluster specific and cannot be used on other TS7700s. For this reason, object partitions will not be included in Backup Settings and, therefore, will not be restorable.

12.Host Commands

This section describes host command modifications and additions for the DS8000 Object Store feature. The following is a list of the commands that will be discussed:

- LI REQ,SETTING,ALERT,RSDTLOW/RSDTHIGH
- LI REQ,SETTING2,ALERT,RSDOLOW/RSDOHIGH
- LI REQ CACHE2
- LI REQ STATUS,GRLNKACT

- **RSDTLOW/RSDTHIGH**

```
LIBRARY REQUEST,<distributed-library>,SETTING,ALERT,[RSDTLOW/RSDTHIGH]
```

RSDTLOW and RSDTHIGH are thresholds, in GBs, at which the TS7700 will generate a message indicating that a logical volume cache partition has exceeded a low or high residential capacity warning limit. These settings will work as they do today but are used not for object partitions when DS8000 Object Store is enabled.

- **RSDOLOW/RSDOHIGH**

LIBRARY REQUEST,<distributed-library>,SETTING2,ALERT,[RSDOLOW/RSDOHIGH]

RSDOLOW and RSDOHIGH work like the RSDTLOW and RSDTHIGH but are used for object partitions. If DS8000 Object Store is enabled, these commands can be used to generate a warning message.

Keyword1	Keyword2	Keyword3	Keyword4	Description	Comp	Dist
SETTING2	ALERT	RSDOHIGH	<value>	The threshold, in GBs of resident data for all object cache partitions, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a high warning limit.	N	Y
SETTING2	ALERT	RSDOLOW	<value>	The threshold, in GBs of resident data for all object cache partitions, at which the TS7700 will generate a message indicating that the amount of resident data has exceeded a low warning limit.	N	Y

Table 1 RSDOLOW/RSDOHIGH keywords

The following is an example of the ALERTS section of the SETTING2 command after RSDOLOW and RSDOHIGH is set:

```
-------------------------------------------------------------------------
ALERTS
PDRVSLOW =         0 PDRVSCRT =         0
LMTDTHR  =         0
CAGALOW  =         0 CAGAHIGH =         0
CAGLOW   =         0 CAGHIGH  =         0
RSDOLOW  =       500 RSDOHIGH =       700
-------------------------------------------------------------------------
```

Figure 16 SETTING2 RSDLOW/RSDOHIGH Values

- **CACHE2**

LIBRARY REQUEST,<distributed-library>,CACHE2

The LI REQ CACHE2 command replaces the LI REQ CACHE command.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Description	Comp	Dist
CACHE2				Requests information about the current state of the cache and the data managed within it associated with the specified distributed library.	N	Y

The main change to the command is a new TCO column. The TCO column maps as follows:

T = Tape migration
C = Cloud migration
O = Partition contains Objects

A 'Y' under the T, C or O indicates the partition supports this option. An 'N' indicates the partition does not support this item. If the O column contains an 'N', the partition is a logical volume partition.

The following is an example output of the CACHE2 command. In this output, Partition 1 is a logical volume partition that supports migration of data to tape and Partition 2 is an object partition:

```
SHOWING RESULTS FOR COMMANDS: CACHE2
TAPE VOLUME CACHE STATE V5 .0
TOTAL SPACE INSTALLED/ENABLED:   162TB/ 162TB
TOTAL ADJUSTED CACHE SPACE USED:  6509GB
CACHE ENCRYPTION STATUS:  NOT CAPABLE
OVERCOMMITTED CACHE PARTITIONS:  NONE
CACHE RESIDENT ONLY PARTITION
 PRIVATE CACHE SPACE USED:   243GB
 SCRATCH CACHE SPACE USED:  6265GB
 CP   ALLOC    USED     PIN     PKP     PRM    COPY    CPYT
  0   158TB  6508GB  2088GB  4419GB   0.0GB   0.0GB       0
 FLASH COPY INFORMATION
 INDEX ENABLED    SIZE
     1      NO   0.0GB
     2      NO   0.0GB
     3      NO   0.0GB
     4      NO   0.0GB
     5      NO   0.0GB
     6      NO   0.0GB
     7      NO   0.0GB
     8      NO   0.0GB
TIERED CACHE PARTITIONS
 CP TCO  ALLOC   USED    PG0    PG1   PMIGR D_PMIGR   COPY  PMT CPYT
  1 YNN 3000GB  0.0GB  0.0GB  0.0GB  0.0GB   0.0GB  0.0GB    0    0
  2 NNY 1000GB  462GB  0.0GB  462GB  0.0GB   0.0GB  0.0GB    0    0
  3 NNN  0.0GB  0.0GB  0.0GB  0.0GB  0.0GB   0.0GB  0.0GB    0    0
  4 NNN  0.0GB  0.0GB  0.0GB  0.0GB  0.0GB   0.0GB  0.0GB    0    0
  5 NNN  0.0GB  0.0GB  0.0GB  0.0GB  0.0GB   0.0GB  0.0GB    0    0
  6 NNN  0.0GB  0.0GB  0.0GB  0.0GB  0.0GB   0.0GB  0.0GB    0    0
  7 NNN  0.0GB  0.0GB  0.0GB  0.0GB  0.0GB   0.0GB  0.0GB    0    0
EXECUTING COMMANDS: CACHE2
```

Figure 17 LI REQ,CACHE2

- **GRLNKACT**

LIBRARY REQUEST,<distributed-library>,STATUS,GRLNKACT

The GRLNKACT command was first introduced in microcode release 8.42.x.x to display point-in-time Grid link activity for all clusters in the Grid but at the perspective of the cluster specified in the distributed library ID field.

In TS7700 microcode level 8.50.x.x, additional fields were added to support DS8000 Object Store and Grid to Grid Migration (GGM) activity.

Keyword 1	Keyword 2	Keyword 3	Keyword 4	Description	Comp	Dist
STATUS	GRLNKACT			Requests information about Grid link point-in-time activity.	N	Y

The command output looks as follows:

```
GRLNKACT STATUS V1 .0
CLUSTER INDEX:  2 LINK COUNT:  4 Time: Thu Apr 11 00:05:13 CUT 2019
GRID LINK ESTABLISHED SOCKET CONNECTIONS-------------------------------
LN INTF IP              C0  C1  C2  C3  C4  C5  C6  C7  MQ  RFA  CLD
L0 en10 9.11.219.178    106 82  0   0   0   0   0   87  268 7    1
L1 en4  9.11.219.190    5   2   0   0   0   0   0   2   3   6    2
L2 en11 9.11.219.191    3   4   0   0   0   0   0   0   0   7    1
L3 en5  9.11.219.192    3   4   0   0   0   0   0   0   0   7    4
LN INTF IP              GGM  OBJ
L0 en10 9.11.219.178    8    6
L1 en4  9.11.219.190    8    6
L2 en11 9.11.219.191    0    0
L3 en5  9.11.219.192    0    0
NET ACTIVITY -----TCP RECV/SEND ADAPTER BUFFER ACTIVITY BYTES---------
LN  TxMBs RxMBs MQ_REC   MQ_SND   GFA_REC  GFA_SND  CLD_REC  CLD_SND
L0  81    225   0        0        0        124792   0        0
L1  58    156   0        0        0        0        0        2089817
L2  2     0     0        0        0        0        0        0
L3  42    0     0        0        0        0        0        1046017
TOT 183   381   -        -        -        -        -        -
LN  GGM_REC  GGM_SND  OBJ_REC  OBJ_SND
L0  0        0        2099835  4179634
L1  1460     0        1050791  1039795
L2  0        0        0        0
L3  0        0        0        0
GRID LINK THROUGHPUT ESTIMATES-MB/s-----------------------------------
DIR C0   C1   C3   C4   C5   C6   C7   GRD_TOT CLD  GGM  OBJ  TOT
Tx  2    3    1    0    0    0    1    7      78   0    98   183
Rx  0    0    0    0    0    0    0    0      247  1    133  381
GRID CLOUD TIER EXPORT AND IMPORT ACTIVITY----------------------------
```

```
ACTIVE EXPORT VOLUME COUNT: 5
ZCL060 ZCL061 ZCL062 ZCL063 ZCL065
ACTIVE IMPORT VOLUME COUNT: 0
```

Figure 18 LI REQ,STATUS,GRLNKACT

The additional fields added in microcode release 8.50.x.x in support of the DS8000 Object Store feature are as follows:

GRID LINK ESTABLISHED SOCKET CONNECTIONS	
OBJ	The number of TCP Sockets currently established between the TS7700 and all DS8000s for each link (L0-L3).
NET ACTIVITY	
OBJ_REC	Amount of data, in bytes, that has been sent for write activity by all DS8000s to this cluster as part of a transparent cloud tiering PUT command in the last 15 seconds. This value is per link.
OBJ_SND	Amount of data, in bytes, that has been read by all DS8000s from this cluster as part of a transparent cloud tiering GET command in the last 15 seconds. This value is per link.
TxMBs TOT	The total transmitted activity of all links in the last 15 seconds. This total will now include OBJ_SND
RxMBs TOT	The total received activity of all links in the last 15 seconds. This total will now include OBJ_REC
GRID LINK THROUGHPUT ESTIMATES-MB/s	
OBJ Tx	Estimated Grid link throughput for object data sent from this cluster to all DS8000s in the last 15 seconds.
OBJ Rx	Estimated Grid link throughput for object data received by this cluster to from all DS8000s in the last 15 seconds.
TOT Tx/Rx	The total estimated transmit and receive grid link throughput activity. These totals will now include OBJ throughput.

Note: The cluster number for which this data was read will no longer be listed in the Grid Link Throughput Estimates-MB/s section of this output to make room for the new OBJ field. The local cluster number would always be zero so has no value. For instance, in Figure 18 there is no C2 column. This is because this data was generated for the distributed library ID belonging to cluster 2.

For detailed information about these Library Request Commands see the z/OS Host Command Line Request User's Guide Version 5.0 or higher located in the IBM Techdocs Library : http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

13.Statistical Data

BVIR output will include object counts and transferred data bytes. Object partition data will also be included as it has in the past.

More information regarding TS7700 statistics can be found in the *IBM TS7700 Series Statistical Data Format White Paper Version 5.x.pdf* found in the IBM Techdocs Library at http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs.

References

White paper - IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide (latest version)

Disclaimers:

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:

September 4, 2019

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 IBM Z.

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