TS7700 Technical Update

What's that I hear about R3.2?

Ralph Beeston
TS7700 Architecture
IBM

Session objectives

- Brief Overview TS7700
- TS7700 Release 3.2
  - TS7720T – TS7720 Tape Attach
    - The Basics
    - Partitions
    - Premigration and Delayed Premigration
    - Odds and Ends
  - Sync Mode Copy Mod option
  - Universal 25 GB volume size support
  - Up 496 Virtual Devices

IBM TS7700 System Z Virtualization Engine

Leveraging Grid as Cloud Storage for System z

Grid Cloud

- 8Gb FICON
- Disk Speeds
- TS1100 Tape Integration
- Transparent Business Continuance
- AES 256 bit Encryption

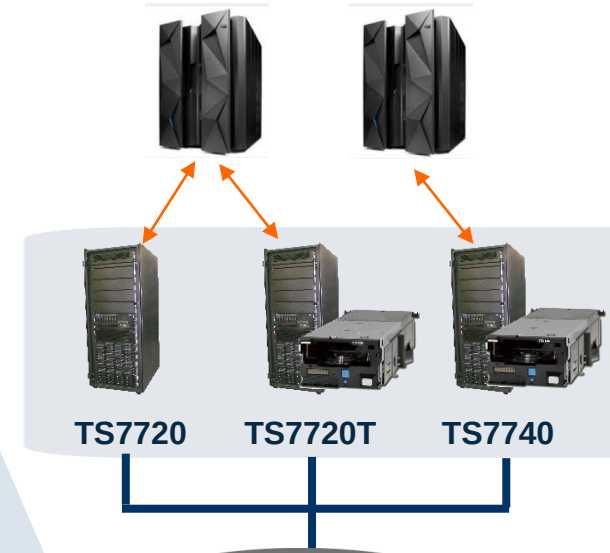

Cumulative FICON throughput of over 13GB/s

System z hosts view up to 2,976 equivalent devices

Grid access to all data independent of where it exists

TS7700 Solutions

- TS7720
  - Access-centric applications (image data, report servers, critical backups, HSM ML2)
  - Cost-efficient applications (HSM, general backups, master-in, master-out, GDGs, archive)
  - Up to 1PB disk cache per system prior to compression. (3 PB with 3:1 compression)
  - Optionally, up to 80PB backend physical tape storage
  - AES 256bit Encryption (disk cache & tape)
  - 3TB NL-SAS 7.2K RAID6

- TS7740 Virtualization Engine
  - Cost-efficient applications (HSM, general backups, master-in, master-out, GDGs, archive)
  - Up to 28TB disk cache per system prior to compression
    - 84TB with 3:1 compression
  - Up to 80PB backend physical tape storage
    - 240PB with 3:1 compression
  - AES 256bit Encryption (disk cache & tape)
  - 600GB SAS 10K RAID6

TS7700 R3.2 – TS7720T Basics

Tape Attach support for TS7720 combines the best of both TS7700 offerings:

- Large cache size for increased or 100% cache hit ratios
- Physical tape for additional copies or added capacity
- Physical tape for seldom or never referenced data

In addition to all the features and functions provided by the TS7720 and TS7740, the TS7720T has additional capabilities:

- Multiple tape partitions
- Delayed pre-migration for tape partitions

TS7700 R3.2 – TS7720T Basics

- Like TS7740, the TS7720T supports the same back end tape attachment.
  - TS3500
  - J1A/TS1120/TS1130/TS1140
- A TS7720T must have at least 1 Tape Partition – but up to 7 independent tape partitions can be defined. These partitions are referred to as CPx in this presentation.
- In addition to the Tape Partitions, a Resident Only Partition exists. This partition is referred to as CP0 throughout this presentation.
  - Only one resident partition exists.
  - This partition acts like a TS7720 disk only cluster.
- Supports all existing functions both on TS7720 and TS7740.
  - Support Flash Copy for DR testing
  - Support Copy Export and Copy Export Restore/Merge

TS7720 Tape Attach (TS7720T)

- **Custom define up to seven tape attached partitions**
  - Each partition utilizes a user defined fixed amount of capacity from the total configured available capacity
  - Each partition is independently hierarchal storage managed to back end tape
    - Contains its own PG0 and PG1 content.
    - Deferred Premigration (retain copy in disk cache only until time from creation/last-access expires)
      - Perfect for Archive!
  - Policy managed as to which partition a workload targets
- **One resident only partition**
  - What space remains after creating one to seven tape managed partitions is viewed as the resident only partition
  - All data in the resident only data partition remains there indefinitely and doesn't offload to tape
  - TS7700 grid removal policies can apply to this partition if it becomes full

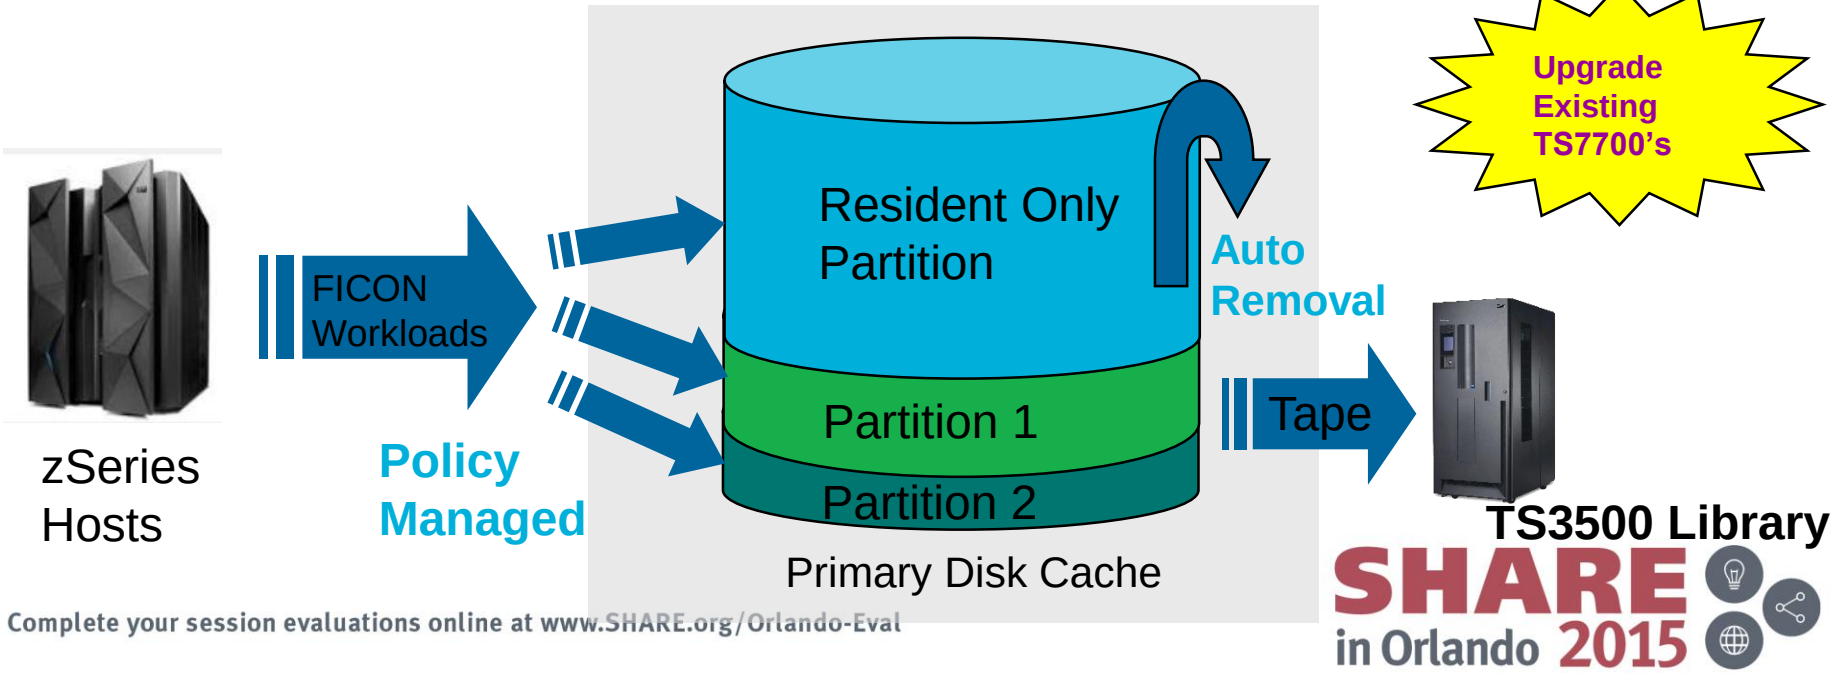

TS7700 R3.2 – TS7720T Partitions

- **Residency only workloads can have dedicated space through CP0**
  - The size of CP0 can be limited in size thus providing a method to limit the amount of residency only data present before auto-removal kicks in
- **Different groups of workloads targeting physical tape may have different cache residency requirements too, thus the tape managed portion of the disk cache can be further divided into partitions too**
  - A given workload may require N days or X TB of content to be resident in cache
  - A dedicated partition can be defined to accommodate such a workload
  - Other workloads with different cache residency requirements can target other partitions without affecting the cache residency of adjacent partitions. For example:
    - Workload A has a 100TB residency requirement, so a 100TB partition A is defined
    - Workload B has a 50TB residency requirement, so a 50TB partition B is defined
    - Workload A creates 10TB of new content a day.
    - Workload B creates 20TB of new content a day
    - Both workloads A and B are setup as PG1.
    - As long as each workload targets their own partition, the 10TB workload for workload A only pushes out workload A content when the partition becomes full. The migration LRU algorithm for partition A is only applicable to workload A data. The same applies for workload B content. Both workloads and their associated daily activity will never alter the adjacent partitions workload's disk cache footprint.

TS7720T Time-delayed Premigration (New Concept!)

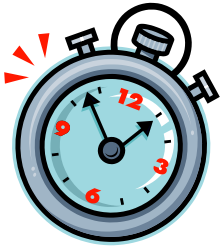

- **Move workloads to tape after a timed delay**
  - Retain data only in disk cache until delay criteria is met
  - Delayed criteria in hours from creation or last access
  - Any content which expires prior to delay period never goes to tape
- **Archive made easy**
  - Have workloads automatically age to tape
  - Limited knowledge needed to determine what is considered archive
  - If it still exists after N hours or hasn't been accessed in N hours, then put it on tape.
- **Once on tape, honor tape based policies**
  - PG0/PG1 honored once copy put on tape
    - PG0 will be removed from disk cache immediately
    - PG1 will be retained using a lease recently used algorithm
- **Limit how much data can be delayed to tape**
  - When sizing a tape partition, a maximum delay premigration limit is defined
  - If delay premigraiton exceeds this limit, it will move to tape early

TS7720T Cache Partitions and Premigration

- Single premigration queue for all tape manage partitions
- Thresholds apply to the global cumulative used disk cache space for all tape manage partitions
- If throttling occurs, it's applied to all workloads targeting a CPx partition.
- Premigration queue size is limited by FC5274 Increments
  - Cumulative post-compressed size of all volumes contained in queue
  - PMTHROT is limited to queue size
- Workloads targeting CP0 are not throttled due to premigration thresholds

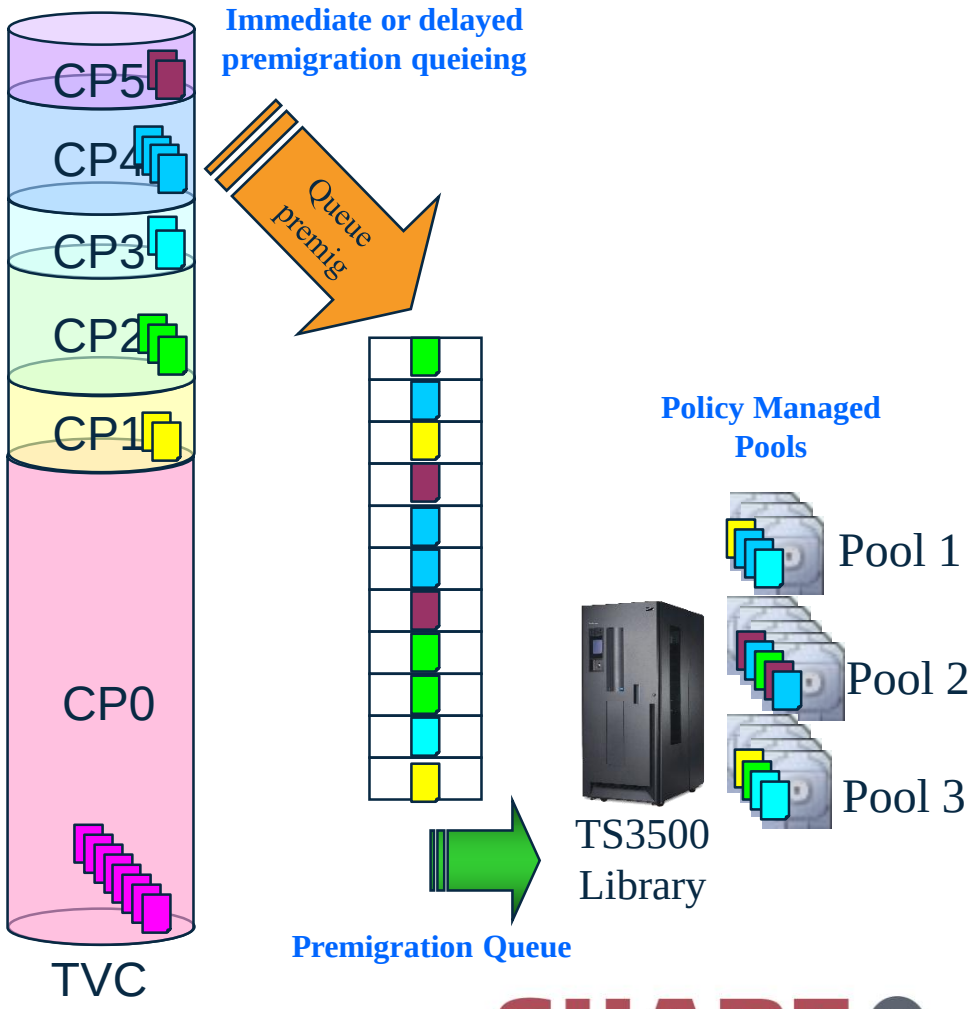

TS7700 R3.2 – TS7720T Odds and ends

- Tape attach features may be ordered as Plant Installed on new systems, or as an MES to an installed TS7720
  - Model VEB server is required
  - Same requirements for the physical tape environment apply to TS7720T as to TS7740
    - TS3500 Tape Library
    - Two fibre channel switches
    - From four to sixteen 3592 tape drives
- TS7720T and Four Disk Cache Strings are Mutually Exclusive
  - If 4 strings of disk cache already exist, TS7720T is not possible
  - Once TS7720T is installed, adding 4th string of disk cache is not possible
- Replace TS7740 with TS7720T and maintain content on physical tape
  - Requires a service outage (non-concurrent for specific TS7740 being targeted)
  - Only TS3500 attached libraries supported.
  - All existing data in TS7740 cache will be flushed to physical tape
  - Multiple backups of DB and configuration data DVDs
  - Push out TS7740, pull in TS7720T and attach existing FC cables from TS3500 to new TS7720T
  - Restore DB and configuration data from backup DVD

TS7700 R3.2 – Sync Mode Copy Mod option

- **Synchronous Copy Mode changes introduced in R3.2**
  - z/OS will now provide information to the TS7700 as part of the mount request as to whether the host is likely to issue an update.
  - The TS7700 has provided a new option in the Management Class setting in the MI. The option states whether the dual open should only occur if a host implied update is detected at volume mount time.

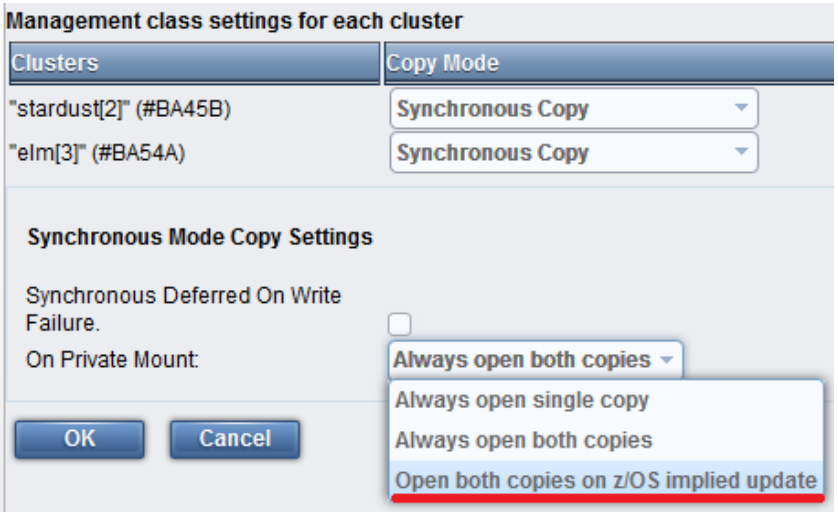

- **Always open single copy**
  Same with unchecked "Open Both Copies On Private Mount" in the previous releases.
- **Always open both copies**
  Same with checked "Open Both Copies On Private Mount" in the previous releases.
- **Open both copies on z/OS implied update**
  When this option is used, a private mount will only open a single 'S' location unless the host implies an update is likely to take place.

Note: Using z/OS DISP information is not guaranteed to always imply an update. Unique applications may not always use DISP as intended.

	Always open single copy	Always open both copies	Open both copies on z/OS implied update
Update implied by JCL and update occurs	A single 'S' location opened (enter synchronous deferred state)	Dual 'S' locations opened (no synchronous deferred state)	Dual 'S' locations opened (no synchronous deferred state)
Update not implied but update still occurs.	A single 'S' location opened (enter synchronous deferred state)	Dual 'S' locations opened (no synchronous deferred state)	A single 'S' location opened (enter synchronous deferred state)

[Table] How DISP and 'S' copy mode options interact.

TS7700 R3.2 – TS7700 25 GB Volume Support

- **25,000 MiB Volumes Now Supported Everywhere!**
  - 25GB volumes previously supported by RPQ 8B3580 in TS7720 only grid configurations
  - Now, any TS7720, TS7740 and TS7720T can utilize 25 GB volumes without RPQ
  - No RUN/Immediate copies supported with 25GB volumes due to long copy times.
    - Any volumes created with a 25GB assigned Data Class and a Management Class with one or more RUN/Immediate mode consistency points will treat the inconsistent RUN consistency points as Deferred within the copy queues.
    - The box will not enter the immediate-deferred state and no warnings will occur.
    - RUN consistency points can still be used for preferred TVC selection, but any inconsistent RUN points at close time will only receive a deferred copy.
  - Appends in deferred copy mode will re-replicate entire volume
    - Workloads which append should stay away from 25GB volumes unless synchronous replication policies can be used.
  - No physical tape recall accelerators
    - If a 25 GB volume only resides on physical tape, a mount will require the entire 25 GB volume to be recalled into disk cache independent of how much data the host actually needs.
  - All clusters in the grid must be at R3.2 if one or more TS7740s are present
    - Otherwise, all clusters must be at R3.2 or contain RPQ 8B3580 on older TS7720 clusters
  - 128 concurrent 25 GB lvol mounts is the default
    - This value can be increased or decreased based on customer use case

TS7700 R3.2 – Support up to 496 Virtual Devices

- **Why more than 256 devices per cluster is required?**
  - The TS7700 has always been limited to 256 devices or 16 LCUs per cluster.
  - With larger TS7720 capacities and faster FICON channels, there is a greater need to consolidate workloads against smaller number of clusters.
  - However, the need remains for large numbers of logical devices.
  - R3.2 adds 15 additional subsystems per TS7700 cluster for a total of 496 devices per cluster.
  - These additional devices are available via a Feature Code (FC 5275) which allows 16 additional devices per purchased feature. A total of fifteen feature codes can be allowed on a single cluster for a maximum of 496 devices.

- **Required configuration to get more than 256 devices per cluster**
  - Only TS7700 clusters with 8Gb FICON supports additional devices due to increased memory requirements.
  - All clusters in the Grid must be at R3.2 or later for this feature to be activated (the feature code itself can be installed with the mixed code configuration, but devices are inaccessible until all clusters are R3.2 or later).
  - The same number of feature licenses do not need to exist on every cluster in the Grid (for example, one cluster can have 256 devices while another has 496 devices).
  - Prior to R3.2, System Z has a maximum device limit of 2048 per grid or composite library. An APAR is available to increase the limit to 4096. Please check the 3957DEVICE PSP-Bucket for all required R3.2 APARs.

Session summary

- TS7700 Release 3.2
  - TS7720T – TS7720 Tape Attach
    - Multiple Partitions
    - Premigration and Delayed Premigration
    - MESable
  - Sync Mode Copy Mod option
  - Universal 25 GB volume size support
  - Up 496 Virtual Devices

Q&A?

- Additional Questions?

Excellent References!

- Techdocs
  - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs
  - Search on TS7700
- TS7700 R3.2 Redbook (1st draft)
  - http://www.redbooks.ibm.com/redpieces/abstracts/sg248122.html
- TS7700 Infocenters
  - Customer - http://pic.dhe.ibm.com/infocenter/ts7700/cust/index.jsp
  - Service - http://pic.dhe.ibm.com/infocenter/ts7700/serv/index.jsp

Thank You