

Washington Systems Center - Storage

Accelerate with IBM Storage: *TS7760 Best Practices* (A View from the Field)

Presenter

Bill Danz

Panelists

Ben Smith

Bob Sommer

Mike Berish



Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2019...

Washington Systems Center – Storage experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

Located in the Accelerate with IBM Storage Blog:
<https://www.ibm.com/developerworks/mydeveloperworks/blogs/accelerate/?lang=en>

Also, check out the WSC YouTube Channel here:
https://www.youtube.com/channel/UCNuks0go01_ZrVVF1jgOD6Q

2019 Webinars:

- January 24** - Discover how IBM Storage Insights and Storage Insights PRO can transform your IBM Storage Support Experience
- January 31** - Elastic Storage Server (ESS) Architecture and use cases
- February 28** - Data Protection and High Availability in Spectrum Virtualize
- March 7** - Overview of IBM's New Spectrum Discover
- March 14** - Why IBM Storage is relevant for IBM Cloud Private deployments
- April 11** - GDPS V4.2 and SPEs: What's new?
- April 18** - IBM DS8880 R8.5 GUI Demo
- April 23** - Copy Services Manager and zHyperLink Update
- April 25** - Data Reduction Pools, Overview and Best Practices
- May 2** - TS7760 Best Practices (A View from the Field)
- May 8** - Discover Latest Capabilities of Spectrum Virtualize Easy Tier Setup(s)
Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=e0ce01b732b304e42100981b38d773dda>
- May 23** - How to use the IBM TS4500 and TS4300 Library Web GUI
Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=edaae73a07257503198d0bdc04d18973e>
- May 30** - New Features with IBM Cloud Object Storage including a File Access Demo
Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=eed0398a8690f50d436c35ba50c0135ac>



IBM TS7700 Leadership and Progression (1997 – 2019)

<p>B16 – Industry First VTS (1997) B18 – Next Gen VTS (1998) P2P VTS – First DR VTS (2000) B10/20 – Next Gen VTS (2001)</p>	<p>TS7720 R1.5 - 12/05/2008 Disk Cache Refresh TS1130 support TS3500 HD support</p>	<p>R2.1 – 12/19/2011 Sync Mode Copy IP Link Failover Grid Merge Copy Export Merge Copy Export Acceleration Customer log collection COPY REFRESH</p>	<p>R3.3 – Sept 2015 TS1150 Support ISKLM for Disk Grid to Grid Migration</p>
<p>TS7740 R1.0 - 09/29/2006 2-way grid TS3500 Support 16 3592 J1A drives</p>	<p>R1.6 - 12/04/09 4-way grid TS7720/TS7740 Hybrid Logical WORM Network Load balancing Cluster Families SNMP</p>	<p>R3.0 – 12/07/2012 Disk Cache Refresh Disk AES256 Encryption Unified GUI 4 Million Logical Volumes Native LDAP Limited IPv6 support</p>	<p>TS7760 R4.0 – June 2016 TS4500 Support Power8 Server Refresh 4x10Gb Grid links Distributed RAID6 Rebuilds</p>
<p>R1.1 - 01/26/2007 AOTM TS1120 Drive Support Larger Disk Cache SW 1Gb Fibre Grid 256 virtual devices</p>	<p>R1.7 - 06/04/2010 Disk Cache Refresh Selective Write Protect Additional Memory FC 6000MB logical volumes LDAP Support</p>	<p>R3.1 - 12/06/2013 1PB TS7720 8x8Gb FICON 512 paths per port Flash Copy for DR Testing Time Delayed Replication</p>	<p>R4.1 – 1Q17 7/8 Way Grid 8TB Drives</p>
<p>R1.2 - 03/09/2007 Tape AES256 Encryption Broadband call home</p>	<p>R2.0 - 06/03/2011 Power7 Server Refresh Server Swap Support 2x10Gb LW Grid 4x1Gb Grid 2 Million Logical Volumes 6-way grid Scratch Allocation Assist TS1140 Support PG0/PG1 Copy Priority</p>	<p>R3.2 – Dec 2014 Tape Attach TS7720 25GB Volume Support 496 UCB/Devices RACF LDAP Support LDAP SSR Bypass</p>	<p>R4.1.2a – 1Q18 16Gb FICON LZ4/ZSTD Compression CUIR Automation Event Redesign (SYSLOG) Resiliency Automation</p>
<p>R1.3 - 08/31/2007 3-way grid LI REQ Support Copy export (standalone) 1M logical volumes Secure data erase ROR Recovery</p>			<p>R4.2 – 3Q18 TS7760C TCT Support</p>
<p>R1.4 - 11/23/2007 Copy export (grid)</p>			

Overview

This collection of TS7760 Best Practices comes from IBM's twenty-two (22) years of experience with mainframe virtual tape. From the original 3494-B16 Virtual Tape Server to today's TS7760 system, we've seen what has worked well for customers and experienced situations where alternate choices would have served everyone's best interests.

By necessity, many of these recommendations are somewhat general in nature and they may not be the right choice in specific instances. **Each customer should carefully evaluate their own environments before implementing any of these suggestions.** Customers can always contact IBM or IBM's Business Partners if there are questions about any of the recommendations contained in this presentation.

The "Best Practices" are organized into these categories:

- General
- Configuration
- Operations
- Network
- Disaster Recovery

BEST PRACTICES – GENERAL

General – KISS Your Grids! ***

*** Keep It Simple and Straightforward

A TS7760 Grid can support up to:

- Eight (8) clusters
- 19.6 PB of useable disk cache storage (2.45 PB x 8)
- 500 PB of customer data in cloud or physical tape storage @ 5-to-1 compression (4M vols x 25 GB vols x 5 comp rate)
- Sixty-four (64) 16 Gb FICON channels (8 channels per TS7760 x 8)
- Thirty-two (32) 10 Gb Ethernet connections for Grid replication (4 connections per TS7760) x 8)
- 256 Scratch Categories, 256 Management Classes, 256 Data Classes, 256 Storage Classes, and 256 Storage Groups



In other words, one can construct very large and extremely complex TS7760 Grids

(Best Practice) Keep each TS7760 Grid as simple as is possible

- Each Grid must still meet its business (availability, performance and data retention) requirements
- Make sure there is a valid business need before introducing a new category or construct (Mgmt Class, Data Class, etc.)
 - Evaluate added complexity versus any potential gains (i.e. amount of data stored, data availability, performance)
- Don't add to the design or configuration complexity in an attempt to "micro-manage" a TS7760 cluster or Grid
 - The software in the TS7760 systems (almost) always makes excellent (optimal) real-time choices
 - Potential performance gains or storage reductions are often temporary and are often not worth the effort
 - Two or three years from now, will anyone remember why a particular "exception" was added?

General – Design for (Planned) Changes



TS7760 Systems may occasionally need to be taken offline to implement planned changes

- System code levels should be upgraded at least once each year
- (Infrequent) There are several hardware upgrades or additions that require a system outage to implement
 - Examples: Physical memory changes, add cloud or physical tape storage support, swap Ethernet or FICON cards
- In-place or “frame roll” replacement of an existing system with a new TS7700 system

Several of the changes listed above can take several hours to implement

- Because they are planned, the change activities can be scheduled during low tape processing times
- One still needs to understand the expected duration of these changes
 - What will be the impact on mainframe operations during the outage?

(Best Practice) Consider having two (2) local TS7760 systems so that mainframe tape processing can continue even if one (1) cluster is unavailable due to a planned change

- A second option is to configure network “FICON channel extension” so that local System Z systems can access a remote TS7760 when the local virtual tape system is not available
- Both options will provide operational redundancy eliminating the need for System Z processing holds

General – Planned Change: TS7720 and TS7740 End-of-Support

IBM announcement letters 918-099 (US) and A18-0183 (Canada) in 2018 gave notice that the TS7720 and TS7740 systems will move to an End-of-Support (EOS) status on April 30, 2020. This includes these models and types:

- Processors: **3957-V07** and **3957-VEB**
- Disk Cache Controllers: **3956-CC9** and **3956-CS9**
- Disk Cache Drawers: **3956-CX9** and **3956-XS9**



End-of-Support (EOS)

- The last date on which IBM will deliver standard support services for a given version/release of a product
- Hardware support extensions are occasionally available. These extensions cover defect support, including existing microcode fixes and patches. The service also covers parts replacement, depending on availability. Extensions do not include:
 - Access to the TS7700 Product Field Engineering (PFE) or TS7700 Development teams for any issues with these machines
 - Corrections, fixes or patches for any new problems including security vulnerabilities

(Best Practice) Develop a plan to replace each TS7720 or TS7740 system with at TS7760, TS7760C or TS7760T system by April 20, 2020

- Most replacements will be done by joining new TS7760 systems into existing Grids, using the Copy Refresh function to move the data off of the TS7720/TS7740 systems, and then removing the older system
- “In-place” replacement (a.k.a. “frame roll”) of TS7740 or TS7720T system with a TS7760T can sometimes be done as well

General – Limit The Use Seven (7) and Eight (8) Cluster Grids



A TS7760 Grid can be:

- From one (1) to eight (8) clusters
- *Note:* Grids that will be more than five (5) clusters require an approved RPQ (Request for Price Quote)
 - Basically an exception process to allow the TS7760 development team to review the proposed configuration

TS7760 Grids with seven (7) and eight (8) cluster actually work very well. However, they can and do:

1. Limit the flexibility to add additional clusters (for data center moves, technology upgrades and application growth)
2. Increase the amount of intra-Grid (metadata) traffic on the replication links (to keep all clusters synchronized)
3. Boost the possibility that a cluster may become overloaded (especially with a poor Grid design)

(Best Practice) Use Seven (7) and Eight (8) Cluster Grids temporarily and sparingly

- For 7 and 8 cluster Grids, use the exception (RPQ) path to engage the TS7760 development team
- Use 7 and 8 cluster Grids only on a temporary basis
 - During technology upgrades (i.e. add new cluster to Grid, copy data, and remove existing cluster from Grid)
 - During data center relocations (i.e. add new cluster to Grid at “to” site, copy data, and removed existing cluster at “from” site from Grid)
- When there is a need to permanently expand to a 5 or 6 cluster Grid due to growth, consider creating a (new) second Grid instead

General – Allow for (Unplanned) Outages



TS7760 Systems do, in rare instances, encounter issues that impact the cluster's normal operations

- A TS7760 system does have a very large amount of built-in hardware redundancy
 - Examples: Twenty (20) processor cores, multiple Ethernet and FICON cards, distributed RAID (RAID-6) in the disk cache, redundant internal communication paths, and mirrored boot drives
- Network, internal issues with system software, and a few hardware components can cause a cluster-wide problem
 - Software: Operating system (AIX), virtual tape software (TS7700 code), and Db2
 - Hardware: System planer board
 - Network: Hardware failures combined with limited to no redundancy

Unplanned outages can occur during the critical System Z processing periods

- What will be the impact on mainframe operations if a tape outage occurs at the worst possible time?

(Best Practice) Consider having two (2) local TS7760 systems so that mainframe tape processing can continue even if one (1) cluster is unavailable due to an unplanned event

- A second option is to configure network “FICON channel extension” so that local System Z systems can access a remote TS7760 when the local virtual tape system is not available
- Both options will provide high availability eliminating the need for System Z processing interruptions

General – Create Two Physical Copies of (Almost) Every Virtual Volume



A TS7760 Grid can support up to:

- Twenty-four (24) physical copies being made of a single virtual volume
 - Requires an eight (8) cluster Grid with each cluster having its own back-end physical tape
 - Each volume is replicated all eight (8) clusters
 - Each cluster keeps one copy of the virtual volume in its disk cache and two (2) more on different physical tapes

While keeping 24 copies is mostly theoretical, keeping a copy in at least two (2) clusters will allow continued access when:

- A problem occurs with a single physical tape cartridge (for those systems using back-end tape)
- A problem occurs with a single TS7760 cluster
- A problem occurs with a site (if the second copy is kept in a TS7760 system in a different data center)

(Best Practice) Use Grid replication or dual physical tape to create (at least) two copies of virtual volumes

- Exceptions are temporary (retention period less than one day) volumes that can easily be recreated

BEST PRACTICES – CONFIGURATION

Config – Exercise Great Care with Grid Copy Consistency Points



Copy Consistency Point (CCP) settings are established within the TS7760 Management Classes

- The CCP settings determine:
 - Where (which clusters) volume copies should be placed
 - When each copy should be consistent with the originating data
- There are five (5) possible values for each CCP: D (Deferred), N (No Copy), R (Run), S (Sync) or T (Time Delayed)
- Establishing the CCP settings correctly is vital to achieving the desired business objectives for the TS7760 Grid and for the optimal functioning of the Grid activities
- As the number of clusters and Management Classes increase, so does the difficulty level of setting the CCPs correctly
 - **An 8 cluster Grid with 256 Management Classes will have 16,384 individual CCP settings**

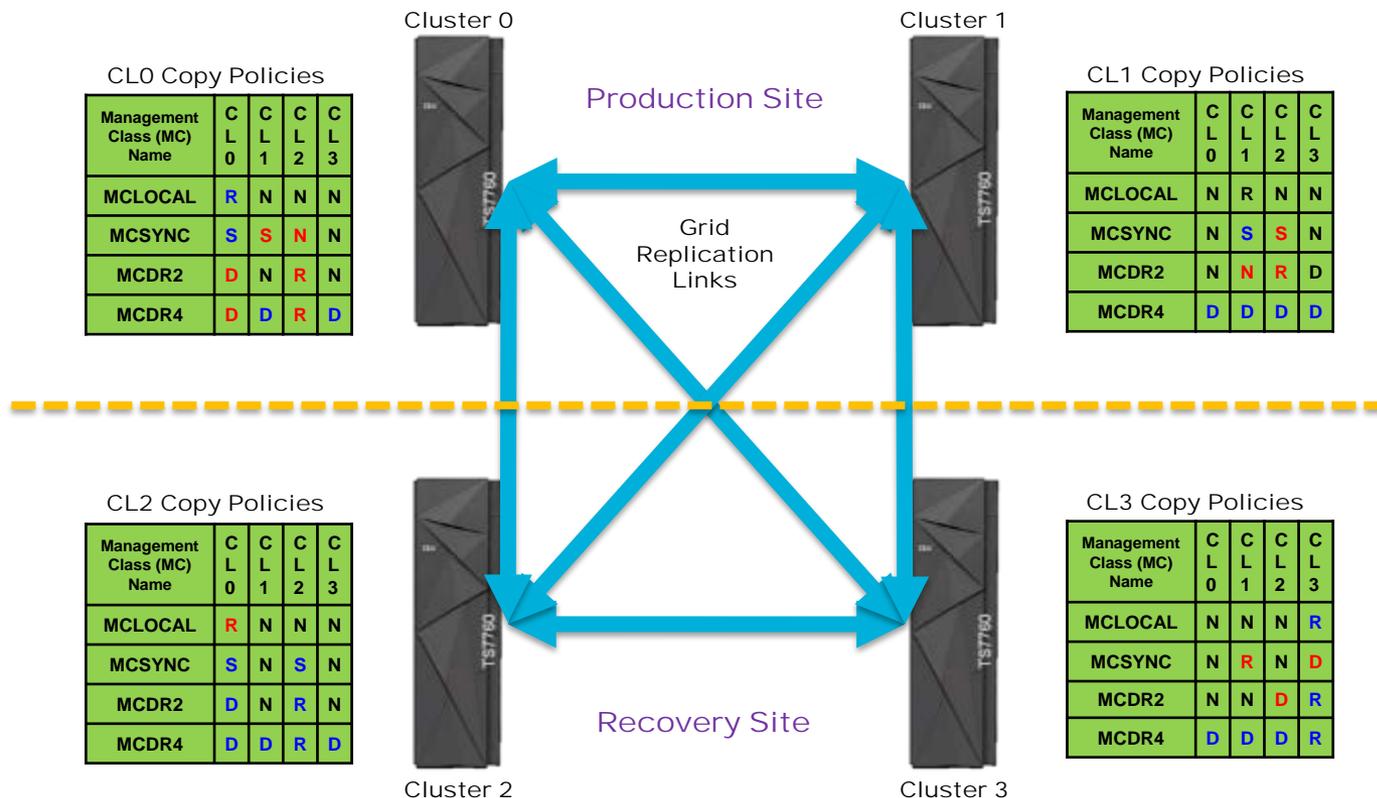
An error made while establishing or entering the CCP settings can be difficult to detect

- The Management Class names are the same across all clusters in a Grid, each set of CCPs must be set based on what actions are desired when volumes are written directly to each cluster
- The clusters will “do what they’re told” and incorrect CCP settings can be a challenge to identify
 - The “made up” example on the next page has at least eight (8) sets of errors

(Best Practice) Pre-develop a CCP plan, carefully review plan before implementing it, and re-check periodically

- Thorough peer reviews of the CCP settings on each cluster should be highly encouraged

Config – Example: Copy Consistency Policies With Multiple Errors



Config – Encrypt All Data-At-Rest (Disk, Physical Tape, and Cloud)

Modern security practices employ multiple layers to ensure only authorized users can access data



- These practices include: user authentication, access control, data-in-flight and data-at-rest protection
- For data-at-rest protection, TS7760 systems use encryption methodologies
 - Encryption is the process of encoding a message or information in such a way that only authorized parties can access it and those who are not authorized cannot
 - Encryption has minimal (<1%) to no impact on read and write throughput rates

TS7760 encryption capabilities

- *(Optional)* **Disk cache – Advanced Encryption Standard (AES)-256 encryption**
 - Internal or external (IBM Security Key Lifecycle Manager [SKLM]) key management
- *(Optional)* **Physical tapes (a.k.a. “stacking” tapes) – AES-256 encryption**
 - External (SKLM) key management
- *(Optional)* **AWS S3 cloud storage – AES-256 encryption**
 - AWS SSE-S3, SSE-KMS, or SSE-C key management
- *(Standard)* **IBM Cloud Object Storage – AES-256 and Secure Hash Algorithm (SHA)-256**
 - Default or SSE-C key management

(Best Practice) Encrypt all disk, tape and cloud data stores (i.e. data-at-rest) used by TS7760 systems

Config – Enable Secure Data Erase in Physical Volume Pools **

(** Applies to TS7720T, TS7740 and TS7760T systems only)

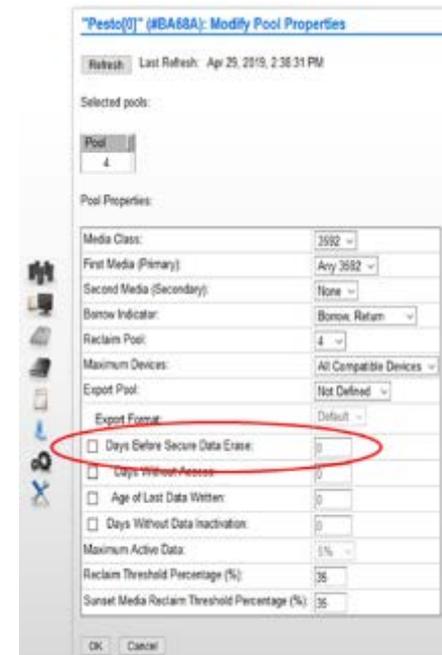
Expired data on a physical volume remains readable until the volume is overwritten with new data. Some clients prefer to delete the content of a reclaimed stacked cartridge, due to security or business requirements. The Secure Data Erase function can be used to delete data that is no longer valid



- The Secure Data Erase (SDE) function is enabled or disabled in each volume pool – the default is disabled
- When SDE is enabled on a pool, one must instruct the TS7700 on how many days to wait to perform the SDE function once a physical volume no longer contains any valid virtual volumes
 - From 1 to 365 days (similar to the Delete Expire Time [DET] setting for virtual volumes)

Two (2) Types of SDE

- For encrypted physical volumes – perform a cryptographic erase
 - Encryption keys used to write the data on the stacking tape are removed
 - Only takes a few minutes to mount the tape and remove the keys
- For non-encrypted volumes – perform a physical, long erase
 - A single full tape pass writes a repeating pattern to eradicate the previous tape contents
 - Takes a long time to complete (e.g. a “JD” cartridge may take four [4] hours or more to complete)



(Best Practice) Enable SDE for encrypted physical volumes. For volumes that are not encrypted, only enable SDE if there if there are enough available tape drives to support it

Config – Use New LZ4 and zSTD Compression Options

LZ3 has been the standard compression type used by all Virtual Tape Servers and TS7700 systems

- Currently known as “FICON compression” – same algorithm was used in ESCON cards
- For each virtual tape volume, the only option was to use it or not use it (DCB=TRTCH=NOCOMP)
- Method typically achieved around **3-to-1** compression (i.e. 300 GB of user data compressed to 100 GB of stored data)



TS7700 R4.1.2a introduced two (2) additional compression options: LZ4 and zSTD – both are done in the TS7700 processor

- All clusters in a Grid must be running the R4.1.2a or higher code level – no additional licenses needed
- New compression options used going forward; existing virtual tape volumes are not recompressed
- **LZ4**: typical compression rate of **4-to-1**, modest CPU overhead, minimal performance impact
- **zSTD**: typical compression rate of **5-to-1**, additional CPU cycles, can decrease single-stream performance by 5 to 10%

(Best Practice) Use zSTD compression for all new virtual tape volumes; for specific instances where performance is critical, revert these volumes back to using LZ4 compression

- Potential positive impacts
 - Decreased disk cache, physical tape and cloud storage usage
 - Lower Grid replication bandwidth requirements (less data to replicate)
 - Improved D/R Recovery Point Objectives (RPOs – fewer or smaller volumes replicate faster)
- When zEnterprise Data Compression (zEDC) is being used on a workload, update the corresponding DFSMS DataClass setting to COMPACTION=NO so that TS7760 doesn't attempt to compress the already compressed data streams

Config – Enable and Use the New Grid Resiliency Functions

Think of a situation where four (4) horses are in a team pulling a very heavy wagon



- The “Best Case” is that all 4 horses are working well and pulling the wagon equally
- The “Not as Good Case” is that one (1) horse has thrown a shoe and has been taken out of the team
- The “Bad Case” is that one (1) horse has thrown a shoe, is still in harness, and is slowing down the rest of the team

The TS7760 R4.1.2a release introduced new functions to address the “bad case” where a cluster is no longer doing its part in the team (i.e. Grid) and is slowing all virtual tape processing down

- The new functions allow the “good” team members to notify your operations team about the “laggard” or to automatically take the cluster having problems out of the Grid (known as “fencing” the cluster) . An automatic reboot of the cluster having problems can be initiated as well
- Most of the new functions are disabled by default (exception: ability to “fence” a local cluster from the Web Interface)
- To use all of the new functions, several decisions about timing (i.e. how long to wait) and autonomy (i.e. notify or automatically take action) must be made
 - The commands to activate the new functions are “on the fly” and can be entered at any time (set and forget)

(Best Practice) Enable and use the new TS7760 Grid resiliency functions

- Even if the decision is just enable them for additional notifications, this should help your operations team identify a “problem” cluster more quickly

Config – Use “Migrations Made Easy” Offering for Grid Changes

TS7700 “Migrations Made Easy” is a new fee-based offering from the IBM Systems Lab Services (LBS) team which substantially automates the Copy Refresh process used to copy or move data between TS7700 clusters



- A new method to simplify hardware refreshes and changes -- alternative to existing host based method
- Runs on the TS7700 Grid, no host access required
- Most migrations will be done remotely by LBS personnel using the remote support (“call-in”) function proved by the TSSC (TS3000 System Console)
- Utilizes existing TS7700 grid functions – same as ones used by the host based method
- Separate from TS7700 code, can be installed on a cluster independently
- Customer can see results of migration with existing interfaces or by asking for updates from LBS
- Customer provides requirements; LBS executes the migration. The requirements/limitations include:
 - Which and how many cluster migrations are running concurrently
 - Priorities may be specified by Management Class, category, volume serial range
 - Rate of migration:
 - ✓ Limit number of copies in queue
 - ✓ Block out windows for no migration copies

(Best Practice) Use the “Migrations Made Easy” offering to automate virtual volume migrations related to Grid changes

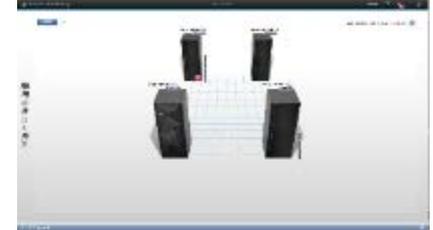
- Contact your IBM Business Partner or IBM Storage Sales representative to get further information including availability, requirements, and pricing

BEST PRACTICES – OPERATIONS

Operations – Use TS7760 Alert Settings and z/OS for Monitoring

Each TS7760 can be set to send console alert messages when certain thresholds are exceeded

- Message: “CBR3750I Message from library distributed library name: Message Text”
- Alerts are set/changed through “**LI REQ,dlibname,SETTING,ALERT,xxxxxxx**” commands
- Thresholds monitored (a subset of what is available)
 - **COPYHIGH** – Uncopied Data High Warning Limit
 - **RESDHIGH** – Resident Data High Warning Limit
 - **PCPYCRIT** – Pending Copy Critical

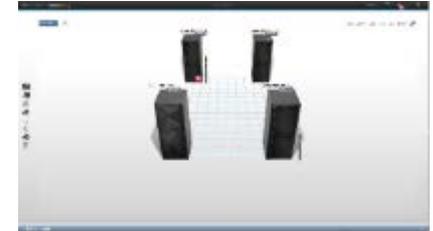


z/OS automation can be used to monitor and alert on even more TS7760 components

- Based on scripts that will issue Library Request commands and interpret the results
- Examples:
 - Replication queue size (by cluster): **LI REQ,dlibname,COPY,SUMMARY**
 - Replication link status and error rates: **LI REQ,dlibname,STATUS,GRIDLINK**

(Best Practice) Determine and set the TS7760 alert settings to insure Grid replication is working properly; consider implementing z/OS automation to provide additional proactive monitoring and alerting

Operations – Library Request Examples for z/OS Monitoring



LI REQ,dlibname,COPY,SUMMARY

```

COPY SUMMARY V3 .0
DISTRIBUTED LIBRARY PESTO CLUSTER 0
COPY COUNT BY ORIGINATING CLUSTER BY COPY QUEUE
CLUSTER 0 1 2 3 4 5 6 7 *
RUN NA 0 NC 0 NC NC 0 NC 0
SYNCDEF NA 0 NC 0 NC NC 0 NC 0
IMEDDEF NA 0 NC 0 NC NC 0 NC 0
DEF NA 1353 NC 0 NC NC 0 NC 0
CPYRFSH NA 0 NC 0 NC NC 0 NC 0
FAMDEF NA 0 NC 0 NC NC 0 NC 0
TDEL NA 0 NC 0 NC NC 0 NC 0

COPY SIZE GB BY ORIGINATING CLUSTER BY COPY QUEUE
CLUSTER 0 1 2 3 4 5 6 7 *
RUN NA 0 NC 0 NC NC 0 NC 0
SYNCDEF NA 0 NC 0 NC NC 0 NC 0
IMEDDEF NA 0 NC 0 NC NC 0 NC 0
DEF NA 1170 NC 0 NC NC 0 NC 0
CPYRFSH NA 0 NC 0 NC NC 0 NC 0
FAMDEF NA 0 NC 0 NC NC 0 NC 0
TDEL NA 0 NC 0 NC NC 0 NC 0

COPY COUNT BY COPY TASK BY COPY QUEUE
TASK COPY RECALL DELAY TDEL
RUN 0 0 0 0 0
SYNCDEF 0 0 0 0 0
IMEDDEF 0 0 0 0 0
DEF 1353 0 0 0 0
CPYRFSH 0 0 0 0 0
FAMDEF 0 0 0 0 0
TDEL 0 0 0 0 0

COPY SIZE GB BY COPY TASK BY COPY QUEUE
TASK COPY RECALL DELAY TDEL
RUN 0 0 0 0 0
SYNCDEF 0 0 0 0 0
IMEDDEF 0 0 0 0 0
DEF 1170 0 0 0 0
CPYRFSH 0 0 0 0 0
FAMDEF 0 0 0 0 0
TDEL 0 0 0 0 0

LONGEST COPY QUEUE AGE : 7191
LONGEST TIME DELAYED TASK : 0
    
```

LI REQ,dlibname,STATUS,GRIDLINK

```

GRIDLINK STATUS V2 .1
CAPTURE TIMESTAMP: 2019-03-05 22:26:33
LINK VIEW
LINK NUM CFG NEG READ WRITE TOTAL ERR LINK STATE
MB/S MB/S MB/S 01234567
0 10000 10000 5.6 0.3 5.9 0 -A A A
1 10000 10000 4.9 0.0 4.9 0 -A A A
2 0 0 0.0 0.0 0.0 0
3 0 0 0.0 0.0 0.0 0

LINK PATH LATENCY VIEW
LIBRARY LINK 0 LINK 1 LINK 2 LINK 3
LATENCY IN MSEC
PESTO 0 0 0 0
SQUINT 0 0 0 0
TOM 0 0 0 0
JERRY 0 0 0 0

CLUSTER VIEW
DATA PACKETS SENT: 152766
DATA PACKETS RETRANSMITTED: 0
PERCENT RETRANSMITTED: 0.0000

LOCAL LINK IP ADDRESS
LINK 0 IP ADDR: 11.12.163.13
LINK 1 IP ADDR: 11.12.164.13
LINK 2 IP ADDR
LINK 3 IP ADDR
    
```

Operations – Collect and Save TS7760 Capacity and Performance Data



Capacity usage and performance data is collected/saved every fifteen (15) minutes by each TS7760

- The last ninety (90) days (on a rolling forward basis) is kept internally within each system
- The internal time of all clusters in a Grid are kept synchronized with each other
 - All clusters will perform their data collection at the same time

Use the Bulk Volume Information Retrieval (BVIR) function to create a copy of the saved usage and performance data

- Data from all (both local and remote) of the clusters in the Grid will be collected with a single job
- Two (2) Options
 - **POINT IN TIME STATISTICS** (data from the 15 minute just before the command was processed)
 - **HISTORICAL STATISTICS FOR xxx-yyy** (collect up to 90 days of data; xxx & yyy are the Julian start and end dates)

(Best Practice) Periodically (once a month) run the BVIR function on a z/OS or z/VSE LPAR to collect the historical statistics for the last thirty-one (31) days

- Save the output data for at least one (1) year; two (2) or more years will be better
- Don't be concerned about the number of days in each month
 - Better to have overlapped days in the data than data gaps
 - The VEHSTATS report program (part of the IBM Tape Tools package) will automatically remove any duplicate data
- Useful for spotting trends or for performing processing comparisons (problem diagnosis) between two (2) points in time

Operations – Establish A Unique Management Class for each TS7760



Capacity usage and performance data is collected/saved every fifteen (15) minutes by each TS7760

- Bulk Volume Information Retrieval (BVIR) command volumes must be processed on a specific cluster
- Creating a specific Management Class (MC) for each cluster allows easy target selection
- MC names are customer selectable (i.e. one doesn't have to use the names shown in the example)
- If a LPAR has access to multiple Grids, create a unique MC name for all of the clusters in these Grids

Example (for a four [4] cluster Grid)

CL0 Copy Policies					CL1 Copy Policies					CL2 Copy Policies					CL3 Copy Policies				
Management Class (MC) Name	C0	C1	C2	C3	Management Class (MC) Name	C0	C1	C2	C3	Management Class (MC) Name	C0	C1	C2	C3	Management Class (MC) Name	C0	C1	C2	C3
MCBVIRC0	R	N	N	N	MCBVIRC0	R	N	N	N	MCBVIRC0	R	N	N	N	MCBVIRC0	R	N	N	N
MCBVIRC1	N	R	N	N	MCBVIRC1	N	R	N	N	MCBVIRC1	N	R	N	N	MCBVIRC1	N	R	N	N
MCBVIRC2	N	N	R	N	MCBVIRC2	N	N	R	N	MCBVIRC2	N	N	R	N	MCBVIRC2	N	N	R	N
MCBVIRC3	N	N	N	R	MCBVIRC3	N	N	N	R	MCBVIRC3	N	N	N	R	MCBVIRC3	N	N	N	R

(Best Practice) Create a specific Management Class for each cluster in a Grid or Grids to allow for easy routing of BVIR command volumes

Operations – Use RSYSLOG for Access and Change Tracking

Rsyslog

RSYSLOG (Remote System Log) Processing Support

- Also known as *Rocket-fast* System for Log processing
- Starting with code level R4.1.2, a TS7760 can send log records to one (1) or two (2) RSYSLOG servers
- RSYSLOG provides for enhanced auditing, reporting and diagnostic capabilities
- A TS7760 only sends the log records – each customer must supply a RSYSLOG server or destination to receive, process and store the log records

Types of log records generated and sent:

- Kernel, user-level and host-destined messages
- Security and authorization (logon/logoff) events
- Call-home activities

(Best Practice) Establish one (1) or two (2) RSYSLOG servers and enable each TS7760 to send log records to it (them)

- Consider using unique Userids for each TS7760 operator and administrator to allow tracking events back to individuals
- The TS7760 systems also support the use of SNMP for logging
 - Customers can use RSYSLOG, SNMP, both RSYSLOG and SNMP, or neither option

Operations – Use Remote Code Load Option for TS7760 Code Upgrades

Remote Code Load (RCL) is a service offering which allows code updates to be performed by remote support engineers, as opposed to an on-site Support Services Representative (SSR)



- Now the preferred code delivery method for providing code upgrades on entitled storage products
- Avoids the need to plan and schedule a SSR to be available at the client's site
- Optional for TS7760 upgrades, on-site SSR code upgrades will still be available
- Client requirements:
 - Confirm which systems require RCL (Remote Code Load); select which code level is intended for the upgrade
 - Schedule remote code update via booking tool or contact the RCL team via IBM Support ticket
 - Arrange for on-site SSR to replace any failed hardware components prior to the upgrade
 - Allow on-site SSR to configure AOS (Assist On-Site) / remote support center function for remote access to the TS7760's TSSC (TS3000 System Console)
 - Ensure that the systems are available for the code upgrade (i.e. host jobs and tasks have been stopped)

(Best Practice) Use the Remote Code Load (RCL) option for TS7760 code level upgrades

- Doesn't have to be all or nothing choice – you can choose which systems use RCL or on-site SSR upgrades

BEST PRACTICES – NETWORK

Network – Provide Enough Network Bandwidth for TS7760 Replication

There are two (2) levels of communications occurring between TS7760 systems in a Grid

- **Messaging:** volume status updates, health check inquiries/responses, and performance checks
- **File Transfers:** pull copies of virtual volumes from one cluster to another

■ Considerations

- Reads and writes to cloud storage will also now use the TS7760 Grid replication ports
- IBM's TS7760 sizing tools assume that the Grid replication links will have access to the full network bandwidth supported by the installed Ethernet adapters (e.g. 2 Gb/sec for twin 1 Gb links, 40 Gb/sec for four 10 Gb links)
- Customers have found that it is possible to provision less than the full supported bandwidth while still meeting the overall business objectives
 - Often a trade-off between bandwidth and D/R Recovery Point Objectives (i.e. volumes take longer to replicate)
- Provisioning too little bandwidth can severely impact production tape mounts (could fail or be very slow)
 - The messaging traffic listed above takes place on a continuous basis between all of the clusters in a Grid
 - The smooth functioning of a Grid depends on all of the clusters being able to quickly communicate with each other

(Best Practice) If a TS7700 is part of a Grid, provide full network bandwidth to each of the Grid links. If this is not possible, design the network to provide a guaranteed amount of reduced bandwidth

- The minimum bandwidth needed will vary based on the number of clusters and the overall mount activity
- A one (1) Gb/sec minimum is a good starting point



Network – Enable Network Flow Control on the Ethernet Grid Ports

Ethernet flow control is a mechanism for temporarily stopping the transmission of data on Ethernet family computer networks. The goal of this mechanism is to ensure zero packet loss in the presence of network congestion



- Most vendors of Ethernet switches and directors DO NOT enable flow control by default – it must be explicitly enabled on each port where the function is desired
- TS7760 Grid replication works best (i.e. fastest throughput and minimized number of data re-transmissions) with flow control enabled

(Best Practice) Work with your data network team to enable bi-directional flow control on each of the Ethernet Ports that will be used for TS7760 Grid replication traffic

- Cisco supports two (2) types of flow control: LLFC (Link Level Flow Control) and PFC (Priority Flow Control)
 - A port can only use one type – use LLFC for TS7760 Grid connections
 - Cisco switch commands to disable PFC and enable LLFC:
 - priority-flow-control mode off
 - no lldp transmit
 - no lldp receive
 - flowcontrol receive on
 - flowcontrol send on

BEST PRACTICES – DISASTER RECOVERY

Disaster Recovery – Use the FlashCopy for D/R Function for Testing

The TS7760's FlashCopy for Disaster Recovery (D/R) function facilitates regular recovery testing

- Periodic tests of D/R capabilities is necessary to verify that the data required for a recovery is offsite and that the recovery processes are still valid
- FlashCopy capability is built into all TS7760/TS7760C/TS7760T clusters – nothing to license



The FlashCopy for D/R function creates a snapshot of the live virtual tape volumes at time zero

- Time zero is when the flash copy is taken. It mimics the time when a real disaster happens. Customers can establish the time zero using a LI REQ command
- Recovery tasks can read, but not change or delete, the production virtual tape volumes that were active at time zero
- Recovery tasks can read, change and delete a temporary set of virtual tape volumes (removed at the end of a test)
- If production processing changes or deletes a live volume the D/R cluster(s) will create a copy of the original volume

Additional disk cache space may be required on the D/R clusters, especially if they are (disk-only) TS7760 models

- To hold temporary data from recovery testing and any volume copies that were made due to production processing

(Best Practice) Setup and use the FlashCopy for Disaster Recovery (D/R) function for regular recovery testing

- Can be setup and used “on the fly” with no outages or replication pauses on any of the clusters in the Grid

Disaster Recovery – Change COPYFSC Setting on D/R Clusters



Copies To Follow Storage Class (COPYFSC) Preference Setting

- Set or changed through a LIBRARY REQUEST command
 - **LIB REQ,xxxx,SETTING,CACHE,COPYFSC,ENABLE** (or **COPYFSC,DISABLE**)
 - Command can be entered through the z/OS console or the TS7760 (Web) Management Interface
- When the ENABLE keyword is specified, logical volumes copied into this cluster's disk cache from a peer TS7700 are managed using the preference group actions defined within this target cluster's Storage Class construct definition
- When the DISABLE keyword is specified, logical volumes copied into this cluster's disk cache from a peer TS7700 are always managed as PGO (prefer to be migrated from cache)
- The default is disabled

For TS7760C (cloud-attach) or TS7760T (tape-attach) clusters being used as targets for disaster recovery copies, staying with the COPYFSC default (DISABLE) can lead to the premature removal of volumes from the clusters' disk cache and slower disaster recovery times

(Best Practice) On any TS7760C/TS7760T cluster being used as a disaster recovery target, change the COPYFSC setting to ENABLE

SUMMARY

Twenty-one (21) TS7760 “Good Things” ...

• General

- KISS your grids!
- Design for (planned) changes
- Remember TS7720/TS7740 End-of-Support
- Limit use of 7 and 8 cluster Grids
- Allow for (unplanned) outages
- Create 2 physical copies of (almost) every virtual volume

• Configuration

- Be careful with Copy Consistency Point (CCP) settings
- Encrypt all data-at-rest
- Enable Secure Data Erase
- Use new LZ4 and zSTD compression
- Use the new Grid Resiliency functions
- Use “Migrations Made Easy” Offering for Grid changes

• Operations

- Use TS7760 alert settings and z/OS automation for monitoring
- Collect and save TS7760 capacity and performance data
- Create a Management Class for each TS7760
- Use RSYSLOG for access/change tracking
- Use the Remote Code Load option for TS7760 code upgrades

• Network

- Provide sufficient network for TS7760 Grid
- Enable network flow control for TS7760 Grid

• Disaster Recovery

- Use FlashCopy for D/R for regular testing
- Change COPYFSC setting on D/R Clusters

Useful TS7700 Reference Documents

- **(Redbook) IBM TS7700 Release 4.2 Guide (SG24-8366-02)**
<http://www.redbooks.ibm.com/abstracts/sq248366.html?Open>
- **(Redpaper) IBM TS7760 R4.2 Cloud Storage Tier Guide (REDP-5514-00)**
<http://www.redbooks.ibm.com/abstracts/redp5514.html?Open>
- **(Knowledge Center) IBM TS7700 – Customer Documentation**
https://www.ibm.com/support/knowledgecenter/en/STFS69/landing/welcome_002.htm
- **(White Paper) IBM TS7700 Series z/OS Host Command Line Request User's Guide**
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101091>
- **(White Paper) IBM TS7700 Full Disk Encryption (FDE)**
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102200>
- **(White Paper) IBM TS7700 Tape Encryption Overview**
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101000>
- **(White Paper) IBM TS7700 Series Grid Resiliency Improvements**
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102742>
- **(White Paper) IBM TS7700 Series Bulk Volume Information Retrieval Function User's Guide**
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101094>
- **(White Paper) IBM TS7700 Best Practices - FlashCopy for DR Testing V1.5**
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102415>

 = Useful for daily operations

THANK YOU!

Disclaimers

- Copyright© 2019 by International Business Machines Corporation.
- No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.
- The performance data contained herein were obtained in a controlled, isolated environment. Results obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. These values do not constitute a guarantee of performance. The use of this information or the implementation of any of the techniques discussed herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into their operating environment. Customers attempting to adapt these techniques to their own environments do so at their own risk.
- 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. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- 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 intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any on-IBM product, program or service.

Disclaimers (continued)

- 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 NONINFRINGEMENT.
- 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 interoperability of any non-IBM products discussed herein.
- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products 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 terms are trademarks or registered trademarks of the IBM Corporation in either the United States, other countries or both.
 - IBM, Power System, Power8, z Systems
 - z/OS, z/VM, VM/ESA, OS/390, AIX, DFSMS/MVS, OS/400, i5, FICON, ESCON
 - VSE/ESA, TPF, DFSMSdfp, DFSMSdss, DFSMShsm, DFSMSrmm
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Other company, product, and service names mentioned may be trademarks or registered trademarks of their respective companies.