



# Washington Systems Center - Storage

## Accelerate with IBM Storage: TS7700 Grid Joins, Merges and Removals

### Presenters

Ben Smith  
Bill Danz

### Panelists

Mike Berish  
Toy Phouybanhdyt



## Accelerate with IBM Storage Webinars

### The Free IBM Storage Technical Webinar Series Continues in 2018...

*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](mailto: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/playlist?list=PLSdmGMn4Aud-gKUBCR8K0kscCiF6E6ZYD&disable\\_polymer=true](https://www.youtube.com/playlist?list=PLSdmGMn4Aud-gKUBCR8K0kscCiF6E6ZYD&disable_polymer=true)



#### 2018 Webinars:

- January 9 - DS8880 Easy Tier
- January 17 - Start 2018 Fast! What's New for Spectrum Scale V5 and ESS
- February 8 - VersaStack - Solutions For Fast Deployments
- February 16 - TS7700 R4.1 Phase 2 GUI with Live Demo
- February 22 - DS8880 Transparent Cloud Tiering Live Demo
- March 7 - Spectrum Storage Management, Control, Insights, Foundation; what's the difference?
- March 15 - IBM FlashSystem A9000/R and SVC Configuration Best Practices
- March 27 - IBM FlashSystem A9000/R Technical Update
- April 12 - Introducing Spectrum NAS - The Newest Member in the Spectrum Storage Family
- April 26 - TS7700 Grid Configuration Changes - Joins, Merges and Removals
- May 8 - DS8880 Technical Update

**Register Here:** <https://ibm2.webex.com/ibm2/onstage/g.php?MTID=e5a4a631f2c718b34c9c3ec4e0f508672>

## Webinar Notice

---

**NOTICE:** This webinar is not meant to be a complete and comprehensive “how to” guide on the steps required to: join a cluster to a Grid, merge two (2) Grids together, or remove a cluster from a Grid

**In practice, the local IBM SSRs (System Service Representatives) and IBM Tape “Top Guns” will perform the actual steps to re-configure a Grid.** When planning for one of these activities, engage them as soon as is practical as they can provide the most accurate and up-to-date information based on their knowledge of your TS7700 environment

The primary purpose of this webinar is to provide a high-level explanation of what occurs during the join, merge and removal operations along with information on the pre-requisites and post-requisites for each activity. Our goals are to:

- Remove some of the mystery about the TS7700 join, merge and removal operations
- Increase your awareness of what is and what is not possible
- Decrease delays due to missing pre-requisites – these are often data network related
- Decrease clean-up activities caused by overlooked updates to TS7700 policy settings

## Webinar Overview

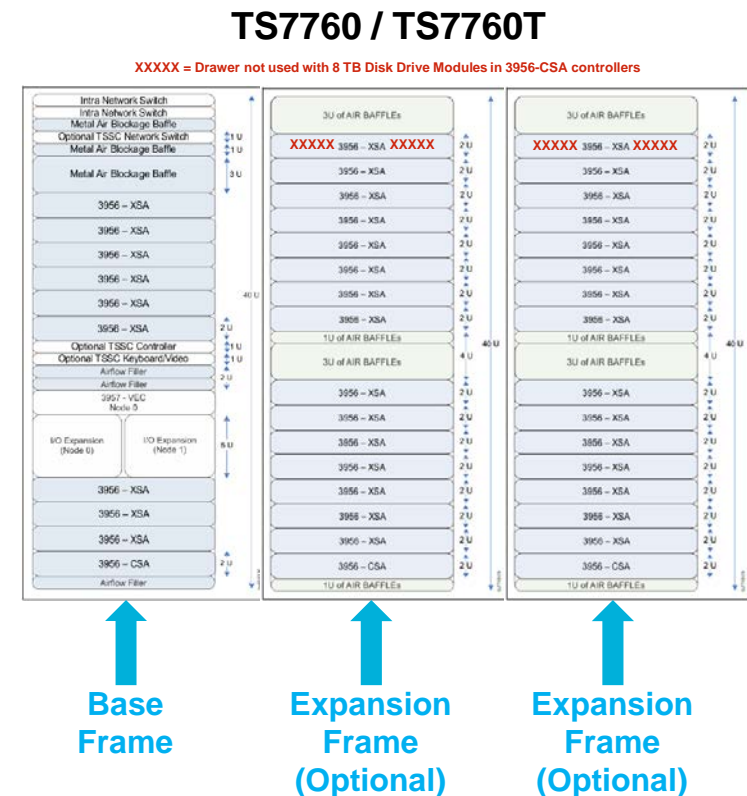
---

Most TS7700 systems that have been installed for few years share a common attribute: at some point in its existence, the cluster has gone through a hardware upgrade or a configuration change. Reconfiguring a TS7700 Grid is an infrequent, but critical change and is the subject of this webinar.

- **Section 1 – TS7700 Grid: Cluster Joins (Bill)**
- Section 2 – TS7700 Grid: Grid Merges (Bill)
- Section 3 – TS7700 Grid: Cluster Removals (Ben)
- Section 4 – TS7700 Grid: Copy Audit and Copy Refresh (Ben)
- Wrap-up and Q&A (Ben, Bill, Mike and Toy)

## TS7700 Grid: Introduction to Clusters

- A **TS7700 cluster** combines the TS7700 server with one or more external (from the server's perspective) disk subsystems. It includes these components:
  - One Tape Base Frame with:
    - One TS7700 Server
    - Two TS7700 Input / Output (I/O) Expansion Drawers (primary and alternate)
    - One TS7700 Cache Controller Drawer with up to nine (9) Cache Expansion Drawers
    - Two Ethernet switches (internal, not for Grid)
    - One TS3000 Total System Storage Console (TSSC)
  - (Optional) One or two optional Storage Expansion Frames with:
    - One TS7700 Cache Controller Drawer with up to fifteen (15) cache drawers
  - (Optional) A TS3500 or TS4500 tape library with:
    - 4 - 16 IBM 3592 tape drives and two Fibre Channel (FC) switches




## TS7700 Grid: Cluster Models

---

- Two (2) current configurations
  - **TS7760** (3957-VEC) – disk only
  - **TS7760T** (3957-VEC) – disk with physical tape library, drives and cartridges
- Three (3) previous (supported) configurations
  - **TS7720** (3957-VEB) – disk only
  - **TS7720T** (3957-VEB) – disk with physical tape library, drives and cartridges
  - **TS7740** (3957-V07) – disk with physical tape library, drives and cartridges
- Note: Two (2) older models TS7720 (3957-VEA) and TS7740 (3956-V06) have reached their End of Support (EOS) dates

## TS7700 Grid: Introduction to Grids

- A **TS7700 Grid** refers to two or more physically separate TS7700 Clusters connected to one another by using a (customer-supplied) Internet Protocol (IP) high-speed data network
  - The TCP/IP infrastructure that connects a TS7700 Grid is known as the Grid Network
  - The term grid refers to the code and functionality that provides replication and management of virtual volumes and their attributes in cluster configurations
  - A Grid can be used to form disaster recovery and high availability solutions
    - A disaster recovery solution is achieved when multiple clusters are geographically distant from one another
    - A high availability solution is achieved when multiple clusters are in close proximity to one another
  - All TS7700 systems begin their “life” as Grid “capable”
    - With the addition of Feature Code 4015 they become Grid “enabled”
    - When joined to one or more other TS7700 systems, each cluster becomes a Grid “member”



## TS7700 Grid: Cluster Requirements for Grid (All clusters)

---

- TS7760, TS7740 and TS7720 systems

- **Feature Code 4015** (Grid Enablement) – enables communication function to allow communication with other TS7700s in a Grid

- TS7760 systems

- Two **Feature Code 1036** (1Gb Grid Dual Port Copper Connection) Ethernet cards, or
- Two **Feature Code 1037** (1Gb Grid Dual Port Optical SW Connection) cards, or
- Two **Feature Code 1038** (10Gb Dual Port Grid Optical LW Connection) cards

And (Optionally – to increase from 2 to four supported connections)

- One **Feature Code 1034** (Enable Dual Port Grid Connection)

- TS7720 and TS7740 systems

- Two **Feature Code 1036** (1Gb Grid Dual Port Copper Connection) Ethernet cards, or
- Two **Feature Code 1037** (1Gb Grid Dual Port Optical SW Connection) cards, or

And (Optionally – to increase from 2 to four supported connections)

- One **Feature Code 1034** (Enable Dual Port Grid Connection)

Or (TS7720 and TS7740 systems do not support four 10 Gb connections)

- Two **Feature Code 1035** (10Gb Single Port Grid Optical LW Connection) cards



## TS7700 Grid: Single Cluster Grids

---


- Every Grid in existence today began its “life” as a single TS7700 cluster
- While a single cluster is not an official Grid, the Grid capability can be activated at any time or the cluster can remain standalone indefinitely
- Most standalone clusters are TS7720T, TS7740 or TS7760T configurations
  - This allows for data to be exported from the TS7700 using physical tape cartridges

### Site “A”




## TS7700 Grid: Multi-cluster Grids

- Up to seven (7) clusters can be added (one at a time) to the original cluster to form a multi-cluster Grid
- If the Grid consists of only (2) clusters and both are located in the same area, it is possible to directly connect the two systems without Ethernet switches
  - Both clusters must have support the same capacity and type of Ethernet connection (1 vs. 10 Gb, copper vs. optical)
  - The direct option is not recommended – must move to switches to add a third cluster



## TS7700 Grid: Growing a Grid Through Joins and Merges


- Grids can grow either by joining an empty cluster with an existing cluster or Grid or by merging two existing Grids together (each with their own data)
- Cluster merge operations are relatively rare
  - The requirements for a merge are quite stringent and one of the Grids must remain completely offline during the merge
  - More on the merge “rules of the road” later on in Section 2 of the presentation



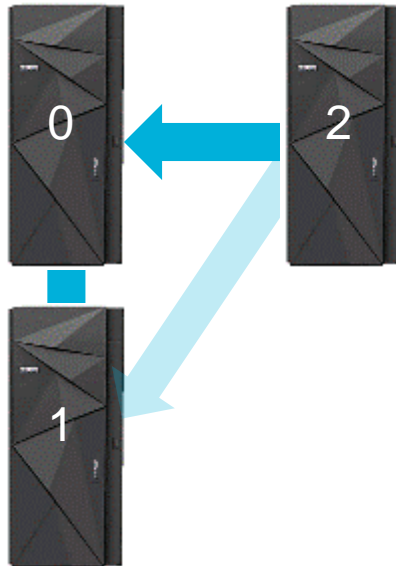
## TS7700 Grid: Cluster Join – Overview

- **Join** - the process that is performed when an empty cluster is joined to another cluster to create or grow a **Grid**. The empty cluster is referred to as the joining cluster. The cluster that is the target of the join is referred to as the existing cluster. The existing cluster might be a new empty cluster, an existing stand-alone cluster, or a cluster that is a member of an existing grid


Join Cluster 1 to Grid



Join Cluster 2 to Grid



Join Cluster 3 to Grid



- When one cluster is joined to a cluster in an existing grid, all other clusters in the existing grid are automatically joined to the new cluster as well

## TS7700 Grid: Cluster Join – Allowable Code Level Combinations

		Go-to code level								
		R1.7	R2.0	R2.1	R3.0	R3.1	R3.2	R3.3	R4.0	R4.1.1/2
Existing code level	R1.7	y	y	y	y <sup>2)</sup>	y <sup>2)</sup>	y <sup>2)</sup>	y <sup>2)</sup>	y <sup>3)</sup>	y <sup>3)</sup>
	R2.0	n	y	y	y	y <sup>2)</sup>	y <sup>2)</sup>	y <sup>2)</sup>	y <sup>3)</sup>	y <sup>3)</sup>
	R2.1	n	n	y	y	y	y	y	y	y
	R3.0	n	n	n	y	y	y	y	y	y
	R3.1	n	n	n	n	y	y	y	y	y
	R3.2	n	n	n	n	n	y	y	y	y
	R3.3	n	n	n	n	n	n	y	y	y
	R4.0	n	n	n	n	n	n	n	y	y
	R4.1.1/2	n	n	n	n	n	n	n	n	y

1) No more than two code levels can be active across all clusters of a grid at any time. Starting with R3.2 up to three code levels are allowed in a grid (with a minimum of R2.1)

2) RPQ/SCORE required; runtime needs to be minimized (e.g. a few weeks only), the intent is to drain the 1.7 cluster(s) and remove from the grid.

3) RPQ/SCORE required to upgrade the R1.7 machine to R2.1 without 8GB Memory MES; runtime should be minimized (e.g. a few weeks only), the intent is to drain the 1.7 cluster(s) and remove from the grid.

Source: TS7700 Code, Hardware and Grid Interoperability V4.1.2.pdf at <http://w3-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD106064>

## TS7700 Grid: Cluster Join – Requirements (1 of 2)

---

- In the documentation used by the IBM SSR for a cluster join operation, the existing, online cluster is referenced as *cluster (x)* and the joining, offline cluster is referenced as *cluster (j)*. Typically, *x* is 0 and the first time through this procedure, *j* is 1, the second time, *j* is 2, and so on
- These are the requirements listed in the join documentation used by the SSRs:
  1. Cluster joins at different TS7700 microcode levels are supported when the existing cluster (x) is at 8.21.x.x or later
  2. The final resulting grid configuration can not have more than three different microcode levels. The microcode level used on a cluster is shown on the login screen
  3. The joining cluster (j) must not contain any logical volumes (i.e. it must be empty)
  4. If cluster (x) is a member of a grid then it must be the latest microcode level of any member in the grid
  5. The joining cluster (j) must be at an equal or later code level than existing cluster (x)
  6. The joining cluster (j) and existing cluster (x) must both have feature code 4015 (Grid Enablement) installed
  7. The joining cluster (j) must have the same number or more of feature code 5270 (Increased Logical Volumes) if cluster (x) has it installed

## TS7700 Grid: Cluster Join – Requirements (2 of 2)

---

- Additional requirements listed in the join documentation used by the SSRs:
  8. The joining cluster (j) must contain feature code 5271 (Selective Device Access Control) if cluster (x) has it installed
  9. If the joining cluster (j) has feature code 1035 (VEB/V07: 10 Gb LW Single Port Optical) or 1038 (VEC: 10 Gb LW Dual Port Optical) installed, the customer infrastructure (i.e. ports on Ethernet switches) must support 10 Gb
  10. If the joining or target cluster has FC 1035 or 1038, then either of the following apply:
    - If the grid will be direct-attached, both clusters must have FC 1035 or 1038 installed
    - If the grid will not be direct-attached, the customer's network infrastructure must support 10 Gb
  11. Schedule this activity during a low activity time on the existing cluster (x) or grid
  12. If the existing cluster has FC 0001 (25 GB Logical Volume Enablement) installed and the joining cluster is at 8.32.x.x or later, the joining cluster does not require the feature code
  13. If the joining cluster has FC 5275 (Additional Logical Devices) installed and it will be joined to clusters below 8.32.0.x, then the additional logical devices on the joining cluster will not be accessible to the host until all clusters in the grid are at 8.32.0.x or later
  14. Prior to the join process, the FICON cables need to be unplugged on joining cluster (j)

## TS7700 Grid: Cluster Join – Preparation is the Key to Success (1 of 2)

---

- If a join operation will result in Grid with more than four (4) clusters in it, an approved no charge, SCORE (Storage Customer Opportunity REquest) RPQ is required. An IBMer or IBM Business Partner can initiate this request. The purpose of this SCORE RPQ is to allow the TS7700 development team to review the proposed Grid design to insure that both the join and normal Grid operations will work correctly
- The new, joining cluster must have the capability for two-way TCP/IP communications across the Grid links to **ALL** of the clusters in the existing Grid. Where copies of logical volumes will be made has NO impact on this requirement. To remain synchronized and to monitor inter-cluster communications, all of the cluster in a Grid regularly communicate with each other
- If network firewalls are being used on any of the Grid links, see the list of the TCP I/P ports in the backup slides that must be open (again, ports must be open for 2-way communications between all of the clusters)
- The IBM SSR team must check that the joining cluster is at the correct code level and they must add any required compatibility fixes (a.k.a. “VTD execs”). Further, the existing clusters in the Grid must be checked and compatibility fixes may need to be added to each of them as well



## TS7700 Grid: Cluster Join – Preparation is the Key to Success (2 of 2)

---

- Determine settings for joining cluster (j)
  - Distributed library-id, distributed library name, cluster number and LIBPORT-IDs
  - FICON connectivity (host channels, director ports, cabling, UCB addresses, etc.)
  - (TS7720T, TS7740, TS7760T) Establish stacked volume ranges and order media
- Update FICON infrastructure (channels, directors, and cables) to allow hosts to be attached to joining cluster (j)
- Establish data network connections and I/P addresses for joining cluster (j)
  - Grid replication links (type, quantity, routing)
  - Management Interface, AOTM (Autonomic Ownership Takeover Manager – if used), and Call Home / Remote Support (TSSC – TS3000 System Console) connections
- Verify Grid replication link connectivity on existing clusters
  - A two-way Grid using direct cabling for their replication links will need to convert to using switches to add a third cluster to the Grid
  - Additional cables may be needed on the existing clusters to use the second pair of Grid links
- Physically install joining cluster (j) (if not already in place)
  - (Customer) Determine floor location (include front and back service clearance space)
  - (Customer) Provide twin power receptacles for each frame (base plus 1 or 2 expansion frames)
  - (IBM SSRs) Install, power-up and activate cluster

## TS7700 Grid: Cluster Join – Fill Out the Worksheets!

For a cluster join to be successful, it is imperative that the IBM SSR team have an accurate and complete list of Grid and cluster settings available to them. These settings are recorded in worksheets (partial example below) and should be kept and updated (when needed) during the life of the Grid. See an IBMer or Business Partner if you need a current (blank) copy of the worksheets

**Table 6. Grid Local Addresses**

Field	Value	Note
IP Address Primary Grid Interface Port A (Slot C1).  Customer Link Name Number: _____	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	The Primary Bay I/O Slot-C1 top port (Grid Link Adapter).  <b>Informational:</b> The Primary Grid Port A Interface at each cluster will connect to the Primary Grid Port A Interface at each of the other clusters in the same Grid  <b>Note:</b> IBM strongly recommends that the primary and alternate Grid interfaces be on separate subnets. If the Grid interfaces are directly connected (without using Ethernet switches) then using separate subnets is required.
Gateway	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	Add Gateway in this section to group information
Network Mask	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	Add Network Mask in this section to group information.
IP Address Alternate Grid Interface Port A (Slot C1).	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	The Alternate Bay I/O Slot-C1 top port (Grid Link Adapter).

## TS7700 Grid: Cluster Join – Scheduling a Join

---

- Copy Export operations are not permitted during a grid join/merge on any member of the grid
- All existing clusters in the Grid will remain online and operational while the join operation is in progress – the joining cluster will remain offline while the join process is running
  - A join should still be scheduled during a non-peak processing period to minimize the time required to complete the join

## TS7700 Grid: Cluster Join – High-Level Steps (IBM SSRs)

---

1. Establish the cluster index number (0 through 7) on joining cluster (j)
  - Must be unique within the Grid. LIBPORT-IDs will be based on this setting. The cluster number to LIBPORT-ID number relationship is detailed in a backup slide
2. Enter the Composite and Distributed Library numbers on joining cluster (j)
3. Perform pre-checks on the existing cluster (x) / Grid
  - All existing clusters must be online and healthy
4. Perform pre-checks on the joining cluster (j)
  - Cluster must be healthy and have a date/time within 5 minutes of existing cluster (x)
5. Configure local I/P addresses for primary and alternate grid paths
  - On cluster (j) and, if needed cluster (x)
6. Configure remote cluster I/Ps on both cluster (j) and cluster (x)
7. Perform a Grid communications test from both cluster (j) and cluster (x)
  - Any problems identified will need to be corrected before a join can be attempted
8. Initiate join operation from cluster (j)
9. Once join is complete, connect FICON cables to cluster (j) and vary it online
10. Run system check-out on cluster (j). When complete, customer administrators can proceed with updating the constructs and library settings

## TS7700 Grid: Cluster Join Times

Joining cluster number of volumes	Existing cluster(s) number of volumes	Approximate elapsed join time
0	0	50 minutes
0	40,000	100 minutes
0	71,000	130 minutes
0	500,000	240 minutes

**Note:** The above times are for guidance purposes only. Your local IBM SSR teams will have access to the latest information from the TS7700 development and support teams. They should always be consulted before a cluster join activity is scheduled

## TS7700 Grid: Cluster Join – Post Join Activities

---

- Check and (if needed) update the TS7700 constructs (Management Class [MC], Data Class [DC], Storage Class [SC], and Storage Group [SG])
  - The construct names and some of the individual custom settings will be copied to joining cluster (j) from existing cluster (x)
  - Some of the settings in joining cluster (j), especially in the MCs, will revert to default values and must be updated. Pay special attention to the Copy Consistency Policy (CCP) and Retain Copy Mode settings as these control where copies of each virtual volume is stored
  - Check the CCP settings in all of the existing clusters to insure that they will continue to direct volume copies to the desired clusters (such as new copies to joining cluster (j))
  - A Best Practice is to check all of the constructs in joining cluster (j) as well as the CCP settings in the existing clusters
- Check the settings controlled through the LIBRARY REQUEST command
  - Issue LIB REQ,distnamx,SETTING and LIB REQ,distname,SETTING2 commands using the distributed library name (distnamx) on existing cluster (x) and joining cluster (j). Save the outputs
  - Compare the two outputs to identify differences. As needed adjust the settings as needed (using the LIB REQ,distnamej,SETTING and SETTING2 commands)
- Check and (if needed) update other settings
  - Such as Physical volume pools, family definitions, cache partition sizes, and inhibit reclaim schedule
- Add distributed library definition for cluster (j) to DFSMS on applicable LPARs

## TS7700 Grid: Cluster Join – Points to Remember

---

- **Volume copies are “extra”** – When performing a join operation, it is important to note that the actual data does not get copied from one cluster to another. This process creates place holders for all of the logical volumes in the final grid
  - Opening a logical volume will cause Copy Consistency Policy (CCP) in the associated Management Class to be re-evaluated and for additional copies to be made (if needed)
  - Entering a COPY REFRESH command for a specific logical volume will cause the CCP for that volume to be examined as well
- In Section 4 (**Copy Audit and Copy Refresh**), we’ll examine how to use the:
  - Copy Audit function of BVIR (Bulk Volume Information Retrieval) to identify volumes that need to be copied into the new cluster
  - LIB REQ,distname,COPYRFSH,volser command to cause copies to be made
  - COPY REFRESH process that is part of the IBM Tape Tools package to help automate the creation and issuing of the COPYRSH commands

## Webinar Overview

---

- Section 1 – TS7700 Grid: Cluster Joins (Bill)
- **Section 2 – TS7700 Grid: Grid Merges (Bill)**
- Section 3 – TS7700 Grid: Cluster Removals (Ben)
- Section 4 – TS7700 Grid: Copy Audit and Copy Refresh (Ben)
- Wrap-up and Q&A (Ben, Bill, Mike and Toy)



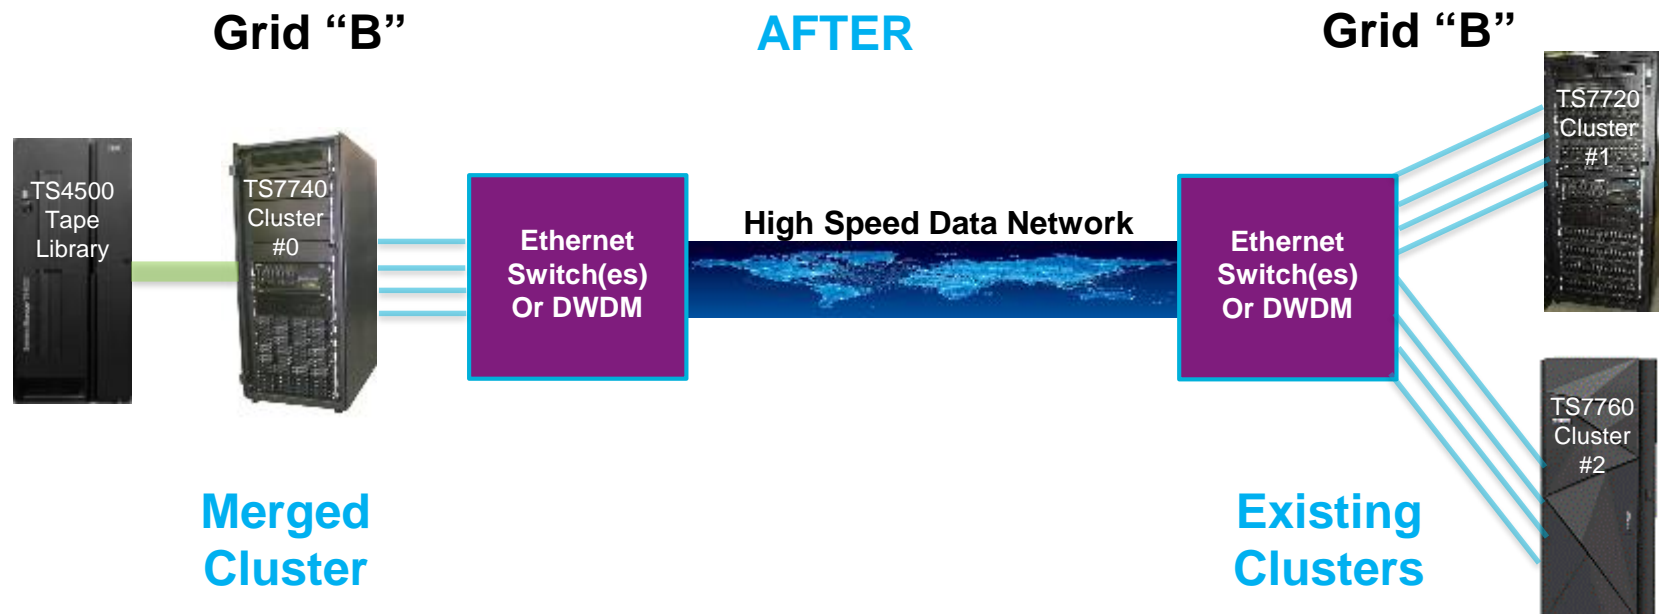
## TS7700 Grid: Grid Merge – Overview (1 of 2)

- **Merge** - the process that is performed when merging a cluster with data to another cluster with data to create a grid. The merging cluster is referred to as the merging cluster. The merging cluster might be a stand-alone cluster or it might be a cluster in an existing grid. Similarly, the existing cluster might be a stand-alone cluster or it might be a cluster in an existing grid
- When one cluster from a Grid is merged with a cluster in a second Grid, all clusters in both Grids are automatically merged into the combined Grid as well



## TS7700 Grid: Grid Merge – Overview (2 of 2)

- **When performing a merge, it is important to note that the actual data does not get copied between clusters.** This process creates place holders for all of the logical volumes in the final grid. When merging with an existing grid, the process is initiated to a single cluster in the grid and the information is populated to all members of the merged Grid
- When the merge operation completes, all virtual volumes can be accessed from any of the clusters in the consolidated Grid (even if no copies have been made)



## TS7700 Grid: Grid Merge – Requirements (1 of 2)

---

- In the documentation used by the IBM SSR for a Grid merge operation, the existing, online cluster is referenced as *cluster (x)* and the merging, offline cluster is referenced as *cluster (j)*.
- These are the requirements listed in the merge documentation used by the SSRs:
  1. If cluster (j) is part of a grid, all clusters in the grid must be offline while the merge process is taking place
  2. **Merges are only supported when all clusters in the resulting grid are at the exact same microcode level**
  3. The merging cluster (j) must have FC 4015 (Grid Enablement) installed.
  4. The merging cluster (j) must support at least the number of logical volumes that are supported by the grid, via FC 5270 (Increased logical volumes)
  5. The merging cluster (j) must contain FC 5271 (Selective Device Access Control) if cluster (x) has it installed
  6. If the merging cluster (j) has FC 1035 or 1038 (10 Gb Grid Optical LW Connection) installed, the customer infrastructure must support 10 Gb.
  7. Grid to grid merging is only supported at 8.21.x.x and later

## TS7700 Grid: Grid Merge – Requirements (2 of 2)

---

- Additional requirements listed in the merge documentation used by the SSRs:
  8. If the merging or target clusters have FC 1035 or 1038 (10 Gb Grid Optical LW Connection), then either of the following apply:
    - If the grid will be direct-attached, both clusters must have FC 1035 or 1038 installed
    - If the grid will not be direct-attached, the customer's network infrastructure must support 10 Gb
  9. Schedule this activity during a low activity time on the existing cluster (x) or grid.
  10. Total combined data classes between clusters should not be over 256.
  11. Total combined categories between clusters should not be over 256
  12. If the microcode level is 8.41.2xx.xx or later and cluster (j) is part of a grid, Control Unit Initiated Reconfiguration (CUIR) must be disabled on that grid prior to initiating the Service Prep action
  - 13. The logical volume ranges used in both Grids cannot overlap. A single volume serial number used in both Grids will cause the merge process to fail**
  14. If back-end physical tape is being used, then the physical volume ranges for any cluster in the merged grid cannot overlap with the merging logical volume ranges

## TS7700 Grid: Grid Merge – Preparation is the Key to Success (1 of 2)

---

- If a merge operation will result in a Grid with more than four (4) clusters in it, an approved no charge, SCORE (Storage Customer Opportunity REquest) RPQ is required. An IBMer or IBM Business Partner can initiate this request. The purpose of this SCORE RPQ is to allow the TS7700 development team to review the proposed Grid design to insure that both the merge and normal Grid operations will work correctly
- The merging clusters must have the capability for two-way TCP/IP communications across the Grid links to **ALL** of the clusters in the existing Grid. Where copies of logical volumes will be made has NO impact on this requirement. To maintain synchronization, all of the cluster in a Grid regularly communicate with each other
- If network firewalls are being used on any of the Grid links, see the list of the TCP I/P ports in the backup slides that must be open (again, ports must be open for 2-way communications between all of the clusters)
- The IBM SSR team must check the fix levels in all of the clusters in both Grids and they must add any required compatibility fixes (a.k.a. “VTD execs”)

## TS7700 Grid: Grid Merge – Preparation is the Key to Success (2 of 2)

---

- Determine settings for merging clusters
  - Distributed library-id numbers, distributed library names, and cluster numbers
  - LIBPORT-IDs (for HCD) will be based on each units cluster number
- Update HCD definitions with new LIBPORT-ID numbers. Defer activates until after the cluster (index) number change operations are complete

## TS7700 Grid: Grid Merge – Fill Out a Consolidated Worksheet!

For a grid merge operation to be successful, it is imperative that the IBM SSR team have an accurate and complete list of Grid and cluster settings available to them. These settings are recorded in worksheets (partial example below) and should be kept and updated (when needed) for the life of the Grid. See an IBMer or Business Partner if you need a current (blank) copy of the worksheets

**Table 6. Grid Local Addresses**

Field	Value	Note
IP Address Primary Grid Interface Port A (Slot C1).  Customer Link Name Number: _____	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	The Primary Bay I/O Slot-C1 top port (Grid Link Adapter).  <b>Informational:</b> The Primary Grid Port A Interface at each cluster will connect to the Primary Grid Port A Interface at each of the other clusters in the same Grid  <b>Note:</b> IBM strongly recommends that the primary and alternate Grid interfaces be on separate subnets. If the Grid interfaces are directly connected (without using Ethernet switches) then using separate subnets is required.
Gateway	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	Add Gateway in this section to group information
Network Mask	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	Add Network Mask in this section to group information.
IP Address Alternate Grid Interface Port A (Slot C1).	CL0: _____ CL1: _____ CL2: _____ CL3: _____ CL4: _____ CL5: _____ CL6: _____ CL7: _____	The Alternate Bay I/O Slot-C1 top port (Grid Link Adapter).

## TS7700 Grid: Grid Merge – Scheduling a Merge

---

- Copy Export operations are not permitted during a grid merge on any member of the grid
- All existing clusters in the existing or target Grid will remain online and operational while the merge operation is in progress – merging cluster (j) and all of the clusters in its Grid must remain offline while the merge process is running
  - A merge should still be scheduled during a non-peak processing period to minimize the time required to complete the operation



## TS7700 Grid: Grid Merge – High-Level Steps (IBM SSRs)

---

1. Perform logical volume duplicate pre-check process (can be done prior to a merge attempt)
  - If Ethernet connectivity is in place, can be done through the network. Otherwise use manual process (VTD Exec 084) The which stores information from cluster (x) on a CD/DVD and the CD/DVD then needs to be transported to the merging cluster (j) for comparison
2. Check to insure that all clusters in both Grids are in a healthy state
3. Vary cluster (j) and all other clusters in its Grid off-line and place them in Service / Service Prep mode
4. Establish the cluster index number (0 through 7) on the merging clusters
  - Must be unique within the Grid. LIBPORT-IDs will be based on this setting. The cluster number to LIBPORT-ID number relationship is detailed in a backup slide
5. Update Composite and Distributed Library Sequence Numbers in cluster (j) and associate clusters
6. Configure local and remote I/P addresses on cluster (j) and cluster (x)
7. Initiate the merge process from cluster (j)
8. Once merge completes, vary merged clusters online and remove Service mode
9. Run system check-out merged cluster (j). When complete, customer administrators can proceed with updating the constructs and library settings

## TS7700 Grid: Grid Merge Times

Merging cluster number of volumes	Existing cluster(s) number of volumes	Approximate cluster renumber time (if required*)	Approximate elapsed merge time
40,000	40,000	45 minutes	100 minutes
71,000	71,000	60 minutes	140 minutes
500,000	500,000	120 minutes	250 minutes

**Note:** The above times are for guidance purposes only. Your local IBM SSR teams will have access to the latest information from the TS7700 development and support teams. They should always be consulted before a Grid merge activity is scheduled

\* In a merged Grid, there can only be one cluster 0, one cluster 1, and so on. If duplicate cluster numbers will result from the merge, one or more of the clusters will need to be changed to a different cluster, unused number. This will impact the LIBPORT-ID settings on those clusters

## TS7700 Grid: Grid Merge – Needed Post Merge Actions

---

- Check and (if needed) update the TS7700 constructs (Management Class [MC], Data Class [DC], Storage Class [SC], and Storage Group [SG]) in all of the clusters in the merged Grid
  - The construct names from both Grid will be merged together and a consolidated set of construct names will appear in all of the clusters
  - Some of the settings, especially in the MCs, will revert to default values and must be updated. Pay special attention to the Copy Consistency Policy (CCP) and Retain Copy Mode settings as these control where copies of each virtual volume is stored
- Check and (if needed) add distributed library definitions each of the clusters in the merged Grid to DFSMS on the applicable LPARs

## TS7700 Grid: Grid Merge – Points to Remember

---

- **Volume copies are “extra”** – When performing a join operation, it is important to note that the actual data does not get copied from one cluster to another. This process creates place holders for all of the logical volumes in the final grid
  - Opening a logical volume will cause Copy Consistency Policy (CCP) in the associated Management Class to be re-evaluated and for additional copies to be made (if needed)
  - Entering a COPY REFRESH command for a specific logical volume will cause the CCP for that volume to be examined as well
- In Section 4 (**Copy Audit and Copy Refresh**), we’ll examine how to use the:
  - Copy Audit function of BVIR (Bulk Volume Information Retrieval) to identify volumes that need to be copied into a cluster
  - LIB REQ,distname,COPYRFSH,volser command to cause copies to be made
  - COPY REFRESH process that is part of the IBM Tape Tools package to help automate the creation and issuing of the COPYRSH commands

## Webinar Overview

---

- Section 1 – TS7700 Grid: Cluster Joins (Bill)
- Section 2 – TS7700 Grid: Grid Merges (Bill)
- **Section 3 – TS7700 Grid: Cluster Removals (Ben)**
- Section 4 – TS7700 Grid: Copy Audit and Copy Refresh (Ben)
- Wrap-up and Q&A (Ben, Bill, Mike and Toy)

## TS7700 Grid: Cluster Removals – Feature Code 4016

---

**Feature Code (FC) 4016 – Remove Cluster from Grid** delivers instructions for a one time process to remove/ unjoin a cluster from a Grid configuration. It can be used for removing one cluster from a two (2) to eight (8) cluster Grid. Subsequent invocations of this feature can be executed in order to remove multiple clusters from the Grid configuration

Use FC 4016 for the following tasks:

- Datacenter consolidation
- TS7700 Reuse
- End of Lease return
- Technology refresh

This topic describes the Prerequisites and Customer Preparation Tasks required to prepare to remove a cluster from the remaining clusters. The following pages provide information on steps that will be performed

## TS7700 Grid: Cluster Removals – Terminology

---

Throughout this section various terms may be used to define the clusters:

- **Remaining cluster** – Refers to a cluster that will still be part of the Grid after the FC 4016 actions are complete
- **Removed cluster** – Describes the cluster that will be taken out of a Grid. All actions for implementing FC 4016 will be performed on the removed cluster from one of the remaining clusters
- **Cluster family** – Clusters in the same grid can be grouped in cluster families
- **Local cluster** – Refers to the cluster that the user is logged into and performing specific actions upon
- **Remote cluster** – Refers a cluster within a Grid that is not the Local cluster. A remote cluster may be in the same room, different building, or even at a location miles away

## TS7700 Grid: Cluster Removals -- Attention

---

**Well before the change window for an IBM SSR to perform FC 4016 is scheduled, the customer should perform a Bulk Volume Information Retrieval (BVIR) copy audit check.** Any volumes shown to be inconsistent (i.e. only copy of virtual volume is in the cluster be removed) should be cleaned up (using Copy Refresh) prior to the cluster removal attempt

Initiating service prep on the cluster you are removing - all clusters in the Grid must be online and available at the start of this procedure. The process to remove a cluster from the Grid **MUST** be performed from a remaining cluster, **NOT** the cluster being removed



## TS7700 Grid: Cluster Removals – Pre-requisites

---

### 1) Ensure that the microcode version prerequisites are met:

- If any cluster in the grid is earlier than microcode level 8.4.1.x (R1.4), then FC 4016 is not supported. In this case all clusters earlier than 8.4.1.x must be upgraded before attempting to remove a cluster from the grid
- If all clusters in the grid are earlier than 8.7.0.134 (R1.7), then all clusters must be at the same microcode level before attempting to remove a cluster using FC 4016
- If at least one remaining cluster in the grid is at microcode level 8.7.0.134 or later, then it is possible to use that remaining cluster to initiate the removal of a cluster using FC 4016 even if the grid contains two different microcode versions

Note: The requirement that all clusters be at 8.4.1.x or later still applies.

### 2) If the microcode level is earlier than 8.5.0.154 (R1.5), the customer must ensure that at least one logical volume has been inserted. Attempting to remove a cluster with no logical volumes in the database will cause the removal to fail

### 3) If a cluster has been removed, the customer might want to modify the host configuration (HCD / IOCP) to remove the LIBPORT-IDs associated with the removed cluster

### 4) Ensure that the customer has performed the “Customer Preparation Tasks” (described on later pages) BEFORE starting the FC 4016 Installation Instruction. Failure to do so might cause the removal of the cluster to fail

Note: For information on code releases see **Backup Slides** section at end of presentation

## TS7700 Grid: Cluster Removals – Requirements (1 of 2)

---

- The remote cluster being removed will be placed in service mode and then taken. All remaining Grid clusters MUST be online and not in service mode or service prep
- During the change, if the SSR receives an error message indicating a count of inconsistent volumes has stopped the removal then the customer will be engage to determine whether to continue the removal. If the count is over 30 volumes the customer should run the BVIR Copy Audit against the remaining clusters to determine what volumes are currently inconsistent
- During the change, the SSR will receive a prompt asking if they should skip the “Copy Consistency Check”. They will ONLY skip the “Copy Consistency Check” if the customer agrees. Extreme CAUTION should be taken, continuing without performing the “Copy Consistency Check” and resolving any issues with logical volumes may result in loss of data
- During the installation time of FC 4016, if the SSR receives an error message that stopped the removal of the cluster because volumes are in the damaged category the customer will be contacted and determine if they would like to repair the volumes before continuing

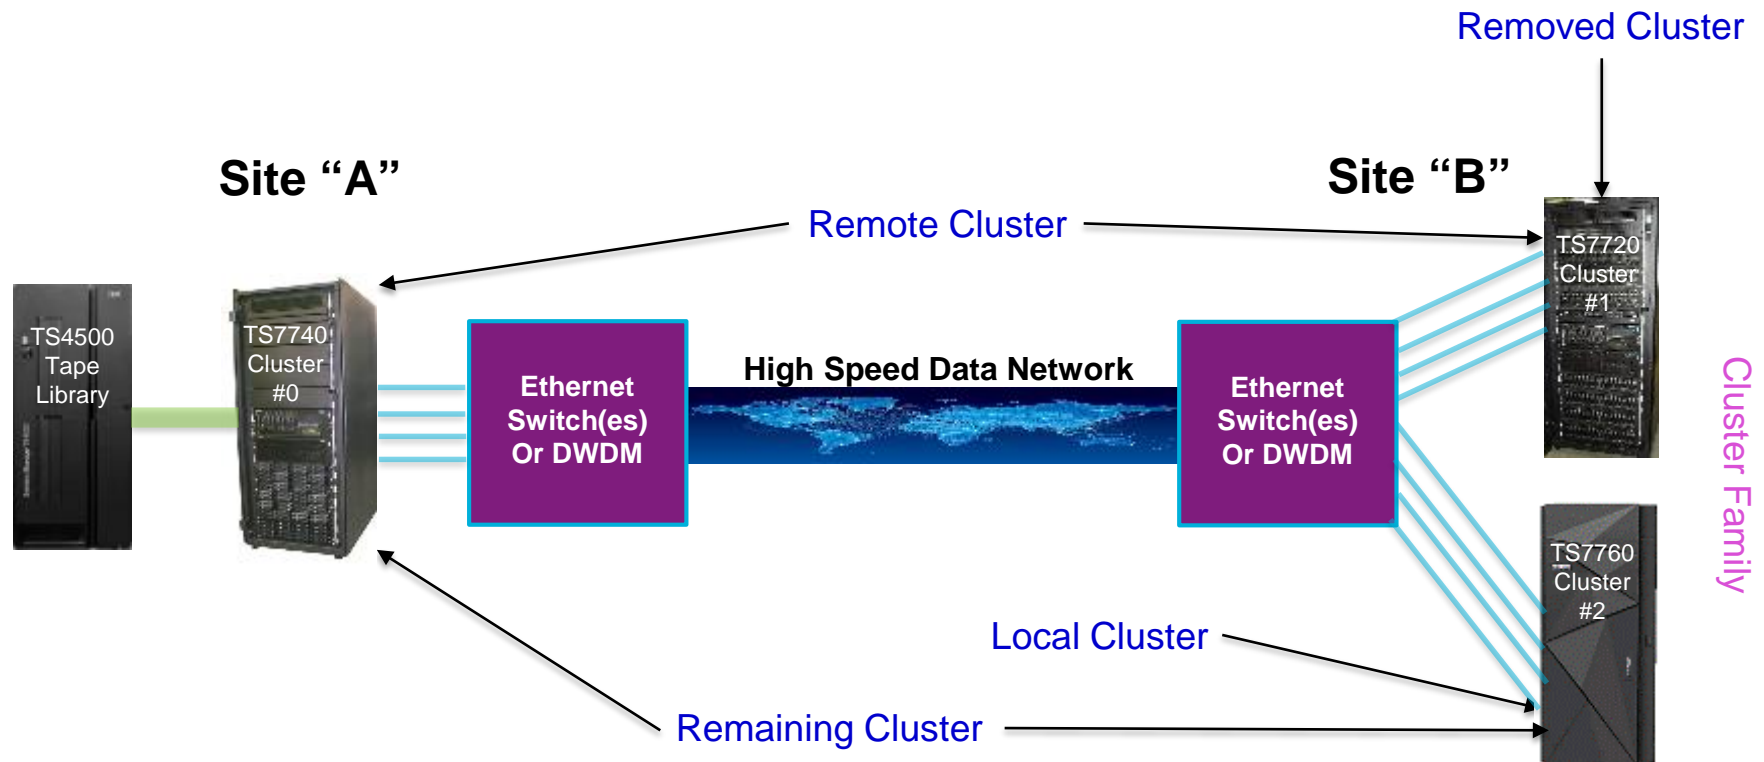
## TS7700 Grid: Cluster Removals – Requirements (2 of 2)

---

- Only ten damaged volumes can be repaired at one time. It takes approximately two minutes to repair a group of ten volumes
- If an error message received that the removal of the cluster failed because the removal process found damaged volumes, the removed cluster **MUST** be varied back online in order for the customer to address the damaged volume issue
- A license key is not required for FC 4016. The FC 4017 license key is required if the cluster is going to be reused
- If the cluster being removing is being used as Grid to Grid Migration (GGM) CSG copy source cluster (vtd\_exec.222 was installed), you must disable GGM and reconfigure GGM to use another CSG copy source cluster
- In order to remove multiple clusters from the Grid configuration the Feature Code 4016 process must be performed in sequential order – only one cluster can be removed at a time

## TS7700 Grid: Cluster Removals – Example

- Existing configuration is a 2-Site, 3-Way GRID
  - In this example, we will remove TS7720 cluster 1 (CL1) from the GRID
  - The Cluster Removal process will be initiated from TS7760 cluster 2 (CL2)



## TS7700 Grid: Cluster Removals – Customer Preparation Tasks (1 of 4)

---

- 1) If the cluster you are removing from the Grid will be reused, you must order FC 4017. If FC 4017 was ordered, install the FC 4017 license key from the TS7700 Management Interface before you run FC 4016
- 2) View the TS7700 Microcode Level. Log in to the TS7700 Management Interface
  - a. From the Grid Summary page, select the cluster that you want to view the microcode level by clicking the cluster image
  - b. When the selected cluster is displayed, click or hover over the cluster name. Information about that cluster will be displayed, including the Microcode Level
  - c. Record the Microcode Level here: \_\_\_\_\_
- 3) Ensure that the cluster you are removing does NOT have the only consistent copy of any volumes. Reference the Bulk Volume Information Retrieval (BVIR) Copy Audit function described in the IBM TS7700 Series Bulk Volume Information Retrieval Function User's Guide white paper (WP101094)
- 4) In order to allow the removal to complete more quickly, the ownership of the logical volumes can be transferred to one of the remaining clusters in the grid. Reference the TS7700 Series z/OS Host Command Line Request Users Guide (WP101091) and review the ownership transfer and verification functions found in command OTCNTL

WARNING: Failure to resolve the inconsistent lvols could result in loss of data.

## TS7700 Grid: Cluster Removals – Customer Preparation Tasks (2 of 4)

---

- 5) If there are volumes that have only a consistent copy on the cluster to be removed, the customer must determine how to handle those volumes by performing one or more of the following tasks:
- Eject the logical volumes by following the steps in IBM Virtualization Engine TS7700: Tape Virtualization for System z Servers
  - Move the volumes to a scratch category. This is a category set up for Fast Ready on the TS7700 Management Interface. Hover over the Virtual icon. Then select Categories
  - Modify the management class of the volume so there will be a copy on one of the remaining clusters. To activate the management class, change either mount/demount the volume or issue a COPYRFSH for the volume. This will result in a new copy on the remaining cluster

Refer to the Host Command Line Request white paper (WP101091) for information on COPYRFSH

**WARNING:** Failure to resolve the inconsistent logical volumes could result in loss of data.

- 6) Ensure that there are no volumes in the damaged category
- From the TS7700 Management Interface, hover over the Service icon. Then select Repair Virtual Volumes.  
Note: You can only display and repair only 10 damaged volumes at a time. To repair each group of 10 volumes takes approximately 2 minutes
- 7) Modify Management Classes so that the removed cluster is no longer the target for copies. Check copies on the removed cluster by using the TS7700 Management Interface
- Hover over the Constructs icon. Then select Management Classes
  - Ensure that “No Copy” is present under the removed cluster ID

## TS7700 Grid: Cluster Removals – Customer Preparation Tasks (3 of 4)

---

- 8) If the cluster being removed is part of a cluster family, it should be removed from the family before the removal of the cluster. Remove the cluster from the family by using the TS7700 Management Interface:
  - a. Select Cluster Families in the Actions tab in the grid summary section.
  - b. If the cluster is part of a family it will be on the bottom of the panel grouped with a family. Clusters that are listed on the top of the panel are not associated with a family. Select the cluster that is to be removed, then drag and drop the cluster to the top of the panel where clusters not associated with a family are listed.
  - c. After you drag out and drop the cluster, if the family is empty, click on the cross to delete the family.
  - d. Click Save
  
- 9) If the cluster being removed is set to write protect, then clear the write protect by using the TS7700 Management Interface:
  - a. The TS7700 Management Interface Grid summary will show a padlock beside the write protected cluster
  - b. To clear the write protect, from the TS7700 Management Interface under Settings > Cluster Settings > Write Protect select disable write protect mode
  - c. If the write protect was set through a Host Command Line Request then it will need to be cleared using Host Command Line Request Commands.  
Reference the DR setup section that is described in the IBM TS7700 Series z/OS Host Command Line Request User's Guide white paper

## TS7700 Grid: Cluster Removals – Customer Preparation Tasks (4 of 4)

---

- 10) If flash was enabled through a Host Command Line Request then it will need to be disabled using Host Command Line Request Commands. Reference the DR setup section that is described in the IBM TS7700 Series z/OS Host Command Line Request User's Guide white paper
- 11) If the cluster is part of a DR family you will need to remove it using Host Command Line Request Commands. Reference the DR Family section that is described in the IBM TS7700 Series z/OS Host Command Line Request User's Guide white paper.

**Note (1):** Be aware that during the removal process that logical volume inserts, ejects, and copy exports are not allowed

**Note (2):** The tasks to remove the cluster from the Grid can be concurrent with hosts jobs running, however, it is recommended that the removal occurs while few or no jobs are running to the remaining clusters



## TS7700 Grid: Cluster Removals – FC 4016 Installation Time

The amount of time it takes to remove a cluster from the Grid depends on how many logical volumes there are in the Grid.

Service Action	Estimated install hours	Estimated machine hours	Estimated system hours	Estimated SSR hours	Number of SSRs
Concurrent <sup>1</sup>	3.5 <sup>2,3</sup>	0	0	3.5	1

### Notes:

- 1) The procedure to remove a cluster from the Grid can be done concurrently with customer operation on the remaining clusters. You must perform the tasks to remove a cluster from a Grid from one of the remaining cluster(s)
- 2) The remaining cluster you use for the removal procedure must be a local cluster.
- 3) The three and one half hours includes the time for service prep. The time is also dependent upon the customer's job processes in the queue. The UNJOIN may need additional time. The process can take 2+ hours per 100K volumes where the owner of the volume was the removed cluster

## TS7700 Grid: Cluster Removals – Feature Code 4017 (1 of 2)

---

**Feature Code 4017 – Cluster Cleanup** (a.k.a. “Factory Reset”) facilitates a one time cluster cleanup, which cleans the database, logically deletes the virtual volumes (LVOL) from the tape volume cache, and removes configuration data for host connections from a TS7700 cluster

Before the cleanup can occur, a cluster with FC 4015 (Grid Enablement) installed, must be removed from the Grid using FC 4016 (Remove Cluster from Grid)

Performing the Cluster Cleanup will:

- a) Clean the TS7700 database
- b) Clear the configuration data from the vital product data (VPD)
- c) Remove the log files
- d) Clean the LM-related data (TS7700 attached to a 3584)
- e) NOT physical delete any data in cache, only removes index information for the LVOL

**Note:** If an additional cleanup is required on this cluster, you must have another license key for FC 4017 (Cluster Cleanup)

## TS7700 Grid: Cluster Removals – Feature Code 4017 (2 of 2)

### Additional Information

- If a cluster is being redeployed, a Factory Reset (FC 4017) may be all that is required
- If a cluster is being disposed of or is being returned to a leasing company, a decision needs to be made on what state the data in cache will be left in
  - Only the index to the data removed (e.g. factory reset)
  - A cryptographically erase, removing the TS7700's encryption keys if feature enabled
  - A Secure Data Overwrite (SDO) which can provide a “Certificate of Completion” document which could be used for auditors
- If utilizing FC 5272 (Disk Encryption - Local Key Manager) or FC 5276 (Disk Encryption – Remote Key Manager) then:
  - During Cluster Cleanup the SSR will receive a prompt early in process on how to handle the encryption keys used by the cache. The SSR can reply YES or NO to the prompt
    - A NO leaves the keys intact and the data could still be accessible after the Factory Reset
    - YES will delete the keys resulting in a cryptographically erase the cache and GPFS file systems.
      - ✓ **Caution:** DATA LOSS WILL OCCUR, the data will no longer be accessible
- If performing an SDO, not handled by way of feature codes
  - An Statement of Work (SOW) will be generated through IBM Lab Services and there will be an additional cost beyond the typical Feature Codes (e.g. 4016, 4017).
  - The cluster will need to be connected to Assist On-Site (AOS, a.k.a. “Call Home / Remote Support”) since the SDO process is handled remotely out of IBM Tucson

## Webinar Overview

---

- Section 1 – TS7700 Grid: Cluster Joins (Bill)
- Section 2 – TS7700 Grid: Grid Merges (Bill)
- Section 3 – TS7700 Grid: Cluster Removals (Ben)
- **Section 4 – TS7700 Grid: Copy Audit and Copy Refresh (Ben)**
- Wrap-up and Q&A (Ben, Bill, Mike and Toy)

## TS7700 Grid: Copy Audit and Copy Refresh

---

Throughout this section we will discuss two processes available to the TS7700's to help insure there are no logical volumes (LVOL) without a copy in at least one the remaining clusters

- **COPY AUDIT** from “Bulk Volume Information Retrieval” (BVIR)
- **Copy Refresh** using a combination of zOS batch JCL and operator commands
  - IBMTOOLS.JCL(COPYRFSH) JCL \*\*\*
  - TS7700 LIBRARY REQUEST command COPYRFSH

### Copy Audit

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

### Copy Refresh

- The COPYRFSH request provides a method to refresh copy policy without mounting/demounting a volume. The copy policy defined in the management class of the volume on the cluster associated with the distributed library in the request is assigned. The request queues a copy job on cluster(s) which need a copy of the volume and the copy job queued using this request will have a lower priority than the newly created workload

\*\*\* The IBM Tape Tools can be downloaded through <ftp://ftp.software.ibm.com/storage/tapetool/>  
The files need to be uploaded to a z/OS LPAR and then restored using RECEIVE commands (see IBMTOOLS.TXT)

## TS7700 Grid: Copy Audit Request (1 of 5)

---

**Purpose:** This request performs an audit of the databases on a set of specified TS7700 distributed libraries to determine if there are any volumes that do not have a valid copy on at least one of them. If the COPYMODE option is specified, whether or not the volume is supposed to have a copy on the distributed library is taken into account (based on the management class settings in the creating cluster) in determining whether that distributed library has a valid copy

- **Request Format:** COPY AUDIT INCLUDE libids or  
COPY AUDIT EXCLUDE libids or  
COPY AUDIT COPYMODE INCLUDE libids or  
COPY AUDIT COPYMODE EXCLUDE libids

where *libids* is a list (one or more) of the 5 character distributed library IDs (separated by a comma) to be included or excluded from the COPY AUDIT operation. INCLUDE or EXCLUDE must be specified or the request will fail

- **Data Scope:** Specified cluster(s)

## TS7700 Grid: Copy Audit Request (2 of 5)

- **Recommended Formatting Approach:** Download file to workstation (PC) and use spreadsheet program (e.g. Microsoft Excel) to format. There are eighteen (18) columns in the report.

- A CSV file with the column headers is included here:



- **Comments:**

- 1) Copy Audit may take more than one (1) hour to complete depending on the number of logical volumes have been defined, how many clusters are configured in the Grid configuration, and how busy the TS7700s are at the time of the request.
- 2) The Copy Audit request is intended to be used for the following:
  - a) A TS7700 is to be removed from a Grid configuration. Prior to its removal you want to ensure that the TS7700s that are to remain in the Grid configuration have a copy of all volumes that were created on the TS7700 that is to be removed
  - b) A condition has occurred (could be a site disaster or as part of a test procedure) where one of the TS7700s in a Grid configuration is no longer available and you want to determine which, if any, volumes on the remaining TS7700s do not have a valid copy
- 3) It is imperative that the correct options and clusters be selected on the COPY AUDIT command. Otherwise the command will run fine, but produce invalid results

## TS7700 Grid: Copy Audit Request (3 of 5)

---

**Example #1:** CL0 (libid 1234A) at remote site

CL1 (libid 1234B) and CL2 (libid 1234C) at production site –  
CL1 will be removed

Single management class (MCRECOV) specifies “D” for CL0, CL1,  
and for CL2.

COPY AUDIT objective, insure that CL2 has a valid copy of all private  
logical (virtual) volumes found in CL1

### **Valid requests include:**

COPY AUDIT INCLUDE 1234A, 1234C

COPY AUDIT EXCLUDE 1234B

COPY AUDIT COPYMODE INCLUDE 1234A,1234C

COPY AUDIT COPYMODE INCLUDE 1234B

COPY AUDIT COPYMODE EXCLUDE 1234A, 123C

### **Invalid requests (no volumes will ever be flagged) include:**

COPY AUDIT INCLUDE 1234A, 1234B, 1234C

COPY AUDIT EXCLUDE 1234A

COPY AUDIT COPYMODE INCLUDE 1234A

COPY AUDIT COPYMODE EXCLUDE 1234B



## TS7700 Grid: Copy Audit Request (4 of 5)

---

**Example #2:** Two (2) cluster Grid – CL0 (libid 1234A) at production sites and CL1 (libid 1234B) at recovery site. Single management class (MCRECOV) specifies “D” for CL0 and “D” for CL1. COPY AUDIT objective is insure that CL1 has a valid copy of all private logical (virtual) volumes

**Valid requests include:**

COPY AUDIT INCLUDE 1234B

COPY AUDIT EXCLUDE 1234A

COPY AUDIT COPYMODE INCLUDE 1234A,1234B

COPY AUDIT COPYMODE INCLUDE 1234B

COPY AUDIT COPYMODE EXCLUDE 1234A

**Invalid requests (no volumes will ever be flagged) include:**

COPY AUDIT INCLUDE 1234A

COPY AUDIT EXCLUDE 1234B

COPY AUDIT COPYMODE INCLUDE 1234A

COPY AUDIT COPYMODE EXCLUDE 1234B

## TS7700 Grid: Copy Audit Request (5 of 5)

---

**Example #3:** Two (2) cluster Grid – CL0 (libid 1234A) at production sites and CL1 (libid 1234B) at recovery site. Two (2) management classes: (MCRECOV) specifies “D” for CL0 and “D” for CL1 and (MCNONE) specifies “R” for CL0 and “N” for CL1. COPY AUDIT objective is insure that CL1 has a valid copy of all private logical (virtual) volumes using the MCRECOV management class

### **Valid requests include:**

COPY AUDIT COPYMODE INCLUDE 1234A,1234B

COPY AUDIT COPYMODE INCLUDE 1234B

COPY AUDIT COPYMODE EXCLUDE 1234A

### **Invalid requests (no volumes will ever be flagged) include:**

COPY AUDIT COPYMODE INCLUDE 1234A

COPY AUDIT COPYMODE EXCLUDE 1234B

## TS7700 Grid: Copy Audit Formatted Data Example (1 of 2)

---

```
VTS BULK VOLUME DATA REQUEST  
COPY AUDIT COPYMODE INCLUDE 04010  
06/01/2017 20:37:33 VERSION 02  
S/N: H5439 LIB ID: 04010
```

```
04010  
NO INVALID VOLUMES ON SPECIFIED CLUSTERS
```

Or

```
VTS BULK VOLUME DATA REQUEST  
COPY AUDIT COPYMODE INCLUDE 04010  
06/01/2017 20:37:33 VERSION 02  
S/N: H5439 LIB ID: 04010
```

```
04010  
L00001 ,R ,D ,D ,N ,N ,N ,N ,N ,N ,R ,R ,R ,N ,N ,N ,N ,R  
VOL004 ,R ,R ,D ,D ,D ,D ,N ,N ,N ,N ,N ,N ,N ,N ,N ,N  
VOL005 ,R ,R ,N ,N ,N ,N ,N ,N ,N ,N ,N ,N ,N ,N ,N ,N
```

## TS7700 Grid: Copy Audit Formatted Data Example (2 of 2)

```

VTS BULK VOLUME DATA REQUEST
COPY AUDIT COPYMODE INCLUDE 04010
06/01/2017 20:37:33 VERSION 02
S/N: H5439 LIB ID: 04010
    
```

04010

NO INVALID VOLUMES ON SPECIFIED CLUSTERS

Volser	Cluster0 Copy Mode	Cluster1 Copy Mode	Cluster2 Copy Mode	Cluster3 Copy Mode	Cluster4 Copy Mode	Cluster5 Copy Mode	Cluster6 Copy Mode	Cluster7 Copy Mode	Cluster0 Removal State	Cluster1 Removal State	Cluster2 Removal State	Cluster3 Removal State	Cluster4 Removal State	Cluster5 Removal State	Cluster6 Removal State	Cluster7 Removal State	Composite Removal State
L00001	R	D	N	N	N	N	N	N	N	D	R	N	N	N	N	N	R
VOL004	R	D	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
VOL005	R	D	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

# TS7700 Grid: Copy Audit – Management Class Usage

- (Best Practice) Establish a Mgmt Class for BVIR Use for Each Cluster

CL0 Copy Policies

Management Class (MC) Name	C0	C1	C2	C3
MCBVIRC0	R	N	N	N
MCBVIRC1	N	R	N	N
MCBVIRC2	N	N	R	N
MCBVIRC3	N	N	N	R

CL1 Copy Policies


Management Class (MC) Name	C0	C1	C2	C3
MCBVIRC0	R	N	N	N
MCBVIRC1	N	R	N	N
MCBVIRC2	N	N	R	N
MCBVIRC3	N	N	N	R

CL2 Copy Policies

Management Class (MC) Name	C0	C1	C2	C3
MCBVIRC0	R	N	N	N
MCBVIRC1	N	R	N	N
MCBVIRC2	N	N	R	N
MCBVIRC3	N	N	N	R

CL3 Copy Policies

Management Class (MC) Name	C0	C1	C2	C3
MCBVIRC0	R	N	N	N
MCBVIRC1	N	R	N	N
MCBVIRC2	N	N	R	N
MCBVIRC3	N	N	N	R



## TS7700 Grid: Copy Refresh Command

---

- **Library REquest COPYRFSH** command format:

LI REQ,*target\_libname*,COPYRFSH,volser,source cluster number

- The Copy Refresh (COPYRFSH) command was introduced with TS7700 R2.1 PGA-3 code level. Instead of mounting a volume, like is done with the PRESTAGE utility, the command instructs the Grid to re-evaluate the logical volume copy status in each cluster. Copies that are no longer needed are marked for removal. Copies that need to be made are run as low priority tasks. The re-evaluation is done using the current Copy Consistency Policy (CCP) settings of the assigned Management Class in the cluster target cluster specified in the command. The (optional) desired cluster to be used as a source for the copies can be specified in the command as well
- The LI REQ command can be used directly if COPY AUDIT identifies a small number of LVOLs needing to be copied. If the number of LVOL's is large then the Copy Refresh process on the next slide may be more appropriate for your GRID
- The LI REQ command can be submitted via batch job to eliminate the need to type and execute multiple LI REQ commands.
  - IBMTTOOLS.JCL(ISSUECMD) -- the ISSUECMD program needs to run from an APF authorized library

## TS7700 Grid: Copy Refresh Process – Steps (1 of 2)

---

You may want to consider setting RECLPG0 to enable on the source cluster in order to have the LVOLS managed as PG0 so they will be removed from cache before PROD volumes

- 1) Be sure the New Management class values are in place
- 2) BVIRMES must be run first to get the volume status file (mes)
- 3) Run VESYNC to get list of virtual volumes to be synced after a new cluster was installed and needs all LVOLS copied or, get a list of some old LVOLS that now need a peer copy
- 4) Run BVIRVTS (TYPE=VOL) to get VOLFILE for TS7740

Run BVIRVTS (TYPE=CACH) for TS7720 and TS7740

Only TYPE=VOL is needed for TS7740 source unless MAXGB=NN used. The commands are issued based on the order LVOLS reside on PVOLS

If CACHFILE is also used for TS7740, then TVC resident LVOLS are grouped together so a physical mount doesn't need to be done

- 5) Review control parameters in SYSCNTL DD statement of this job
- 6) Look at ISSUECOMMANDS comments to list or issue commands
- 7) Submit this job and wait for Copy Refresh to complete

Note: NO host mounts are done

If the process gets interrupted before all COPYRFSH commands are issued, you should start with step 2 again. Once the COPYRFSH has been issued, the MESFLAG field is changed so the LVOL won't be selected again by the VESYNC tool.

## TS7700 Grid: Copy Refresh Process – Steps (2 of 2)

---

- The following provides the name of the JCL files in the IBM Tape Tools package:
  - IBMTOOLS.JCL(PRESTAGE)
  - IBMTOOLS.JCL(VESYNC)
  - IBMTOOLS.JCL(BVIRVTS)
  - IBMTOOLS.JCL(COPYRFSH)
- The files above do not have to be run. It is best to read the comment box located in each of the JCL members to fully understand the need for that job. A REMOVE (aka UNJOIN) can be as simple as the BVIR Copy Audit and then either the LI REQ command or just the IBMTOOLS.JCL(COPYRFSH) with the volume serial numbers (LVOL) of those that need to be copied to one of the remaining clusters.

**Note:** There is a member, IBMTOOLS.JCL(LICENSE), that contains an expiration date or valid through date. If it's been a long time since the last execution of JCL / programs from the IBMTOOLS datasets then you may want to look at refreshing the datasets. There are download and installation instructions in the "ibmtools.txt" file found on the IBM TAPETOOLS FTP site



## TS7700 Grid: Copy Refresh Pre-Stage Process (1 of 2)

---

- The following provides the name of the JCL files in the IBM Tape Tools package:
  - IBMTOOLS.JCL(PRESTAGE)
- Pre-stage processing is similar to copy refresh processing in as much as you need to provide a list of volumes to recall
- In the following list CLx is an existing cluster, CLn is the new cluster that has been added
  - 1) Turn off reclaim on CLx
  - 2) Run BVIRVTS on CLx to get VOLUME MAP
  - 3) Run BVIRMES on CLn to get VOLUME STATUS
  - 4) Run VESYNC to get current list of logical volumes needing copies
  - 5) Run PRESTAGE to build IEBGENER jobs to orderly recall logical volumes.  
Go back to step 3 for next set of recalls.
- A way of running PRESTAGE would be to use the MAXDR= 1 parameter so that only one recall at a time would be occurring, but you could recall the logical volumes from multiple physicals and more or less be continually doing multiple recalls at a time. To be more aggressive, MAXDR=3 would allow 3 at a time to be recalled, but that would be three times the link traffic. Yes, there are ways of limiting how many recalls are submitted at one time to minimize TVC flushing and the recalled volumes can be set to PG0 also by using the STORECLASS= parameter. If it isn't used, then logical volumes will use their original PGn attribute when recalled
- After all of these recalls have completed, you would need to start the process over again from the beginning because the previous list of logical volumes would now have a bunch of them recalled

## TS7700 Grid: Copy Refresh Pre-Stage Process (2 of 2)

---

- The customer will need to restrict when reclaim is being done on a cluster or it might affect the efficient recall of the logical volumes. PRESTAGE orders the logical volume recalls based on where they reside on the physical volume. Reclaim might change where the logical volumes reside and cause too many physical volume mounts
- PRESTAGE will allow the user to assign a temporary MGMTCLAS value (like RDD) to influence recalls to be in the desired cluster. It also provides for changing it back to the original MC name. Both of the changes are done in the job actually doing the recalls
- The actions performed by a management class can be changed so a copy will be made by the next dismount, but the NAME of the management class can only be changed via the LIBRARY LMPOLICY command which PRESTAGE will do if requested

## Webinar Overview

---

- Section 1 – TS7700 Grid: Cluster Joins (Bill)
- Section 2 – TS7700 Grid: Grid Merges (Bill)
- Section 3 – TS7700 Grid: Cluster Removals (Ben)
- Section 4 – TS7700 Grid: Copy Audit and Copy Refresh (Ben)
- **Wrap-up and Q&A (Ben, Bill, Mike and Toy)**

## TS7700 Grid: Summary

---

During this session, we've covered three (3) major operations used to change the "landscape" of a TS7700 Grid. These include:

- 1) Cluster joins – add a cluster to a Grid
- 2) Grid mergers – combine two (2) Grids together
- 3) Cluster removals – delete a cluster from a Grid

In addition, we covered the use of the COPY AUDIT function (part of the BVIR utility) to identify logical volumes needing to be copied within a Grid before a cluster can be removed. Further, we discussed the use of the COPYRFSH sub-command (of LIBRARY REQUEST) to cause these copies to take place along with the Copy Refresh process that can be used to semi-automate the submission of the COPYRFSH sub-commands

Finally, due to time and space constraints, we did not cover other related functions or operations such a GGM (Grid-to-Grid Migration) or the direct cluster replacement operation (a.k.a. "frame roll" – only applicable when physical tape is part of a cluster's configuration). If there is interest, we can cover these topics in a future webinar

# Washington Systems Center - Storage

Thank You



## Backup Slides

---

# Backup Slides

## TS7700 Grid: Bibliography

---

- TS7700 Release 4.1 and 4.1.2 Guide Redbook (SG24-8366-01)  
<http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248366.html?Open>
- TS7700 Knowledge Center – Customer Documentation 4.1.2.0  
[https://www.ibm.com/support/knowledgecenter/STFS69\\_4.1.2/hydra\\_c\\_ichome.html](https://www.ibm.com/support/knowledgecenter/STFS69_4.1.2/hydra_c_ichome.html)
- TS7700 Installation Instructions FC 1034 (Enable Dual Port Grid Connection) – PN 2726141 EC M13543A
- TS7700 Installation Instructions FC 1035/1036/1037 Grid Link (Adapter Card Conversion) – PN 46X9160 EC M13327
- TS7700 Installation Instructions FC 1036/1037/1038 Grid Link (Adapter Card Conversion) – PN 2726842 EC M13923
- TS7700 Customer Installation Instructions FC 4015 (Install Grid Enablement Feature) – PN 38L6368 EC M13543A
- TS7700 Installation Instructions FC 4016 (Remove Cluster from Grid) – PN 38L6369 EC M13923
- TS7700 Installation Instructions FC 4017 (Cluster Cleanup) – PN 2727199 EC M13923
- TS7700 Installation Instructions for Frame Replacement Migration – PN 2727108 EC N99426
- TS7720 v8.33 PGA2 Installation Roadmap -- PN 01PL100 EC M13937
- TS7720 Tape Attach v8.33 PGA2 Installation Roadmap -- PN 01PL101 EC M13937
- TS7740 v8.33 PGA2 Installation Roadmap – PN 01PL102 EC M13937
- TS7760/TS7760 Tape Attach v8.41.2a Installation Roadmap – PN 02CE656 EC N99425

## TS7700 Grid: Major Code Releases

---

<b><u>TS7700 Code Versions</u></b>		<b><u>1<sup>st</sup> Code GA Date</u></b>
Version 8.41.2x.x	(R4.1.2)	2018
Version 8.41.x.x	(R4.1)	2017
Version 8.40.x.x	(R4.0)	2016
Version 8.33.x.x	(R3.3)	2015
Version 8.32.x.x	(R3.2)	2014
Version 8.31.x.x	(R3.1)	2013
Version 8.30.x.x	(R3.0)	2012
Version 8.21.x.x	(R2.1)	2011
Version 8.20.x.x	(R2.0)	2011
Version 8.7.x.x	(R1.7)	2010
Version 8.6.x.x	(R1.6)	2009
Version 8.5.x.x	(R1.5)	2008
Version 8.4.x.x	(R1.4)	2007
Version 8.3.x.x	(R1.3)	2007
Version 8.2.x.x	(R1.2)	2007
Version 8.1.x.x	(R1.1)	2007
Version 8.7.x.x	(R1.0)	2006



## TS7700 Grid: Cluster Joins – Open Ports Needed on Grid Links

TCP/IP Port	Role
<b>ICMP</b>	Check cluster health
<b>9</b>	Discard port for speed measurement between grid clusters
<b>20</b>	Recommended to remain open for FTP data
<b>21</b>	Recommended to remain open for FTP control
<b>80</b>	Access the remote management interface when clusters are operating at different microcode levels (HTTP)
<b>123</b>	Network time protocol (NTP) time server
<b>350</b>	TS7700 file replication (distributed library file transfer)
<b>443</b>	Access the TS7700 Management Interface
<b>500</b>	IPsec Key Exchange (TCP and UDP): must remain open when grid encryption is enabled
<b>1415</b>	WebSphere message queues (grid-to-grid)
<b>1416</b>	WebSphere message queues (HDM-to-HDM)
<b>5988</b>	CIM Server
<b>8500</b>	IPsec Key Exchange (TCP and UDP): must remain open when grid encryption is enabled

Source: [https://www.ibm.com/support/knowledgecenter/STFS69\\_4.1.2/ts7740\\_infrastructure\\_requirements\\_network\\_switches\\_tcp\\_ip\\_ports.html](https://www.ibm.com/support/knowledgecenter/STFS69_4.1.2/ts7740_infrastructure_requirements_network_switches_tcp_ip_ports.html)

**Red** = Required; **Black** = Highly Recommended; Grey = Optional

## TS7700 Grid: Cluster Numbers and LIBPORT-ID Numbers

- LIBPORT numbers are assigned to each string of 16 devices
- 256 virtual tape drives are supported in a standard TS7700 cluster (256 = 16 strings with 16 devices per string)
- Up to 496 virtual tape devices can be supported thru the addition of one or more (up to 15) 5275 Feature Codes (496 = 31 strings x 16 devices)

For Cluster 0 use 01-1F	For Cluster 4* use 21-3F
For Cluster 1 use 41-5F	For Cluster 5* use 61-7F
For Cluster 2 use 81-9F	For Cluster 6* use A1-BF
For Cluster 3 use C1-DF	For Cluster 7* use E1-FF

*\* Grids with more than four (4) cluster are only supported with an approved RPQ*

- This is true for the IOCP on any of the hosts

## Disclaimers

---

- Copyright© 2018 by International Business Machines Corporation.
- No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.
- 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.