



ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG Webinar: Storage Virtualize 8.6.3 and 8.7.0 Update

Byron Grossnickle

IBM Storage Virtualize/FlashSystem SME

byrongro@us.ibm.com



Meet the Speakers



Byron Grossnickle is an IBM Storage Technical Specialist concentrating on Storage Virtualize software. This includes FlashSystem, SVC, and Storage Virtualize for Public Cloud. Byron has been with IBM 19 years exclusively in storage. Prior to working for IBM, Byron spent 6 years engineering storage in the Telcom Industry. Prior to that he worked 8 years in healthcare IT. Byron lives in the Kansas City area and is available to travel to customer engagements.

Release Schedule

- eGA – 8.6.3 – March 1, 2024 – Non-LTS
- eGA – 8.7.0 – June 14, 2024
- pGA – There is no hardware associated with this release

8.7.0 is an LTS Release. This means that it will get planned patches and updates until 8.7.0.x code goes EOS. Non-LTS releases will follow until the next LTS release

Storage Virtualize was previously known as Spectrum Virtualize



8.6.3 and 8.7.0 Updates



- Release Schedule
- Flash Grid – 8.7
- Storage Partition Migration – 8.6.3
 - With SI assist – 8.7
- Replication and HA Updates
- Snapshot Updates
- Ethernet Updates
- GUI Updates
- Code Update Changes
- Auto-space Reclaim in Standard Pools
- Ransomware Threat Detection – FS Awareness
- Misc.
 - 64Gb FC Support FS5300 – 8.7.0

Flash Grid – 8.7.0

- A Flash Grid is a collection of single I/O group FlashSystem or SVC systems that ‘looks and feels’ like a single storage solution, providing single pane of glass management and non-disruptive migration

Flash Grid

- 2Q24
- Storage Insights only
- 4Q24+
- Onboard GUI

IBM Storage Insights

Search

Alerts Tickets Messages Scope Global

Overview / Storage Systems / Pepsi_Hursley

Pepsi_Hursley

Flash Grid ID ABC1234

A Flash Grid is a flexible and scalable federated cluster of storage devices that allows for simple administration and advanced data services.

Capacity

4 Storage Systems

The combined or summed-up storage capacity of all storage resources within an Grouped Systems

Used Capacity **28.54 TiB**

Available Capacity **72.46 TiB**

100 TiB Total

Total Provisioned **67.5 TiB**

Data Reduction **28.54 TiB**

Total Savings **72.46 TiB**

Partitions type

Total partitions 12

Review the distribution of storage partitions across various configurations.

Local	Disaster Recovery	2-site High Availability	3-site High Availability
01	02	09	0

01 Storage Partition under active migration

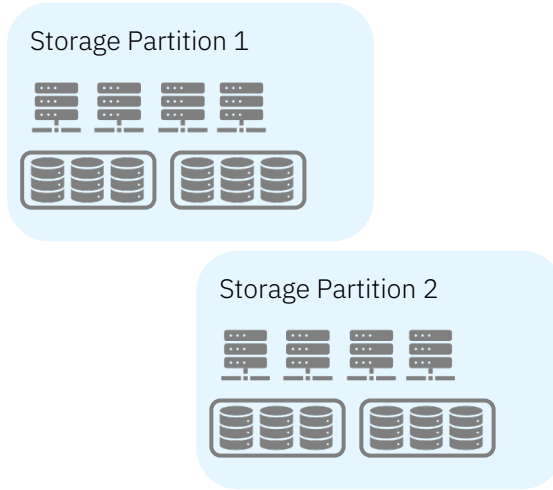
Summary of migration status

In progress: 01 | Migration error: 01 | User action required: 01 | Request submitted: 02 | In queue: 04

Storage systems	Partitions	Active migration	Volumes groups	Host	Used Capacity (GiB)	Snapshots	Contr
tpcflash9100	04	Pepsi_loc_all	01	03	28,927 GiB	Yes	-
tpcflash9000	04	Pepsi_loc_all	02	02	50.88 GiB	Yes	-
tpcflash9500	01	-	02	01	23.88 GiB	Yes	-
tpcflash7200	03	-	01	01	23.88 GiB	Yes	-


100 1 - 100 of 100 items

What is a Flash Grid?

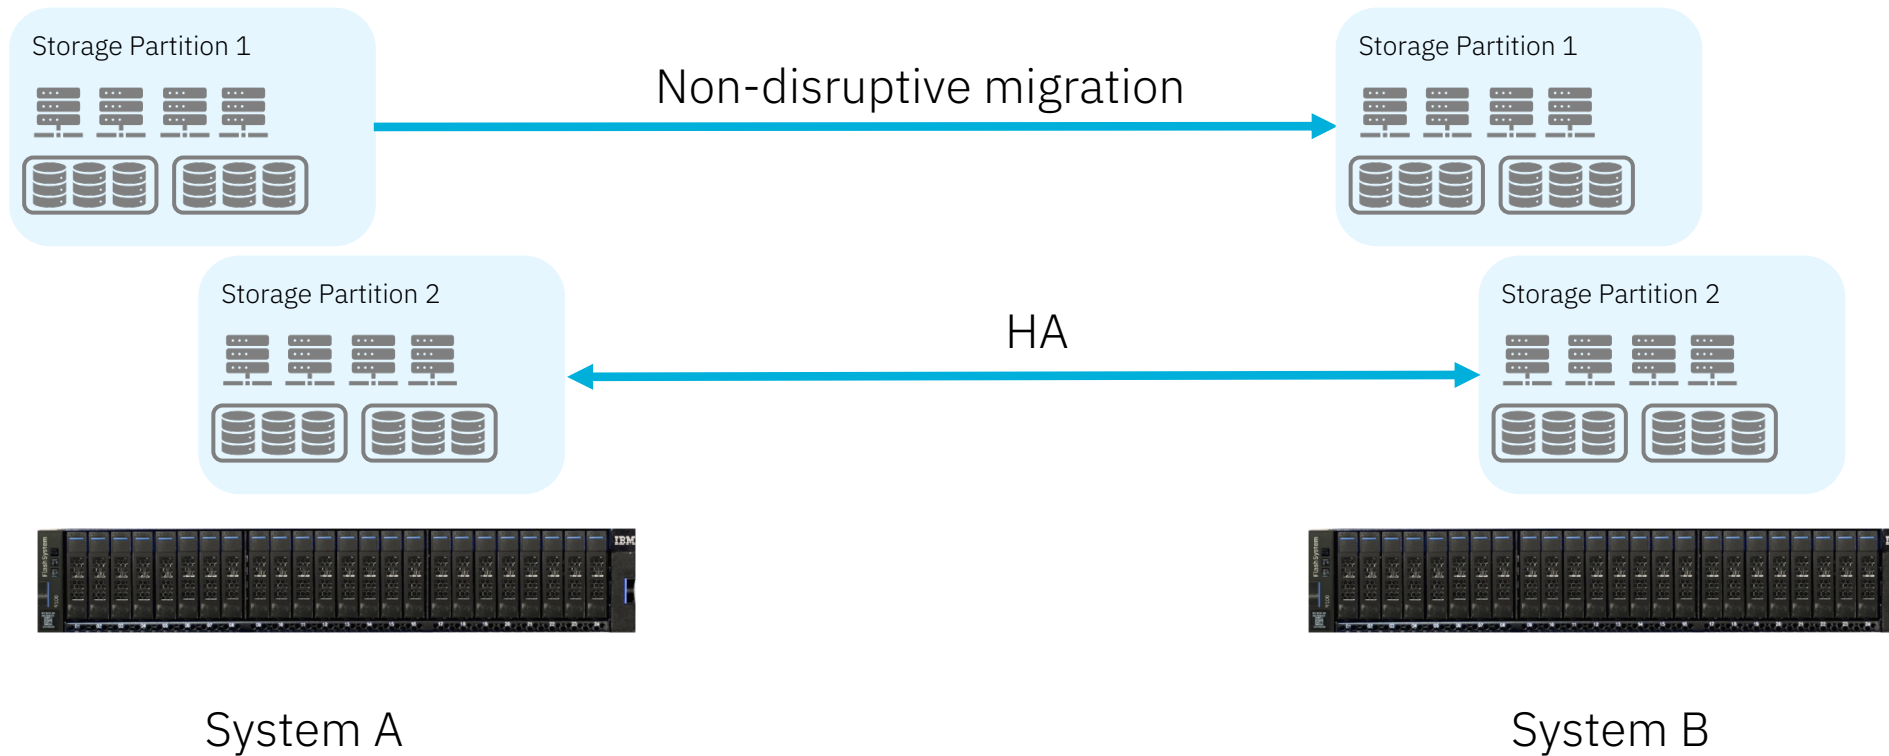


System A

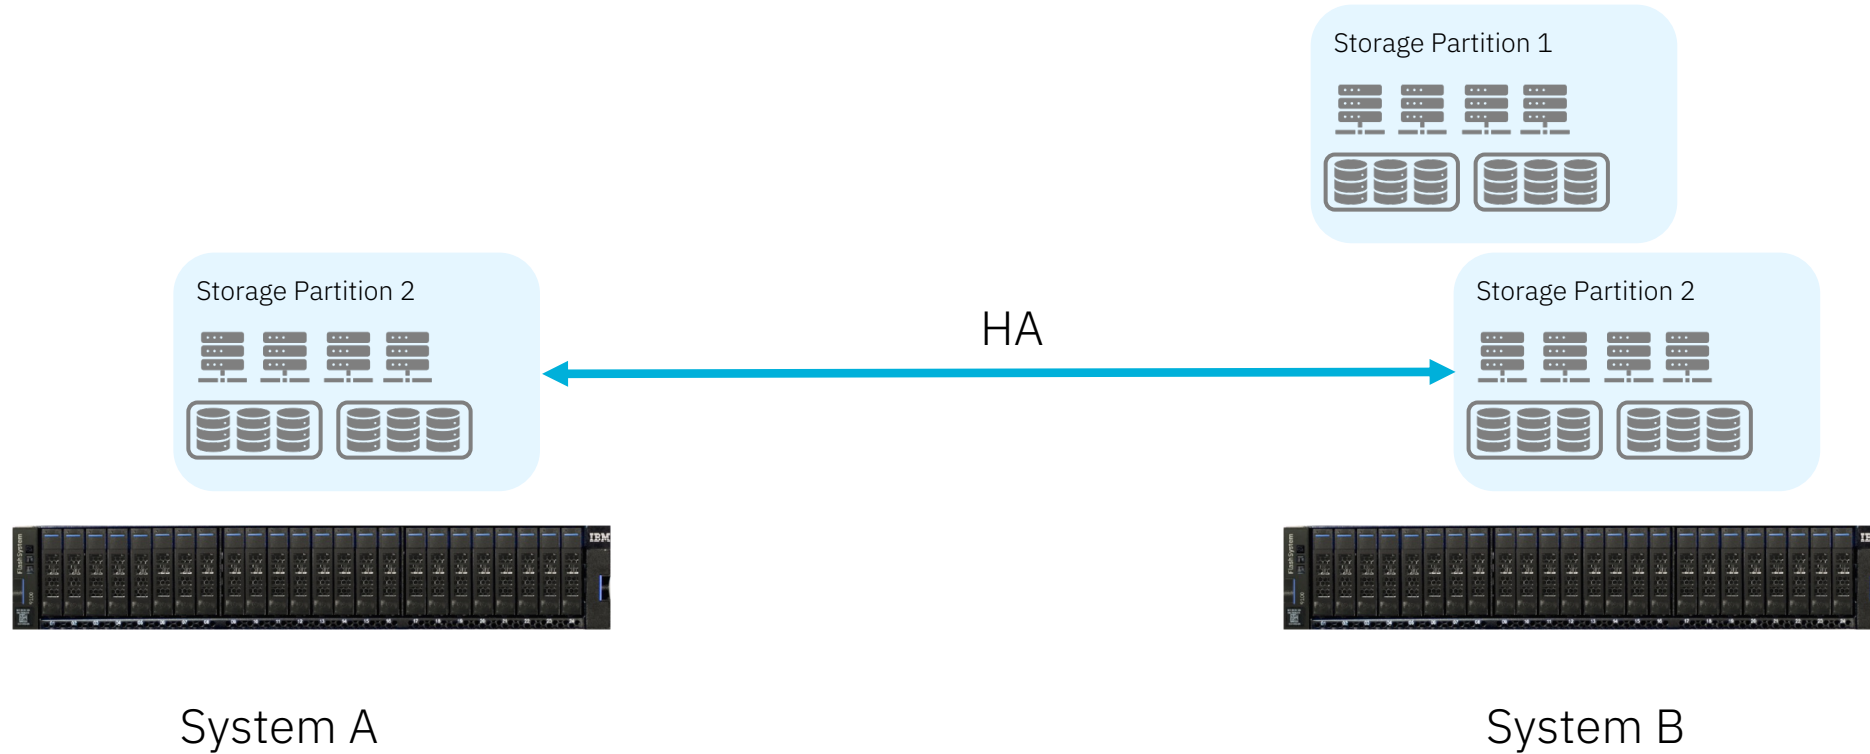
What is a Flash Grid?



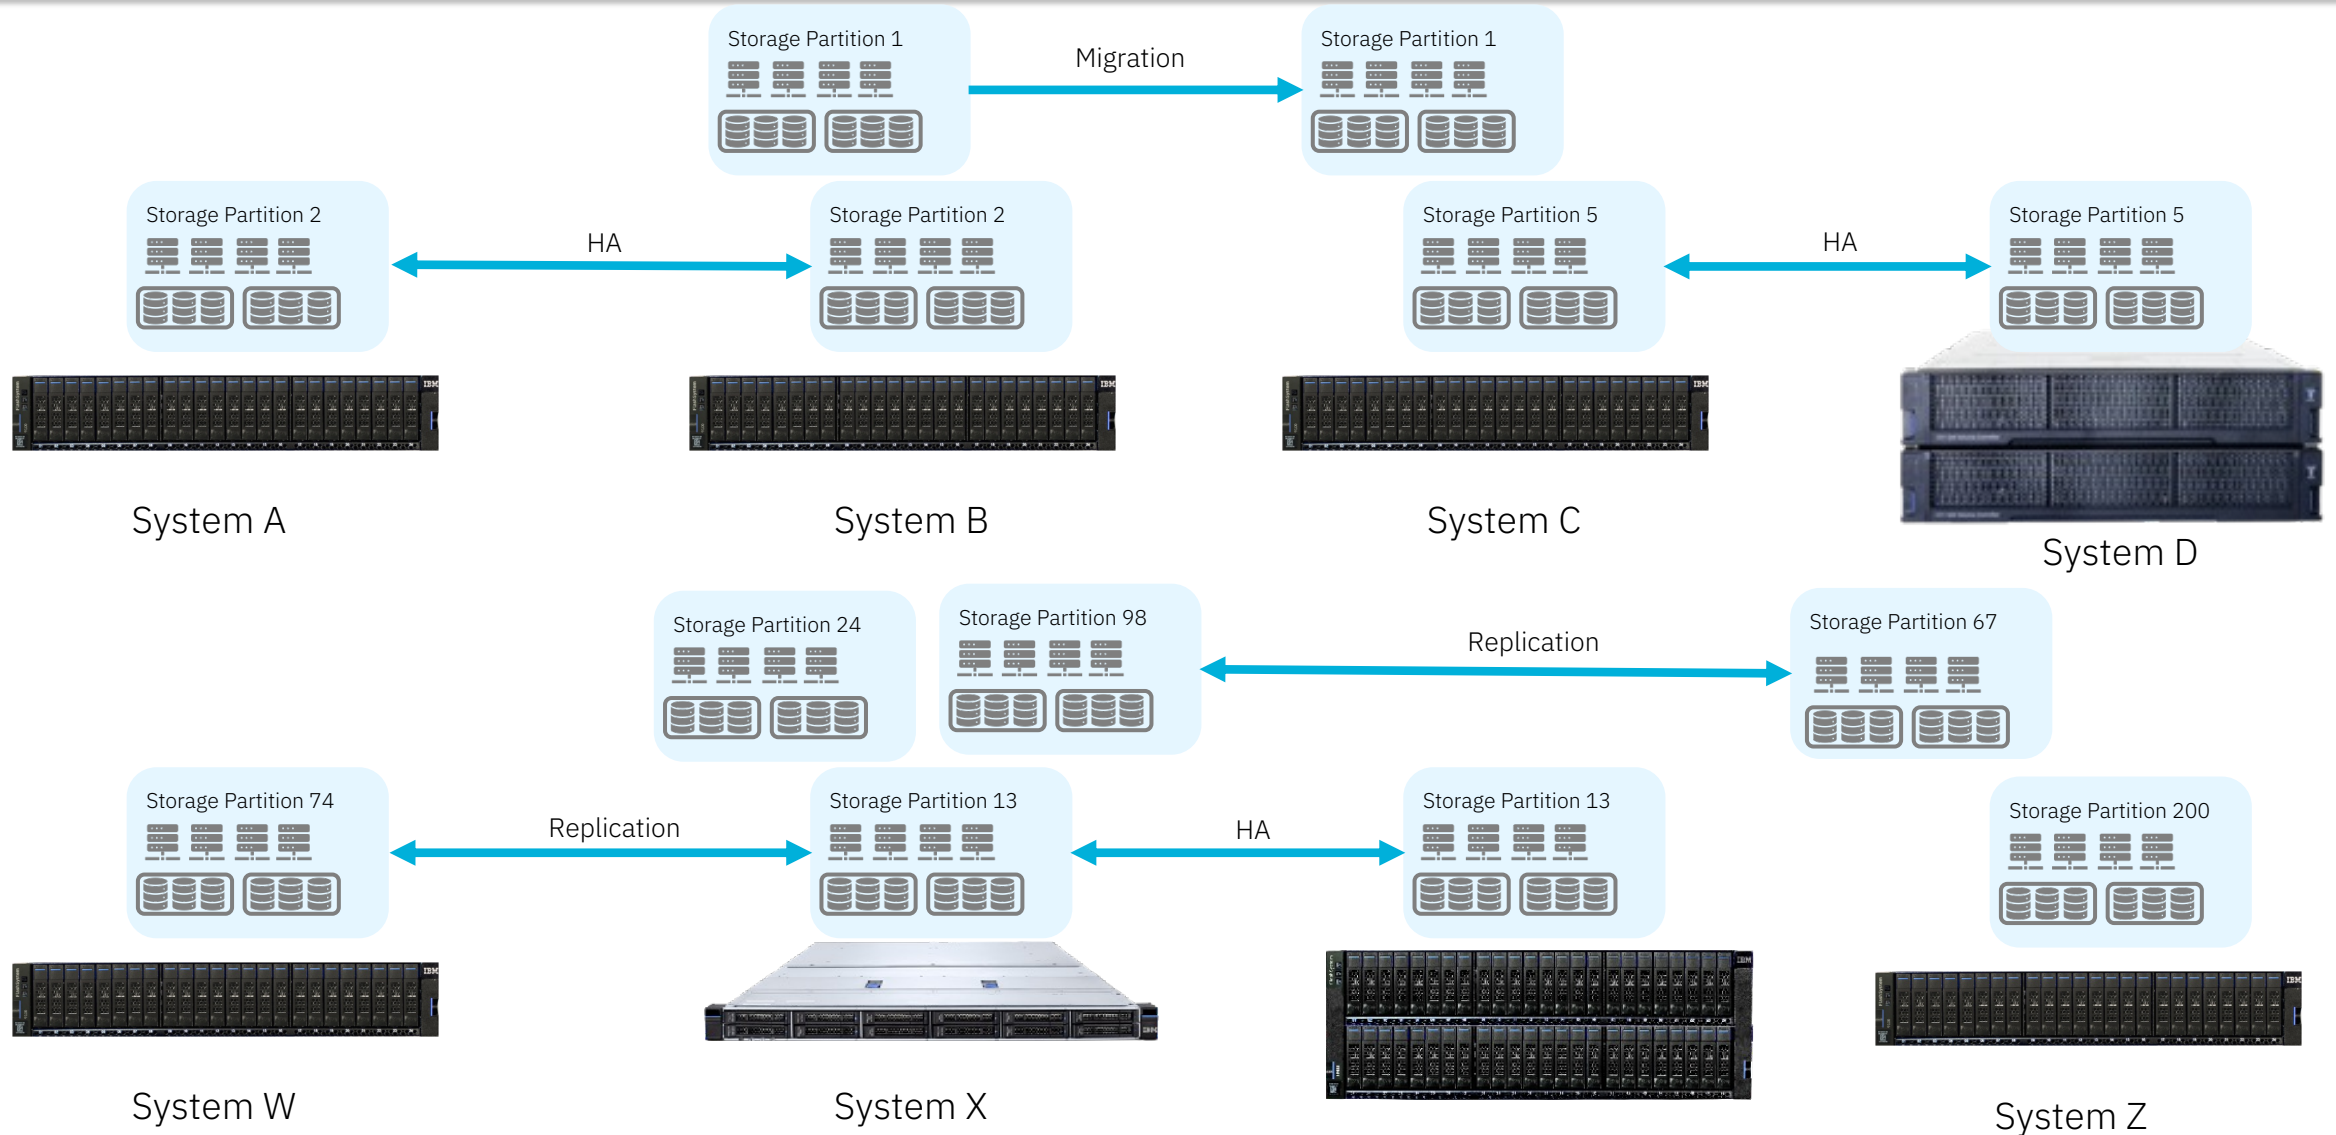
What is a Flash Grid?



What is a Flash Grid?



What is a Flash Grid ?



Traditional scale-out using I/O groups vs Flash Grids

• Clustering

- Up to four I/O groups per system.
- Compatible hardware required
- Single software version.
- Single failure domain (clustering software).
- Object counts are per system, e.g. 32000 production volumes, regardless of I/O groups.
- Non-disruptive migration of volumes between I/O groups.

• Flash Grids

- Up to eight systems in a Flash Grid.
- Mix-and-match from across the portfolio.
- Independent software updates.
- No single point of failure.
- Object counts are per system, allowing for linear scalability of objects.
- Non-disruptive migration of storage partitions between systems.

• Comparison

- Up to 2x increase in capacity and performance with linear increases.
- Simpler hardware upgrades and tiering of performance within a Flash Grid.
- Upgrades can be staged to allow for risk-mitigation.
- Increases availability when using policy-based HA compared to single-system HA options.
- Up to 8x increase in the number of volumes and snapshots, with linear increases in limits.
- Migrations of entire workloads in a single click, powered by AI recommendations.

In 8.7.0...

- CLI-only management of a Flash Grid, consumed by Storage Insights to group systems together and providing AI-assisted migrations of storage partitions.
- Up to 8 systems in a Flash Grid.
- Supported on all NVMe FlashSystem products and SVC.
- The membership of a Flash Grid is managed from a single system (the Flash Grid owner).
- Tiering of performance and capabilities, by mixing hardware within a single Flash Grid.
- Independent software updates for each member system.

```

• manageflashgrid -create -name <name>

manageflashgrid -join -ip <owner ip> -truststore
<id/name>

manageflashgrid -accept -ip <new ip> -truststore
<id/name>

• manageflashgrid -remove <ip>

manageflashgrid -leave

$ lssystem
...
• flash_grid_system_uuid 0EEB2BD0-4A6C-5504-94AC-
0C91682CE84D
• flash_grid_uuid 1EDAB66A-F6F0-5481-88BF-1AAA133E480E
• flash_grid_name IBM Storage Flash Grid
• flash_grid_owner_uuid 0EEB2BD0-4A6C-5504-94AC-
0C91682CE84D

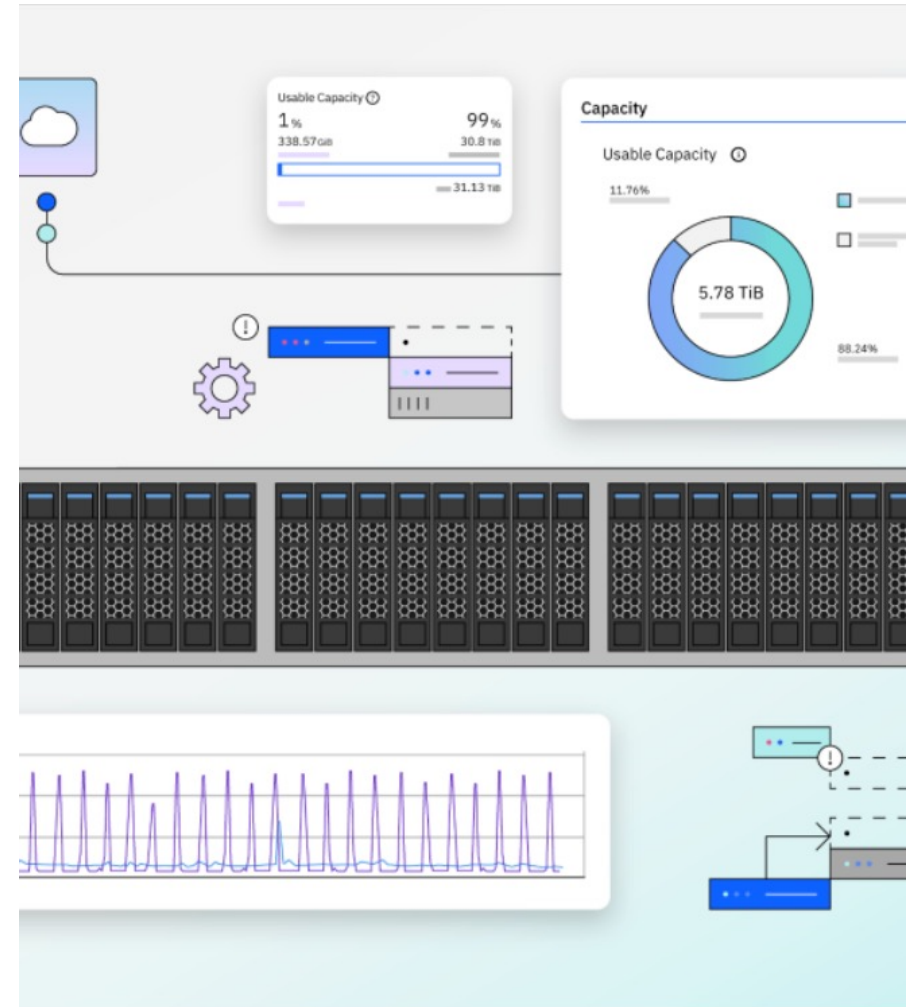
```

In the future...

- Onboard GUI support for managing the entire Flash Grid from a single pane of glass.
- Seamless interaction between Storage Insights and the Flash Grid allowing for AI-based migrations for load-balancing and data placement recommendations.
- Fully automated migration of storage partitions between systems
 - Including the volumes, volume groups, snapshots, policies, hosts, host clusters, mappings, access controls, replication, HA...
- 32, maybe 64, systems per Flash Grid, offering:
 - 1M+ production volumes
 - 2M+ snapshots
 - 1+ Exabyte of manageable capacity
- *Disclaimer: IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.*

IBM Storage Virtualize

Storage Partition Migration



Storage Partition Migration – 8.6.3

- The 8.6.3 release delivers Storage Partition Migration. With this feature IBM® Storage Virtualize provides nondisruptive migration of the storage partition across Storage Virtualize devices.
- Think of this feature as Volume Mobility for an entire partition.

Storage Partition Migration Provides:

- Non-disruptive Upgrade of system hardware from older systems to newer ones regardless of model (e.g. FS5200 to FS7300)
- Load-balancing by migrating Storage partitions from overloaded systems to other systems
- Use case-specific scenarios which requires a storage partition migrated from a source system to a specified target system

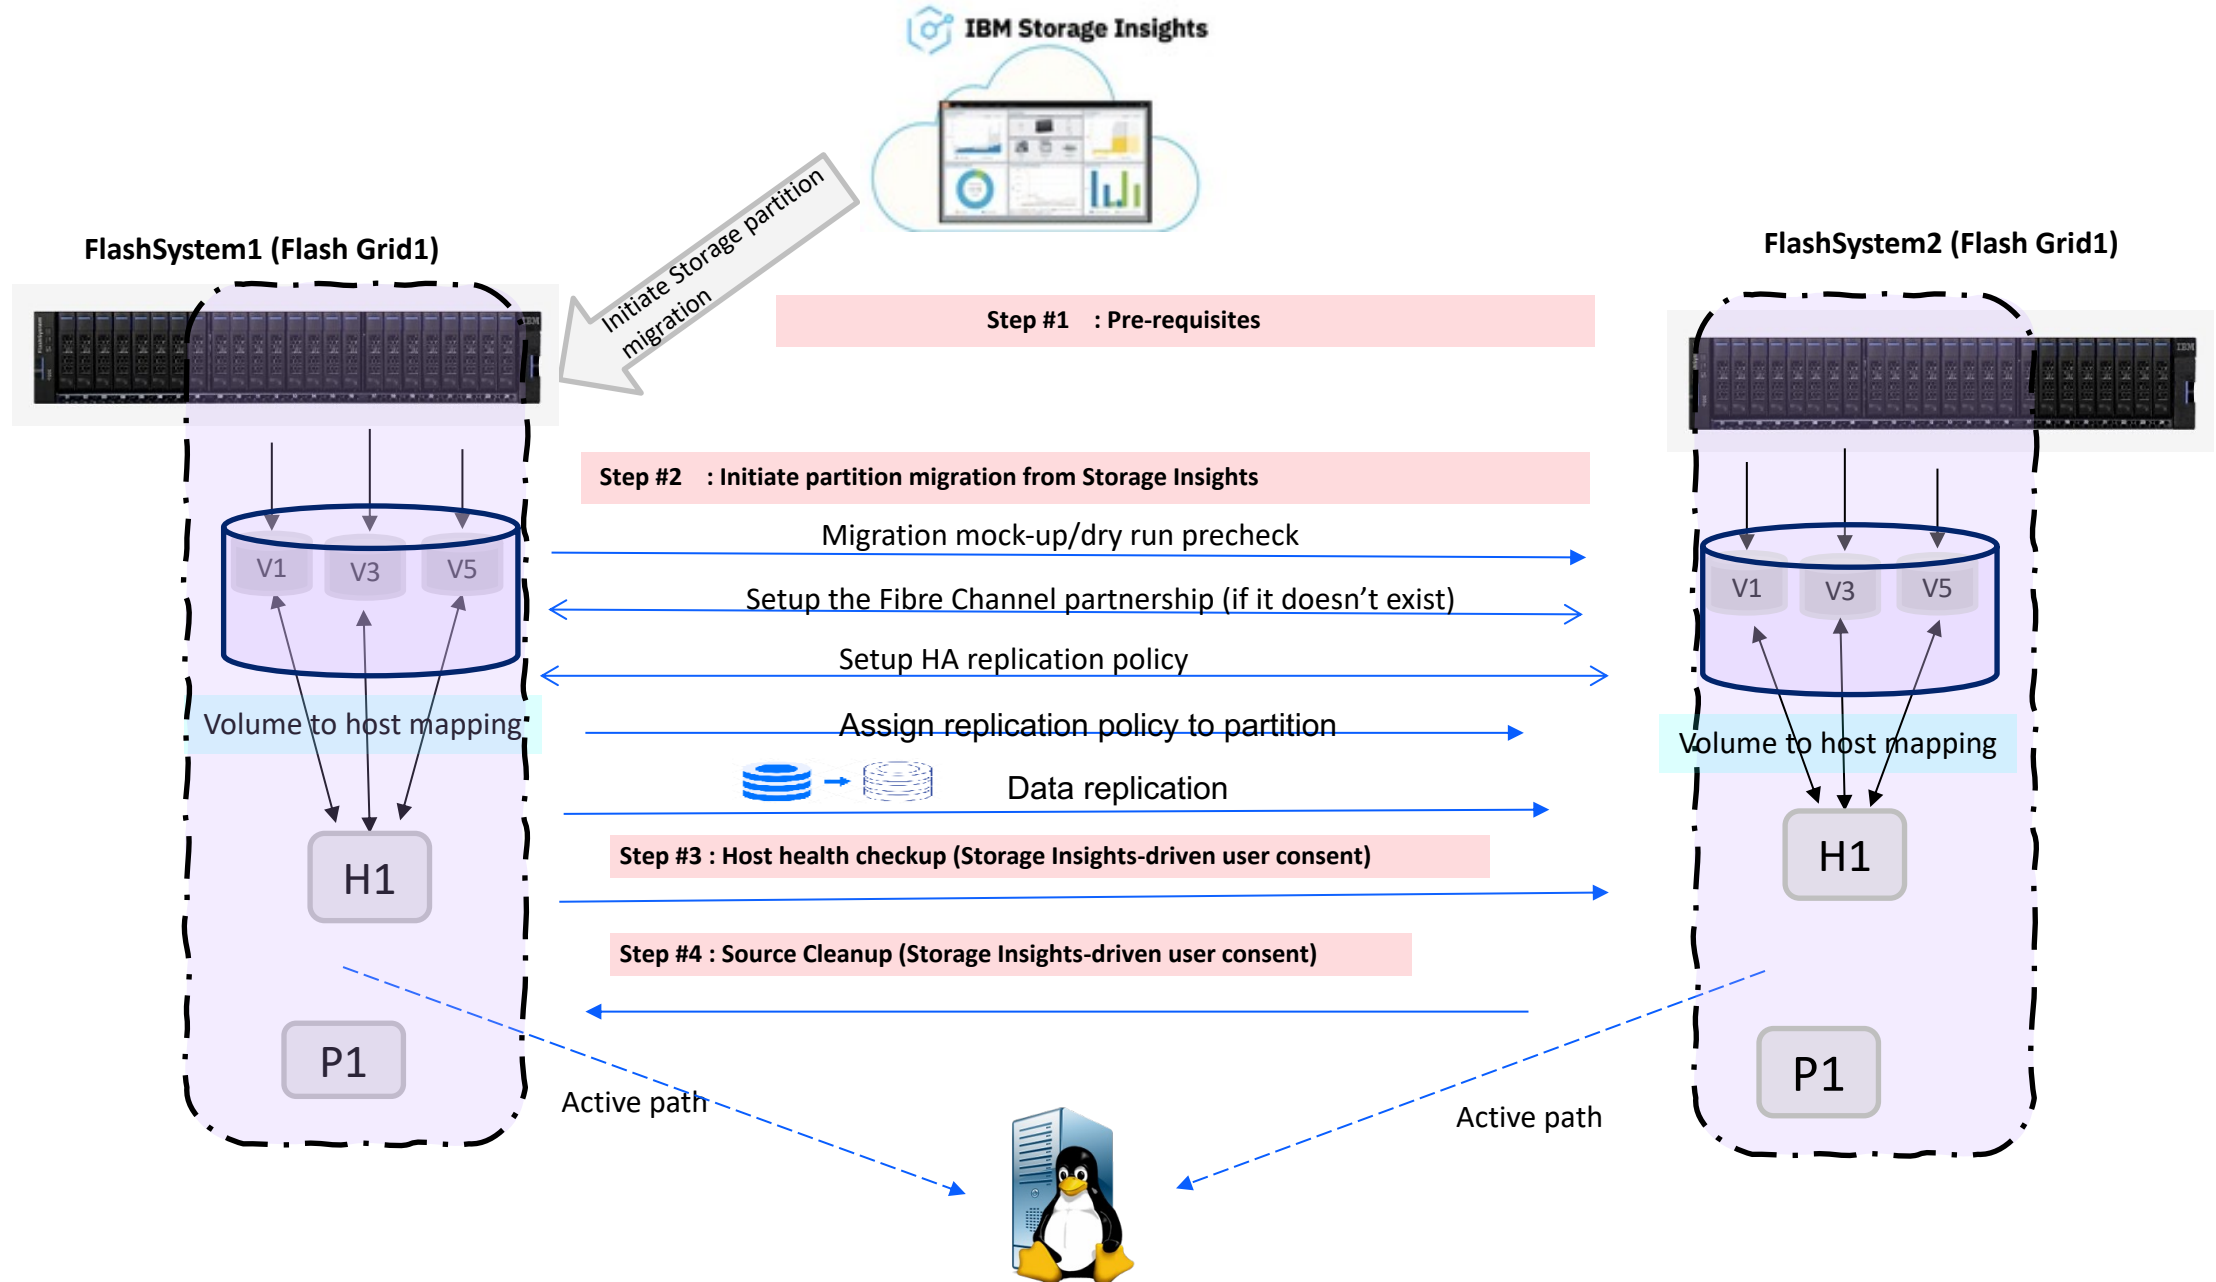
Storage Partition Migration Workflow

- All Storage Partition associated objects hierarchy such as volumes groups, volumes, hosts, host mappings are relocated to the target IBM® Storage Virtualize system non-disruptively
- Host IO shall be now served from the target IBM® Storage FlashSystem after a manual rescan of Host IO paths from the target
- The CLI “svctask chpartition” is the sole CLI to trigger a migration. This makes triggering the migration Uncomplicated and direct.
- As this is a CLI only release we have a developed event-based notification for various migration stages involving manual interventions (e.g. Host IO path detection for new Storage, Confirm Source Storage Partition Copy deletion)
- Events are used to trigger pauses until ready to continue
- The event-based notification will also inform the Storage Admin of any faults during the Storage Partition Migration process
- Using the event-based notification feature we have also allowed the Storage Partition Migration rollback option before we can confirm migration completion

Limitations and Restrictions

- Currently draft storage partition objects, snapshots, HA, or DR relationships are not supported for automated migrations...yet
- Only one storage partition can be migrated at a time from a system
 - Any consecutive migrations that are attempted get queued and are scheduled automatically as per the sequence of invocations when the earlier migrations complete
- This 8.6.3 (1Q24) release of Storage Partition Migration is accessible only via the CLI
- Automatic IO Path detection from the Host to the target is not available in this release
 - Manual rescan of paths are needed based on event raised during migration.
- The Flash Systems GUI is not able provide complete control and visualization of the feature
- IBM Storage Insights provides complete visualization and control of the feature in the 8.7.0 release
- Currently snapshots and policies are not transferred. The snapshots remain on the source system until they expire.
- You may have orphaned snapshots

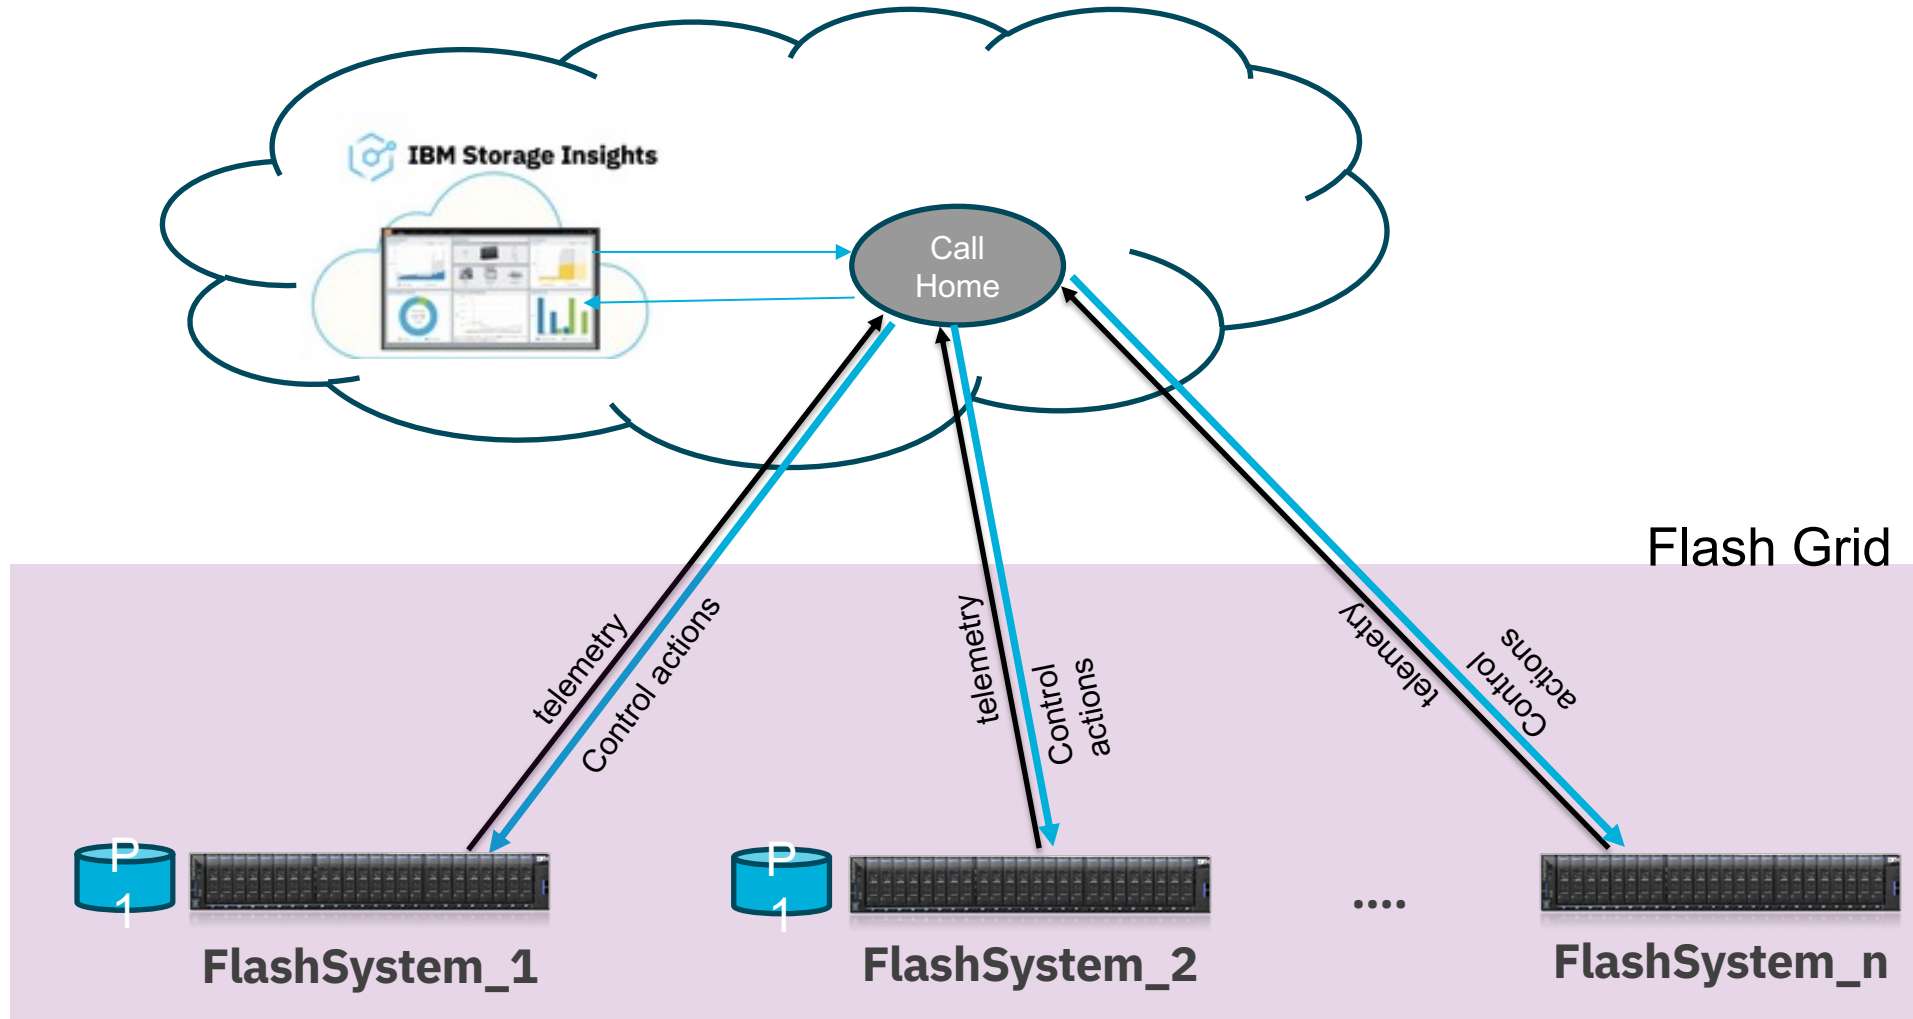
Storage Partition Migration Workflow – 8.7



Storage Insights Integration Highlights – 8.7.0

- Control actions provided to Storage Insights to manage partition migrations:
 - Initiate (Initiate partition migration)
 - Abort (Abort on-going/queued partition migration)
 - Host rescan fix (User consent for host path rescan)
 - Commit (User consent for deleting source copy)
 - Rollback (User consent for switching control/PML to source cluster)
- Integration of storage partition migration within Flash Grid with Storage Insights
- Adds control support in Storage Virtualize to allow control actions from Storage Insights for storage partition migration
 - If desired – This option is not mandatory
- Enhanced storage partition migration telemetry for Storage Insights
- Storage Insights to provide migration advisory for storage partition migrations

Storage Environment



Prerequisites For Storage Partition Migration Using SI

- Participating systems should be added in Storage Insights
- Participating systems must be in same Flash Grid
- Pool linking between participating systems must be configured
- Storage Insights control access must be enabled on participating systems
- Storage Insights Users must be given permissions to migrate partitions

Examples

Storage systems	Health	Alerts	Events	Data collection	Used capacity	Partitions	I/O rate read	I/O rate write	
ATG_FlashGrid... 02 systems	✔	✔	i 05	✔ 02/02	2.79 GiB (0%)	-	<div style="width: 10%;"></div>	<div style="width: 10%;"></div>	⋮
FS7300-2	✔	✔	i 04	✔ Running Refreshed A few seconds	2.67 GiB (0.01%)	3	<div style="width: 20%;"></div>	<div style="width: 20%;"></div>	⋮
FS9100-10	✔	✔	i 01	✔ Running Refreshed A few seconds	0.12 GiB (0%)	3	<div style="width: 20%;"></div>	<div style="width: 20%;"></div>	⋮

```

IBM_FlashSystem:FS7300-2:byrongro>lssystem |grep control
storage_insights_control_access yes
IBM_FlashSystem:FS7300-2:byrongro>
    
```

User Access Management

Manage additional privileges and REST API key access

Name	Role	Additional Privileges
<input type="checkbox"/> BYRON GROSSNICKLE	ADMIN	Partition Management

Items per page: 20 | 1-20 of 39 items

Storage Partitions

Total partitions: 03
A storage partition is a logical grouping of volume groups, volumes and hosts that can have high availability configured or migrated non-disruptively between systems.

[Why can't I perform migration actions?](#)

Partition name	Partition type	Migration status	Host ...	Snapshots	Hosts	Volume groups	
<input type="checkbox"/> FS7300-2_to_FS9100-10	2-site-ha		0.00 GiB	Yes	00	01	⋮
<input type="checkbox"/> PB_HA_1	2-site-ha		0.00 GiB	Yes	02	01	⋮
<input type="checkbox"/> Non-HA Partition	local		48.62 ...	Yes	01	01	⋮

Items per page: 10 | 1-3 of 3 items

1 of 1 page

- Block Storage (4)**
- Device Name
- cluster_Cabaro_216
 - cluster_Cabaro_219
 - Source_FS5200
 - Target_FS5200

Overview

Block Storage

Overview of 4 Storage systems

Resources
All resources are normal

Call Home with cloud services 2


Storage Systems
4

Fabrics
0

Hosts
3

2 6

Capacity
Usable Capacity, Data Reduction and Provisioned Data



Data Reduction
0 Bytes

Total Savings
0 Bytes

Data Provisioned
120.00 GiB

Performance Deviations

This experience will be live soon

Sales_records 4 Hours Ago	Increased Write IOPS 200% higher than usual	200 Ops/s Usual is 600 Ops/s	
Lab_storage 2 Hours Ago	High Read Throughput 52% higher than usual	345 MBps Usual is 200 MBps	
Oracle_store 4 Hours Ago	Decreased Write IOPS 82% higher than usual	100 Ops/s Usual is 245 Ops/s	

Workload Protection
Derived from 4 of 4 applicable systems

Volumes Mapped to a Host 6 Total

System Protection
Derived from 4 FlashSystems

Storage Systems 4 FlashSystems


Environment
Derived from 4 of 4 applicable systems

Power | Temperature

Storage Insight Dashboard

IBM Storage Virtualize

Replication/HA Updates



Partition Concepts

- 8.6.1 introduced the concept of partitions to a system
 - Partitions are a way of sectioning off a system for certain functions to be available on part of a system
 - Such as Policy Based HA (PBHA)
 - A system can have up to 4 partitions
 - A partition can be HA or non-HA
 - There was no way to get existing things into a partition
- **New concepts in 8.6.3**
 - Draft Storage Partition
 - Merging Partitions

Draft Storage Partitions – 8.6.3

- Used to get existing configurations into partitions so that features like Policy-based HA or Storage Partition Migration can be used without production interruption

Draft Storage Partition Requirements

- Volume Groups/Hosts are moved to the draft partition
 - You must have at least 1 VG
 - Volumes that are not in a VG that will be affected by the move can be dealt with in the wizard during the process
- VGs cannot have a replication policy on them

Starting configuration, example: 3 volume groups, 4 volumes

The screenshot shows the management interface for an IBM FlashSystem 7200. The header indicates the system name 'IBM FlashSystem 7200', the controller 'gobibear-c', and the current view 'Volumes'. Below the header, there are navigation options: 'Create Volumes', 'Actions', and 'All Volumes'. A table lists the configured volumes with their respective states, pools, volume groups, and protocol types.

Name	↑	State	Pool	Volume Group	Protocol Type
vol0		✓ Online	pool0	volgrp0	SCSI
vol1		✓ Online	pool0	volgrp1	SCSI
vol2		✓ Online	pool0		SCSI
vol3		✓ Online	pool0	volgrp2	SCSI

Starting configuration, example: 5 volume mappings to 3 hosts

The screenshot shows the 'Mappings' page for an IBM FlashSystem 7200. The page title is 'IBM FlashSystem 7200 gobibear-c Mappings'. There is a dropdown menu for 'Private Mappings' and an 'Actions' menu. Below is a table with 5 rows of volume mappings.

Host Name	SCSI ID	Volume Name	UID
host0	0	vol0	6005076810BC840778000000000000000
host1	1	vol1	6005076810BC840778000000000000001
host1	0	vol0	6005076810BC840778000000000000000
host1	2	vol2	6005076810BC840778000000000000002
host2	0	vol3	6005076810BC840778000000000000003

Create Draft Partition

IBM FlashSystem 7200

- Dashboard
- Monitoring
- Pools
- Volumes
- Hosts
- Copy services
- Policies
- Access
- Settings
- Storage partitions**

Create new storage partition

Complete the steps to create new partition

- Storage partition basics
- Add existing resources
- Summary Review

Storage partition basics

Select the type of partition you would like to create and enter the required information in the form fields below.

High Availability Disabled

Storage partition name
Enter partition name

Add existing resources to the new storage partition. This cannot be done at a later time

Cancel Skip Continue

Select Existing Volume Group(s)

Select volume groups

Select from the list of existing volume groups in the system to add to this new storage partition.

Search

<input type="checkbox"/> Volume group	Volume Count	Host Count	Policies
<input checked="" type="checkbox"/> volgrp0	1	2	
<input type="checkbox"/> volgrp1	1	1	
<input type="checkbox"/> volgrp2	1	0	

Items per page: 25

1-3 of 3 items

1 1 of 1 page

Volume Group Selected

Create new storage partition

Complete the steps to create new partition

Storage partition basics

Setup partition resources
Optional

Select and review volume groups

Assign unattached volumes

Summary Review

Select existing volume groups (optional)

Selected volume groups (1) Automatically included volume groups (1)

Search

Select volume groups +

Volume group	Volume Count	Host Count	Policies
volgrp0	1	2 mapped hosts	

Items per page: 25 1-1 of 1 item 1 1 of 1 page

Create new storage partition

Complete the steps to create new partition

Storage partition basics

Setup partition resources
Optional

Select and review volume groups

Assign unattached volumes

Summary Review

Select existing volume groups (optional)

Selected volume groups / **Automatically included volume groups (1)**

Search

Volume group	Volume Count	Host Count	Policies
volgrp1	1	1 mapped host	

Items per page: 25 1-1 of 1 item 1 1 of 1 page

Volumes not in Volume Group must be added to new VG

Create new storage partition

Complete the steps to create new partition

- ☑ Storage partition basics
- ☑ Setup partition resources
Optional
- ☑ Select and review volume groups
- 🔵 Assign unattached volumes
- ⊙ Summary Review

Create volume group for existing volumes

Partitions require volumes to be attached to a volume group.

i All unassigned volumes must be included in a volume group
 The volumes listed below are mapped to hosts attached to the new partition, but are not in a volume group. These volumes will be added to the volume group created in this step. Cancel and save a draft of your progress if you do not want to create a default volume group.

Volume group name

volgrp3

Volume	Pool
vol2	pool0

Items per page: 10 ▾ 1-1 of 1 item

Partition is Published

Create new storage partition

Complete the steps to create new partition

✔ Storage partition basics

✔ Setup partition resources
Optional

✔ Select and review volume groups

✔ Assign unattached volumes

🔄 Summary Review

Summary review

Review all the components and configurations that have been setup for policy-based replication below.

✔ **Partition basics created**
myPartition1 created

✔ **Existing volumes group selected**
Expand to view all selected volumes groups

✔ **Unattached volumes assigned**
volgrp3 created

✔ **Local storage partition has been setup!**
You can now view your partition and its components after exiting this process.

Make HA (via View Partition)

The screenshot displays the IBM FlashSystem 7200 management interface. The top navigation bar shows 'IBM FlashSystem 7200', 'gobibear-c', and 'Storage partition: myPartition1'. The left sidebar contains a navigation menu with options like 'myPartition1', 'Storage partition overview', 'Volumes', 'Hosts', 'Copy services', 'Policies', and 'IP quorum'. The main content area is titled 'myPartition1 overview' and features a 'Manage partition' button in the top right. A central message states 'High availability replication is not configured.' with a red circle highlighting the 'Add high availability replication +' button below it. To the right, the 'Storage components' section lists 'Hosts (2)' (Online), 'Volumes (2)', and 'Volume groups (2)' (Online). The bottom status bar shows performance metrics: Latency 0 ms, Read 0 ms, Write 0 ms, Bandwidth 0 MBps, Read 0 MBps, Write 0 MBps, and IOPS 0 (Read 0, Write 0).

Merging Storage Partitions – 8.6.3

- Used to combine storage partitions
- Can be used with HA and non-HA partitions
 - Partitions MUST have the same properties
 - Same HA policy, etc

Before Merge 2 Storage Partitions Exist:

The screenshot shows the IBM FlashSystem 7200 management interface. The top navigation bar includes 'IBM FlashSystem 7200', 'gobibear-c', and 'Storage partitions'. The main content area is titled 'All storage partitions' and displays two partition cards: 'myPartition1' and 'myPartition2'. Each card shows a table of 'Hosts' and 'Volumes', and a 'Volume groups' section. A 'View partition' button is at the bottom of each card. For 'myPartition2', a context menu is open, showing 'Merge partition' and 'Delete' options, both of which are circled in red. The 'Merge partition' option is also circled in red.

Hosts	Volumes
2	3

Hosts	Volumes
1	1

Volume groups: 3

Volume groups: 1

Select Target Storage Partition:

Merge partition ×

Select a partition that you would like to merge with. The objects in the current partition will be moved into the selected target partition and the current partition will be deleted.

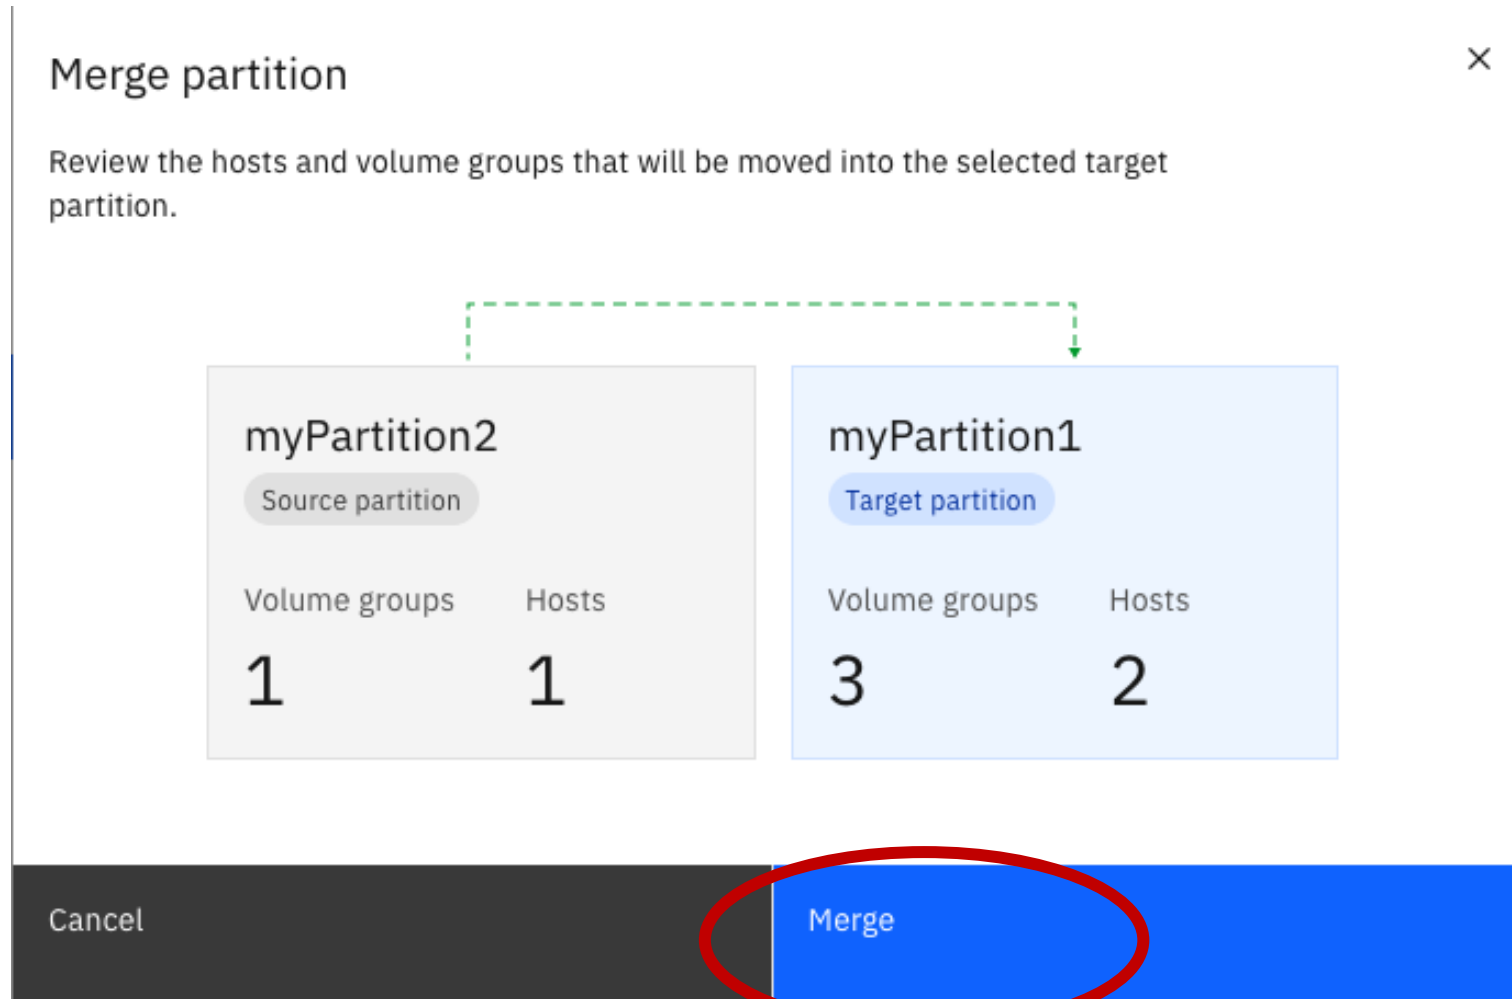
Select target partition

myPartition1 ∨

This is the partition that objects will be moved into

Cancel Next

Diagram:



Merge Result (select View Partition to make High Availability):

The screenshot shows the IBM FlashSystem 7200 management interface. The top navigation bar includes a hamburger menu icon, the text 'IBM FlashSystem 7200', the identifier 'gobibear-c', and the current page title 'Storage partitions'. On the left side, there is a vertical sidebar with several icons representing different system components. The main content area is titled 'All storage partitions' and displays a card for 'myPartition1'. This card contains a table with the following data:

Hosts	Volumes
3	4
Volume groups	4

Below the table, there is a blue button labeled 'View partition', which is circled in red. A vertical ellipsis menu icon is visible to the right of the partition name.

Notes

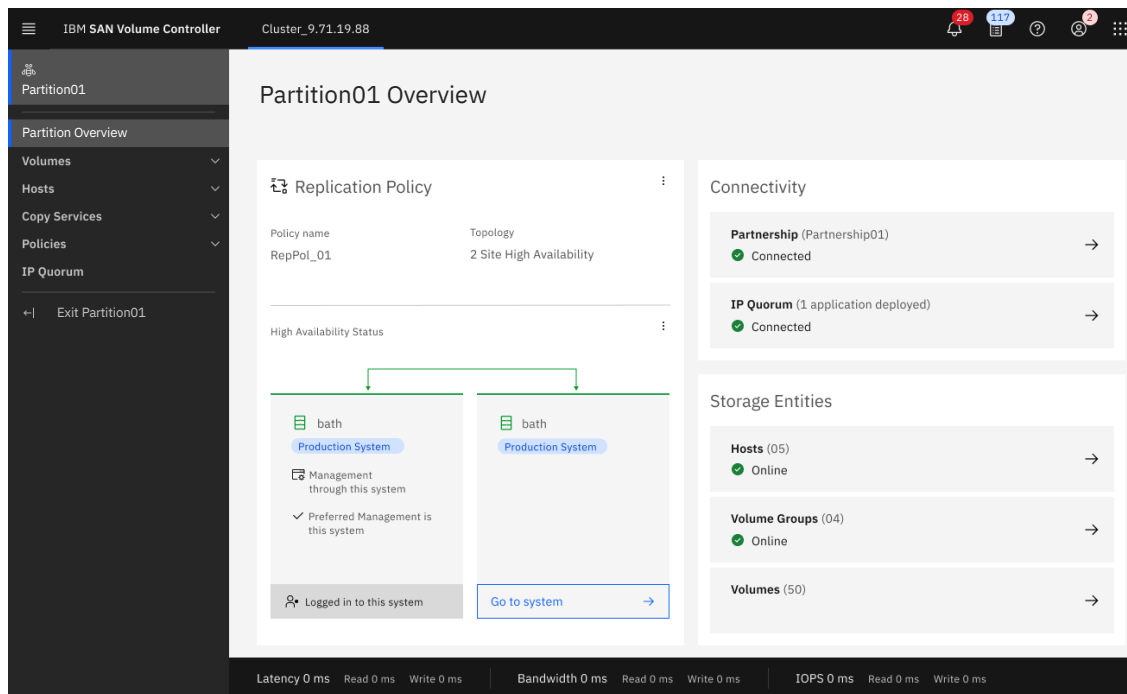
- There is no way to take existing unpartitioned resources and put them into an existing partition
 - Create a draft partition with the resources with the same partition type of the desired partition
 - Merge the partitions
- There is no way to move resources between partitions
 - Future plan
- Today there are only 4 partitions available
 - Future plans for more

Byron's Recommendations

- Only use permanent partitions for resources needing PB-HA
- Leave Non-PB-HA resources out of partitions
 - No way to extract or move individual resources between partitions
 - Only option is to delete the whole partition
- Use temporary partitions to migrate storage resources
 - Migrate resources into a partition
 - Migrate storage partition to new system
 - Delete storage partition

Storage Partitions (recap)

Storage partitions simplify the configuration, management and monitoring, with a single point of control for HA.



- Create and populate

- Create a storage partition in a few clicks, and create or add any number of hosts, volume groups and volumes.

- Policy-based

- Add a HA replication policy to the partition, and everything within it will automatically be configured to be highly available, with all remote provisioning handled automatically.

- Highly available


- In a disaster, the storage partition automatically manages its availability ensuring that applications remain accessible, with per-partition control over which system is preferred in the event of a loss-of-connectivity.

- Scoped

- HA problems have zero impact on any non-HA volumes: HA and non-HA storage can happily co-exist on the same system.

PB-HA Ground-up Implementation (recap)

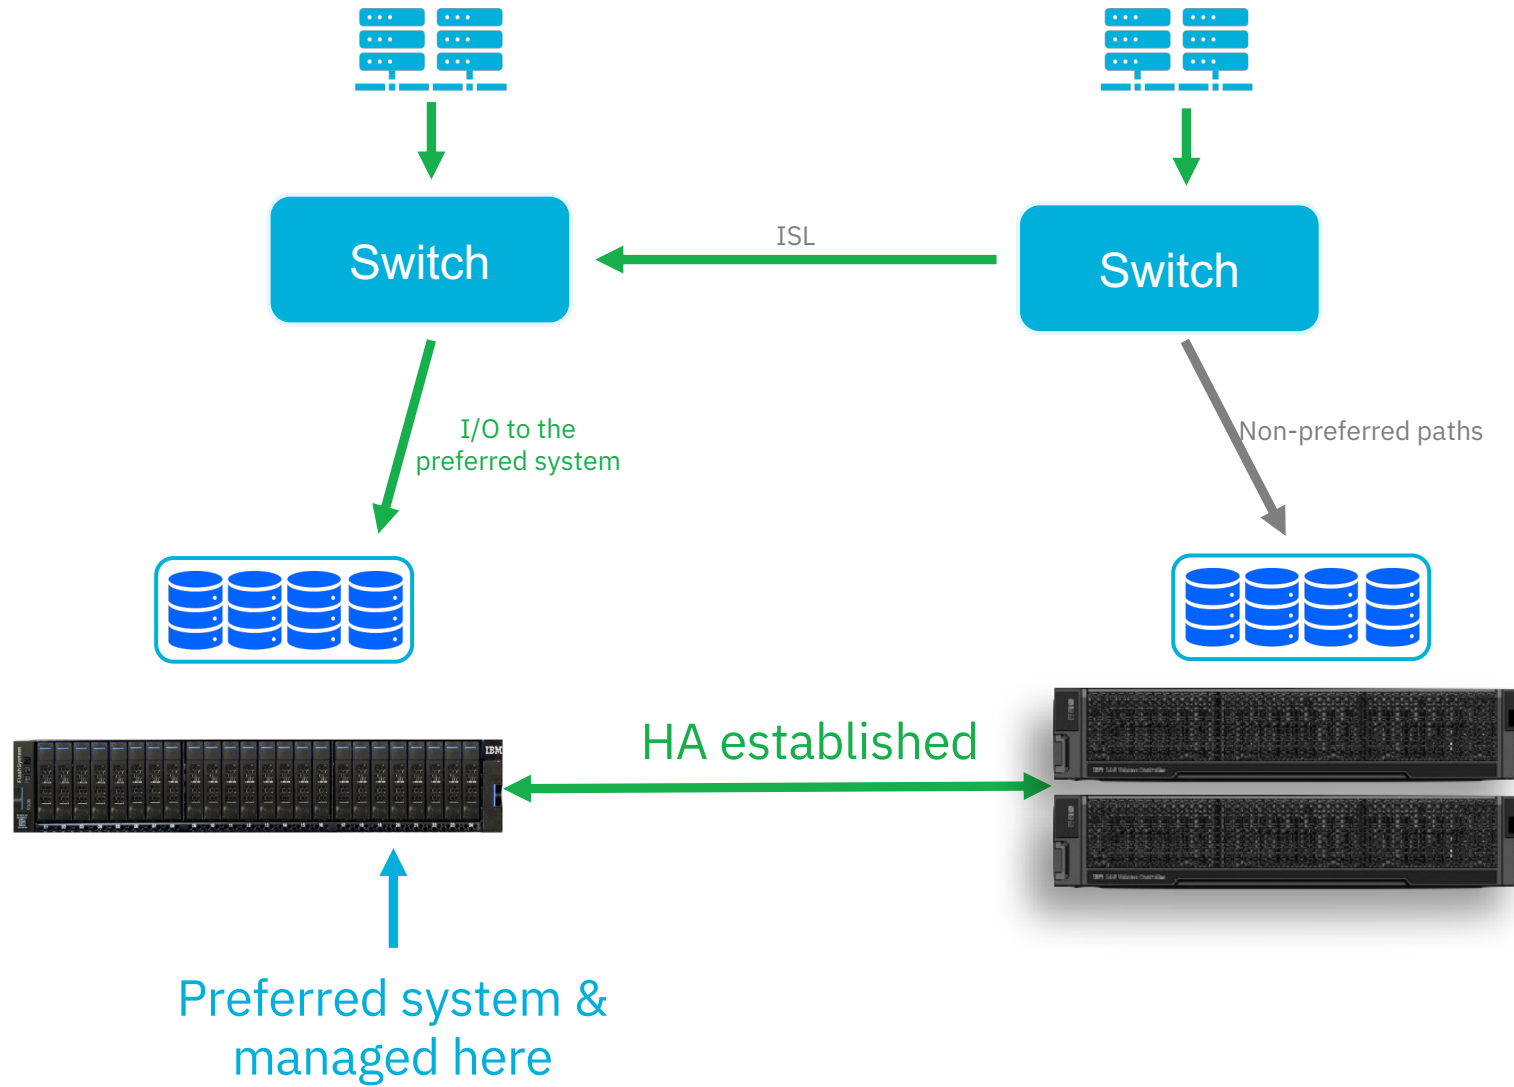
- Similar to policy-based replication, a ground-up re-write provides a significant boost in performance vs HyperSwap.
- HA performance is close to the not-HA performance, maximizing the investment.
- Two FlashSystem 9500 systems (one I/O group each) connected via FC, 16KiB 50%r/50%w/50% hit at negligible distance



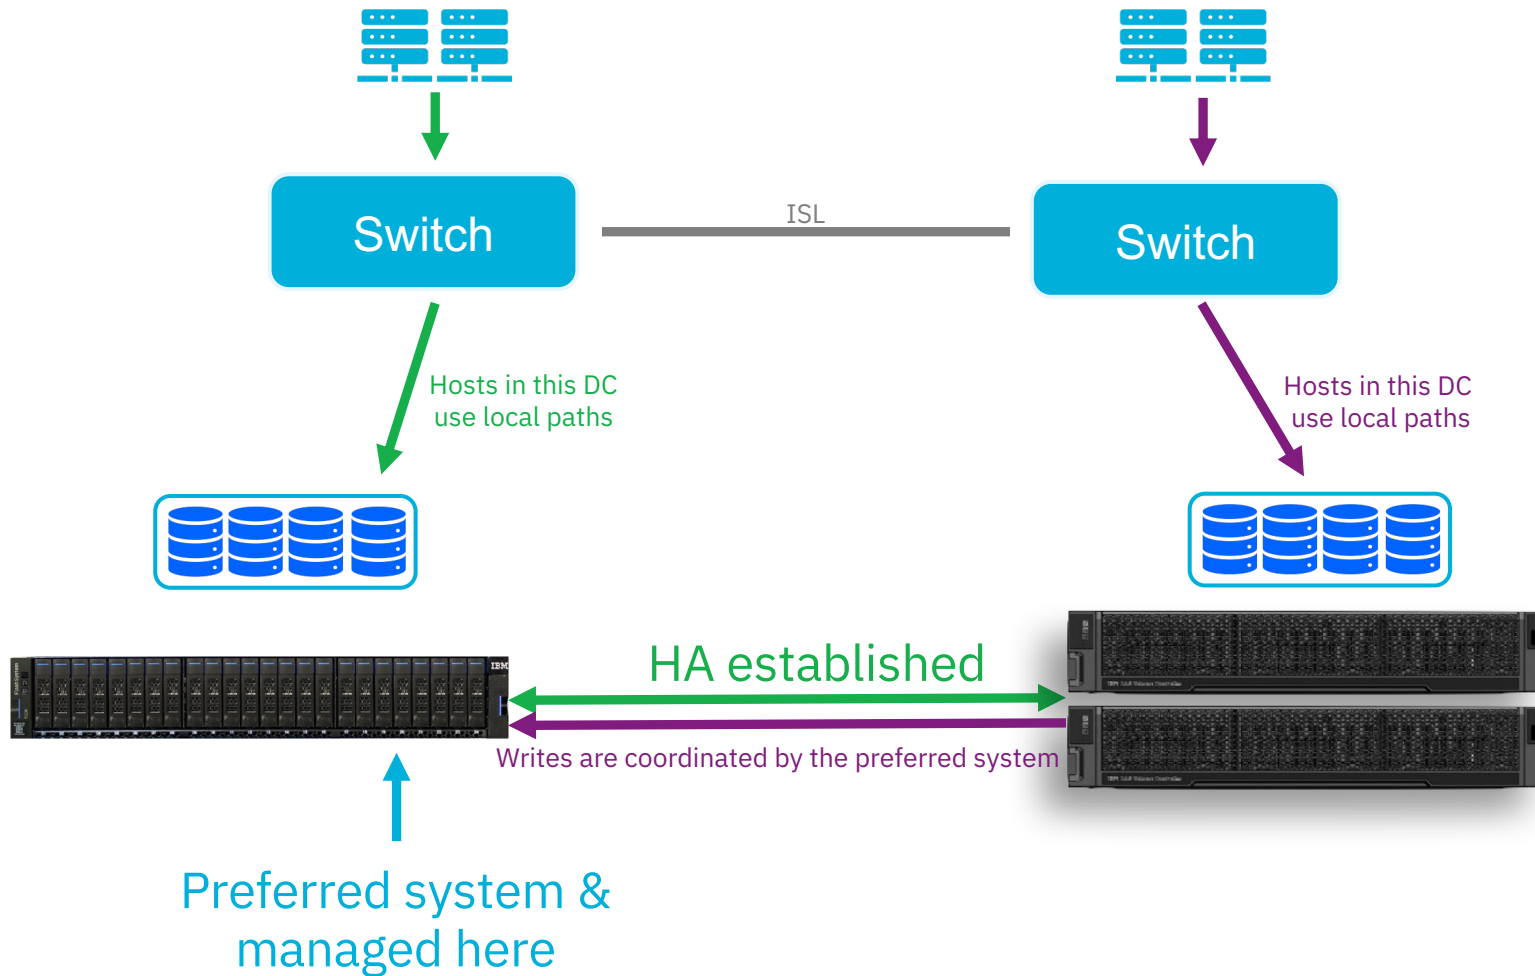
Active/Active HA – 8.7.0

- All HA volumes will automatically transition from Active/Passive to Active/Active when BOTH systems have completed upgrade to 8.7.0
- Volumes are active/active only when HA is established (initial copy done)
- Writes to the non-preferred system have an additional control round trip, but the data only crosses the ISL once
- If a location is set for a host, the volumes report preferred access to the system in the same location
- Allows for the hosts located at a specific physical location to read and write to the locally connected system, reducing ISL traffic and latency
- Allows for both ‘uniform’ and ‘non-uniform’ setups
- Supports the following Fibre Channel SCSI operating systems:
 - Red Hat Enterprise Linux 7 and later
 - VMware ESXi 7 and later
 - AIX 7.2 and later (**new in 8.7.0**)
 - Post-GA support planned for these operating systems using FC SCSI (submit SCORE requests as required):
 - Windows servers/Hyper-V
 - IBM i
 - SCSI persistent reservations

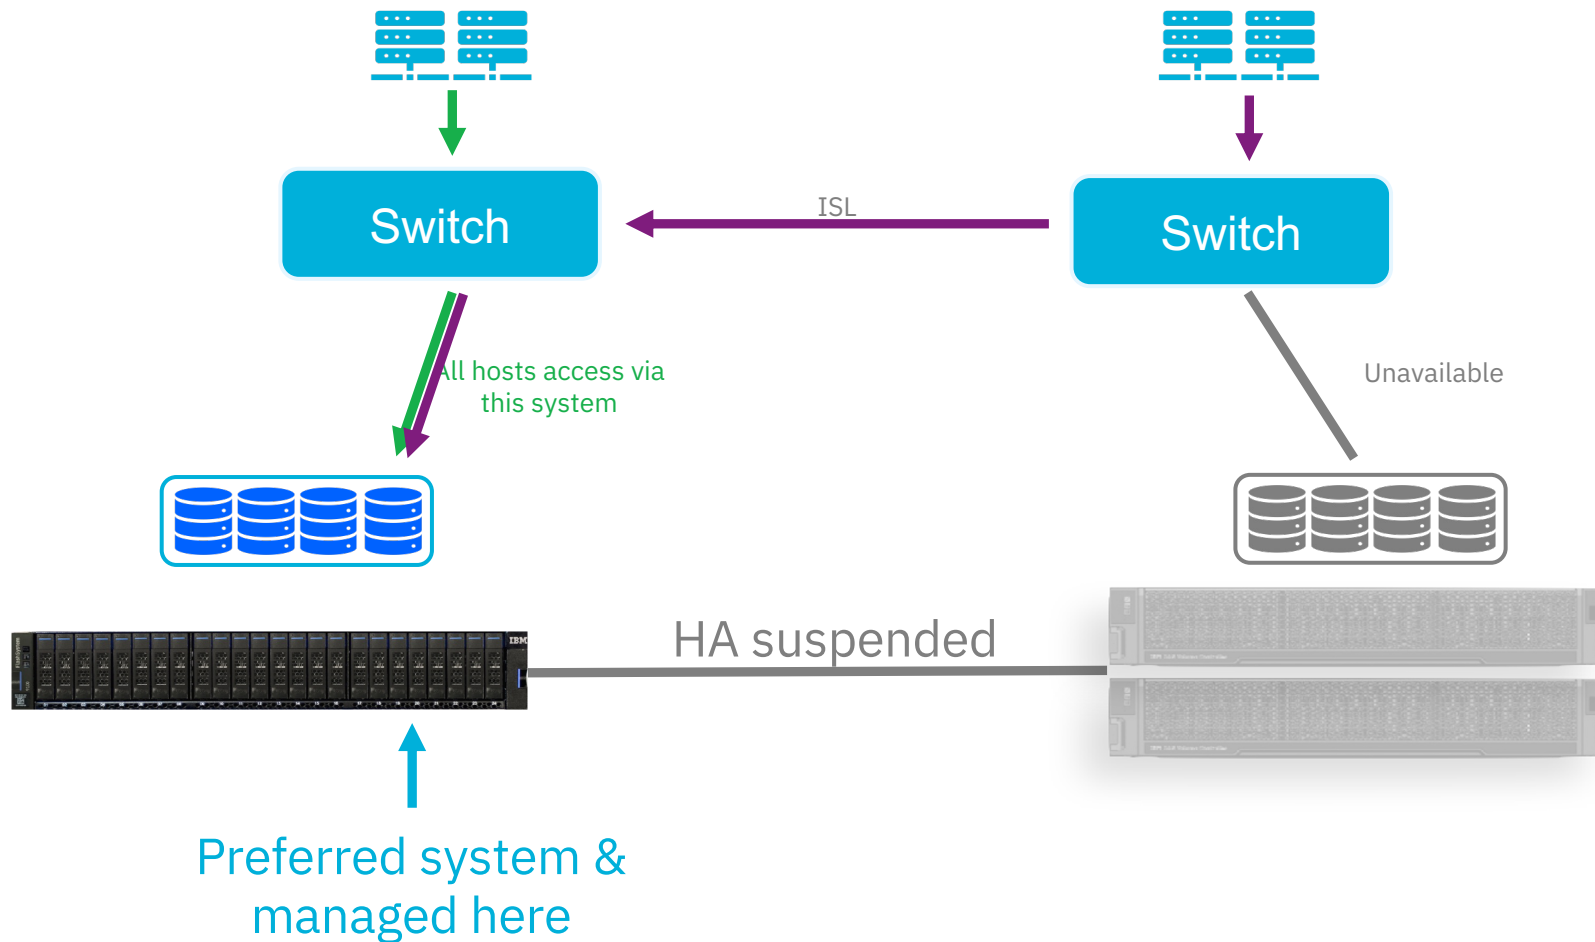
HA Without Host Locations



HA With Host Locations



HA With Host Locations – Storage Failure



Managing Host Location

IBM FlashSystem 7200 | pandabear | Storage partition: myPartition - Hosts

myPartition

Storage partition overview

Volumes

Hosts

Copy services

Policies

IP quorum

Exit myPartition

Hosts

+ Add Host | Actions | Download

Default | Contains | Filter

Name	Status	Host Type	# of Ports	Host Mappings	Host Cluster ID	Host Cluster Name	Location	Protocol Type
pandabear-h0	Online	Generic	1	Yes			pandabear	SCSI
pandabear-h1	Online	Generic	1	Yes				SCSI

Edit and rename host [X]

Name: pandabear-h1

Assign location

Host location: sunbear

Advanced ^

Cancel Save

Partnership Support – 8.7.0

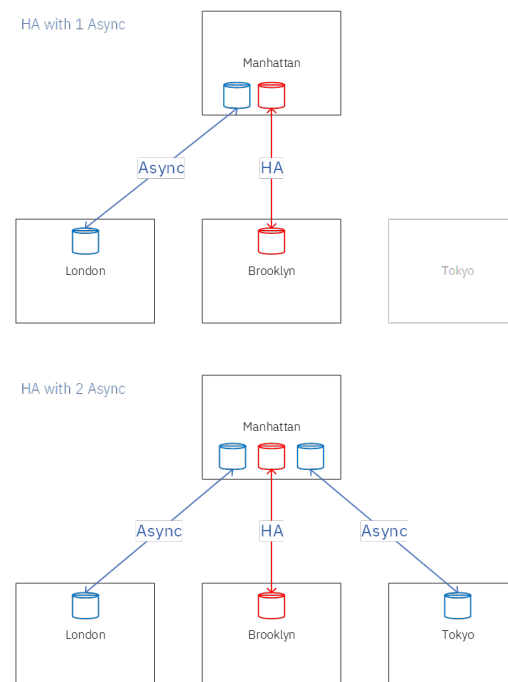
- Policy-based HA can be configured using either:
 - Fibre Channel partnerships, or
 - Short-distance partnerships using RDMA (High Speed Ethernet).
- Fibre Channel and short-distance RDMA partnerships have equivalent performance.
- If sharing an ISL for node-to-node and host traffic, QoS should be configured to prioritise replication traffic.
- Dedicated ports are required for both types.
- Max 1ms RTT (SCORE requests can be submitted for higher).

• HA + async on different partnerships

- HA to 1 remote system with Async to 1 or 2 other remote systems, or 3 partnerships using async policy-based replication.

Note: A Volume Group can either be participating in policy-based HA or Async policy-based replication, but not both.

- Note: Limit of 3 partnerships per system, of which at most 2 can be short-distance Ethernet using RDMA partnerships.



FlashSystem 5045 Support – 8.7.0

- HA and PBR async support available from 8.7.0
- Single I/O group systems only
- Must have 32GB per node
- No support for vVol replication or storage partition migration
- Cannot use Data Reduction Pools (DRP) and replication policies on the same system.

Model	Replicated volume count	Replicated capacity (per I/O group)
FlashSystem 5015/5035	Not supported	Not supported
FlashSystem 5045	4096	200 TiB
FlashSystem 5200/5300	7932	1024 TiB
FlashSystem 7200	7932	2048 TiB
FlashSystem 7300	16050	2048 TiB
FlashSystem 91x0/9200	7932	2048 TiB
FlashSystem 9500	32500	4096 TiB
SAN Volume Controller (SA2/SV2)	7932	2048 TiB
SAN Volume Controller (SV3)	7932	4096 TiB

Miscellaneous – 8.7.0

- UUIDs and policy reporting
 - Partitions now have RFC4122-compliant UUIDs
 - e.g `8124F74C-99CE-5F3D-9841-396DC5527D55`
 - Static for life of partition
 - HA partitions share UUIDs
 - HA partitions generate UUID once both systems have upgraded to 8.7.0
 - Useful for automation (track partitions as they're migrated between systems)
- Volume groups now report DR and (inherited) HA replication policies separately

Support Statement

- **8.7.0 is the final version that supports Remote Copy**
 - **8.6.3 does NOT**
- **Remote Copy will be supported long-term on 8.7.0 for as long as the hardware has a valid support contract.**
- Includes:
 - Global Mirror
 - Global Mirror with Change Volumes
 - HyperSwap
 - Metro Mirror
 - Migration relationships
 - HyperSwap and Metro Mirror 3-site solutions
- Note: Entry-level FlashSystem 5015 and 5035 will not have replication capabilities if upgraded beyond 8.7.0.
- Global Mirror and Global Mirror with Change Volumes are replaced by policy-based replication. A migration procedure is available in the product documentation.
- HyperSwap is replaced by policy-based HA.
- Migration is using storage partition migration.
- Policy-based HA + async DR (3-site) is planned for 2H24.
- Synchronous policy-based replication is planned for 2025.
- *Disclaimer: IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.*

Compatibility Mode – 8.7.0

- There are too many ways to perform the same functions on a Storage Virtualize system confusing users
 - We want to prompt new users to adopt the new ways of doing things to simplify deployment and reduce problems
 - When 8.7.0.x start to ship on new units, the systems will have compatibility mode off, meaning that only the new features are allowed from both the GUI and the CLI/REST API
- **Examples**
 - **Replication:** Global Mirror, Global Mirror with Change Volumes, Metro Mirror, Policy-based replication.
 - **High availability:** Policy-based HA, HyperSwap, Stretched Cluster.
 - **Snapshots:** FlashCopy, Volume Group Snapshots
 - **Migration:** Image-mode import/export, Clustering and NDVM, Volume Mobility (migration relationships), Partition Mobility
 - **Multiple I/O groups:** Newer features will only be allowed on single I/O group systems. Although not a feature that is going away, users will have to use informed consent

Compatibility Mode – 8.7.0

- Changes the default behaviour of newly installed systems to provide clearer choices to our users
- Encourages the user to use the newer/strategic solutions by disabling legacy/non-strategic functions
- Disabled functionality can be re-enabled, if needed, by enabling compatibility mode
- Upgrading systems have no changes of behaviour as compatibility mode is enabled as part of the upgrade

Remote Copy functions in Copy Services:

By default, Remote Copy functionality that is related to Remote Copy relationships and Remote Copy consistency groups is disabled. To use this functionality, set this value to Enabled.

Remote Copy functions in Copy Services:

Disabled

Multiple I/O Groups

By default, support for multiple I/O groups is disabled on new systems as some newer features are only supported on single I/O group systems. These features include storage partitions, VMware vVol replication and policy-based high availability. Adding additional I/O groups to the system will prevent the use of these features. To enable support for multiple I/O groups, set this value to Enabled.

Multiple I/O Groups


Disabled

Default Layering Change – 8.7.0

- Newly configured FlashSystems will default to the 'replication' layer. The layer will not change on upgrade of existing systems.
- This is to improve the interoperability of replication and HA between FlashSystem and SVC
- Using the GUI setup wizard to configure a FlashSystem to be backend storage for SVC will correctly adjust the layer

IBM Storage Virtualize

Snapshot Updates



Internal TCT Scheduling – 8.6.3

- There are 3 types of snapshot policies allowed on a VG (max of 1)
 - Local – Local Snapshot Policy
 - Cloud – TCT Policy
 - Hybrid – Combination of the above
- Snapshot Policy
 - Increase to a max of 40 policies
 - 5 Predefined – Cannot be deleted/changed
 - 3 Local
 - 1 Cloud
 - 1 Hybrid

Screen Shot

The screenshot displays the IBM FlashSystem 5200 GUI for the 'tct-check' system, specifically the 'Volume groups' section. The browser address bar shows '9.193.230.107/gui#volumes-volumegroups'. The page title is 'tct-check - Volume groups'. The breadcrumb navigation is 'Volume groups / volumegroup0'. The main heading is 'volumegroup0' with a '20.00 GiB Total Group Capacity' indicator. Below the heading are tabs for 'Volumes (2)', 'Policies (1)', 'Local snapshots (0)', and 'Cloud Snapshots'. The 'Policies (1)' tab is active, showing a 'Snapshot Policy' configuration. The policy details include: Name: 'my_cloudbackup_policy_1', Target: 'Cloud', Cloud schedule: 'Every day at 03:00 AM, retained for 7 days', and Next at: '30/1/2024 12:00 AM'. A message states 'Policy-based replication is not configured.' with a 'Configure policy-based replication' button. A footer bar shows performance metrics: Latency 0 ms, Read 0 ms, Write 0 ms, Bandwidth 0 MBps, Read 0 MBps, Write 0 MBps, IOPS 0, Read 0, Write 0.

Volume groups / volumegroup0

volumegroup0

20.00 GiB Total Group Capacity

Volumes (2) **Policies (1)** Local snapshots (0) Cloud Snapshots

Snapshot Policy

Name	Target
my_cloudbackup_policy_1	Cloud

Cloud schedule
Every day at 03:00 AM, retained for 7 days

Next at: 30/1/2024 12:00 AM

Policy-based replication is not configured.

[Configure policy-based replication](#)

To use External Safeguarded Backup Policy, [Learn More](#)

Latency 0 ms Read 0 ms Write 0 ms Bandwidth 0 MBps Read 0 MBps Write 0 MBps IOPS 0 Read 0 Write 0

Notes

- As with local snapshots, no expiry date is attached to manually created cloud backups
- CLI's "chvdisk" & "chvolume group" used with "-retainbackupenabled" can be used to retain cloud backup enabled on the volumes while removing/moving vdisk from volume group or while un-associating snapshot policy from the volume group. This option is only available through CLI.
- A hybrid policy cannot be safeguarded
 - Since no mechanism exists today to safeguard TCT, a hybrid policy cannot be safeguarded
- TCT will not work with SGCv1
- To have a TCT backup of a Safeguarded Copy, a thin clone must be created first
 - Refer to the following Redpaper: [Offloading Safeguarded Copies with TCT](#)
- To get a combination of TCT and SGCv2, use a cloud policy and have CSM take the SGC copies
 - CSM does not support TCT today
- Internal code improvements to increase backup/restore speed

Restore Timestamps – 8.7.0

- lsvolumegroup shows the last time that a volumegroup was restored
- lsvolumegroupsnapshot shows the last time that snapshot was used to restore a volumegroup
- Format: YYMMDDHHMMSS
- A blank field means there has never been a restore

```
[11:11:41] sq1-fab3-3site-n1:~ # lsvolumegroupsnapshot -delim :
id:name:volume_group_id:volume_group_name:time:state:matches_group:parent_uid:expiration_time:protection_provisioned_cap
acity:protection_written_capacity:operation_start_time:operation_completion_estimate:owner_id:owner_name:auto_snapshot:s
afeguarded:last_restored_from
0:snapshot0:0:vg24:240212111009:active:yes:0::30.00GB:2.25MB::::no:no:
1:snapshot1:1:vg25:240212111034:active:yes:1::20.00GB:1.50MB::::no:no:240212111109
2:snapshot2:1:vg25:240212111034:active:yes:1::20.00GB:1.50MB::::no:no:240212111139
3:snapshot3:1:vg25:240212111034:active:yes:1::20.00GB:1.50MB::::no:no:240212111139
4:snapshot4:0:vg24:240212111039:active:yes:0::30.00GB:2.25MB::::no:no:
[11:11:46] sq1-fab3-3site-n1:~ # lsvolumegroup -delim :
id:name:volume_count:backup_status:last_backup_time:owner_id:owner_name:safeguarded_policy_id:safeguarded_policy_name:re
plication_policy_id:replication_policy_name:volume_group_type:uid:source_volume_group_id:source_volume_group_name:parent
_uid:source_snapshot:snapshot_policy_id:snapshot_policy_name:partition_id:partition_name:restore_in_progress:owner_type:
draft_partition_id:draft_partition_name:last_restore_time
0:vg24:3:off:::::0::::no:none:
1:vg25:2:off:::::1::::yes:none:240212111139
2:vg26:0:empty:::::2::::no:none:
```

Last Snapshot Restore

[Volume groups](#) / BCTest

BCTest


200.00 GiB Total Group Capacity

[View properties](#)

Volumes (4) Policies (0) **Local snapshots (1)** Cloud Snapshots

Capacity for Snapshots

[Hide details](#) ^



Search

[Take Snapshot](#) +

Snapshot name	State	Safeguarded	Time Created	Last restore time	Expiration Time	Matches Group	Auto Snapshot
snapshot0	Active	No	5/10/2024 10:22 AM	5/18/2024 12:02 PM	-	Yes	No

Items per page: 25 1-1 of 1 item 1 1 of 1 page

Last Volume Group Restore

Volume groups / BGTTest

BGTTest 200.00 GiB Total Group Capacity [View properties](#)

Volumes (4) Policies (0) **Local snapshots (1)** Cloud Snapshots

Capacity for Snapshots Hide details ^

3.00 MiB / 200.00 GiB (0.00%)

Written Capacity

Search

Take Snapshot +

Snapshot name	State	Safeguarded	Expiration Time	Matches Group	Auto Snapshot
snapshot0	Active	No	-	Yes	No

Items per page: 25 1-1 of 1 item 1 1 of 1 page

Volume group properties

Last Restored: 5/18/2024 12:02 PM

Default Grain Size for FlashCopy/Snapshots – 8.7.0

- The default grainsize for snapshot mappings on a Storage Virtualize system is 256KB
- Some applications (such as EPIC) are sensitive to the I/O amplification that this larger grainsize can cause
- chsystem can be used to change the default globally to 64KB or 256KB
- When changed, the new default grainsize setting will affect new mappings only
- If the default is changed to 64KB, the target capacity allowed will be decreased by a factor of 4
- Applies to both FlashCopies and Volume Group Snapshots
 - Instead of 40PB on a FS9500, it will now be 10PB

```
host_unmap on
backend_unmap on
quorum_mode standard
quorum_site_id
quorum_site_name
quorum_lease short
automatic_vdisk_analysis_enabled on
callhome_accepted_usage no
safeguarded_copy_suspended no
protection_provisioned_capacity 200.00GB
protection_written_capacity 3.00MB
flashcopy_gui_enabled no
snapshot_policy_suspended no
snapshot_preserve_parent no
anomaly_detection on
anomaly_detection_event off
flashcopy_default_grainsize 256
storage_insights_control_access no
flash_grid_system_uuid D2C03E7B-CB49-5D84-991E-7F3D516D4D73
flash_grid_uuid
flash_grid_name
flash_grid_owner_uuid
auto_drive_download on
```

Converting a Thin Clone to a Clone – 8.7.0


- In some circumstances it may be desirable to convert a thin clone into a full independent clone copy
 - A thin clone always requires the source volume(s), so it/they cannot be truly deleted
- The command: `converttoclone` has been added to convert a thin clone into a free-standing cloned copy
- The command can either be used to convert an entire volume group or a subset of the volumes
- The GUI can only be used to convert the whole volumegroup
- If necessary, the speed of the background copy can be managed from the command line

Miscellaneous – 8.7.0

- Can now restore/clone snapshots into non-HA partitions.
- By default, snapshots are restored/cloned into the same partition.
- Option to restore outside of a partition, or into a different partition.

IBM Storage Virtualize

Ethernet Updates



iSCSI Security Enhancements – 8.6.3

- Introduction of SHA1 support in the authentication framework
 - Seamless integration within the CHAP (Challenge-Handshake Authentication Protocol) process
 - Ensures compatibility and adherence to industry-standard cryptographic practices
- Stronger Cipher Suites
 - Introduction of stronger cipher suites in iSCSI
 - Support for SHA1, SHA256 and SHA3-256
 - Options for customers to use these cipher suites for FIPS (Federal Information Processing Standards) compliance

Management Port Changes – 8.7.0

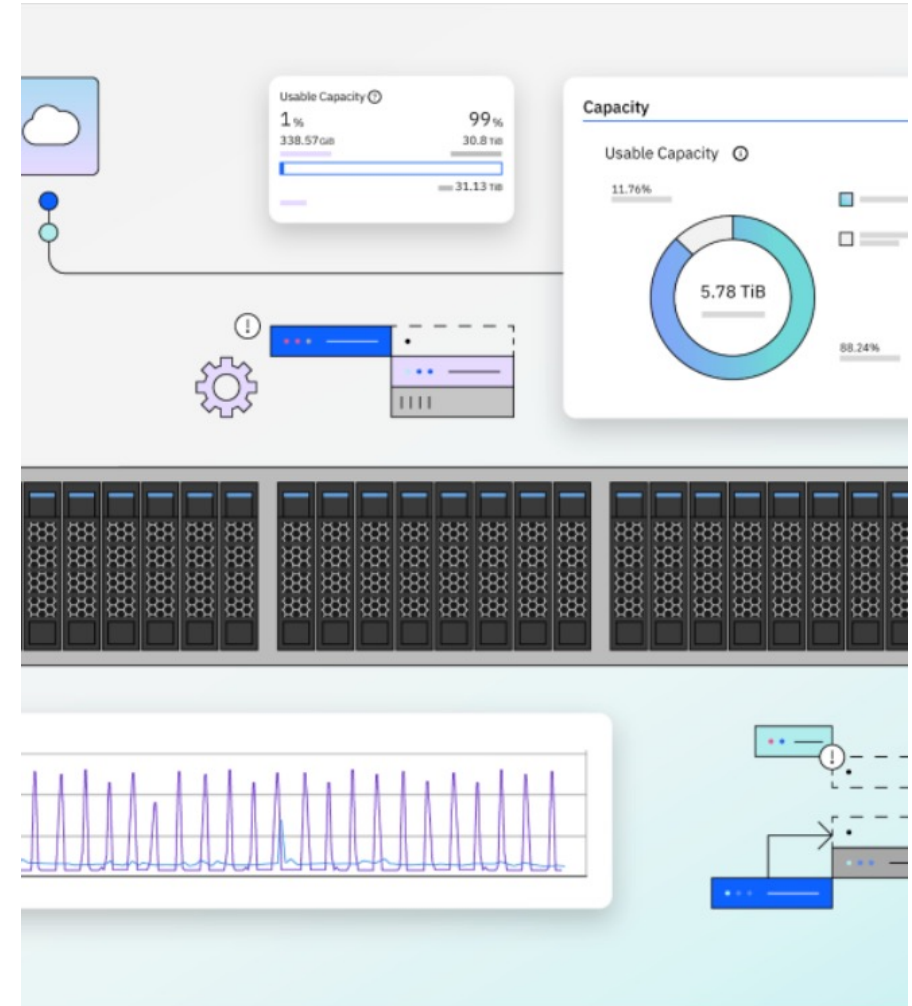
- Management addresses can be configured on any Ethernet ports on the system
 - Can be configured with VLAN
 - Common cli commands for Ethernet and management IP addresses
 - mkip/rmip/lsip/chip
 - Supported on the GUI
 - For now, service IPs stay on logical port 1
- New portset: SystemManagement
 - Maximum of 2 ports
 - Portset ID 72
- 4 routable data IP addresses per port per node
- Config node failover time reduced by 10%

Management Port Commands

- mkcluster - New optional param port_id and vlan, new syntax like mkip
- mkip - to configure second system IP(cluster IP) address
- rmip - to remove system IP(cluster IP) address
- lsip - displays both data and system IP addresses
- lspportset - displays new default management portset (SystemManagement)
- chip – to change system IP address
- chssytemip – No changes. It is recommended to use mkip/rmip/chip
 - Left for compatibility with existing scripts

IBM Storage Virtualize

GUI Updates



New GUI Setup Encryption Default – 8.6.3

System Setup
✕

- ✔ Welcome
- ✔ License Agreement
- ✔ Change Password
- ✔ System Name
- ✔ Licensed Functions
- ✔ DNS
- ✔ Date and Time
- Encryption
- Call Home
- IBM Storage Insights
- Support Assistance
- Summary

Encryption

Was the encryption feature purchased for this system?

No
 Yes

Activate the encryption licenses for each enclosure.

☰ Actions ▾
⬇

Type	↑	M/T-M	S/N	Licensed	⬇
Control Enclosure		2076-824	78E3006		

Showing 1 row | Selecting 0 rows

🔍 Need Help
Cancel
◀ Back
Next ▶

GUI Updates – 8.7.0

- The GUI has been updated to allow giving a host in a PB-HA partition a location
- The location is based on the storage unit and the name cannot be changed
- The location is used to meter locality for PB-HA to keep traffic off the distance link wherever possible

GUI Updates – 8.7.0

- Volume Groups has been added to the Logical Components section of the Dashboard

System Health Version: 8.7.0.0 (build 177.23.2405081401000) | Cluster ID: 00000204246003CA

Hardware Components ✓ | **Logical Components** ⊕ | Connectivity Components ✓

Nodes 2 ✓ All Online More Details	Array MDisks 2 ✓ All Online View Page	I/O Groups 1 ✓ All Online View Page	Volumes 8 ✓ All Online View Page	Pools 2 ✓ All Online View Page	Volume Groups 3 View Page
----------------------------------------------------------------------	--------------------------------------------------------------------------	------------------------------------------------------------------------	---------------------------------------------------------------------	-------------------------------------------------------------------	----------------------------------------------

GUI Updates – 8.7.0

- Volume Groups can now be assigned to an Ownership Group

Create Volume Group

Select how to assign volumes to a new volume group. You can specify from existing volumes or select a snapshot of volumes in another volume group.

Enter name (optional)

VMware_VG

Ownership group (optional)

VMware

Assign volumes (optional)

Choose existing volumes

Choose an existing snapshot from a volume group

Cancel

Create Empty Group

Volume Groups

Search table...

Name	Volume Count	Snapshot Policy	Snapshot Count	Safeguarded	Ownership group
BGRecovery	4	-	0	No	
BGTest	0	-	0	No	
BGTest1	4	Daily7	8	Yes	
VMware_VG	0	-	0	No	VMware

Create Volume Group +

IBM FlashSystem 9100 FS9100-1 Dashboard

Dashboard

Pools 1

Volume Groups 1

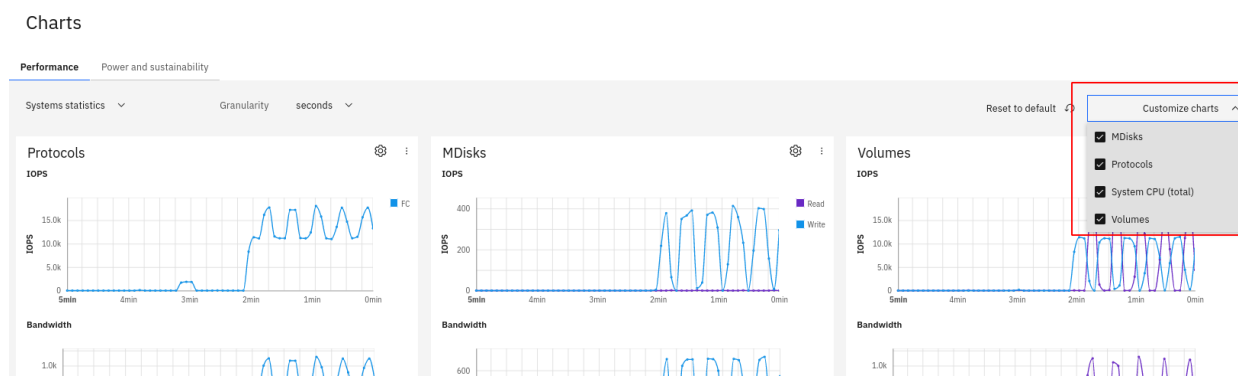
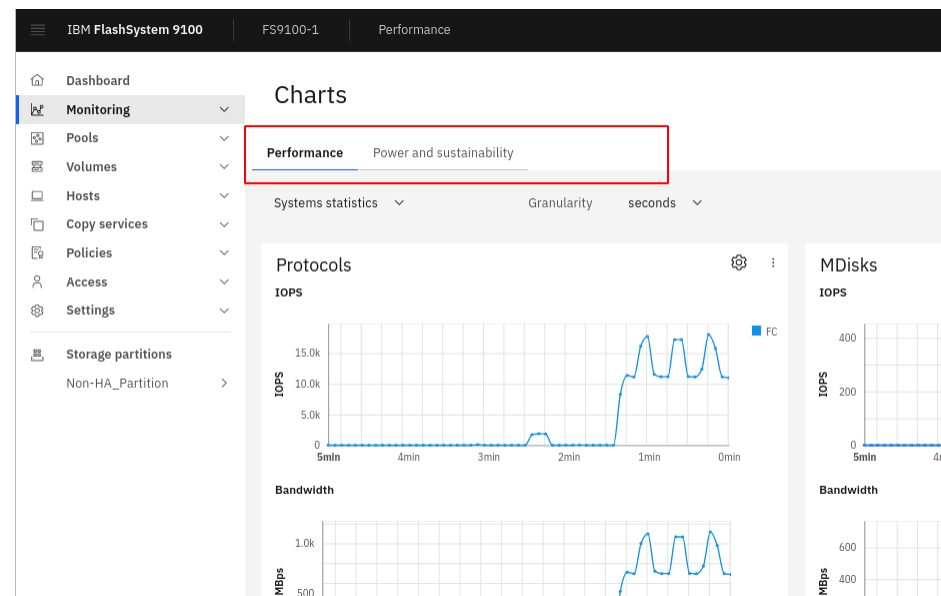
All Online

View Page

View Page

GUI Updates – 8.7.0

- GUI performance panel modernization using Carbon v11 components and Carbon Charts
- Tabs based implementation for future scalability of the charts to incorporate a growing number of statistics
- Charts by user customization across the user sessions
- Responsive flexible Layout of 1 column, 2 column or 3 columns
- Improved usability and accessibility through drag and drop feature
- New option for ‘Power’ and ‘Temperature’ charts to display statistics at systems, node or enclosure level
- Easy to compare the IOPS, Bandwidth and Latency through single view
- ‘Restore to default’ feature to remove all user customizations and restore default view




GUI Updates – 8.7.0

Charts


Performance Power and sustainability

Systems statistics Granularity seconds


Protocols




Bandwidth




MDisks



Bandwidth




Latency




- Move up
- Move down
- Move left
- Move right


Volumes



Bandwidth




Latency



System CPU (total)



GUI Updates – 8.7.0




Charts



IBM Storage Virtualize

Code Update Changes



Current Updates

IBM FlashSystem 7300 FS7300-TestDrive System

Update System

Current software version: Version 8.6.0.3 (build 169.25.2401311801000) [Long Term Support](#)

Available Update Versions

8.6.1.0
8.6.2.0

To learn more about available updates and to download firmware, visit [support](#).

There are no update related files currently on the system

Provide files to the system using one of the options below. If you have manually transferred files to the system using the CLI, click Check for Files.

[Check for files](#)

Provide a package

Obtain the package directly

Use Call Home with cloud services to directly transfer the update package to the storage system.

[Transfer](#)

Provide the package manually

Upload the update package or patch file to the storage system manually.

[Upload](#)

byrongro@FlashSystem-TestDrive... FS7300-TestDrive - System — Mo... 1 / 2

Latency 0 ms 0 ms 0 ms Bandwidth 0 MBps 0 MBps 0 MBps IOPS 0 0 0

Current Updates

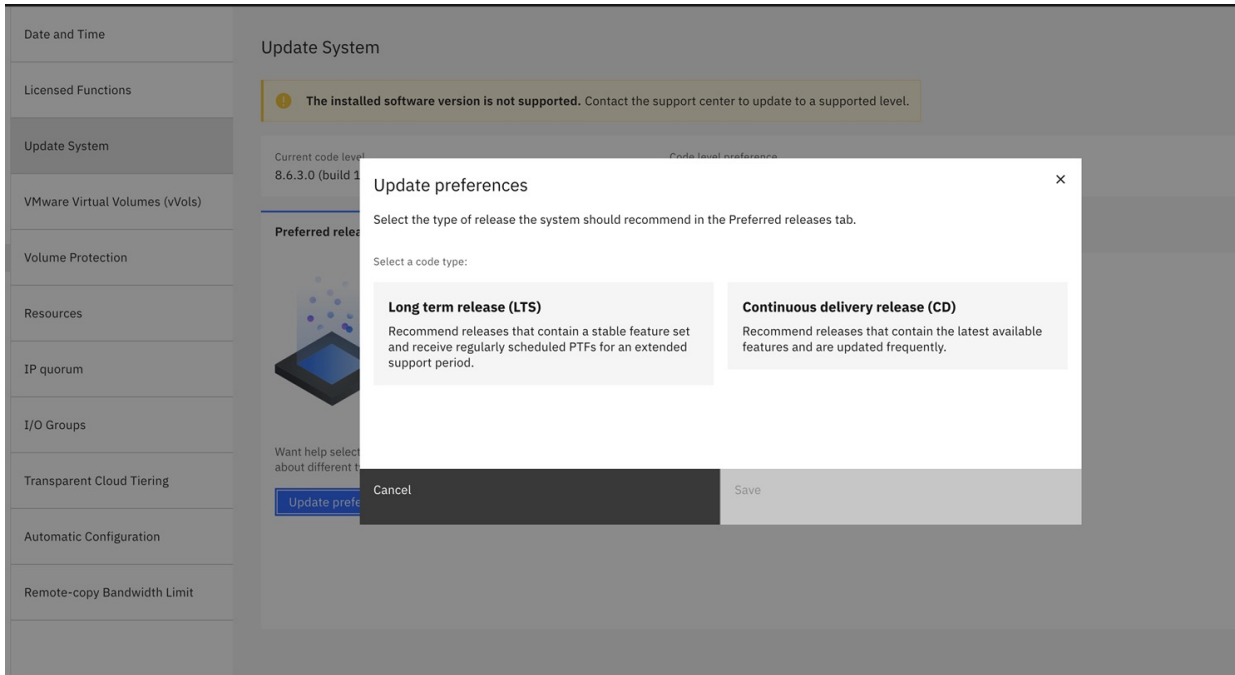
The screenshot displays the IBM FlashSystem 7300 management interface. The top navigation bar includes the system name 'IBM FlashSystem 7300', the environment 'FS7300-TestDrive', and the user 'byrongro Security Administrator'. The main content area is divided into three sections:

- Performance:** A line chart showing latency, bandwidth, IOPS, and CPU utilization over a 5-minute period. The y-axis for latency ranges from 0 to 100k μ s. The x-axis shows time intervals from 5min to 0min. Summary statistics are shown as follows:

Metric	read	write
Latency	0 μ s	0 μ s
Bandwidth	0 MBps	0 MBps
IOPS	0 IOPS	0 IOPS
CPU Utilization	12 %	
- Capacity:** A section with a loading indicator: 'Loading the capacity data, please wait...'
- System Health:** A section showing the status of hardware, logical, and connectivity components. All are reported as 'All Online'. Each component has an 'Expand' button.


At the bottom right, a notification box titled 'Update Available Suggested Task' is highlighted with a red border. It contains a close button (X), a 'Not Now' button, and a 'Learn More' button. The system version is noted as 'Version: 8.6.0.3 (build 169.25.2401311801000)' and the cluster ID is '0000020420814428'.

New Update Preference – 8.6.3



Current code level 8.6.3.0 (build 175.11.2312191502000)	Code level preference Continuous delivery release
------------------------------------------------------------	------------------------------------------------------

Preferred releases **Alternative update options**

 **There are no update related files currently on the system**

Provide files to the system using one of the options below. If you have manually transferred files to the system using the CLI, click Check for Files.

[Check for files](#)

Provide a package

Obtain the package directly

Use Call Home with cloud services to directly transfer the update package to the storage system.

[Transfer](#)

Provide the package manually

Upload the update package or patch file to the storage system manually.

[Upload](#)


Auto Patch Updating - What is a Patch

- A patch is a lightweight update to a function or service, which can be installed on a user's system:
- It shall not change any file directly used by the IO stack
- It shall never require a node reboot or reset when installed
- It may restart a Linux service when installed
- It shall install on all platform types
- Is generally small in size
- Patches are stored on FixCentral
- Cloud Call Home is used to access patches
- Currently new versions of existing patches would require users to manually install them

Automatic Patch Updating – 8.7.0

- Automatic Patch Updating can be set up, on a user's system, either using the GUI, or CLI commands
- The User can choose which patches to automatically update when a new version comes out
 - A good example may be the inference engine
 - The user may want that patch to be updated automatically, but not others
- A check is made once a day via cloud call home to see if updates are available and automatically updates the patches selected

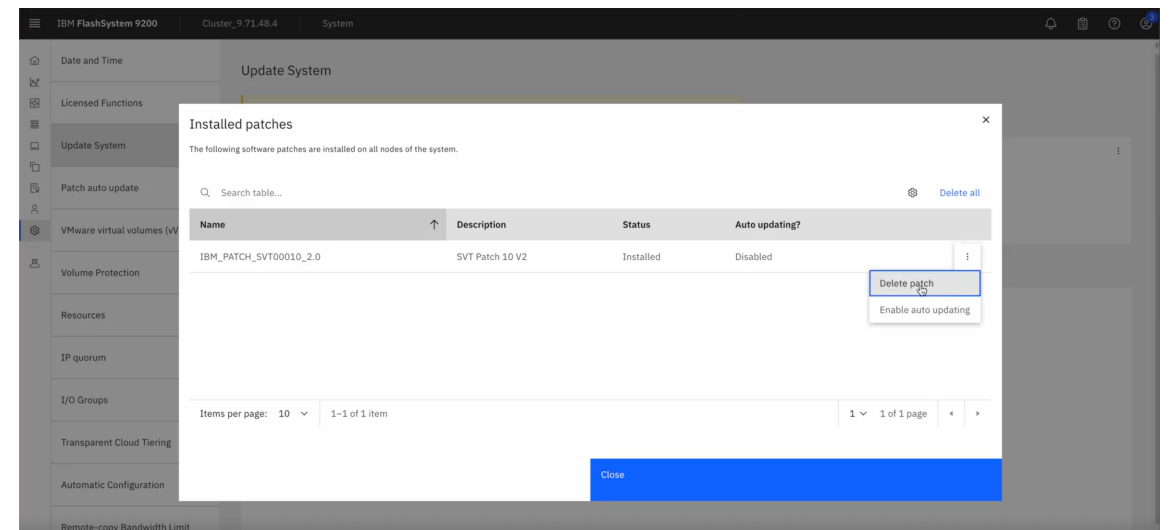
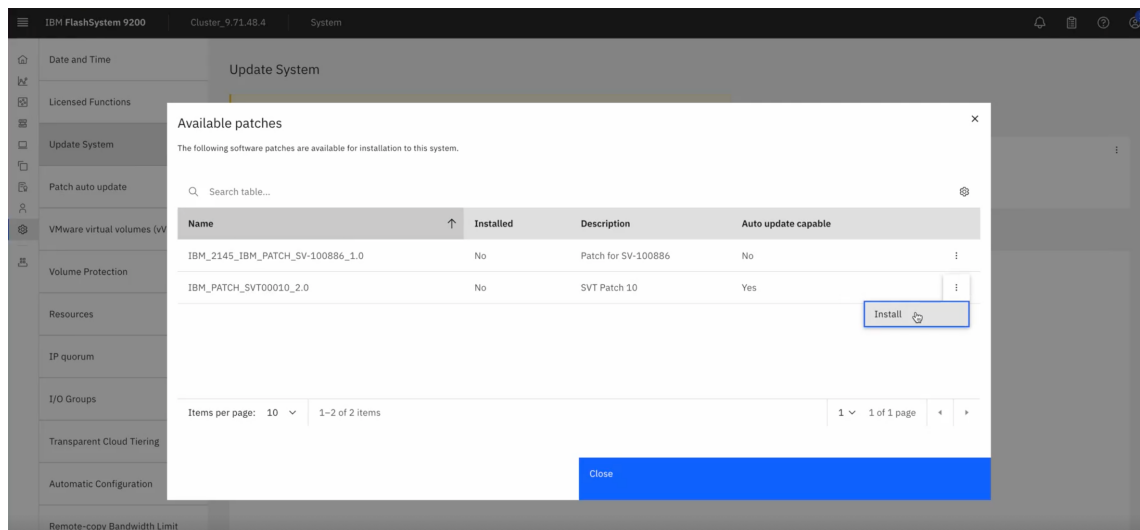
- Starting the patch updating mechanism from the GUI:



Auto Patch Update - GUI Continued

- List available patches and install one

- Enable/Disable patch to automatically update



Auto Patch Update - CLI

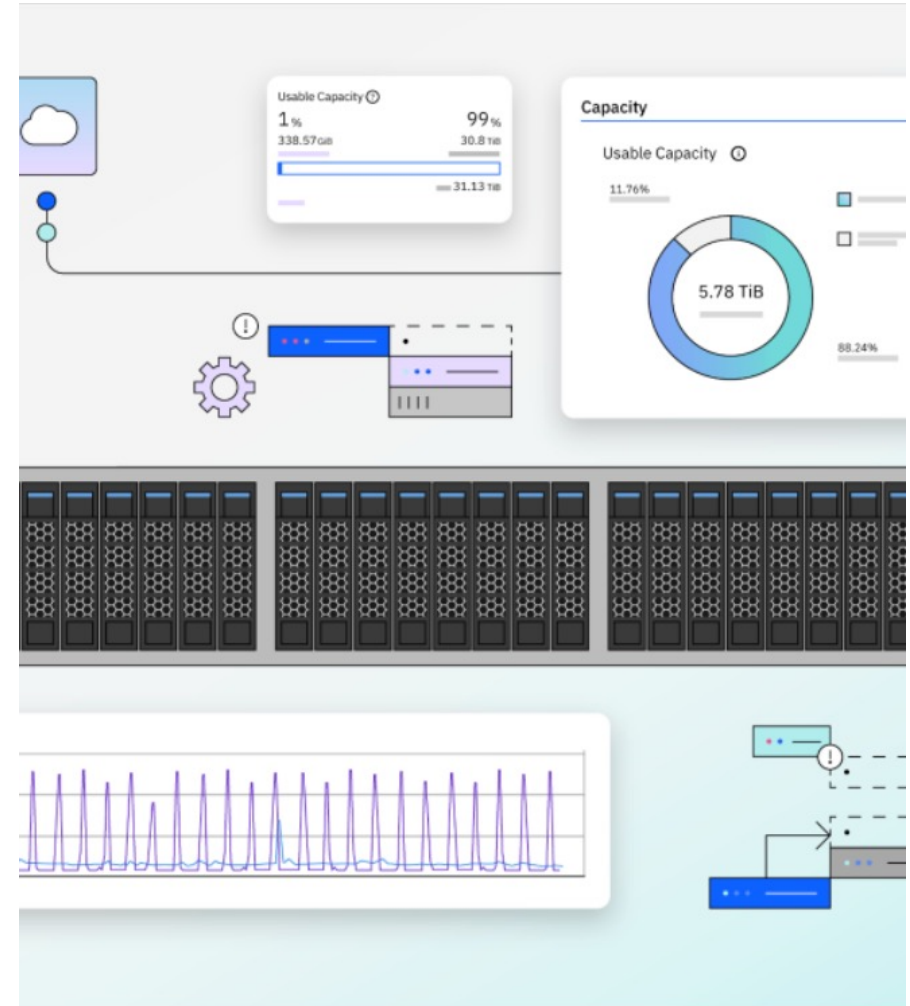
- `svctask chsecurity -patchautoupdate yes` (starts the automatic patch updating mechanism)
- `svcinfo lsavailablepatch` (gets the list of available patches from FixCentral)
- `satask downloadsoftware -patchid <patch_name>` (retrieves a patch from FixCentral)
- `svctask applysoftware -file <patch_name>` (installs a patch)
- `svctask chsystempatch -patch <patch_name> -autoupdate yes` (sets a patch to automatically update)

Automatic Drive Download – 8.7.0


- Automatic Drive Download is a mechanism that utilizes the new patch infrastructure to enable drive firmware to be stored on a system/cluster and ensure a standard drive firmware level is maintained. **This change is for FCM drives only**
- FCM drive firmware is now built as a patch
- Applysoftware <firmware patch> will now copy the drive firmware to every node in the system/cluster
- Any time an FCM is configured or replaced, it will check the system for newer firmware and download it if available
- Patches are stored on FixCentral
- The automatic update only applies to new drives in the system. Existing drives will need to be updated with the normal drive upgrade process

IBM Storage Virtualize

Auto-Space Reclaim in Standard Pools – 8.7.0



Thin Provisioned Volumes in Standard Pools Until Now



What does this feature deliver?

A mechanism to automatically reclaim thin-provisioned volumes in standard pools.

What is the value of this feature?


- Automatically recovers space in standard pools after hosts submit unmap IO (or overwrite with zeros)
- Customers using a utility implementation asked for this to be added to the product
- Previously this required careful monitoring and manual actions

What you need to know


- Feature is disabled by default
- The **reclaim_se** CLI can be used to enable / disable / configure / monitor
- A daemon runs on the config node and checks for eligible VDisks and starts reclaiming
- Reclaiming works by adding a VDisk copy to the existing VDisks using the **-autodelete** flag
- **reclaim_se --help** to get started

Reclamation Process



1. Initial State



2. Reclamation in Progress



3. Reclamation Complete



Control Commands

- **Enable the feature**

```
reclaim_se --start [optional configuration]
```

- **Monitor the feature**

```
reclaim_se --status [--dry-run] [optional configuration]
```

- **Change configuration**

```
reclaim_se --reload [optional configuration]
```

- **Disable the feature (undocumented: optionally discard the options file)**

```
reclaim_se --cancel [--delete-options-file]
```

Configuration

- Configuration parameters can be passed:
 1. When enabling the feature using --start
 2. When reloading configuration using --reload
- Parameters not specified either use the previously configured values or the defaults

Parameter	Default	Description
interval	60	Interval in minutes (5-10080)
reclaimable (capacity)	500 GB	Reclaimable capacity threshold
reclaimable (percentage)	N/A	Reclaimable percentage threshold (0-100)
syncrate	100	Syncrate for volume copy synchronization (1-150)
reclaims-per-pool	1	Number of concurrent reclamations per parent pool
rsize	N/A	Override of the real size
rsize_fallback	2%	Real size to fallback to if it cannot be guessed
autoexpand-off-mode	disabled	Mode for volumes with autoexpand disabled

Example

```
IBM_FlashSystem:Cluster_9.71.48.127:superuser>reclaim_se --start --reclaimable 500MB
```

```
IBM_FlashSystem:Cluster_9.71.48.127:superuser>reclaim_se --status  
reclaim_se is running (monitoring) with the following parameters:
```

```
status:                enabled  
interval:             60 minutes  
reclaimable:         500.0MB  
reclaims-per-pool:   1  
rsiz:                 same as existing vdisk  
rsiz-fallback:       2%  
autoexpand-off-mode: disabled  
syncrate:            100 (64.0MB/s)
```

```
Next check in approximately 3.1 minutes
```

```
Reclaim in progress on vdisk 1362 (start_time=2024-03-11 17:38:36 mdisk_grp_id=0  
parent_mdisk_grp_id=0 reclaimable=1.9GB reclaimable_percentage=0.2 progress=0  
estimated_completion_time=240311210539)
```

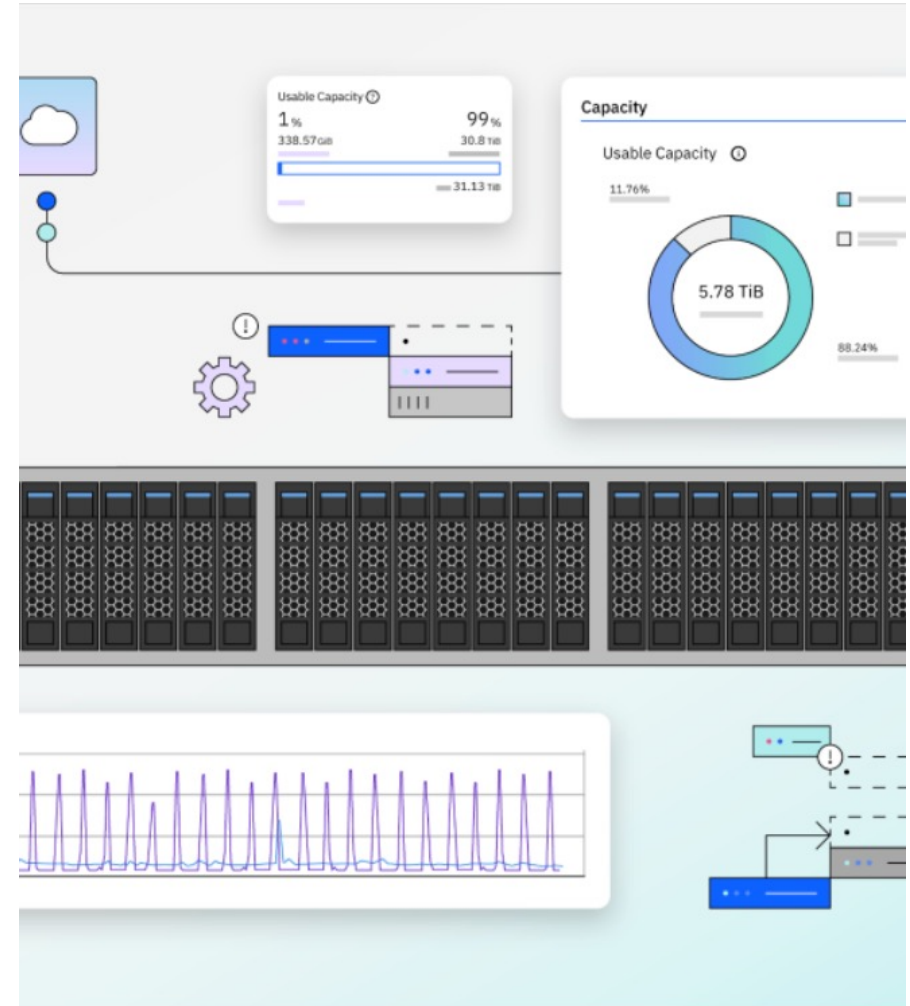
```
IBM_FlashSystem:Cluster_9.71.48.127:superuser>catauditlog | tail | grep addvdiskcopy  
7828          240311173835 system svctask addvdiskcopy -mdiskgrp 0 -unit b -rsiz  
28568211456 -warning 80% -autoexpand -grainsiz 256 -syncrate 100 -autodelete 1362
```

Limitations

- No GUI support
- Feature requires free space
- Not supported on mirrored volumes
- No new reclamations started while upgrade is in process
- No new reclamations started on offline or degraded VDisks
- Reclamations are only started if there if lsvdiskanalysis has been run within the last 24 hours
 - System defaults to running lsvdiskanalysis every 12 hours
- Sparse VDisks (empty) will only be reclaimed if lsvdiskanalysis has been run while the daemon is running
 - The command could be run manually to kick off the reclamations

IBM Storage Virtualize

Ransomware Threat Detection File System Awareness – 8.7.0




Why Add Volume File System Awareness?

- Volumes can be used by many different applications, operating systems, and file systems
- AI for ransomware detection can infer IO patterns better if we know what file system is used in each volume
- Storage admins are not always aware of what volumes are used for, by what application, and different teams use the storage
 - Which is why we want to infer/determine rather than making it manual
- Support can benefit from understanding what file systems are in which volumes in some recovery scenarios
- Can infer operating systems, and sometimes applications like VMware

Logistics

- Every 12 hours file system is automatically updated for each volume
 - Will also update by analyzevdisk or analyzevdiskbysystem CLI commands
- Background reads are sent to a volume
- Open source libraries used to determine file system
- Output is displayed in lsvdiskanalysis's file system field
 - 15 character max for field
 - Can display multiple file systems
- File system used by inferencing engine to improve ransomware detection
- Data read on the volume is only used to determine file system type and is NOT sent to SI for processing



CLI Example


```
[17:30:30] fs9500:~ # analyzevdisk 0
```

```
[17:30:36] fs9500:~ # lsvdiskanalysis 0
```

```
id 0  
name vdisk0  
state sparse  
started_time 240425173030  
analysis_time 240425173030  
capacity 100.00GB  
thin_size 0.00MB  
thin_savings 0.00MB  
thin_savings_ratio 0  
compressed_size 0.00MB  
compression_savings 0.00MB  
compression_savings_ratio 0  
total_savings 0.00MB  
total_savings_ratio 0  
margin_of_error 0  
file_system ext4,xfs
```

IBM Storage Virtualize

Miscellaneous



64 Gb FC Adapter in FS5300 – 8.7.0

- Feature ALB1 – 64Gb FC Adapter Pair
- Availability
 - FS5300 announced in April
 - Available August 9

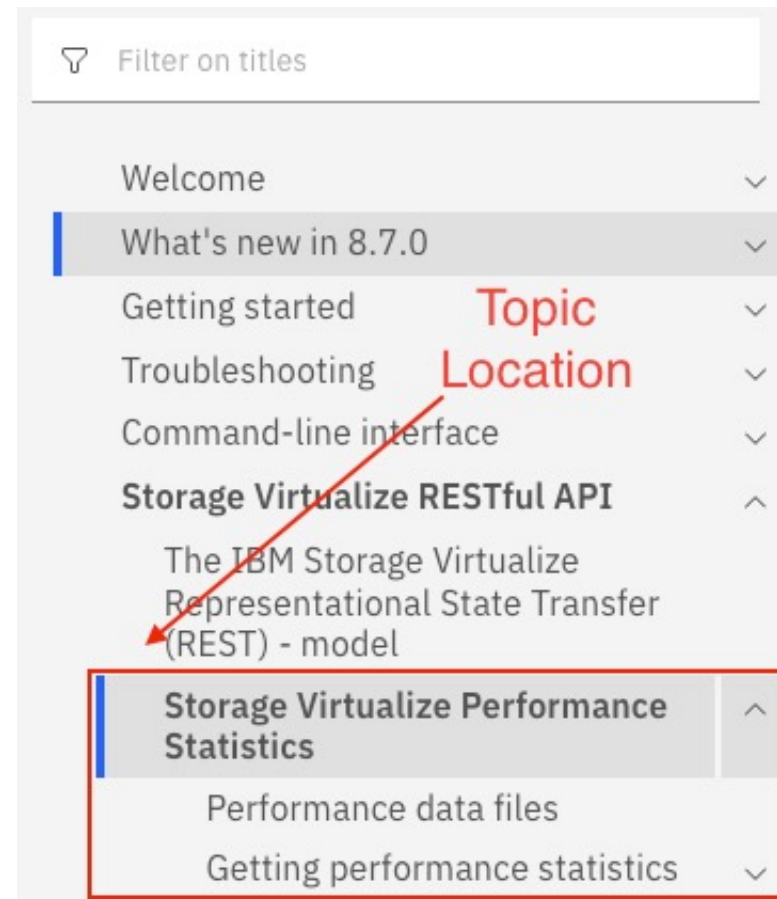
Node Join Performance – 8.7.0

- Decreased I/O pause when a node rejoins the cluster
- Node join used to be one large block of processing
- Node join in 8.7.0+ is 2 blocks of processing
 - Extends overall rejoin time slightly, but in the example given in the DST the I/O pause was less than $\frac{1}{2}$
- Results may vary based on system configuration and circumstance, however overall, I/O pause should be smaller than today

Performance Stats on the REST API

• Timeline

- 8.6.1.0 – Disabled CIMOM
- 8.6.1.0 – Enabled REST API file download
 - Drop-in replacement
- 8.7.0.0 – Documentation update
 - Migration guide



Performance Stats on the REST API

Steps



Limits

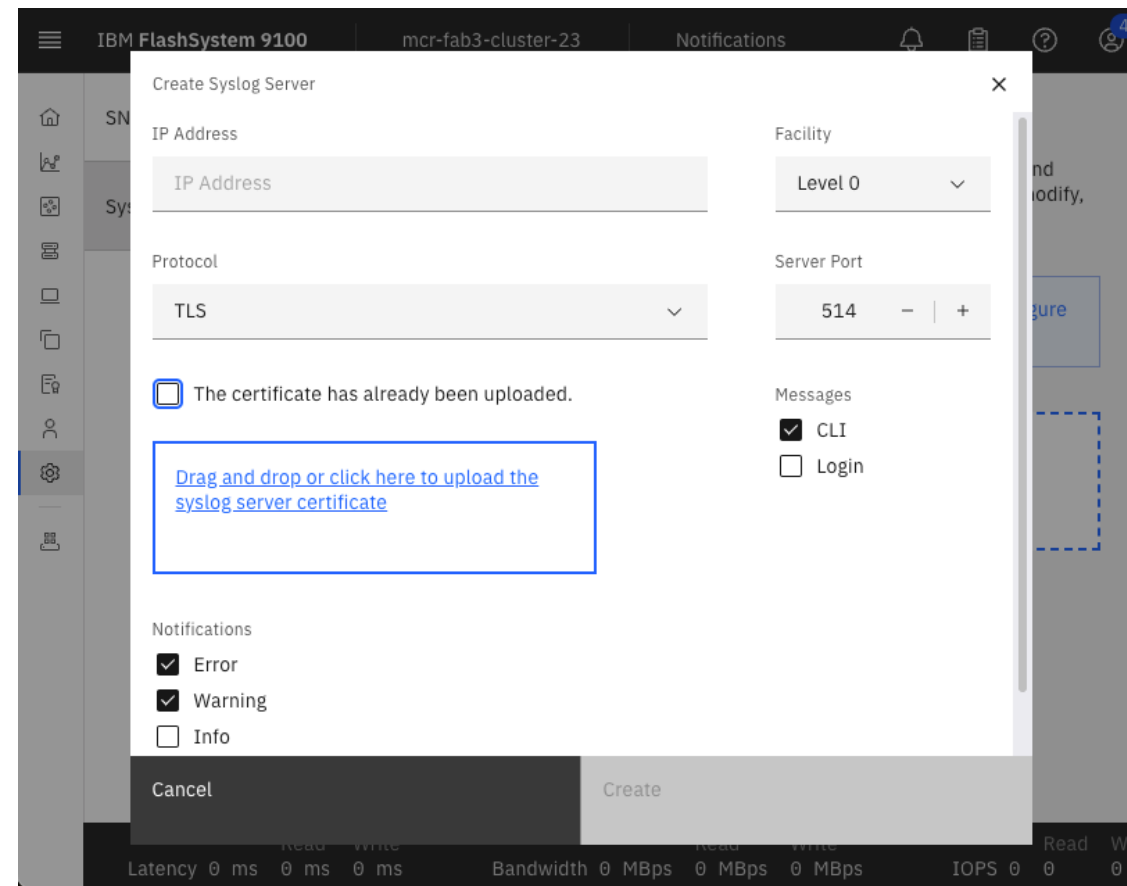
- Existing limits apply
- 3 req/sec auth
- 10 req/sec command

Syslog over SSL

- System level security configurations
- Cipher selection
- Server-side certificate validation



- GUI Support



Syslog over SSL

• Steps

1. Prepare the CA bundle

1. `scp <path-to-ca-bundle-used-by-the-syslog-server>
superuser@<system>:/tmp/<ca-bundle>`

2. Create a truststore entry for the CA bundle, and enable the “syslog” flag

1. `mktruststore -syslog on -file /tmp/<ca-bundle>`

3. Create a syslog object for the encrypted transfer of system events

1. `mksyslogserver -name <name> -login on -ip <remote-syslog-server-ip> -protocol tls`

4. Verify that the logs are received

1. Check the system has a truststore with “syslog” flag enabled

- `lstruststore`

1. Check the system has a syslog object with “tls” flag enabled

- `lssyslogserver`

Thank you!

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](https://www.menti.com/join/17086924) with code 1708 6924 or

Direct link <https://www.menti.com/alwhyze7z1gz>

Or

QR Code

