ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar: Storage Virtualize 8.7.2 Technical Update

Byron Grossnickle
IBM Storage Virtualize/FlashSystem SME
byrongro@us.ibm.com

Accelerate with ATG Technical Webinar Series

Advanced Technology Group 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.

2025 Upcoming Webinars – Register Here!

Unlocking the Future of AI and Application Modernization- IBM Fusion 2.9 - January 21st, 2025

Take Your CSM to the Next Level: What's New in IBM Copy Services Manager 6.3.12 & 6.3.13 - January 28th, 2025

IBM Storage Ceph RADOS Gateway Deep Dive - February 6th, 2025

Important Links to bookmark:

ATG Accelerate Site: https://ibm.biz/BdSUFN

ATG MediaCenter Channel: https://ibm.biz/BdfEgQ

Offerings

Client Technical Workshops

- IBM Storage Scale & Storage Scale Functions
- Cyber Resilience with IBM Storage Defender
- IBM DS8000 G10 Advanced Functions
- IBM Fusion & Ceph
- IBM FlashSystem Deep Dive & Advanced Functions

TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Representative or Business Partner for more information.

***IMPORTANT* The ATG team serves clients and Business Partners in the Americas, concentrating on North America.**

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 with code 5151 0447 or

Direct link https://www.menti.com/alhsf3bgvxu6
Or

QR Code

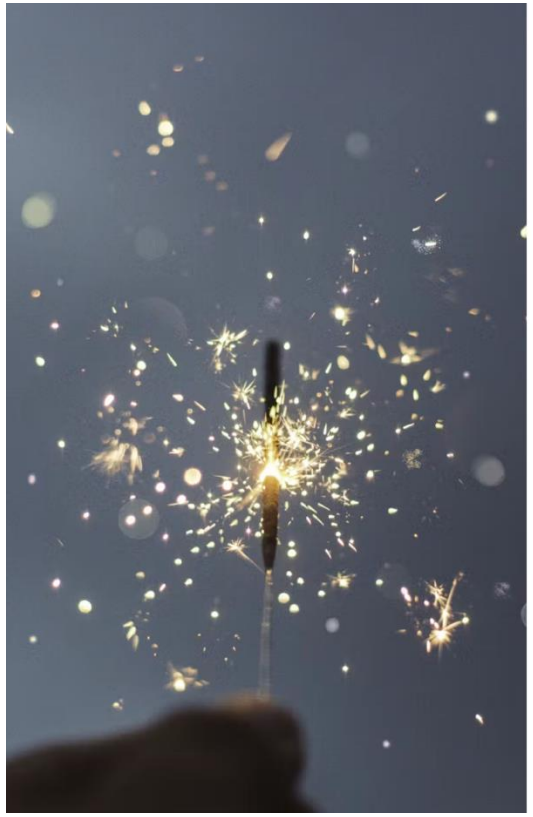

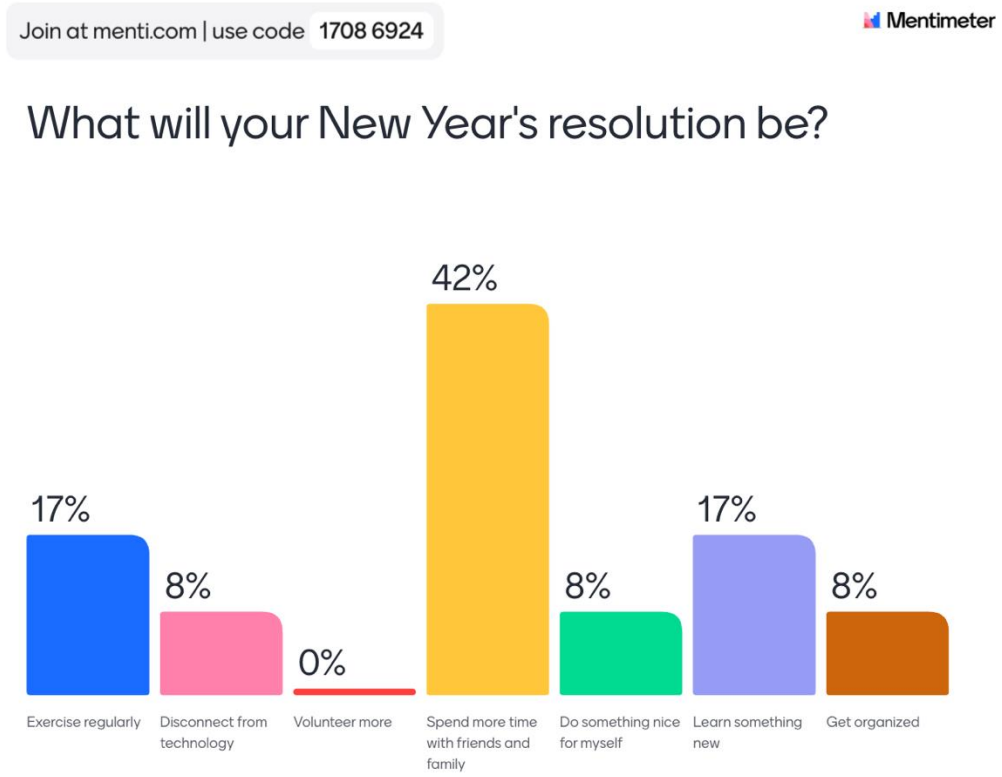

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar: Storage Virtualize 8.7.2 Technical 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.

IBM Storage Virtualize 8.7.2 Agenda

- Release Schedule
- Security Updates
  - Internal Key Management
- Host Attach Enhancements
  - Using Fabric (FDMI) Info for Host Management
  - Jumbo Frames
  - CHAP Secret Enhancements
  - Scaling NVMe/FC
- Replication and HA Updates
  - iSCSI HA with Portset Linking
  - CHAP Secret Changes
  - Interop Update for PBHA
  - Restoring/Refreshing Snapshots PBR/PBHA
- Flash Grid Updates
  - GUI
  - Storage Partition Migration Automation
- Misc Updates
- Plugin Updates
  - VMware
  - CSI

Release Schedule

- RFA Announce – 8.7.2 – November 12, 2024
- eGA – 8.7.2 – November 29, 2024
- pGA – There is no hardware associated with this release

8.7.2 is a Non-LTS Release. Non-LTS releases are tested identically to LTS releases. Non-LTS releases will not get any patches and will not be maintained long term. Those patches will be available in the next Non-LTS or LTS release.

– Release FAQ

IBM Storage Virtualize

Encryption
Internal Key Management

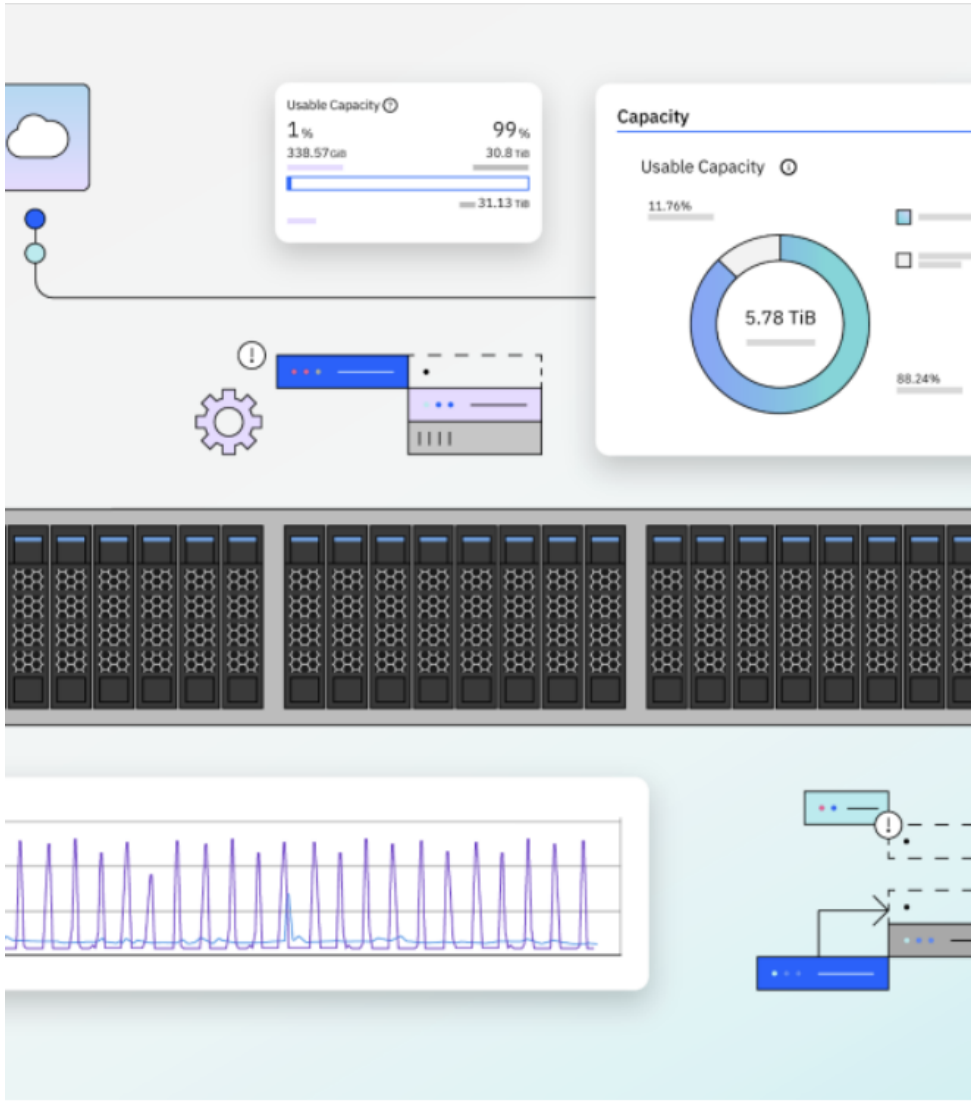

Internal Key Management (IKM) Overview

- Existing External Key Management
- USB
- Key Servers- GKLM, CypherTrust (Thales)

- External Key Challenges
- Storing the main keys
- Redundant server infrastructure for key manager - complexity
- Managing the main key for encryption enabled system
- How many copies of keys and where to store them
- Can lose keys, security risk if lost, access to encrypted data lost forever if all keys destroyed

- Other Drivers
- USB constraints, Security Issues
- Key server software cost and maintenance overhead

- Internal Key Management
- Stores a copy on the internal boot drive(s) of every node
- AES-256 bit key, sealed by the node's unique TPM chip
- Automatic system, unlocks, no user intervention
- Designed for secure site environments

Internal Key Management (IKM)

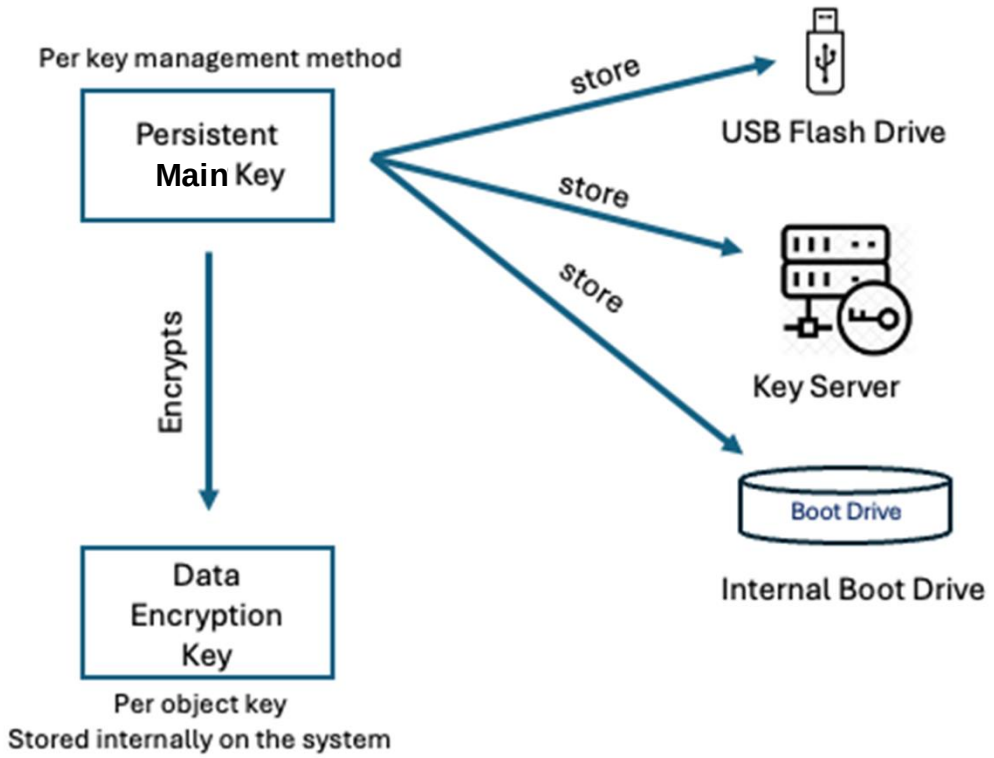

- **How it works**

- IKM stores keys secured internally
- Keys are secured using TPM Hardware
- TPM Sealed Keys are stored on all the nodes for redundancy
- System automatically attempts to unlock on power up
- IKM Keys are rekeyed automatically every 24 hours
- Encryption Recovery Key is recommended to be enabled

- **Security**

- User shall not have access to IKM Keys
- Keys are stored on internal drives within systems and so it is more secured
- Each node stores encrypted IKM Keys with its own unique TPM hardware

Internal Key Management (IKM)

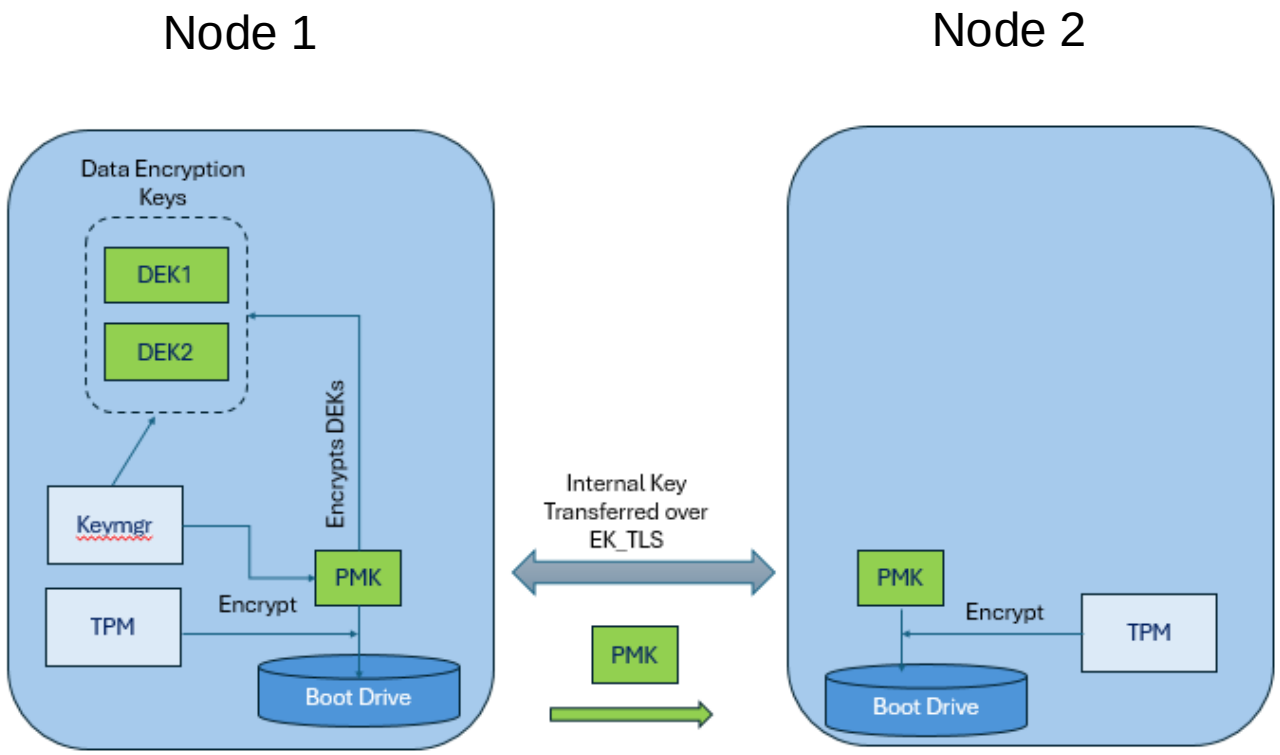

- **Recovering Data**

- **IKM Keys are stored on all the nodes in the cluster for redundancy**
- **With any node surviving in the cluster, system is unlocked automatically**
- **During T1, T2, T3 recovery, system can automatically unlock using TPM hardware, internally stored key and recover data**
- **In case of multiple hardware failures (TPM, Drive, Node..) , Encryption Recovery Key allows us to recover data**

CLI Changes

- A new CLI option “-internal” in existing “chencryption”

```
#chencryption -internal -enable
                        -disable
```

- A modification in ”lsencryption” CLI to display the internal key management status and last rekey time

```
#lsencryption

status licensed
error_sequence_number
usb_rekey no_key
usb_key_copies 0
usb_key_filename
usb_rekey_filename
keyserver_status licensed
keyserver_rekey no_key
keyserver_pmk_uid
keyserver_pmk_rekey_uid
recovery_key_status enabled
recovery_key_rekey no
recovery_key_name recoverykey_000002043CA0D058_3CA0D05800000001_Fab3Plus
recovery_key_rekey_name
internal_key_status enabled
usb_last_rekey_time
keyserver_last_rekey_time
recovery_key_last_rekey_time 240808111327
internal_key_last_rekey_time 240808111247
```

Things to Think About

- Reasons to Use IKM

- Easy to configure
- Prevents “losing” the key
- Automatic rekey
- No intervention upon coming up from a power down

- Reasons to Think About the Deployment of IKM

- If your location is not secured and the physical array is stolen your data is compromised
  - This can be prevented with external key management
- When decommissioning a FlashSystem, you need to remember to do a secure erase on the unit before it goes out the door
  - If not, again, your data is compromised
  - While external key management does not replace the best practice of a secure erase, if the external keys are not with the storage array when decommissioned, the data is not compromised
- Internal and external key management can be used simultaneously, but other forms will need to be configured from the command line

IBM Storage Virtualize

Non-HA Host Attach Updates

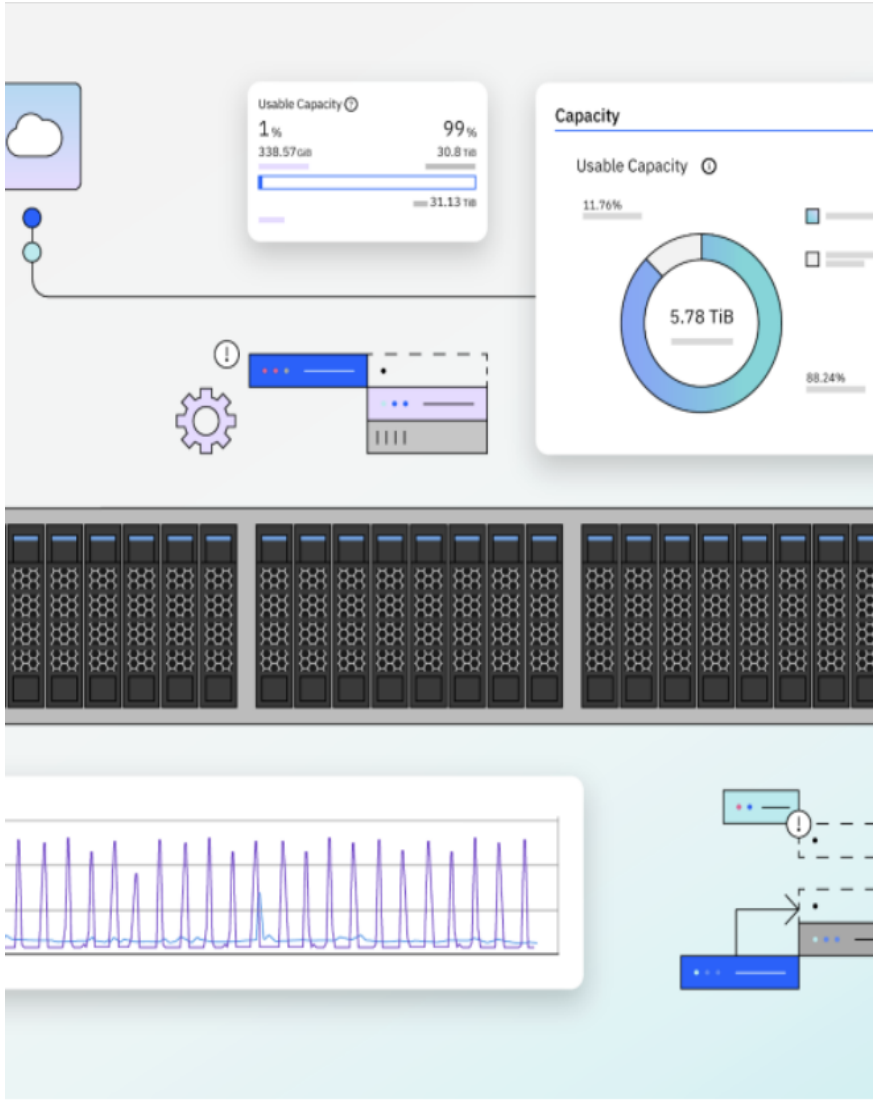

Improved FC Host Management with Fabric Device Management Interface (FDMI)

- Current Challenges
  - Only zoned WWPNs are listed in the GUI
  - No grouping of WWPNs based on the host server
  - Hard to correlate listed WWPNs with the real host server
  - Hard to configure unzoned WWPNs
  - Addition or replacement of host port requires burdensome management steps

- Benefits
  - Understand fabric better
  - Simplified host configurations
  - Know host connectivity
  - Easier Interop Checking

Fabric Insights Help with Host Management

- FC Fabric Insights
- All hosts connected in the fabric (zoned / unzoned)
- Host name configured on the server
- Host OS
- Host Port details
- Host Paths to FlashSystem
- 1 click fabric view refresh is available

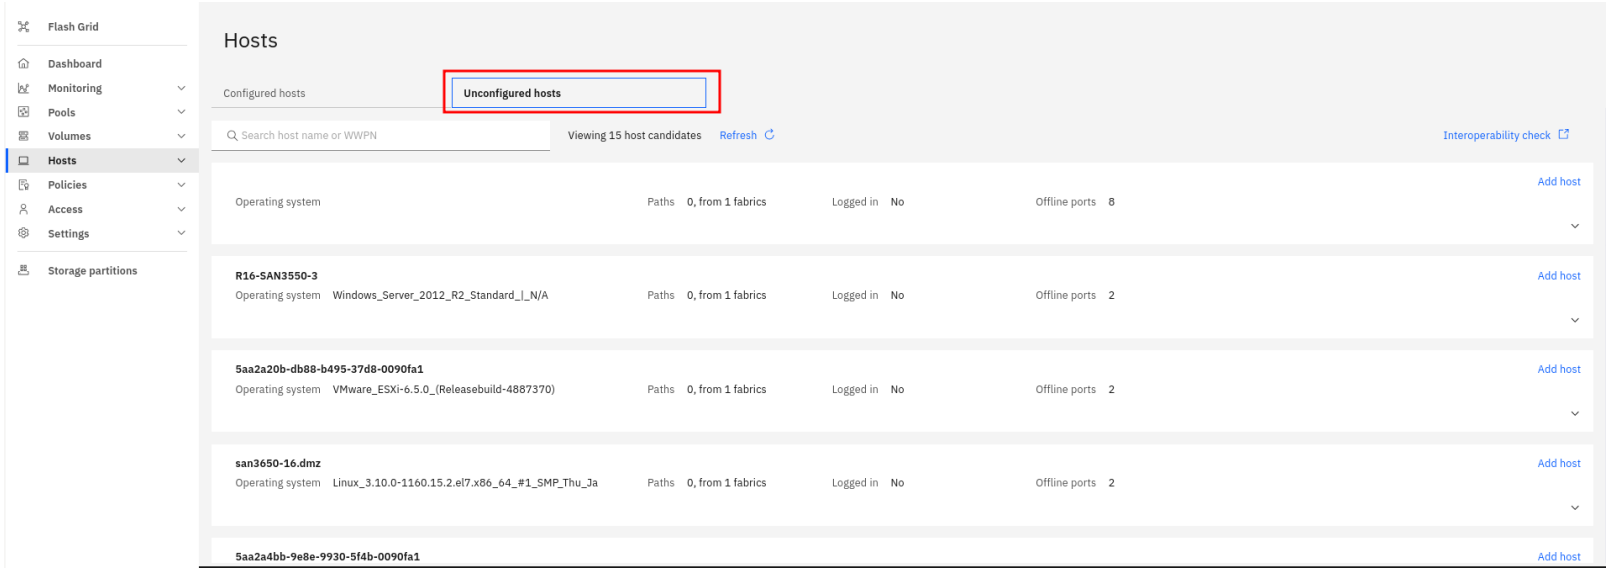

One Click Host Configuration

- The discovered host can be added with names configured on the server
- Host Server name
- Host OS
- WWPNs from the server (zoned / unzone)
- Host connectivity details to FlashSystem
- Simplified CLI to configure a host with the server's name
- System automatically discovers all WWPNs on the server

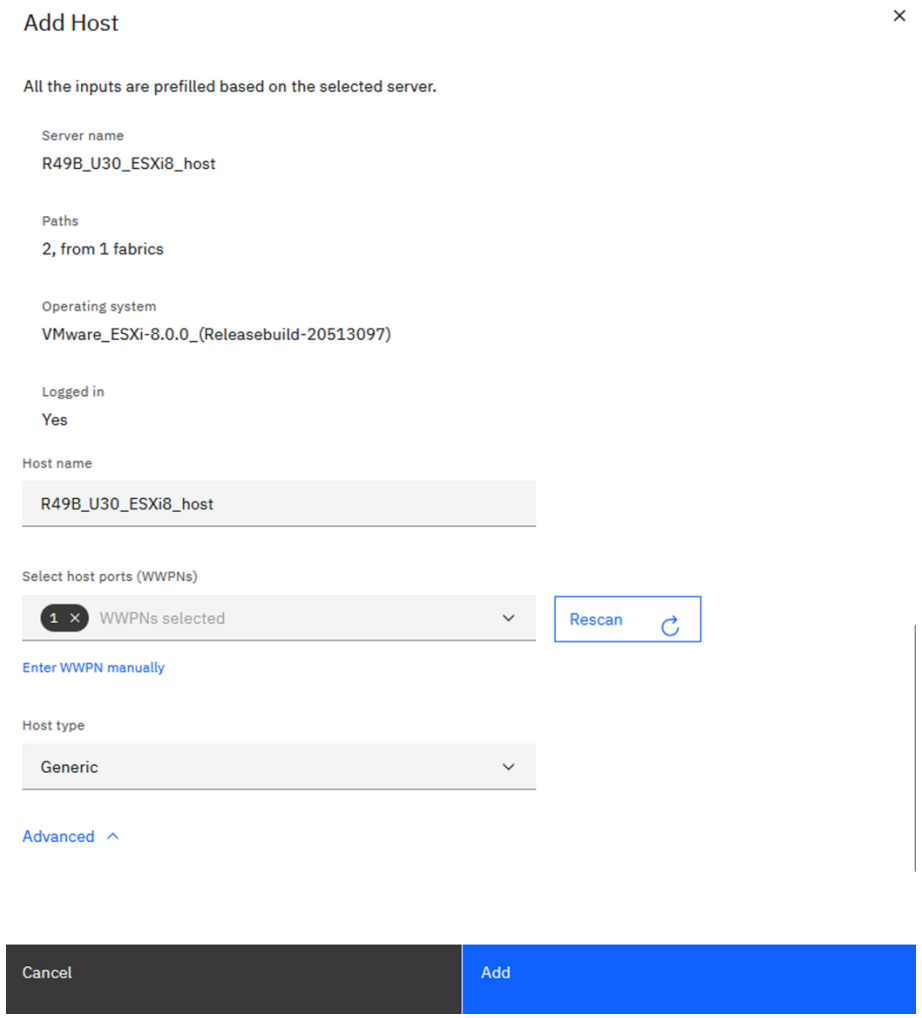

Host Connectivity Insights

- Login Status
- Symmetric/ Asymmetric /Offline
- Login Counts
- Total logins to the FlashSystem
- Valid Login Count
- Logins serving valid SCSI paths
- Speed
- Host port speed for troubleshooting

rhel9.4-pune Add host

Operating system: Linux_5.14.0-427.26.1.el9_4.x86_64_#1_SMP_PREEM
Paths: 30, from 5 fabrics
Logged in: Yes
Offline ports: 1

WWPNs	Login status	Speed	Login count	Valid login count	Fabric WWN	I/O group count	Serial number
10000090FAD16058	Asymmetric	16Gb	1	1	500507680B53AC4A	1	11S01CV842Y650HY67L070
10000090FA941EDF	Symmetric	16Gb	4	4	200100DEFBDA4C21	1	11S01AC345Y450HY58L00B
10000090FA9421CC	Symmetric	16Gb	4	4	100050EB1AD959B9	1	11S01AC346Y450HY58L03G
21000024FF16C47D	Asymmetric	32Gb	1	1	500507680B5A5D2E	1	RFD1707R30419

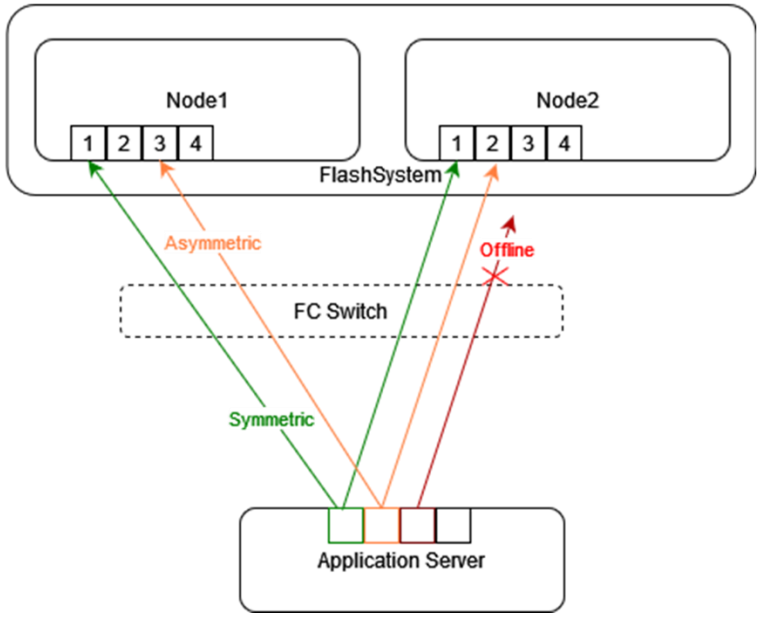

Symmetric vs Asymmetric Host Access

- Symmetric
- Nodes have same number of valid logins and are to the same set of ports

- Asymmetric
- Nodes have different login counts from the host port OR logins are not on the same ports
- If a host shows as asymmetric, it generally means there is a connectivity or zoning issue

Easier Interop Checking

- Operating System
- OS version
- HBA Details
- FW version
- Driver version
- Vendor Name
- These details can be used to verify interop support using SSIC [https://www-03.ibm.com/systems/support/storage/ssic]

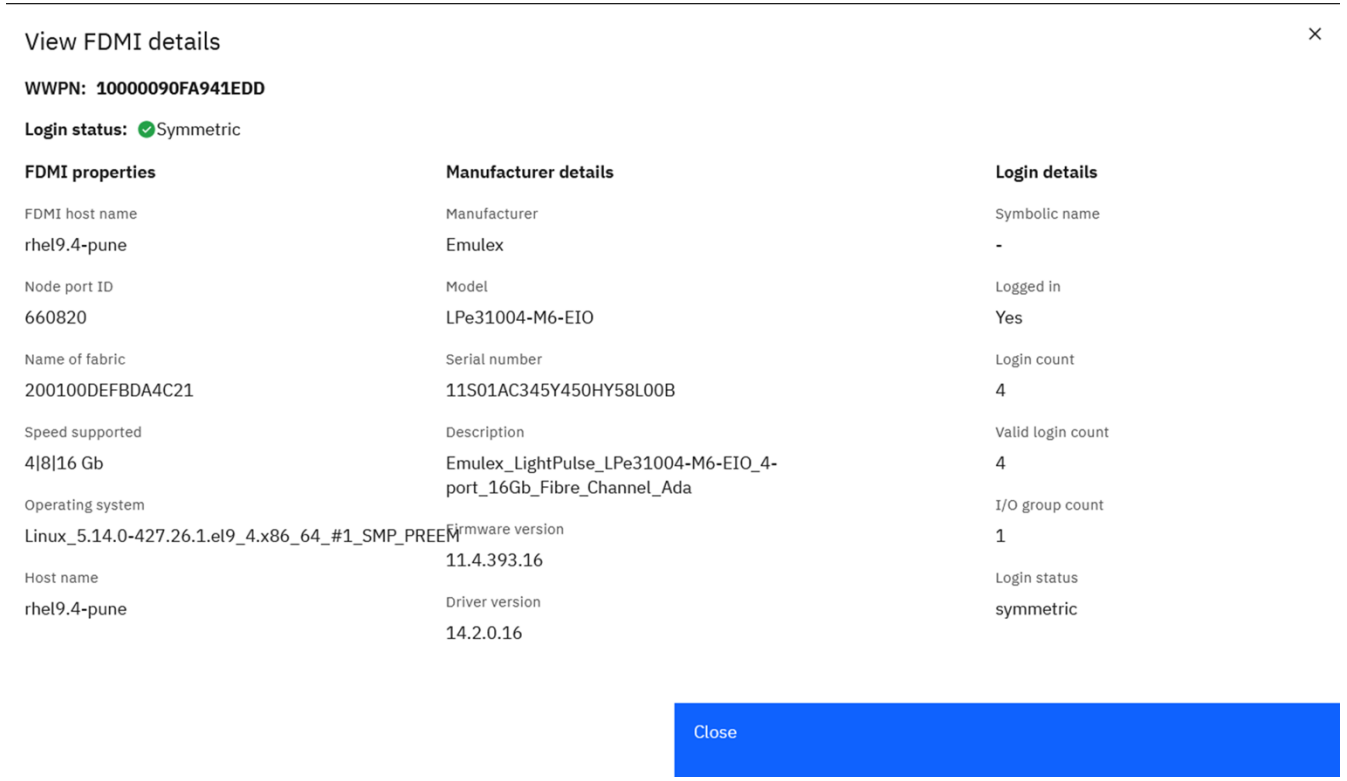

– Note: FDMI detail will vary based on how well host OS communicates with FDMI

Jumbo Frames – MTU9000

- Prior to 8.7.2, MTU 9000 was supported only for iSCSI host attachment.
- Starting from 8.7.2, MTU 9000 can now be used for NVMe/TCP host attachment as well.
- There are no changes to the CLI/GUI

Scaling NVMe/FC

- Why -
  There can be a benefit to NVMe/FC, particularly with VMware where their NVMe stack is more scalable/performant than their SCSI stack. Our current limits were insufficient for most of VMware environments

- What -
- We have increased FC-NVMe/FC host limit to 256 (previously 32) and number of per FS-node port logins to 64

- How -
- User will have to make use of non-default FC port-set while creating NVMe/FC host using GUI/CLI

iSCSI Authentication Updates

- 8.7.2 Onward

- Deprecates system-wide Bidirectional CHAP, previously configured through the chsystem CLI
- Introduces Host-wide Bidirectional CHAP, configurable through the chhost CLI
- Adds an option to enforce FIPS-compliant hashing algorithms for CHAP in the chsecurity CLI
- Eliminates the display of cleartext CHAP secrets in both the CLI and GUI

- Benefits

- This request was primarily driven by a customer (RFE SCSI-I-1247)
- Provides better control over hashing algorithms used in CHAP authentication
- Offers granular control of Bidirectional CHAP, making it more suitable for multi-tenancy
- Eliminates system-wide disruptions. Previously, changing or removing system-wide CHAP was a disruptive process, potentially causing LUN outages if not properly coordinated with host configurations.
- The new model aligns well with PBHA, facilitating smoother volume mobility

Hashing Algorithm Control

- Introduces a new security setting called iSCSI Host Authentication Mode
- When **disabled**:
  No restrictions on hashing algorithms; all currently supported algorithms MD5, SHA1, SHA2, and SHA3-256 are allowed.
- When **enabled**:
  Only FIPS-compliant hashing algorithms are permitted, specifically SHA2 and SHA3-256 from the currently supported options.

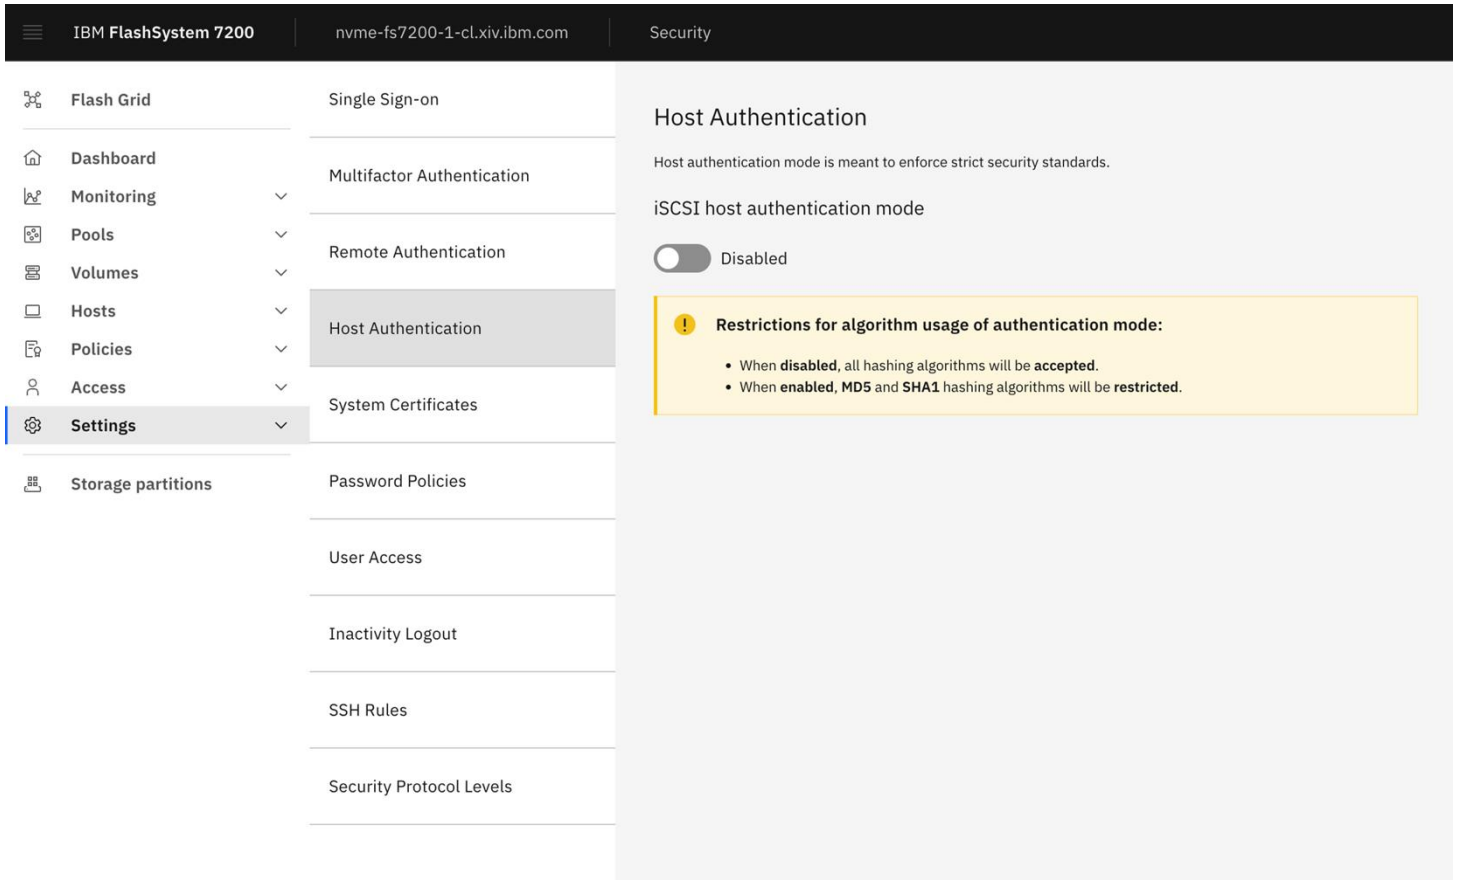

Per Host Bi-directional CHAP

- From 8.7.2 onwards at host level we will also have bidirectional CHAP configuration option along side unidirectional
- Secrets are no longer displayed
- When upgrading, the system wide CHAP will automatically be moved to the host definitions so there will be no disruption when migrating from one CHAP model to the other

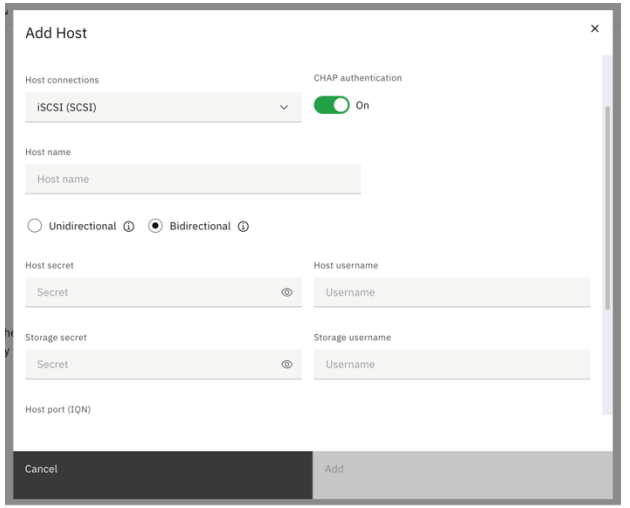

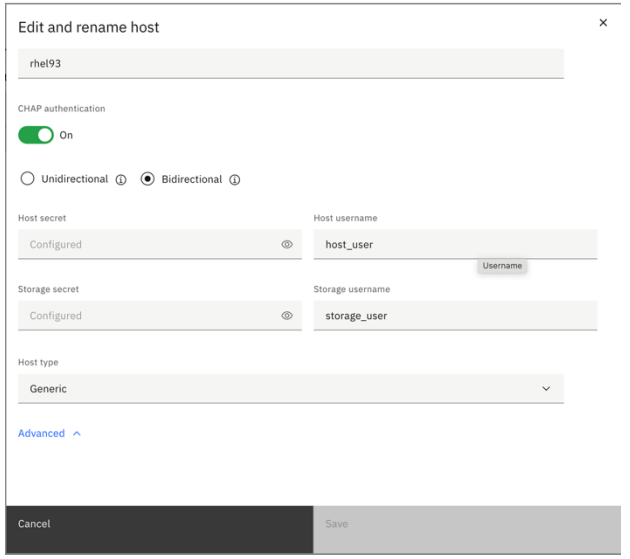

IBM Storage Virtualize

Replication and HA Changes

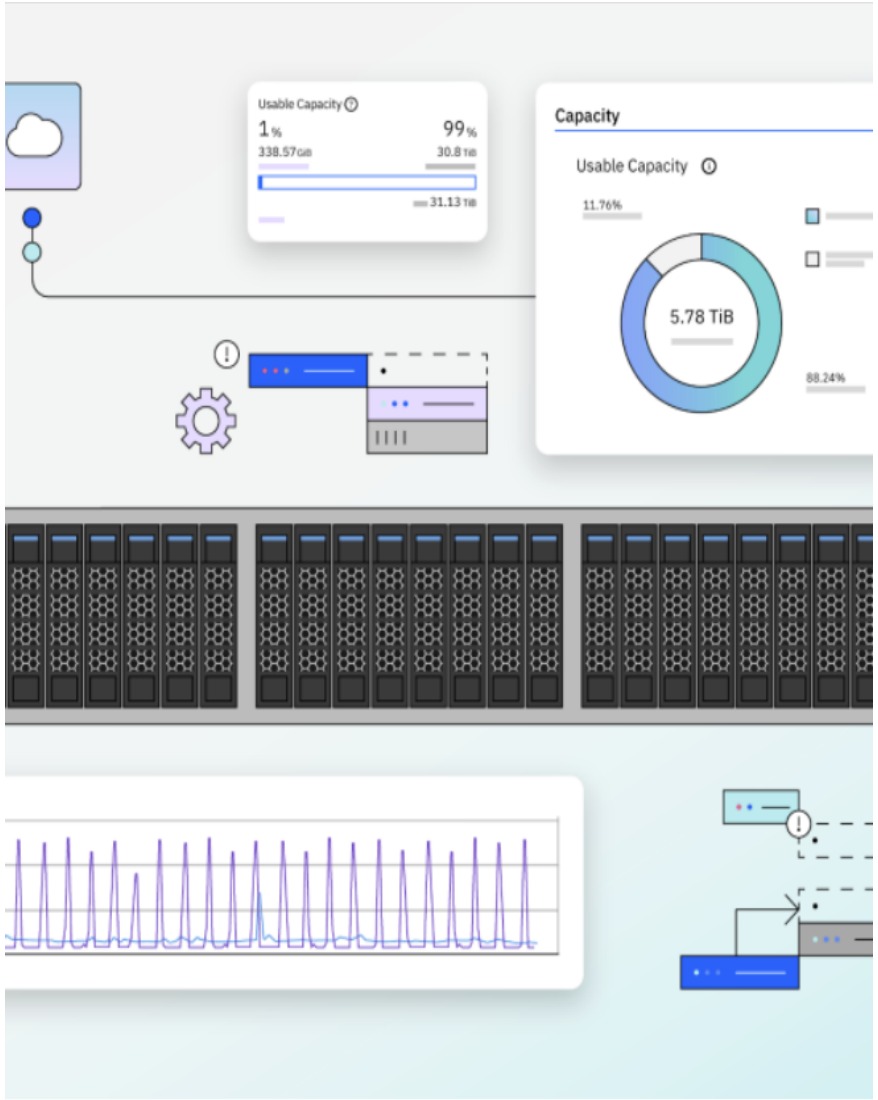

Restore/Refresh Replicated Volumes

- Support for restoring or refreshing from snapshots for volume groups/volumes using in policy-based replication, policy-based HA or HA+DR.
- Allows instant restore from regular snapshots and safeguarded snapshots.
- With asynchronous replication, restoring and refreshing can only be run against production or independent volume groups/volumes.
- HA volume groups/volumes can only be restored or refreshed using snapshots at the active management location – remember that the management location can be changed!
- Restoring or refreshing HA volume groups/volumes will cause HA to be **temporarily** lost while the contents of the snapshot is synchronized to the partner. This is automatic and requires no intervention.

Policy-based High Availability

- A warning is displayed indicating that restoring will temporarily cause HA to be suspended until the resynchronization completes

- While the partition is resynchronizing, a snapshot is automatically taken at the other site to maintain a point in time should a failure occur

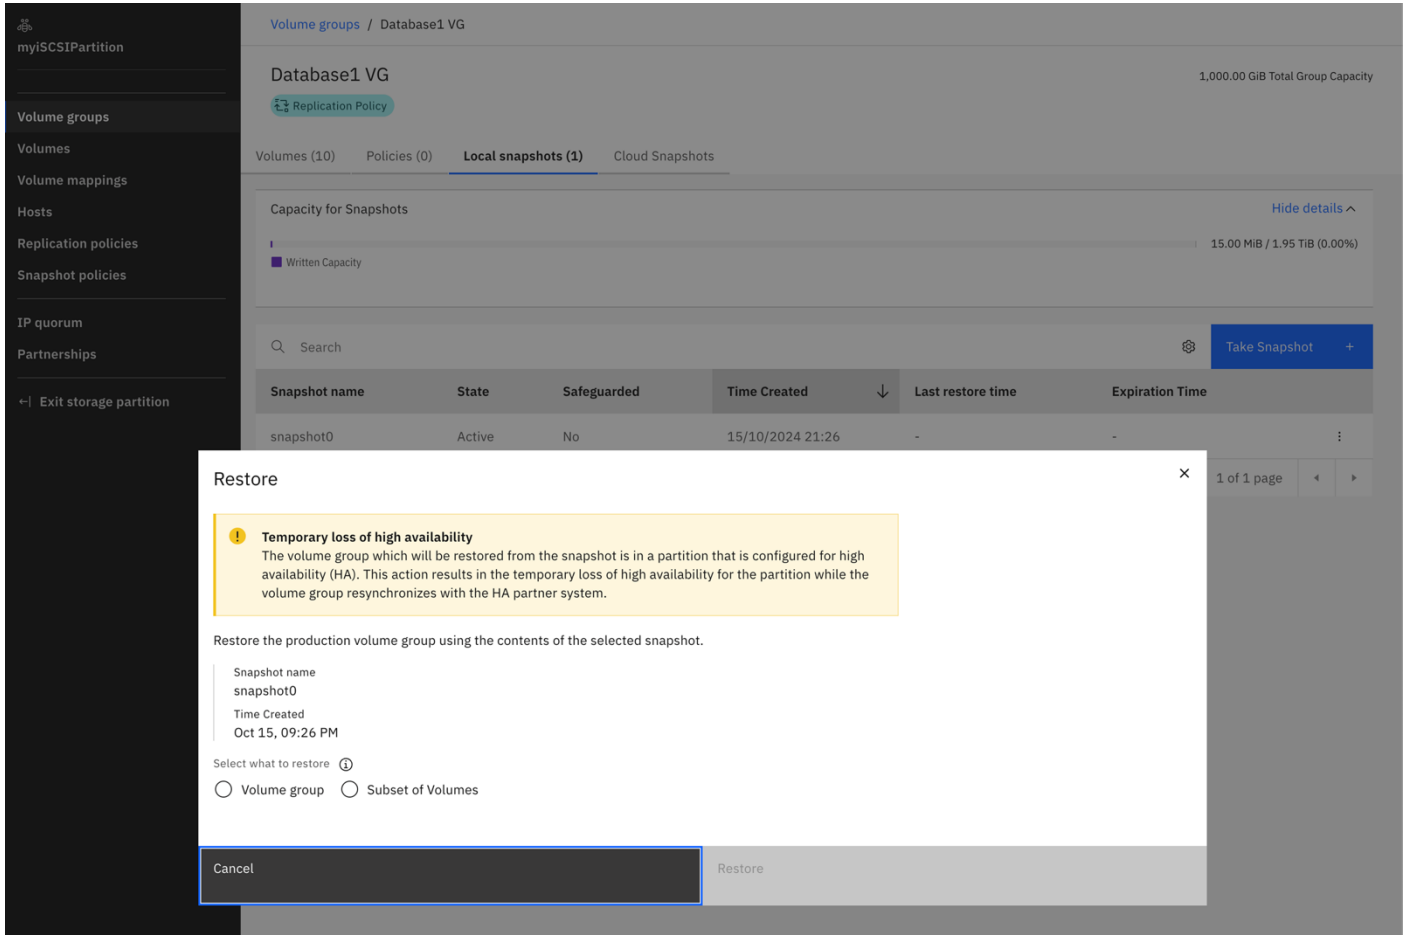

Policy-based High Availability

- Refreshing replicated volume groups or volumes works the same as non-replicated volume groups
- To refresh a volume or volume group, it must be a thin-clone
- Replication or HA will incrementally resynchronize the changes from the snapshot

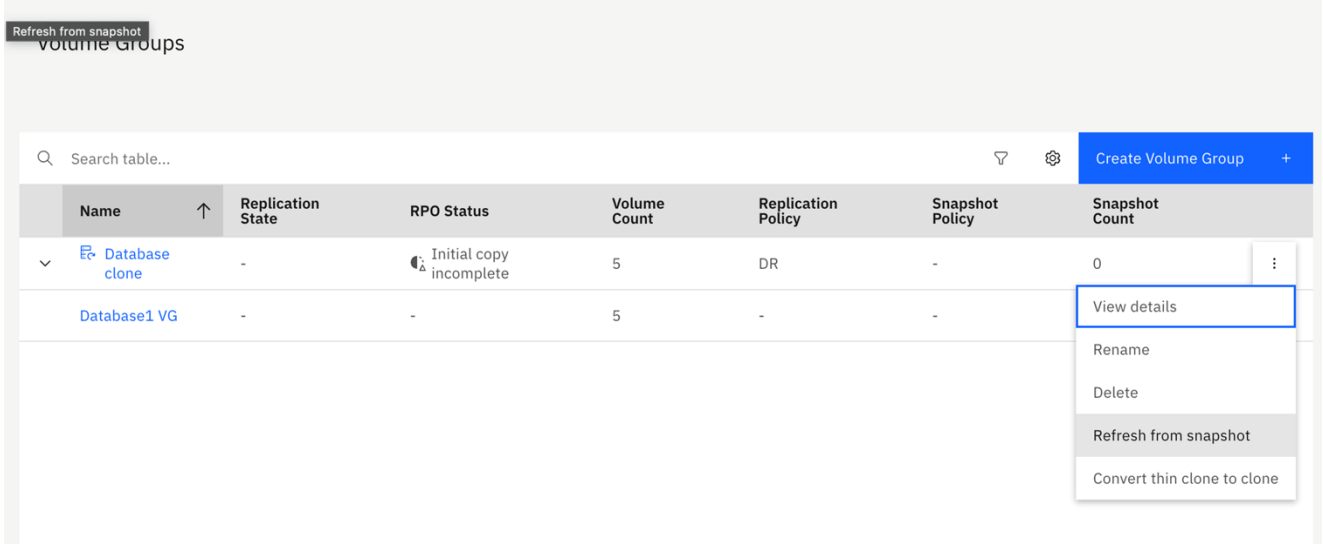

iSCSI Support for Policy-based High Availability

All-Ethernet HA deployment

- RDMA-based Ethernet partnerships can be used for HA, with optionally either RDMA or TCP/IP partnerships for async replication to create a HA+DR solution.
- Combined with iSCSI host support, allows for an Ethernet-only deployment.
- The user-experience is *very* similar to configuring a Fibre Channel setup, with the only difference being linked portsets.
- CHAP secrets are supported for HA hosts, but only the new 8.7.2 per-host username/secret implementation as it is more secure.

Linked portsets

- When configuring HA with iSCSI hosts, portsets must be linked between the systems. Portsets are optional for Fibre Channel, but this applies to Fibre Channel too.
- Reminder: Portsets are a group of IP addresses or WWPNs assigned to ports on the system.
- Using the same model as pool linking, portsets are linked between the two systems which defines which portset will be used on the partner when hosts are created.
- Default host portsets are auto-linked after the partnership is created.
- Hosts cannot be associated with an empty portset! This applies to non-HA hosts as well, to avoid unintentional access issues.

Interoperability support

- All supported versions of:
  - Red Hat Enterprise Linux
  - SUSE Linux Enterprise
  - VMware ESXi
- Other operating systems will be added as testing completes.
- PBHA Host Interop

Portset Linking in PBHA

- Portsets can be linked from the GUI network settings menu after creating portsets…

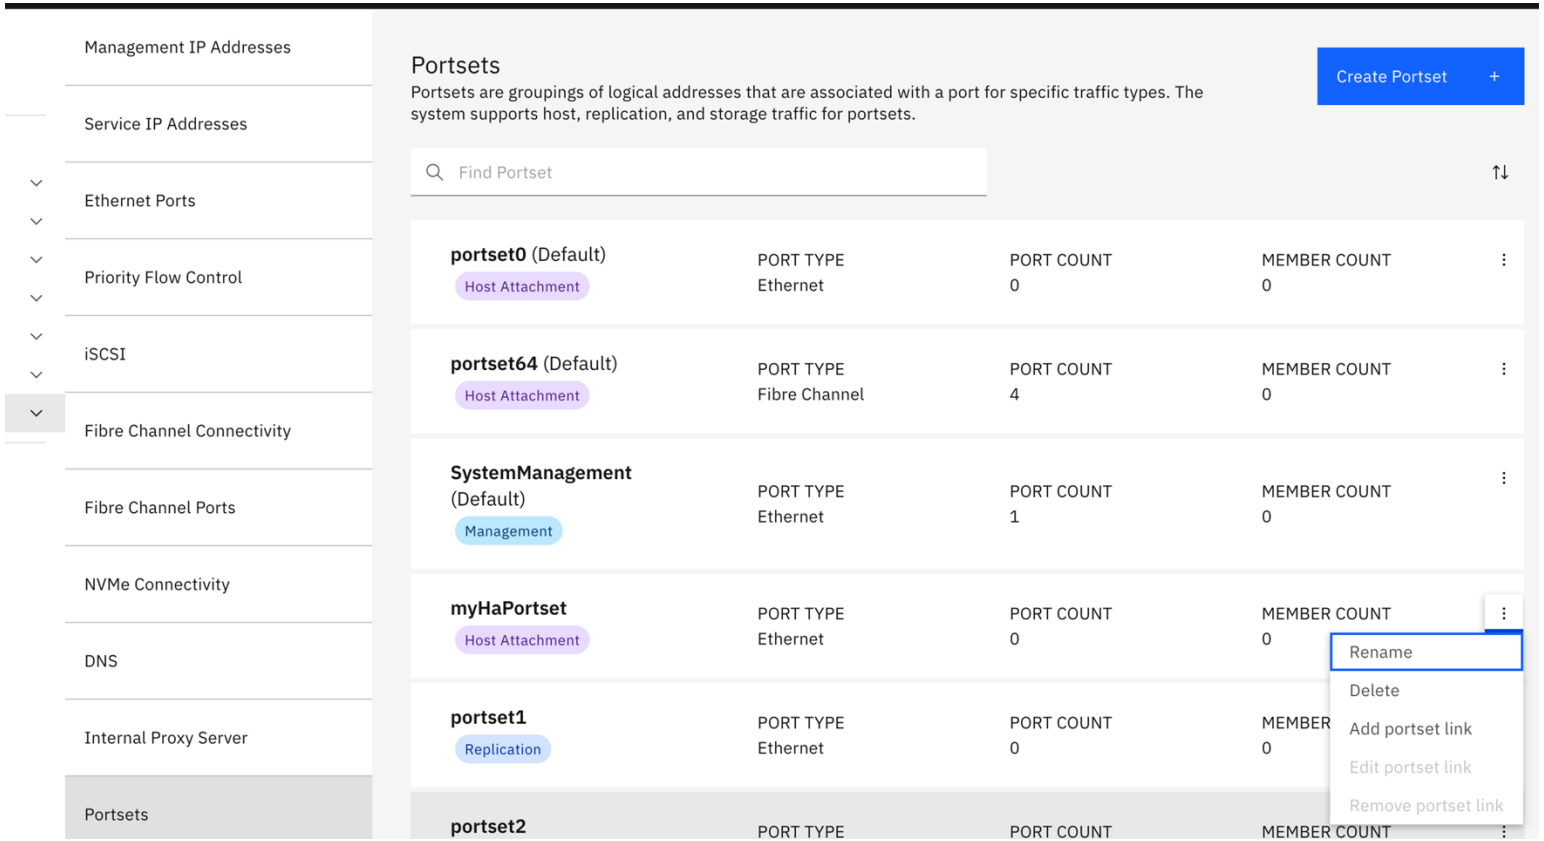

Portset Linking

- Or, using the partition HA setup wizard…

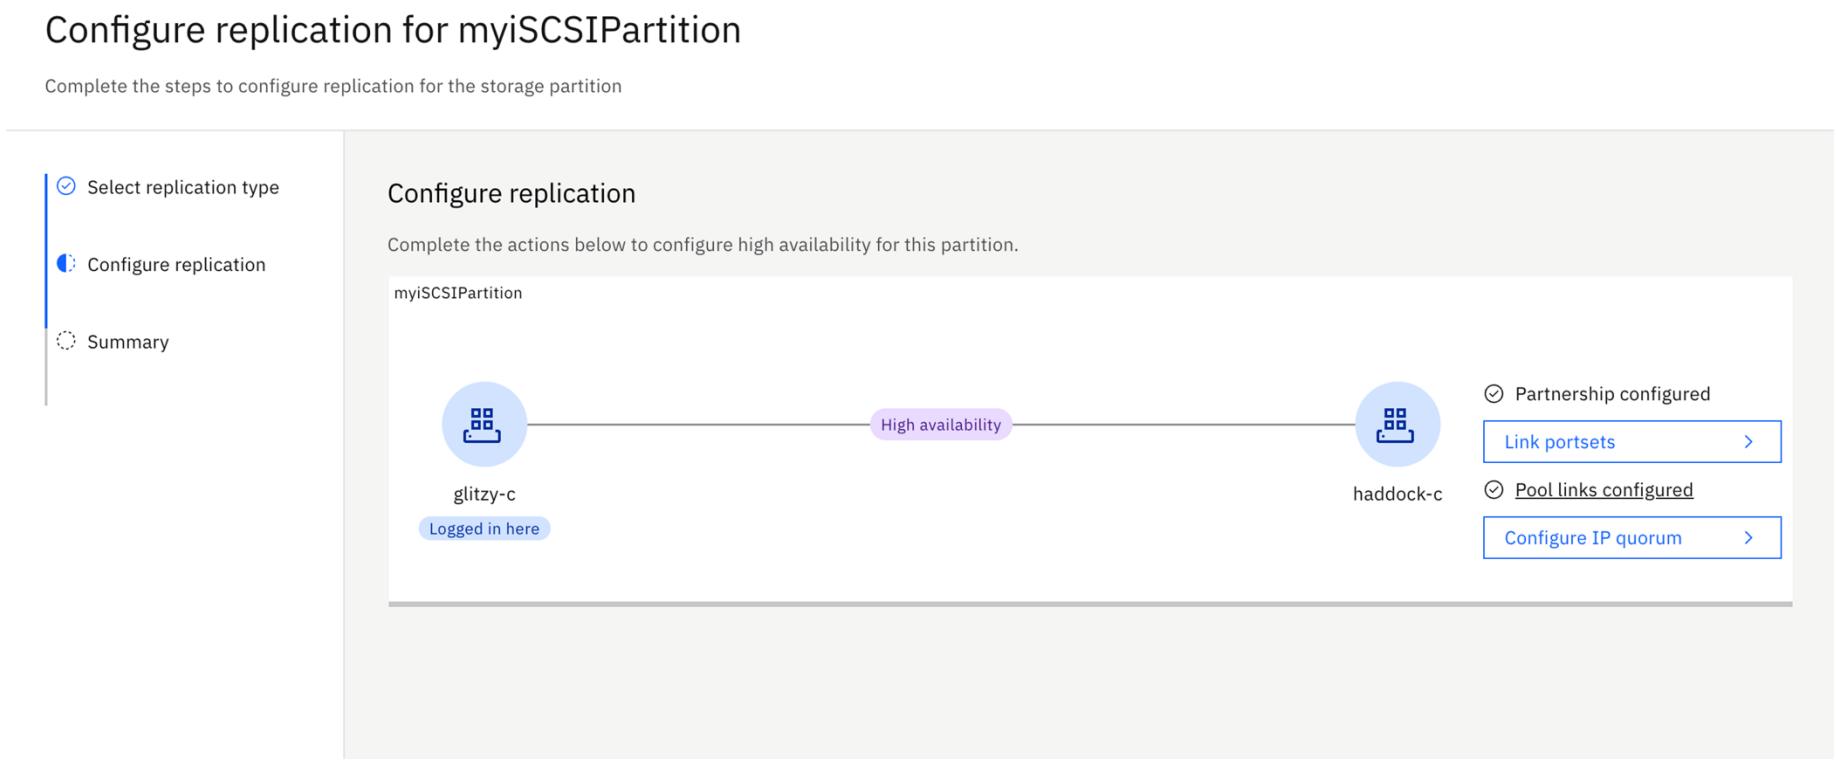

Portset Linking

- Or, using the partition HA setup wizard...

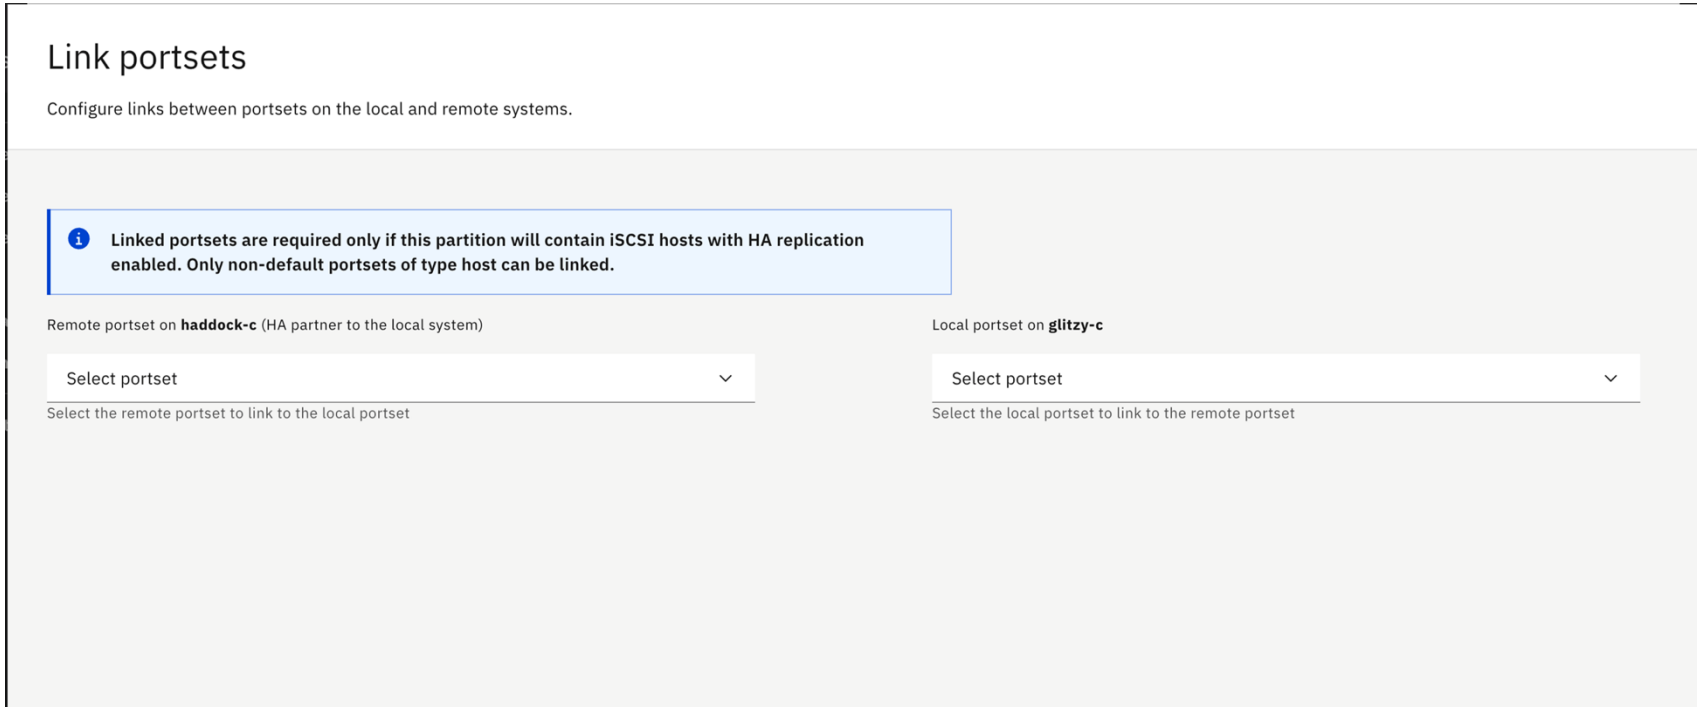

Import from external storage while maintaining existing DR

- 1. Create image-mode volumes with cache-disabled from the external storage controller
- The external storage can maintain its existing replication. If a disaster occurs before the new replica is established, the storage admin can access the DR copy maintained by their old array

- 2. Add the image mode volumes to a volume group and assign a replication policy
- While the sync progresses, add a copy of the volume to the internal arrays

- 3. When the synchronization completes, remove the replication from the external array, and convert the volumes to be cache-enabled on the production system
- Delete the image mode volume copy once its copied to the internal array

- Notes:
- Cache must be changed to read/write mode once the synchronization completes as the replication direction cannot be reversed while cache-disabled volumes exist in the volume group.
- Only image mode volumes may be added with cache disabled, and cache cannot be disabled on replicated volumes.
- DR recovery volumes should not be exported using image mode as write cache is always enabled.

Misc. Replication Items

- Automatic Pool Linking
  - The first pool created in the system will automatically be linked to partnered systems if both systems are running 8.7.2+.
  - On upgrade, if there is only a single pool in the system and it's not currently linked, it will configure itself to be auto-linked.
- Interoperability for HA (and HA + DR)
  - The following page will be updated as new operating systems are qualified:

PBHA/PBHA+DR Interop

Supported OS's

- FC SCSI

  - Red Hat Enterprise Linux, including with KVM (all supported versions)
  - SUSE Linux Enterprise (all supported versions).
  - IBM AIX (version 7.2 and later)
  - Linux on IBM Z (all supported versions and distributions)
  - VMware ESXi (all supported versions)
  - Microsoft Windows Server, including Hyper-V (all supported versions). This requires both systems to be running IBM Storage Virtualize 8.7.2 or later

  – Hoping to have IBM-I support soon (expected this month)

- iSCSI

  - Red Hat Enterprise Linux, including with KVM (all supported versions).
  - SUSE Linux Enterprise (all supported versions).
  - VMware ESXi (all supported versions).

IBM Storage Virtualize

Flash Grid GUI

Partition Migration Automation

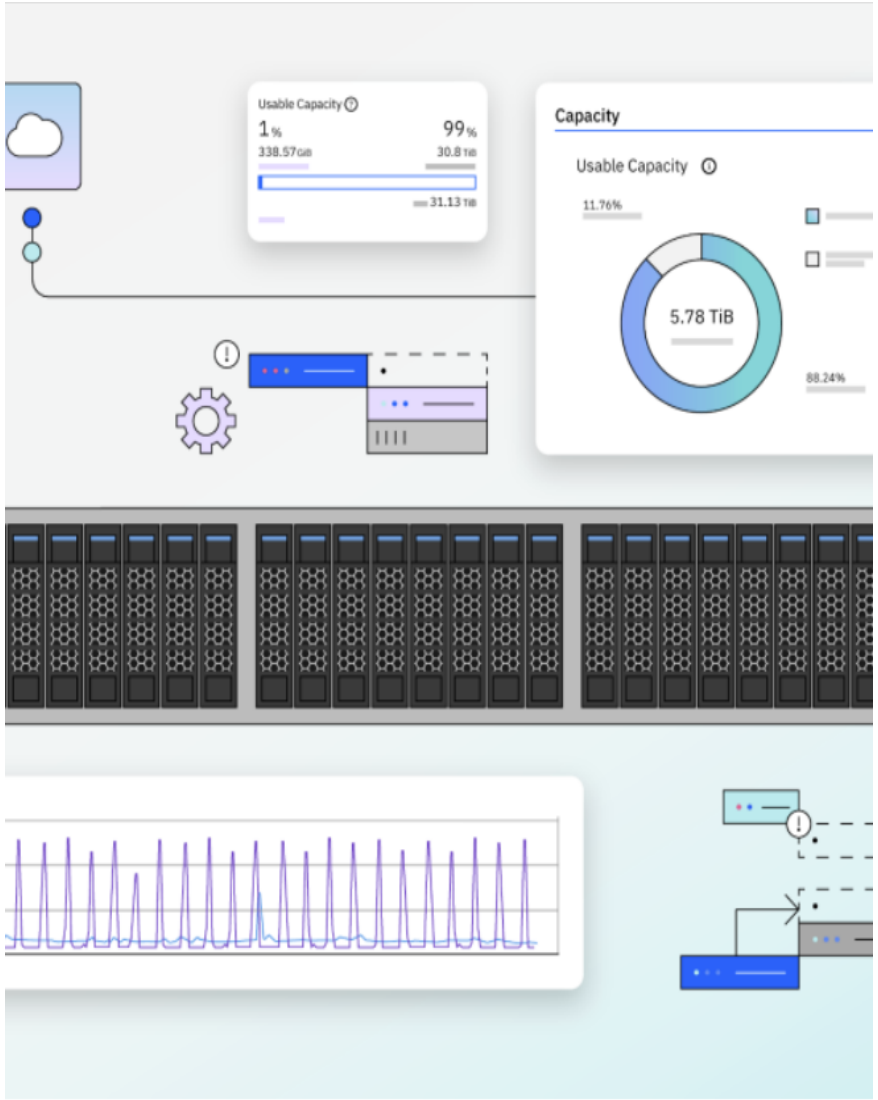

Flash Grid Configuration, Monitoring and Partition Migration

Before 8.7.2	In 8.7.2
Configuration (Create Flash Grid, Add Member) CLI Only	**Configuration (Create Flash Grid, Add Member)** **CLI and** GUI
Monitoring Storage Insights Only	**Monitoring** Storage Insights and **GUI**
Partition Migration CLI/Storage Insights	**Partition Migration** CLI/Storage Insights

Create a New Flash Grid

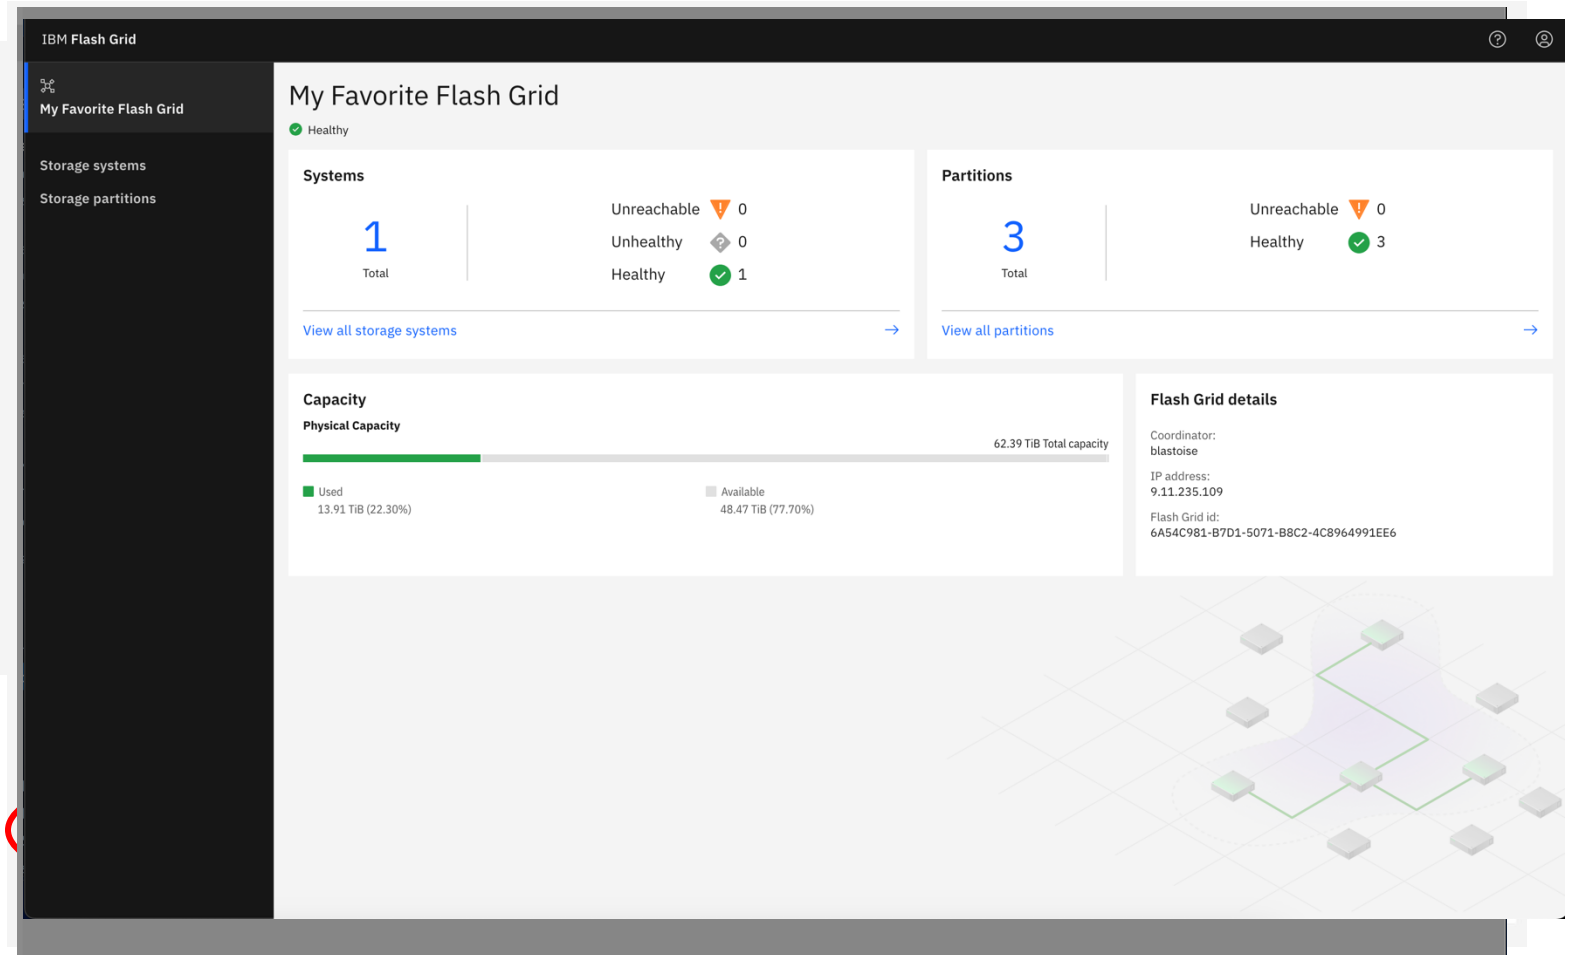

- 8.7.2 systems that support Flash Grid will have a new notification
  - Notification can be disabled

Steps

1. Click on Create a new Flash Grid
2. Name your Flash Grid
3. Successful creation will take you to the Flash Grid dashboard
   - Single system added
   - Flash Grid Coordinator

Flash Grid Dashboard

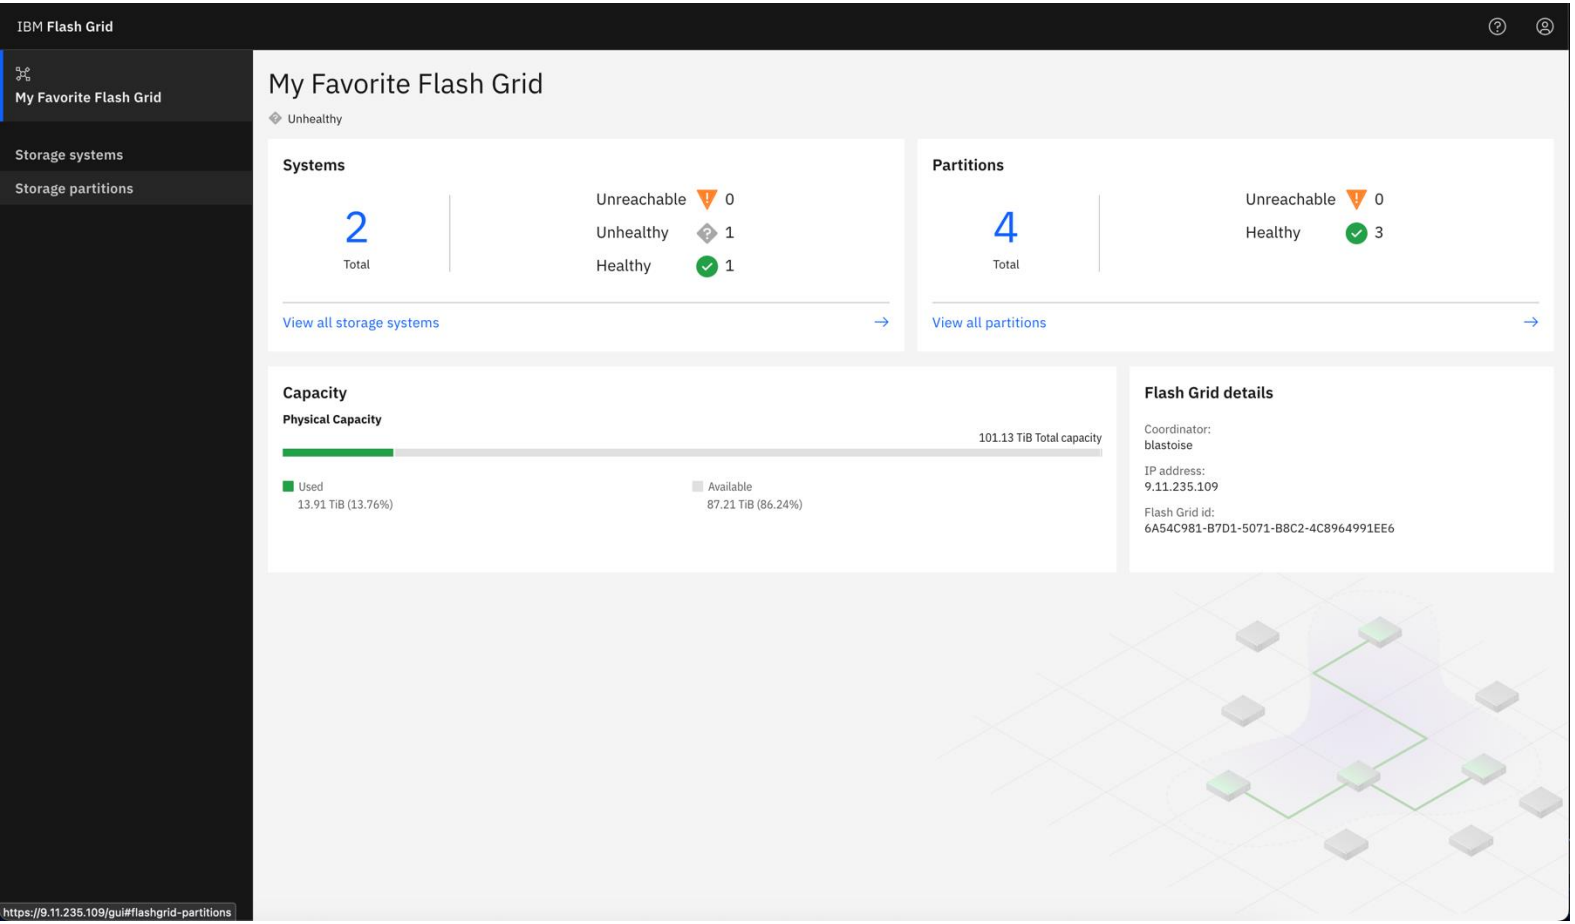

Number of systems and partitions

System / Partition Health

- Healthy, Unhealthy, Unreachable

Aggregated Capacity

- Physical capacity only in this release

Flash Grid details

Side menus for Storage systems and Storage partitions

Flash Grid Login Screen

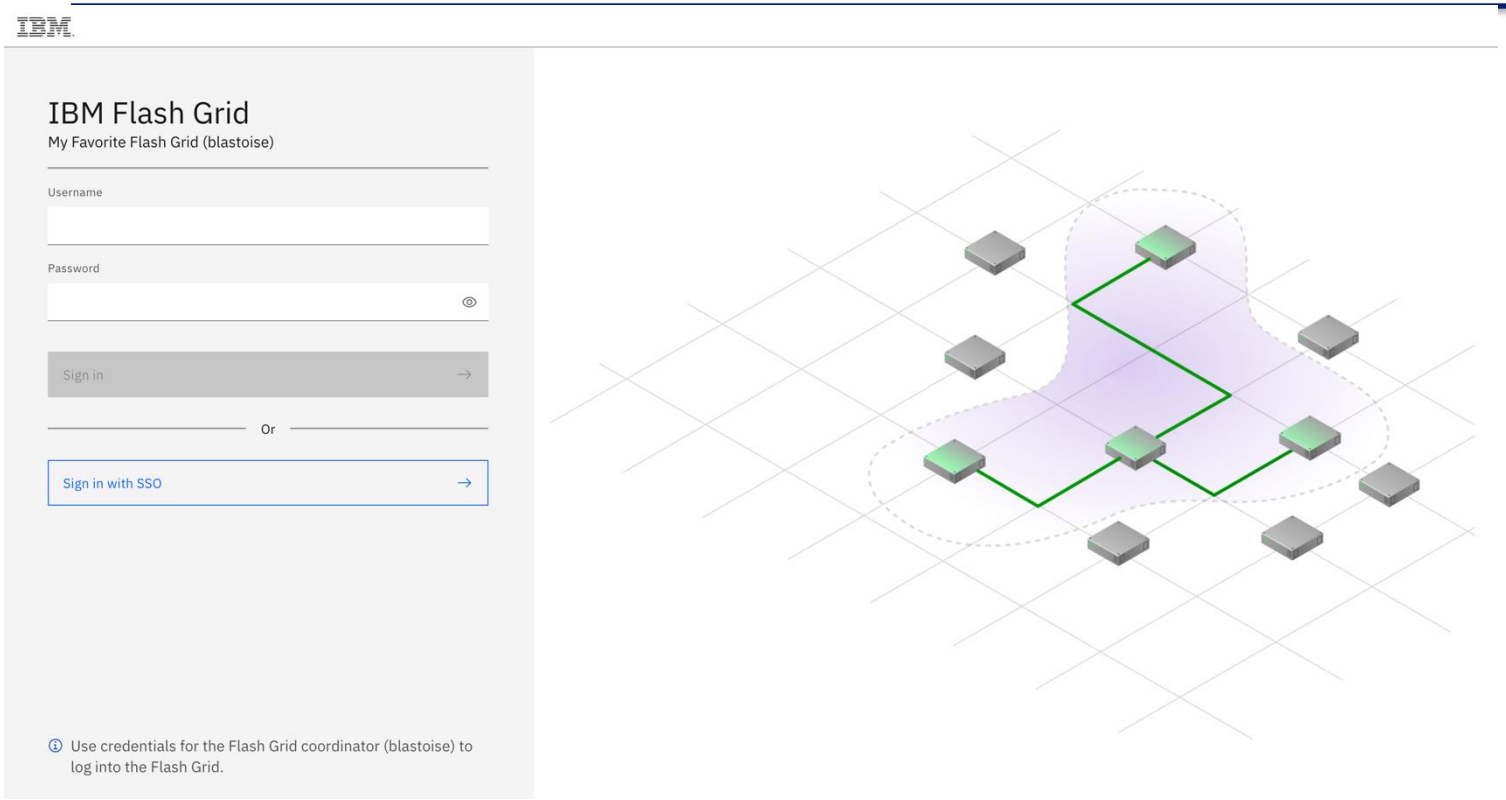

The Flash Grid Coordinator will have a new Flash Grid login screen

User credentials to log in are those of the system you connect to

- Supports LDAP, SSO

Flash Grid Login Screen

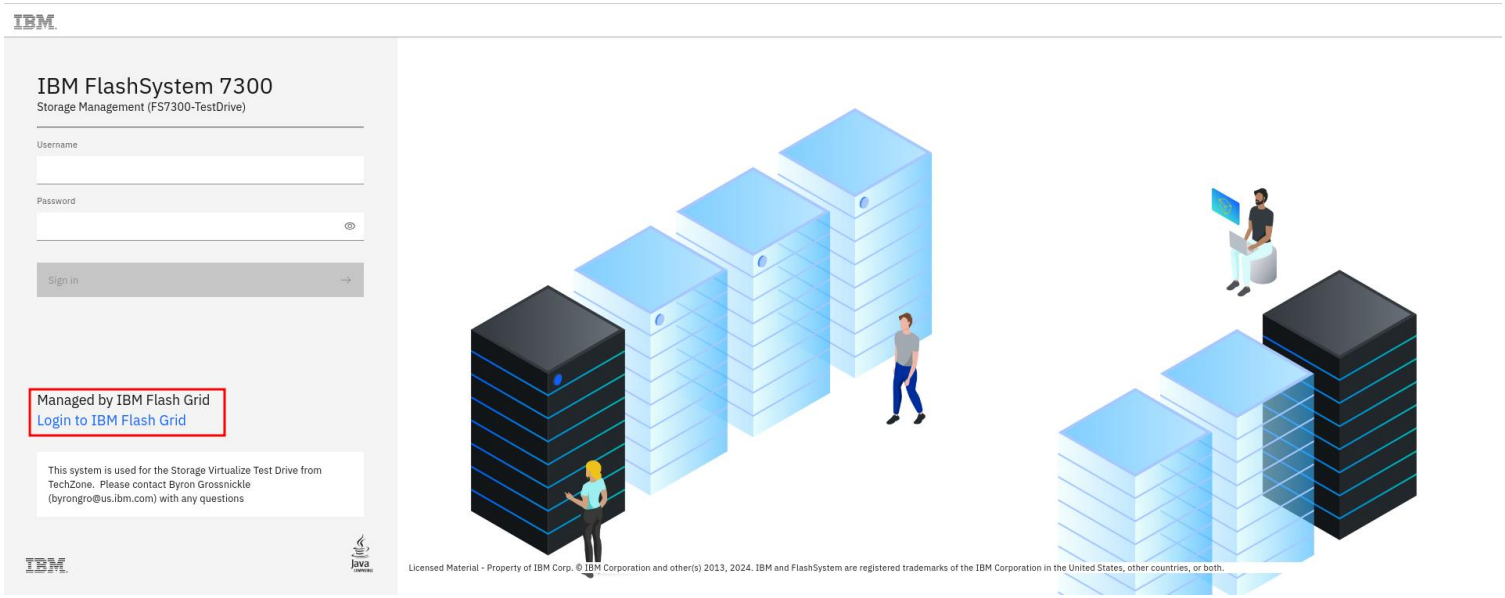

A member of the Flash Grid will have a normal login screen with the option to connect to the Flash Grid Coordinator that is running the Flash Grid GUI front end

Add a member to the Flash Grid

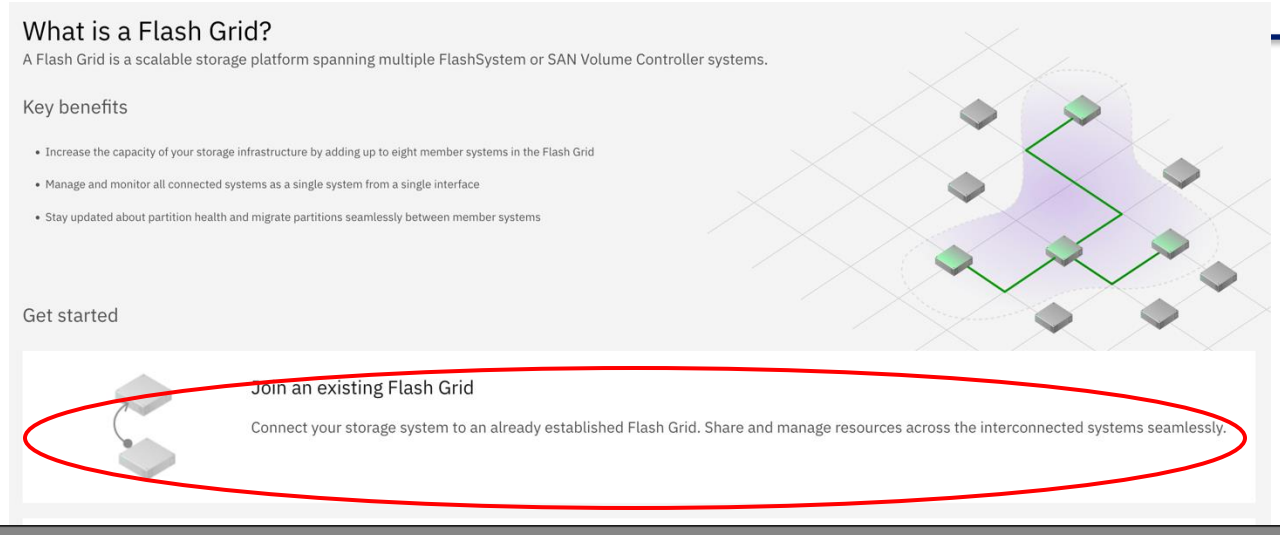

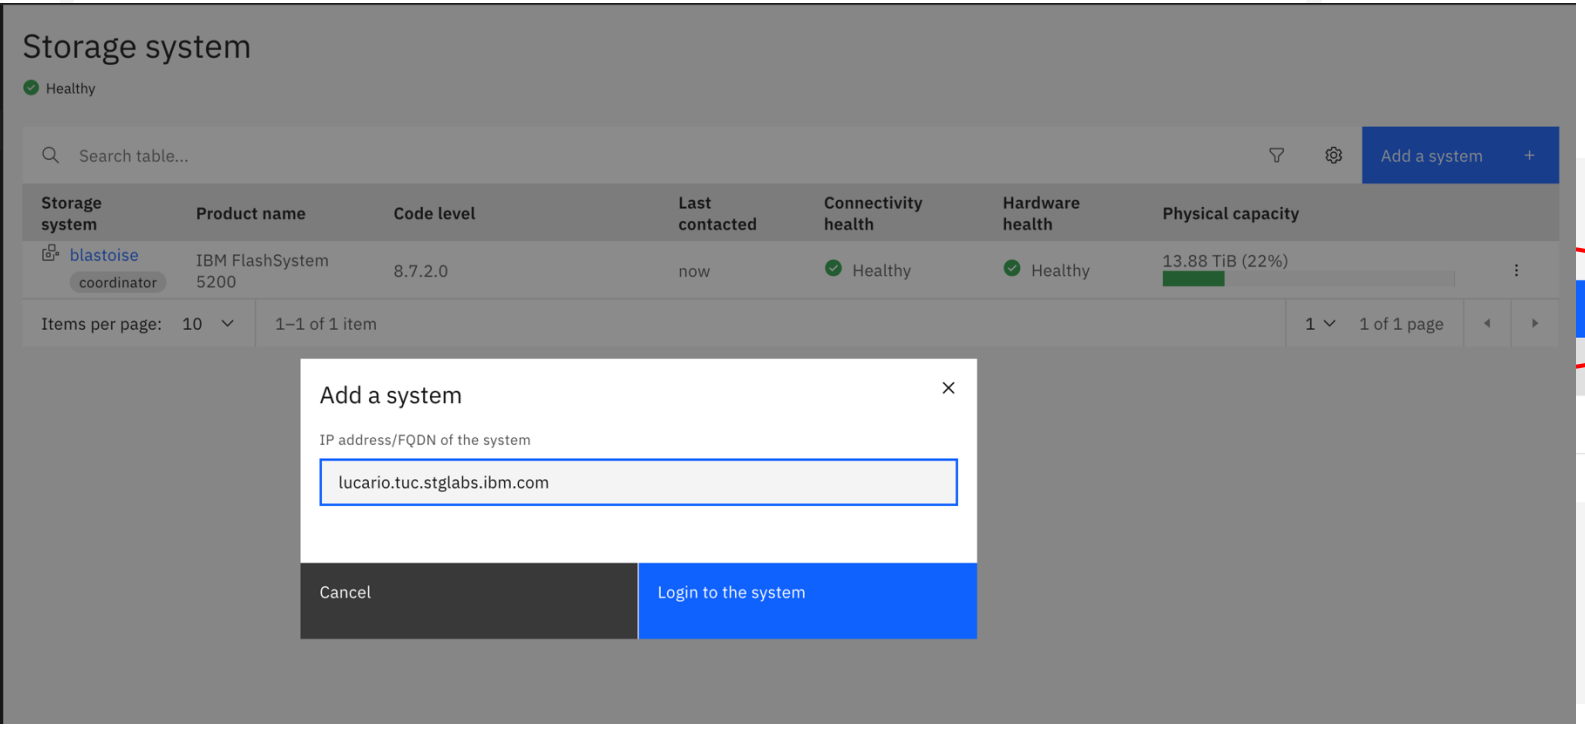

Add a member from the Storage systems page

OR

From the new member system's Flash Grid configuration screen

- Select Join an existing Flash Grid

Add a member to the Flash Grid – Cont'd

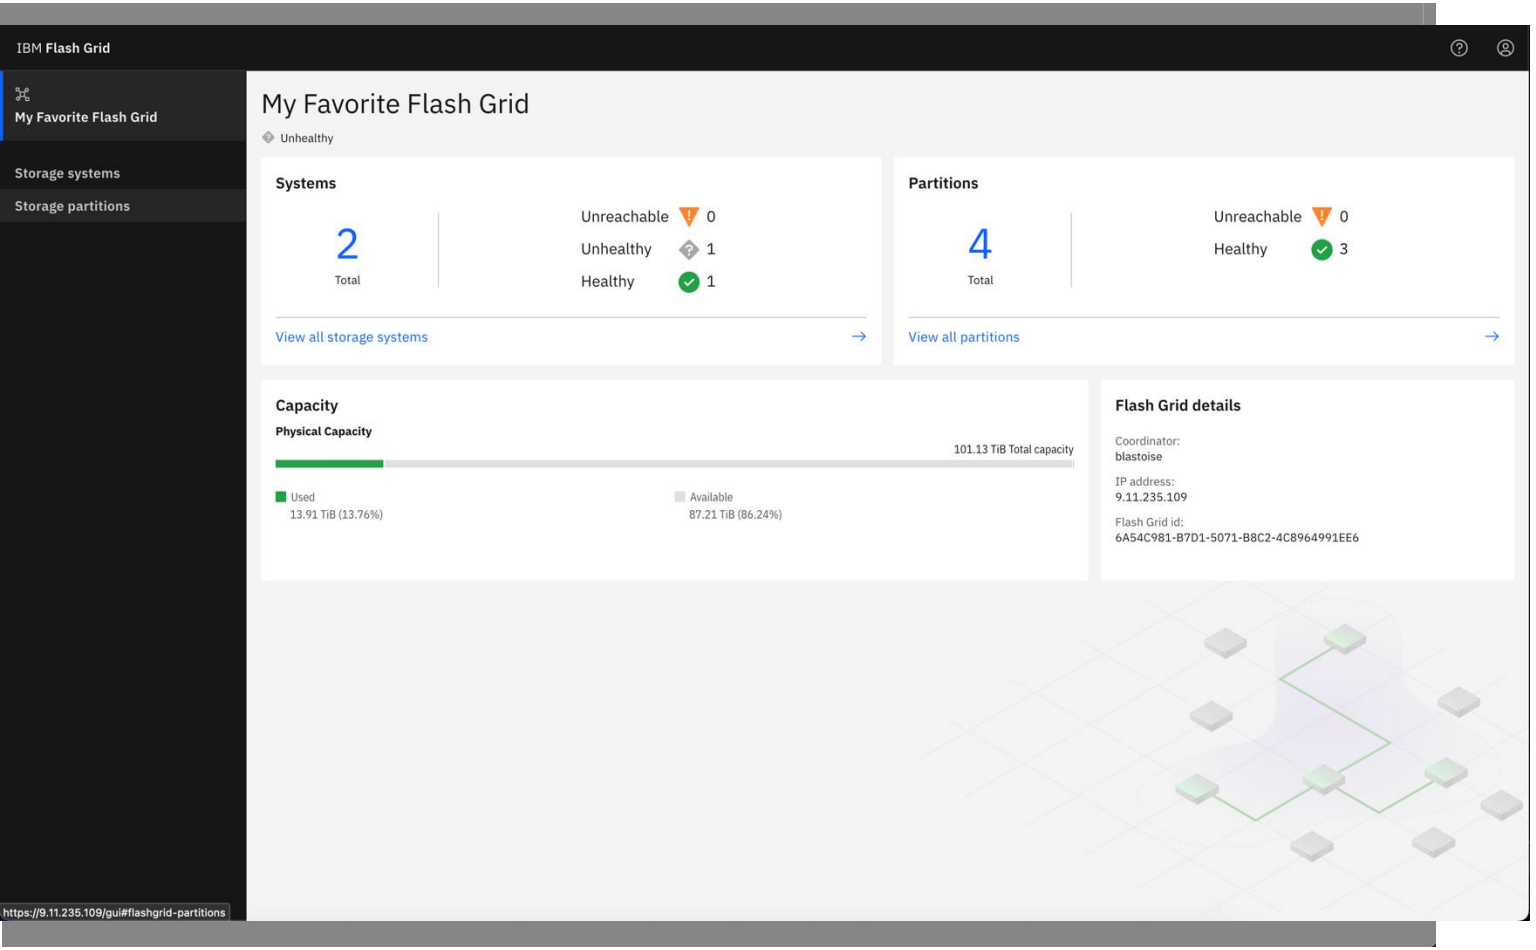

GUI automatically navigates you to the member system

- May require logging in

Verify and exchange certificates

GUI automatically navigates back to the Flash Grid coordinator to complete join process

Approve Flash Grid join request and new member is added

Dashboard now shows additional systems and new aggregate capacity

Flash Grid – Storage systems

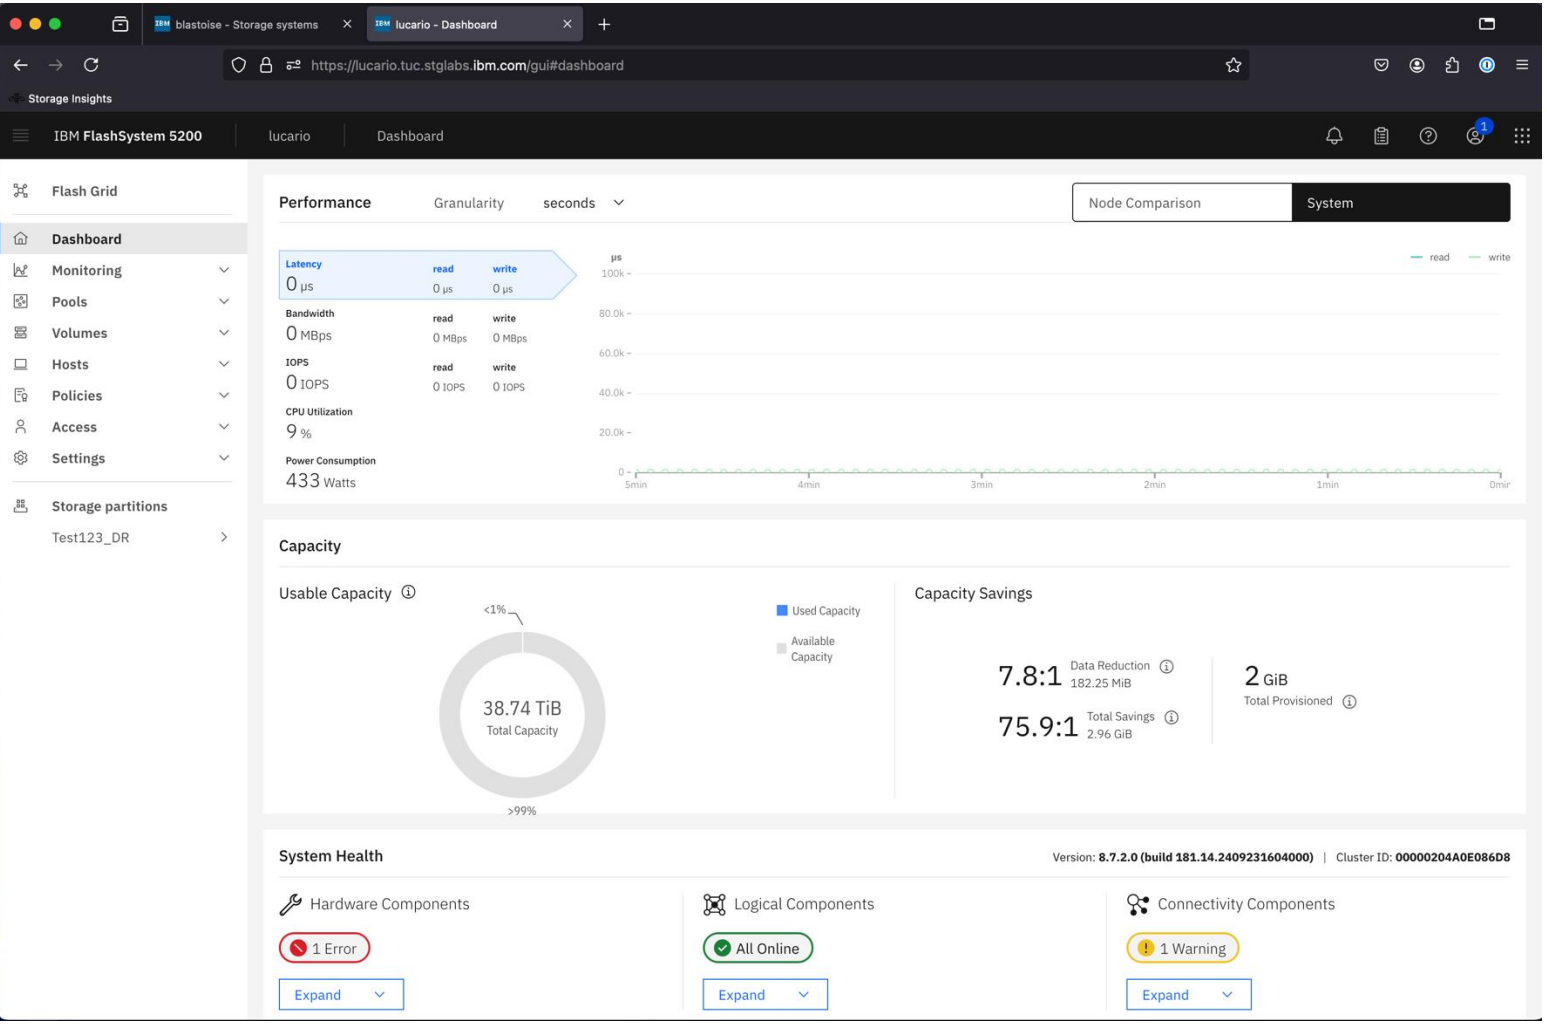

All storage systems in the Flash Grid are managed here

Product name, Code level details, Physical capacity

Last contacted shows last time the system was queried from the Flash Grid

Connectivity, Hardware health provides similar attributes to system health boxes

Clicking on storage system name will pop up system dashboard in a separate tab

Flash Grid – Storage partitions

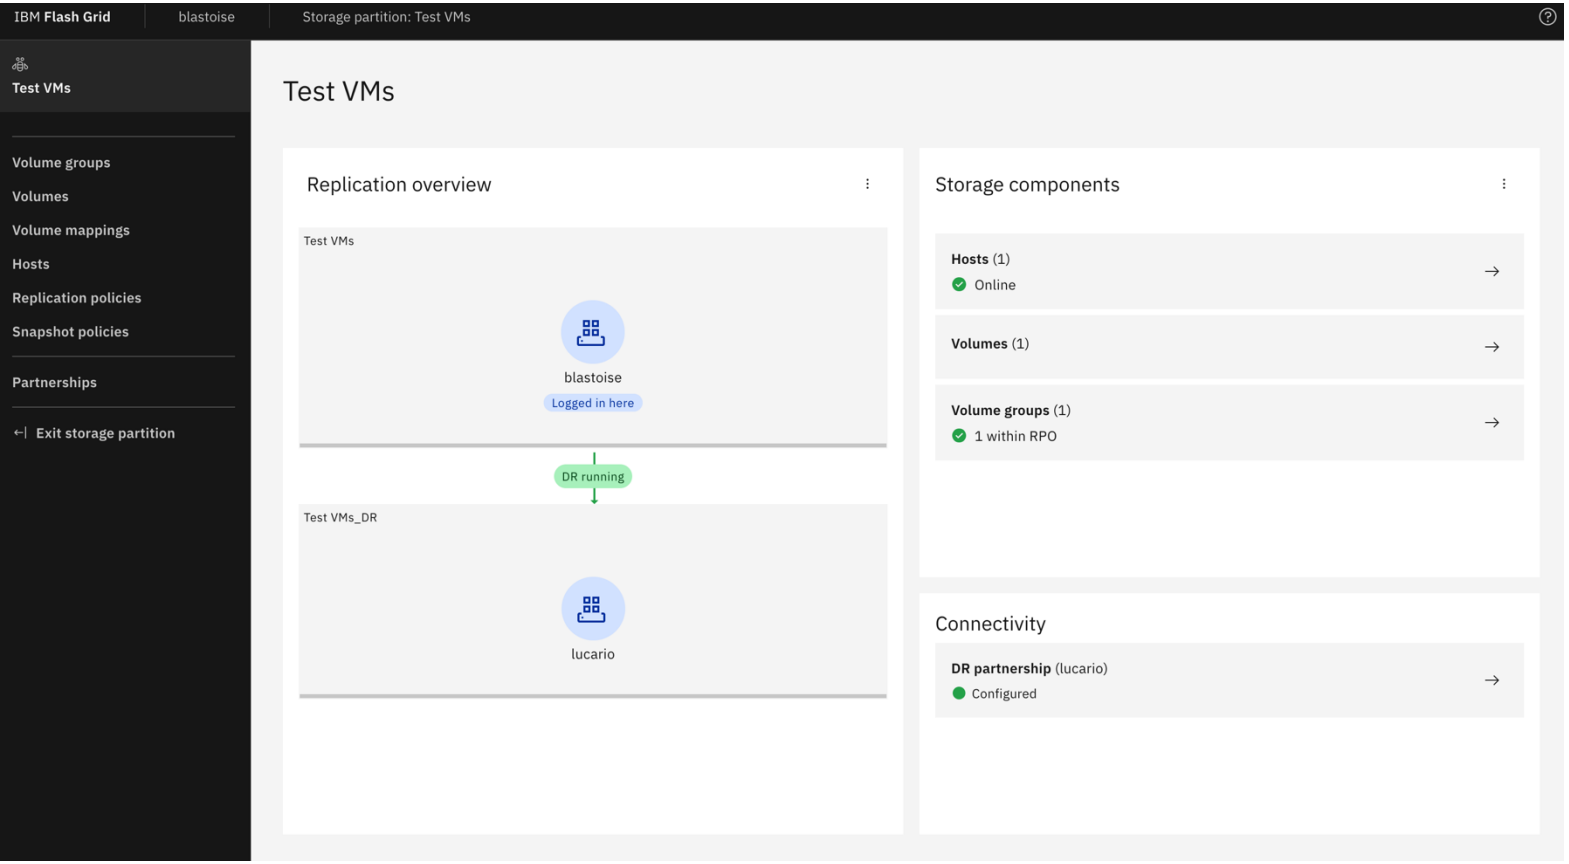

Aggregated list of all storage partitions in the Flash Grid

Provides replication and migration status and number of volume groups

Partition and Volume Group links

- Clicking on links will launch to the partition or volume group
  - Launches within the window
    - No additional tab
  - May require login
    - SSO configured systems will auto log in

Exiting storage partition brings you back to the Storage partitions list

Removing a member from a Flash Grid

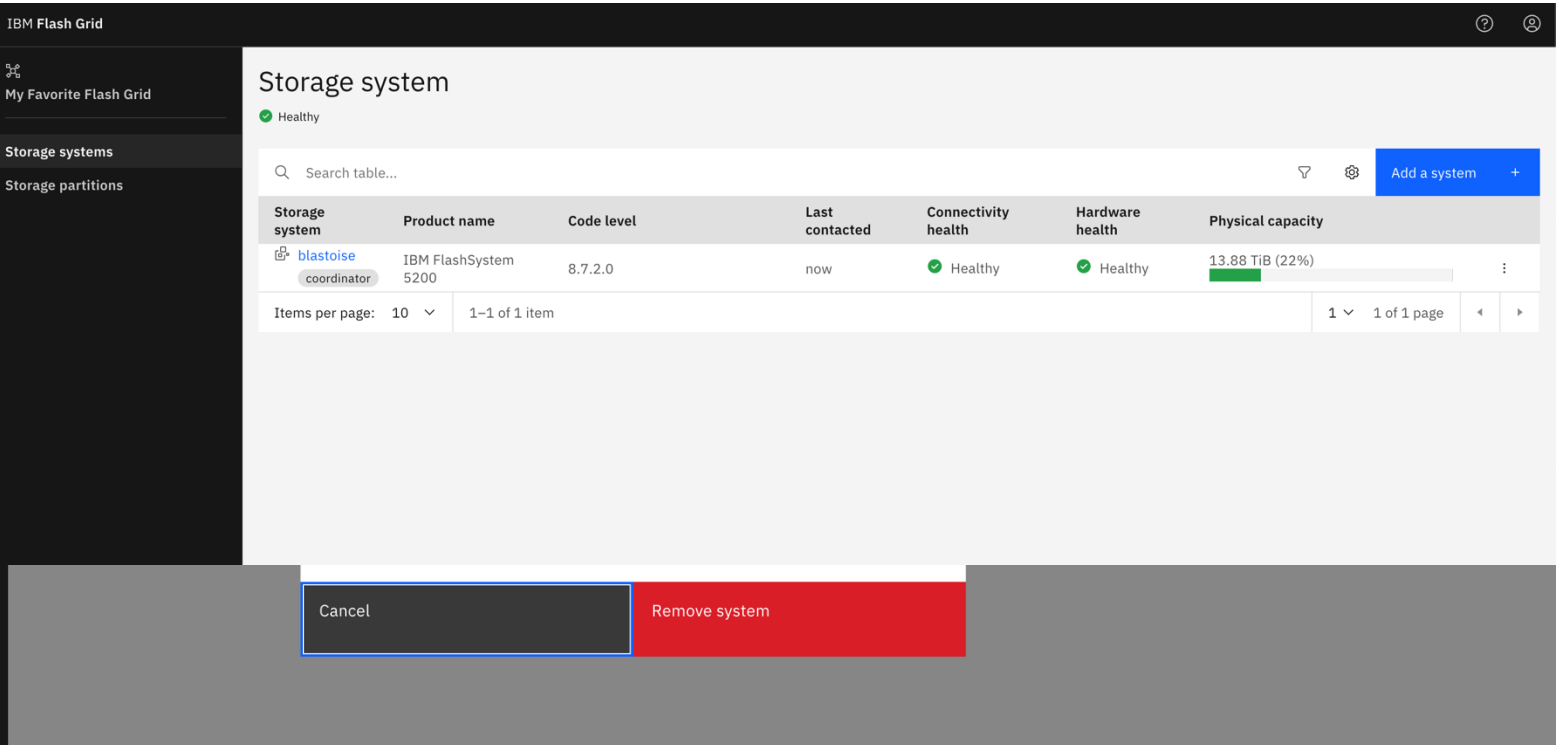

Accessible from Storage systems page

Select Remove system from the action menu of the system you want to remove

Confirm and system is removed

Removing the LAST system (coordinator) from the Flash Grid will dissolve the Flash Grid entirely

Storage Partition : New Unit for Managing Workload Storage

– New unit of management

– Consists of multiple volume groups, hosts

– Typically, assigned to one workload

– Currently limited to four partitions per system

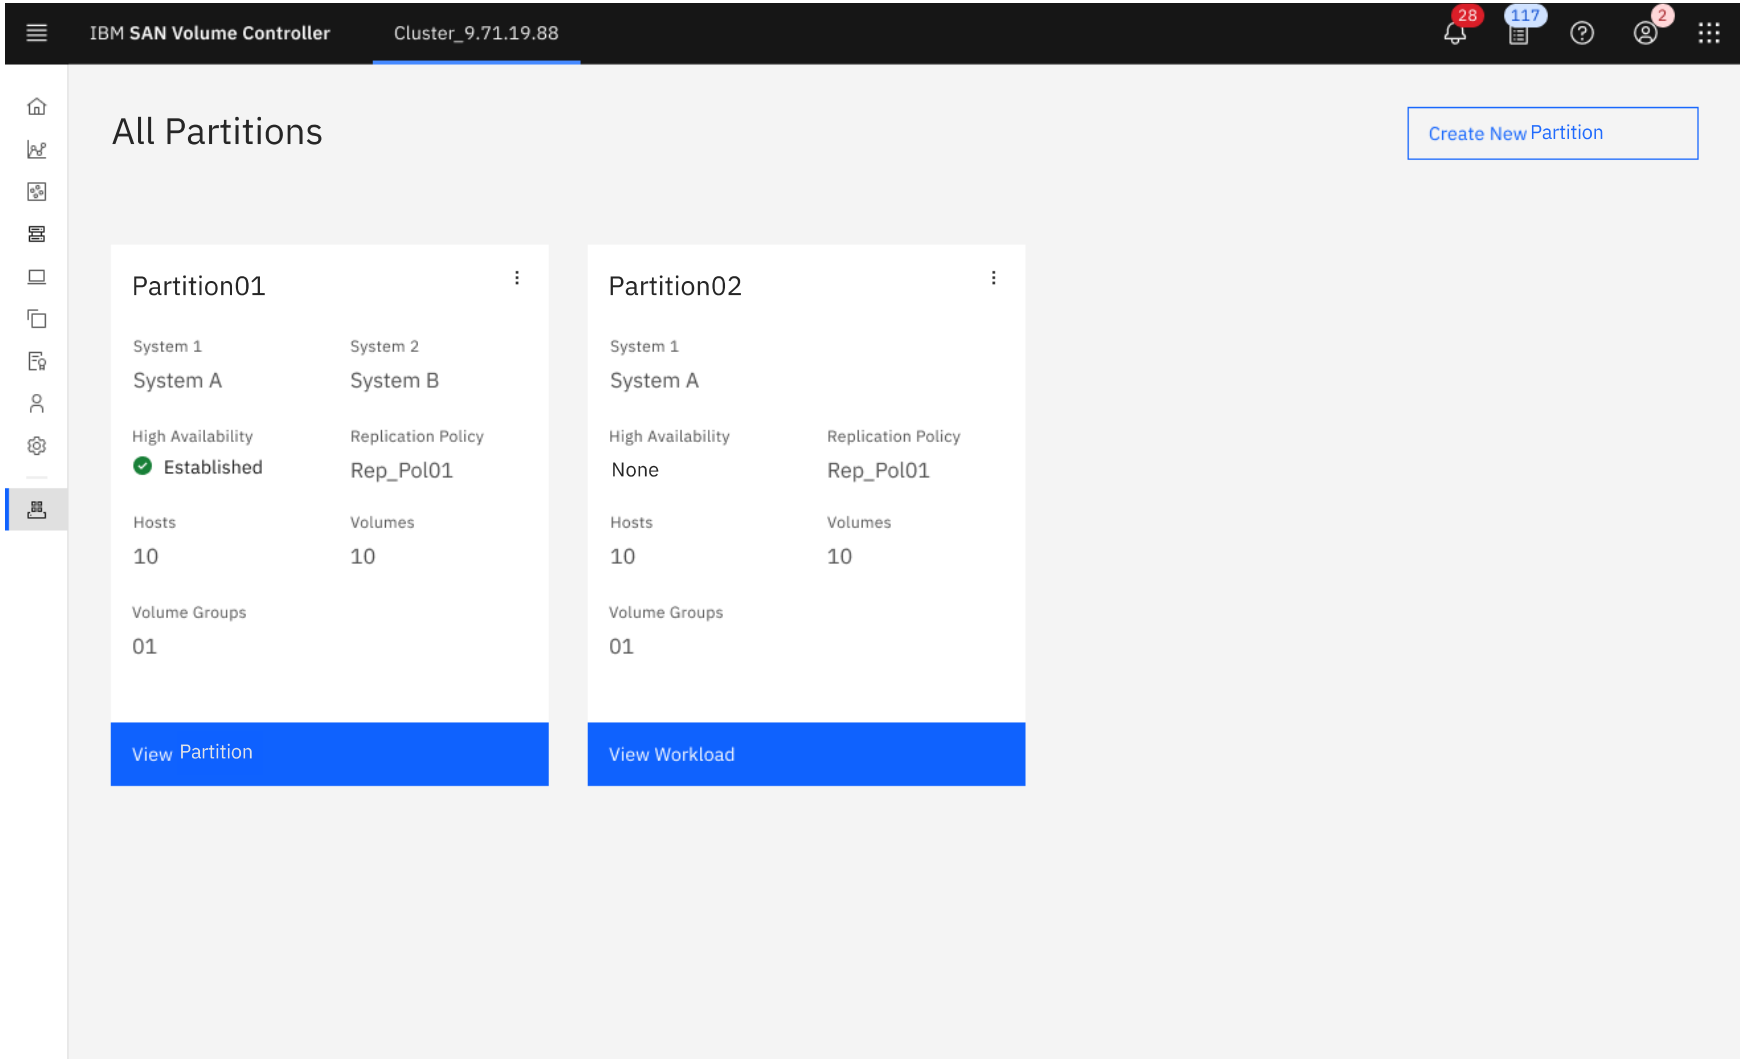

Storage Partition Migration:

Mobility Drivers

1. **Migration Use case:** Storage Partition Movement
2. **Flash Grid Optimization:** Load-Balancing across Fleet of Storage
3. **Hardware Refresh:** Non-disruptive System Retirals

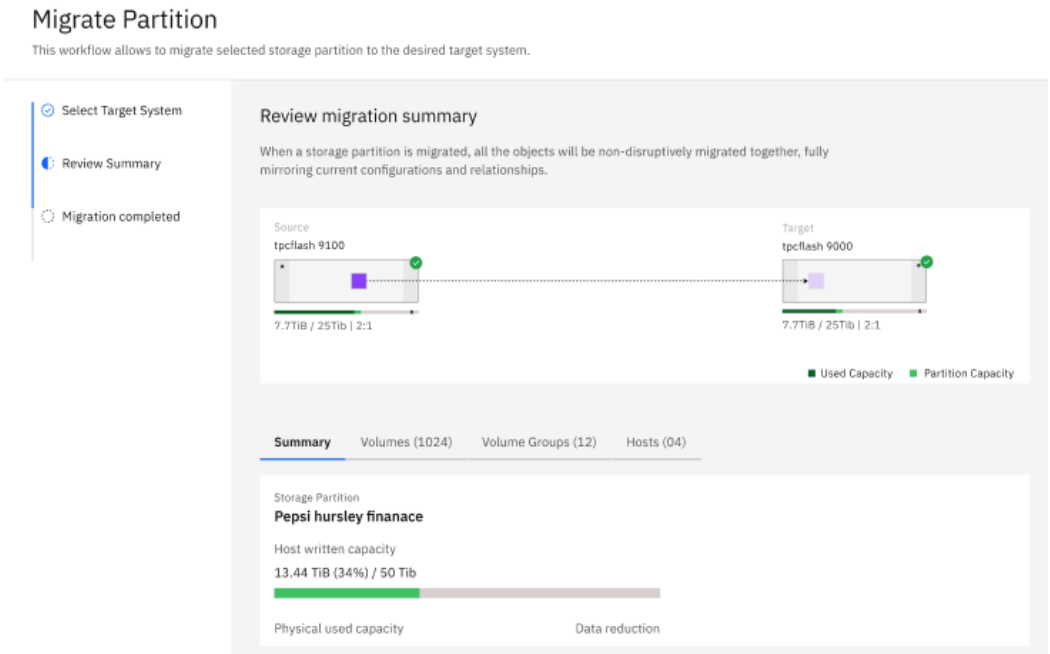

Solution and Roadmap

1. Single-click Automation for Migration (24Q1)
2. Zero downtime (Non-disruptive) (24Q1)
3. AI driven automated migration through Storage Insights (24Q2)
4. Integration with VMWare Host environments (24Q4)
5. Automated Network, Host actions (8.7.2/Future)

User Interaction during Storage Partition Migration

Before 8.7.2	In 8.7.2
User initiates storage partition migration	User initiates storage partition migration
User Action#1 Request Host Rescan - Scan paths at host level and continue	**User Action#1 Request Host Rescan - Scan paths at host level and continue (not required for VMWare environments)**
User Action#2 Copy Delete Consent - Delete source copy and complete the migration	**User Action#2 Copy Delete Consent - Delete source copy and complete the migration**

- Ease of use / Minimum intervention
  - Providing a way to detect environment that has automated rescan support
  - Providing a way for user to configure if given host has periodic automated rescan capability
    - CLI Only: chhost –autostoragediscovery yes

– Note: FDMI data is used to set this value for hosts known to auto-rescan

IBM Storage Virtualize

Misc. Topics

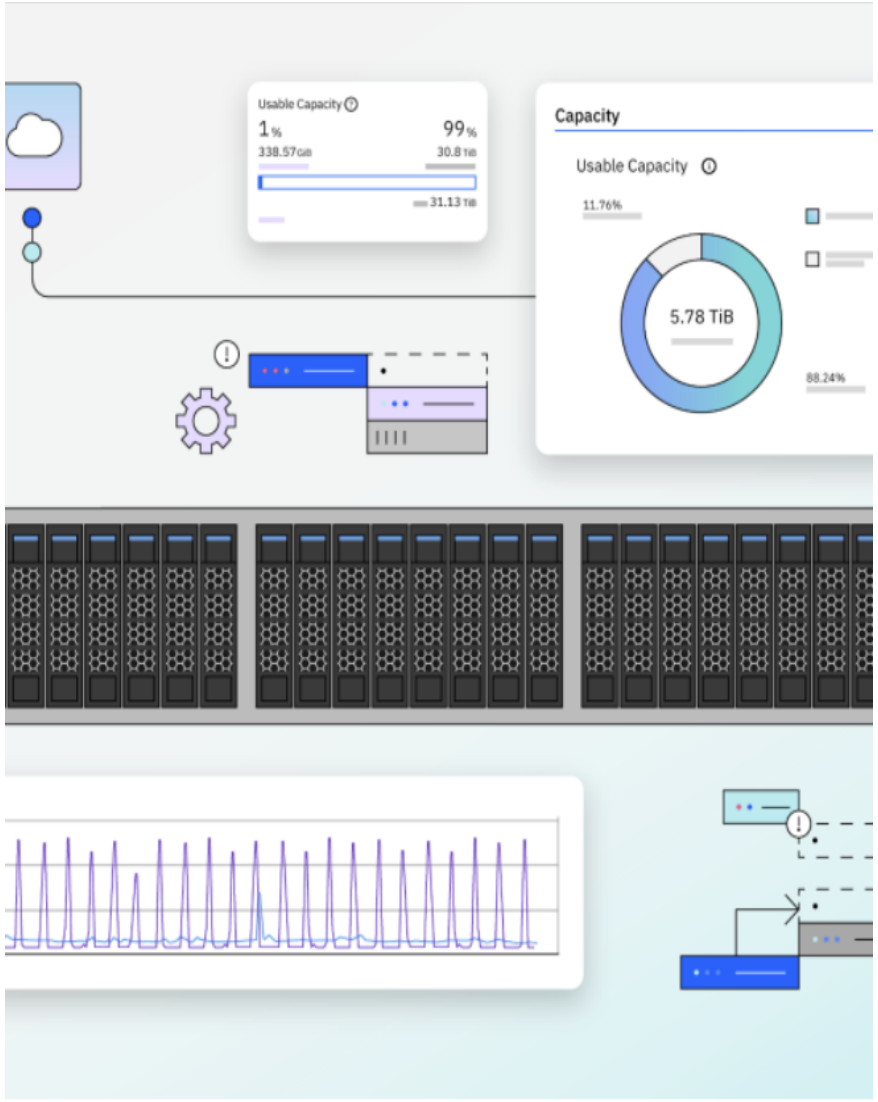

Service IP Updates

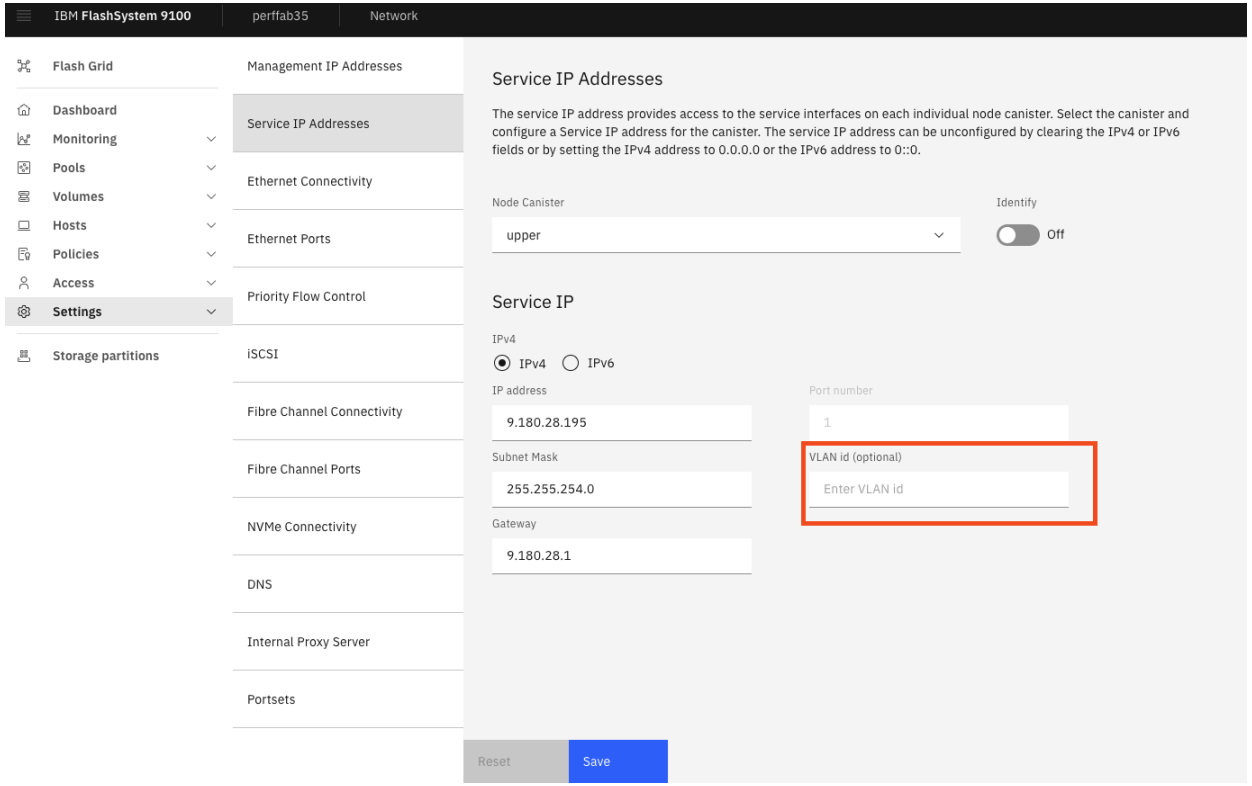

- To align with the overall improvements of IP Addresses, we have added VLAN tagging support to Service IPs that are managed from the service GUI/CLI or the System GUI/CLI

- Note – we will only support either an IPv4 OR an IPv6 address for the service IP, not both

Additional IP Replication Statistics

Stats	Explanation	SI Metric Counter	SI Metric Name
ipbz	Indicates the average size (in bytes) of data that is being submitted to the IP partnership driver since the last statistics collection period.	1264	IP Replication Transfer Size (Receive)
iprc	Indicates the total bytes that are received before any decompression takes place.	1280	IP Replication Compressed Data Rate (Recieve)
ipre	Indicates the bytes retransmitted to other nodes in other clusters by the IP partnership driver.	1278	IP Replication-to-Remote Node Data Rate (Resend)
iprt	Indicates the average round-trip time in microseconds for the IP partnership link since the last statistics collection period.	1265	IP Replication Latency
iprx	Indicates the bytes received from other nodes in other clusters by the IP partnership driver.	1277	IP Replication-to-Remote Node Data Rate (Receive)
ipsz	Indicates the average size (in bytes) of data that is being transmitted by the IP partnership driver since the last statistics collection period.	1263	IP Replication Transfer Size (Send)
iptc	Indicates the total bytes that are transmitted after any compression (if active) takes place.	1279	IP Replication Compressed Data Rate (Send)
iptx	Indicates the bytes transmitted to other nodes in other clusters by the IP partnership driver.	1276	IP Replication-to-Remote Node Data Rate (Send)

- Provide XML IO Statistics for individual IP Partnerships for multi-site configurations
- These are not in SI right now, but available in the XML statistics

```
<partnership_ipperf index="1" cluster="auxnear" remote_cluster_id="0x0000020439602218" port_id="4"
iptx="680288517951" iprx="12495722729" ipre="10366232" ipsz="1210"
ipbz="1147" iprt="46" iptc="673275635991" iprc="6473577404"/>
```

Quotaless Child pools

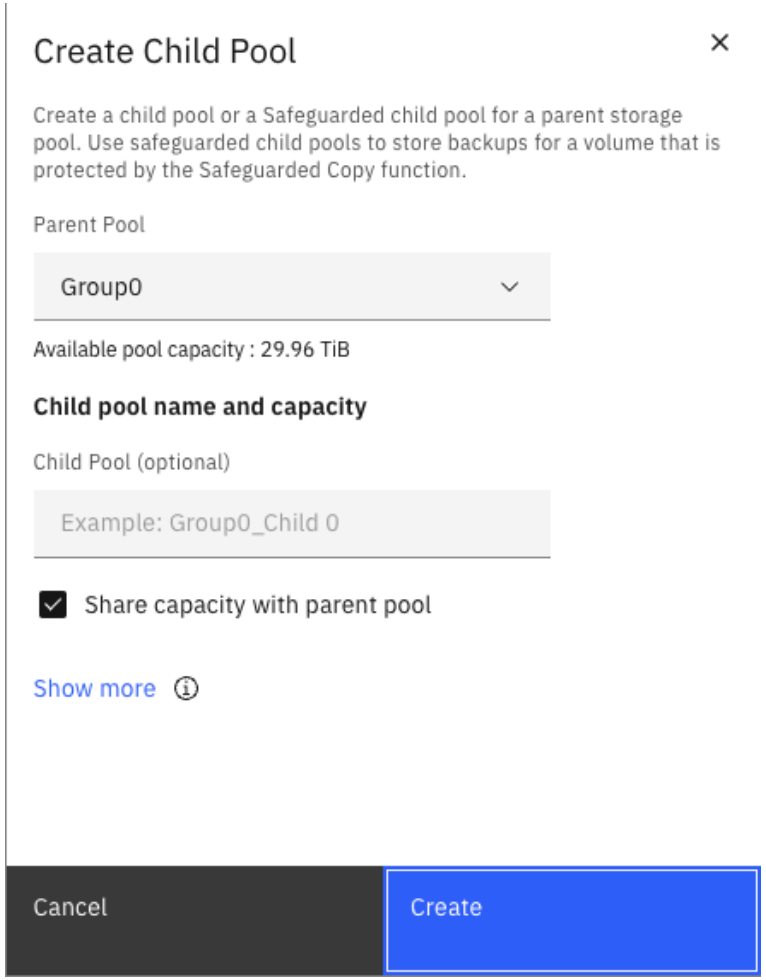

- Provide a new mechanism for creating a Standard Child Pools without a capacity quota for easier maintenance, Object Based Access Controls and IO throttling
- `mkmdiskgrp –noquota –parentmdiskgrp {id} –datareduction no`
- VMware vVols still require a quota-controlled standard pool(s)

Host Offline Alert Changes

8.6.1 introduced a new alert when a host changes status to `offline`

Feedback from customers was that sometimes hosts going offline was expected.

- New option to ignore certain hosts when they go offline:
  - CLI/API option on `chhost`
  - DMP-driven approach when resolving the host-offline event to select certain hosts to ignore status changes.

Updates to Config Backup

- Modernized tools in order to:
  - use less memory
  - use JSON instead of XML
  - Significantly improves the performance of config backup, T3 and T4 prepare steps – down to **20s** from 3 hours
  - T4 No longer restores
    - Cloud accounts,
    - Host iSCSI Auth,
    - Cloud Backup
  - **No longer in /tmp**

/dumps/svcconfig/svc.config.backup.json

Capacity reporting CLIs

Terminology changes for capacity

- Physical Capacity (Usable)
- Logical Capacity (Effective)

Note: Logical capacity metrics are only available for systems and pools which are backed by FCM-only arrays

- Provides consistency across GUI, SI and external APIs
  - Prevents each piece of software doing complex math separately for capacity metrics

lssystemcapacity

```
total_physical 38.74TB
used 12.45GB
warning_exceeded no
total_logical 178.62TB
allocated 602.00GB
written 482.44GB
written_snapshots 0.75MB
provisioned_host 600.00GB
reserved_host_volume 117.56GB
reserved_contingency 2.00GB
reserved_childpools 0.00MB
```

lspoolcapacity

```
id 0
name swimming_pool
total_physical 38.74TB
used 12.45GB
warning 8000
warning_set no
total_logical 178.62TB
allocated 602.00GB
written 482.44GB
written_snapshots 0.75MB
provisioned_host 600.00GB
reserved_host_volume 117.56GB
reserved_contingency 2.00GB
reserved_childpools 0.00MB
```

New Login Panel

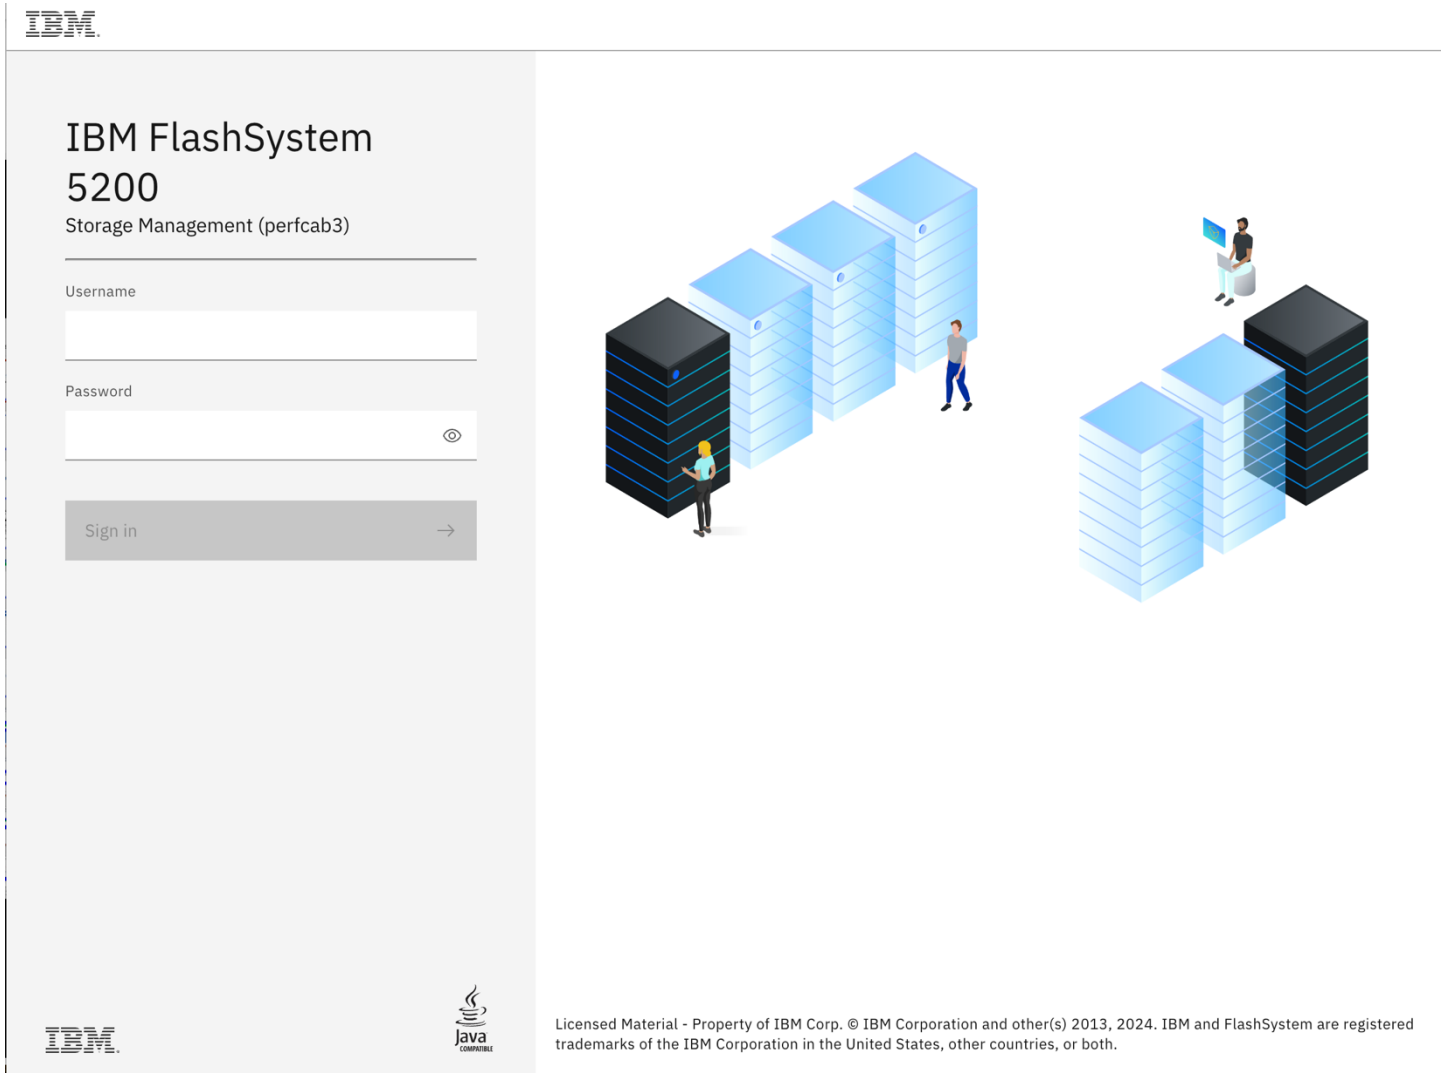

IBM Storage Virtualize

Plugin Updates

vSphere Plugin
CSI Driver

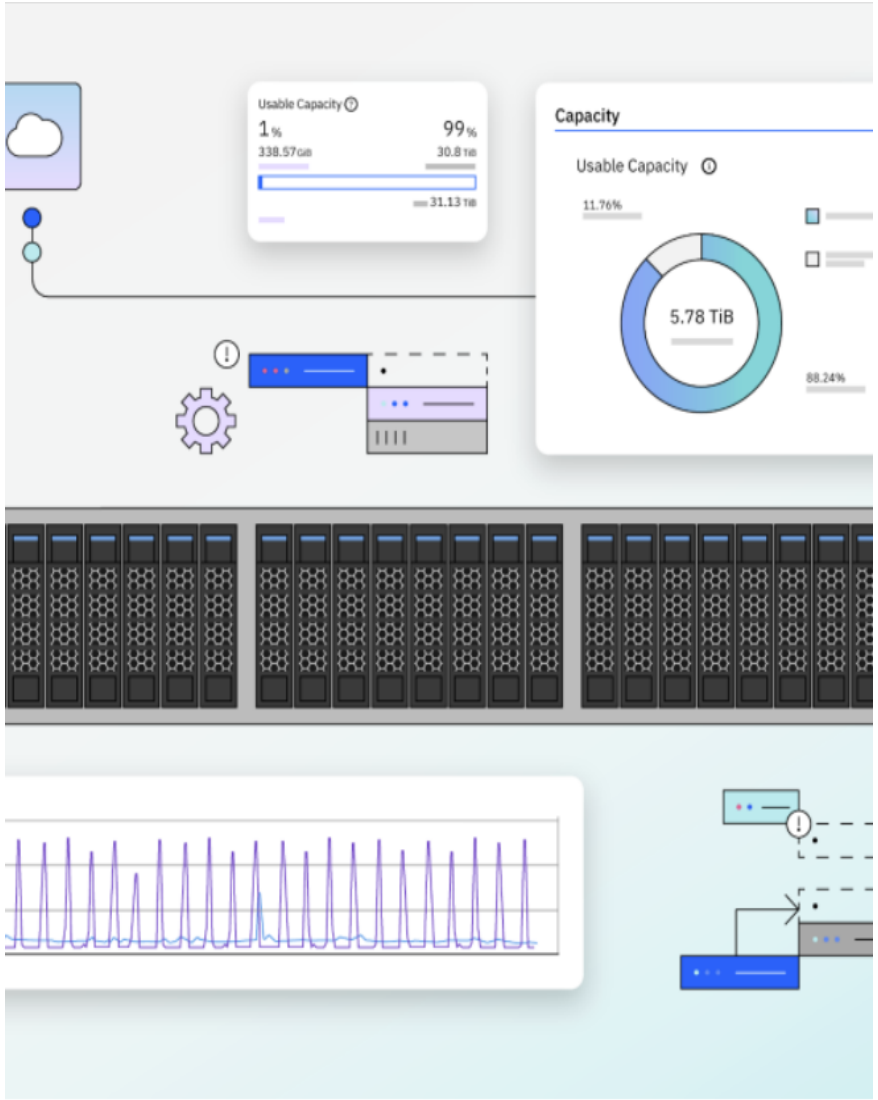

IBM Storage Virtualize

Plugin Updates

vSphere 2.0.0
December

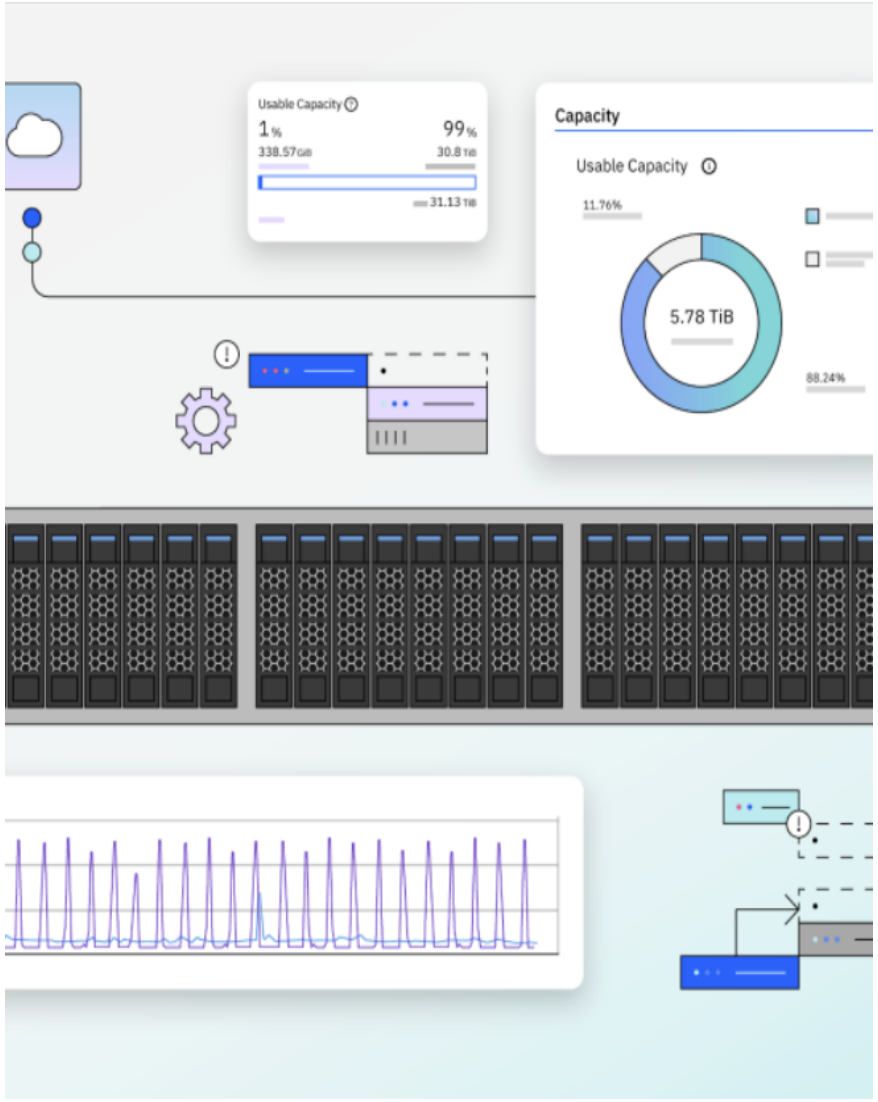

IBM Storage Virtualize plugin for vSphere

Automates vSphere and storage tasks together All from the vCenter UI

- Supports all Storage Virtualize platforms:
  FlashSystem (8.5.0+), SAN Volume Controller and Storwize products)
- Manage multiple storage systems from a single plugin instance
- Simplifies Storage Provisioning & Management
- Provides storage context for vSphere objects
- OVA can be Download from IBM Fix Central
- Backed by crisp documentation
  https://www.ibm.com/docs/en/svpfv

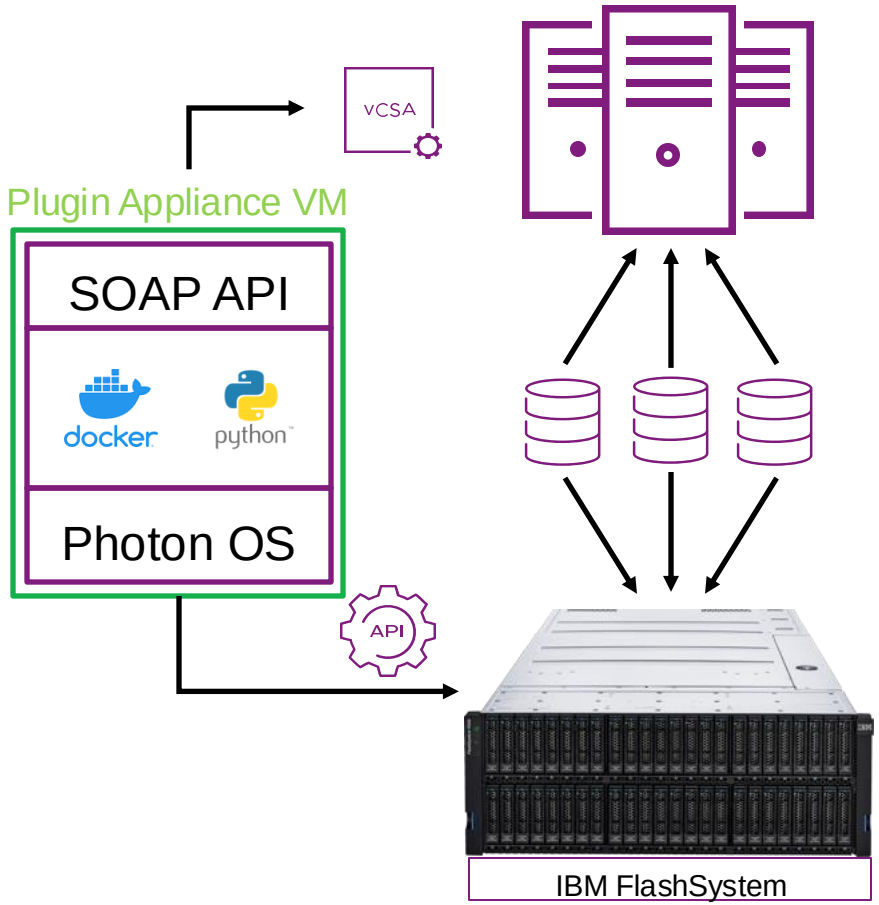

Current vSphere Plugin Capabilities

- Manage FC host
- Manage FC host cluster
- vSphere privileges check
- Storage system details
- Manual reconciliation

- Manage VMFS datastores
  - Create datastores
  - Expand datastore
  - Delete datastore

- Manage datastores volume group
  - View volume group details
  - Move datastore to a volume group
  - Remove datastore from a volume group

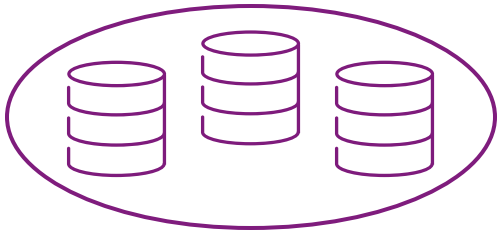

- Manage datastores snapshot
  - Take snapshot
  - View snapshots
  - Copy to new datastore
  - Delete snapshot

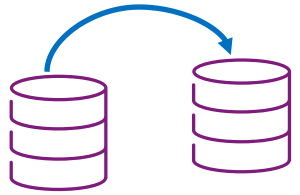

What's New in vSphere Plugin 2.0 Overview

Performance Monitoring Dashboard

- Integration with Storage Insights
- Provision to connect SI through proxy server(For dark site users)
- List top 5 storage systems by IOPS, Bandwidth or latency
- List top 5 datastores by IOPS, Bandwidth or latency
- View recent CPU utilization of each registered storage system
- Doughnut chart to categorize storage systems based on version, platform, call-home enabled, and status
- Aggregate view of number of datastores and VMs with focus on publishing status with error if any
- View system(storage system) level statistics with respect to IOPs, bandwidth and latency
- View volume(datastore) level statistics with respect to IOPs, bandwidth and latency

Plugin settings view

Storage Insights Integration

- From vSphere Plugin dashboard summary, add SI instance by providing SI tenant ID, API Key and API Key alias.
- This result into availability of performance statistics automatically at system and datastore level for those storage system which are part of this SI inventory.
- This also enables the ransomware threat detection alert feature right in the vSphere UI

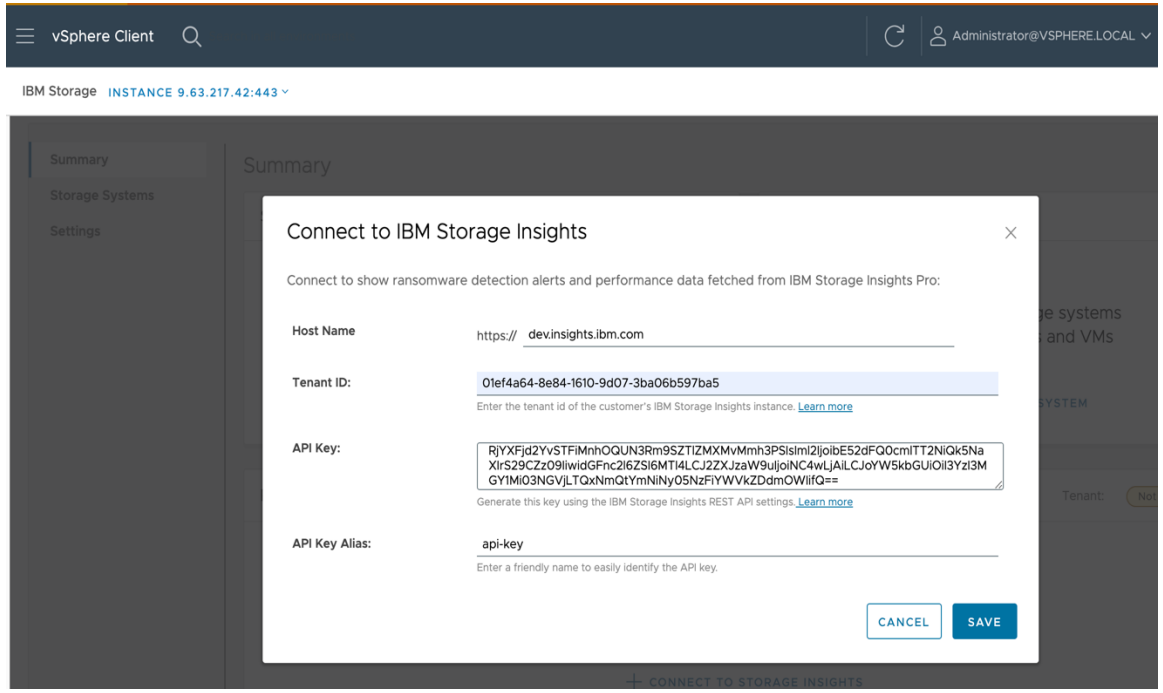

Add Storage Insights

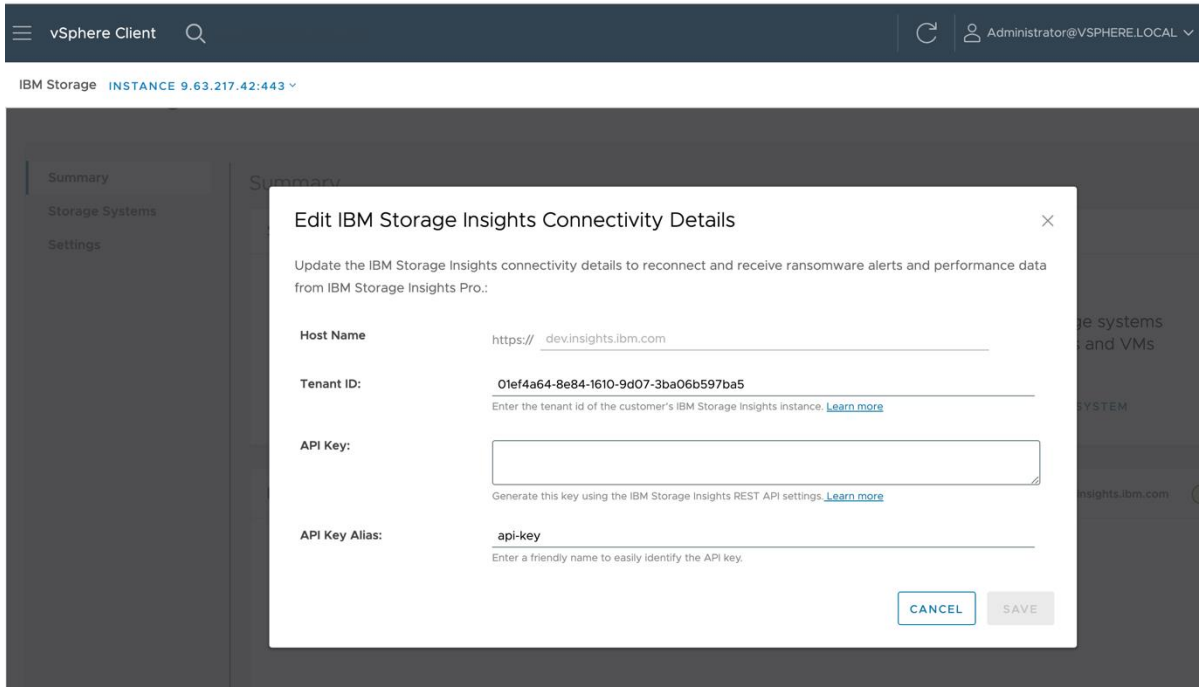

Edit Storage Insights

Proxy Server and plugin settings

- For internet connectivity purpose if required, from settings page, proxy server can be added by providing URL, port, username and password based on proxy type selection

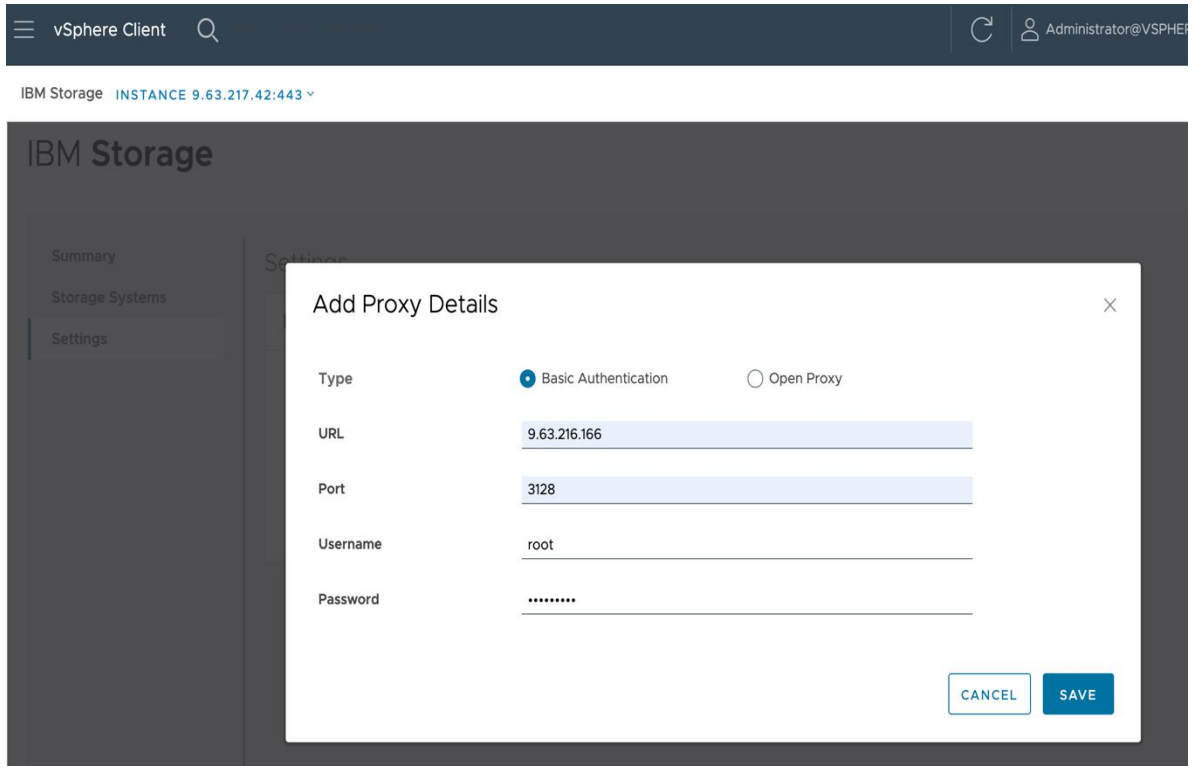

Add Proxy server

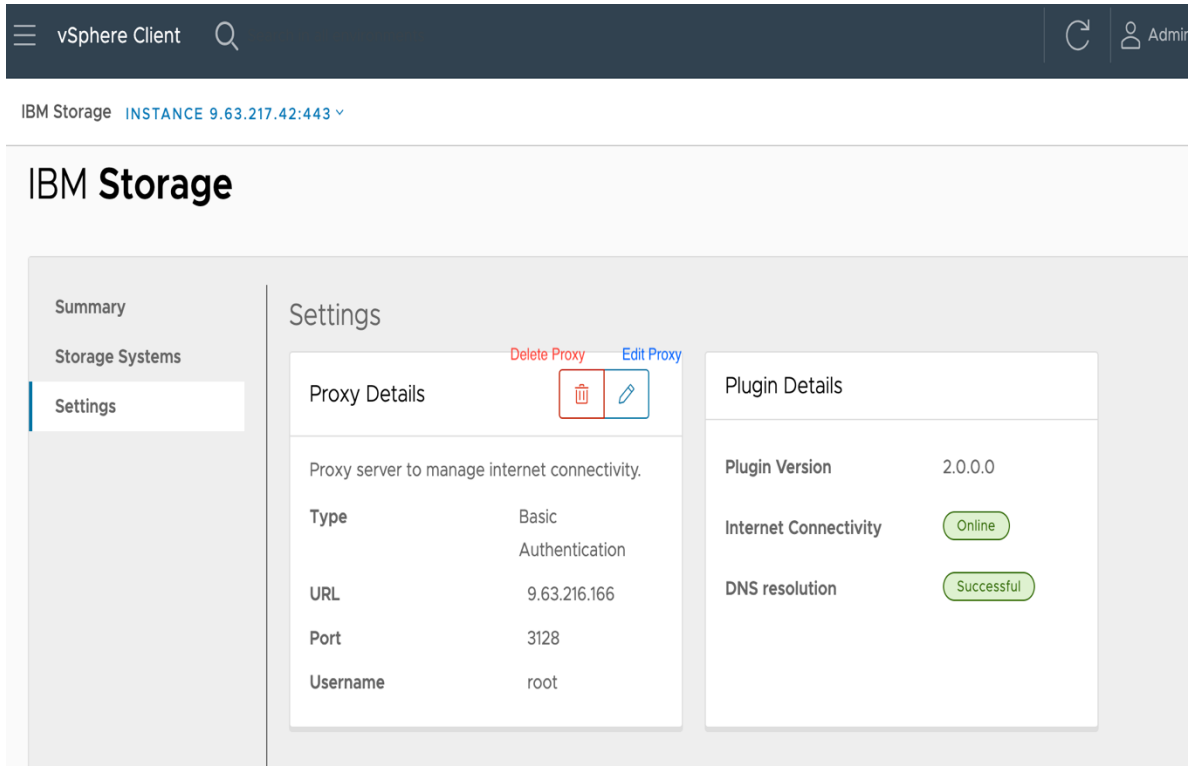

Edit/Delete Proxy server

Sample performance monitoring dashboard

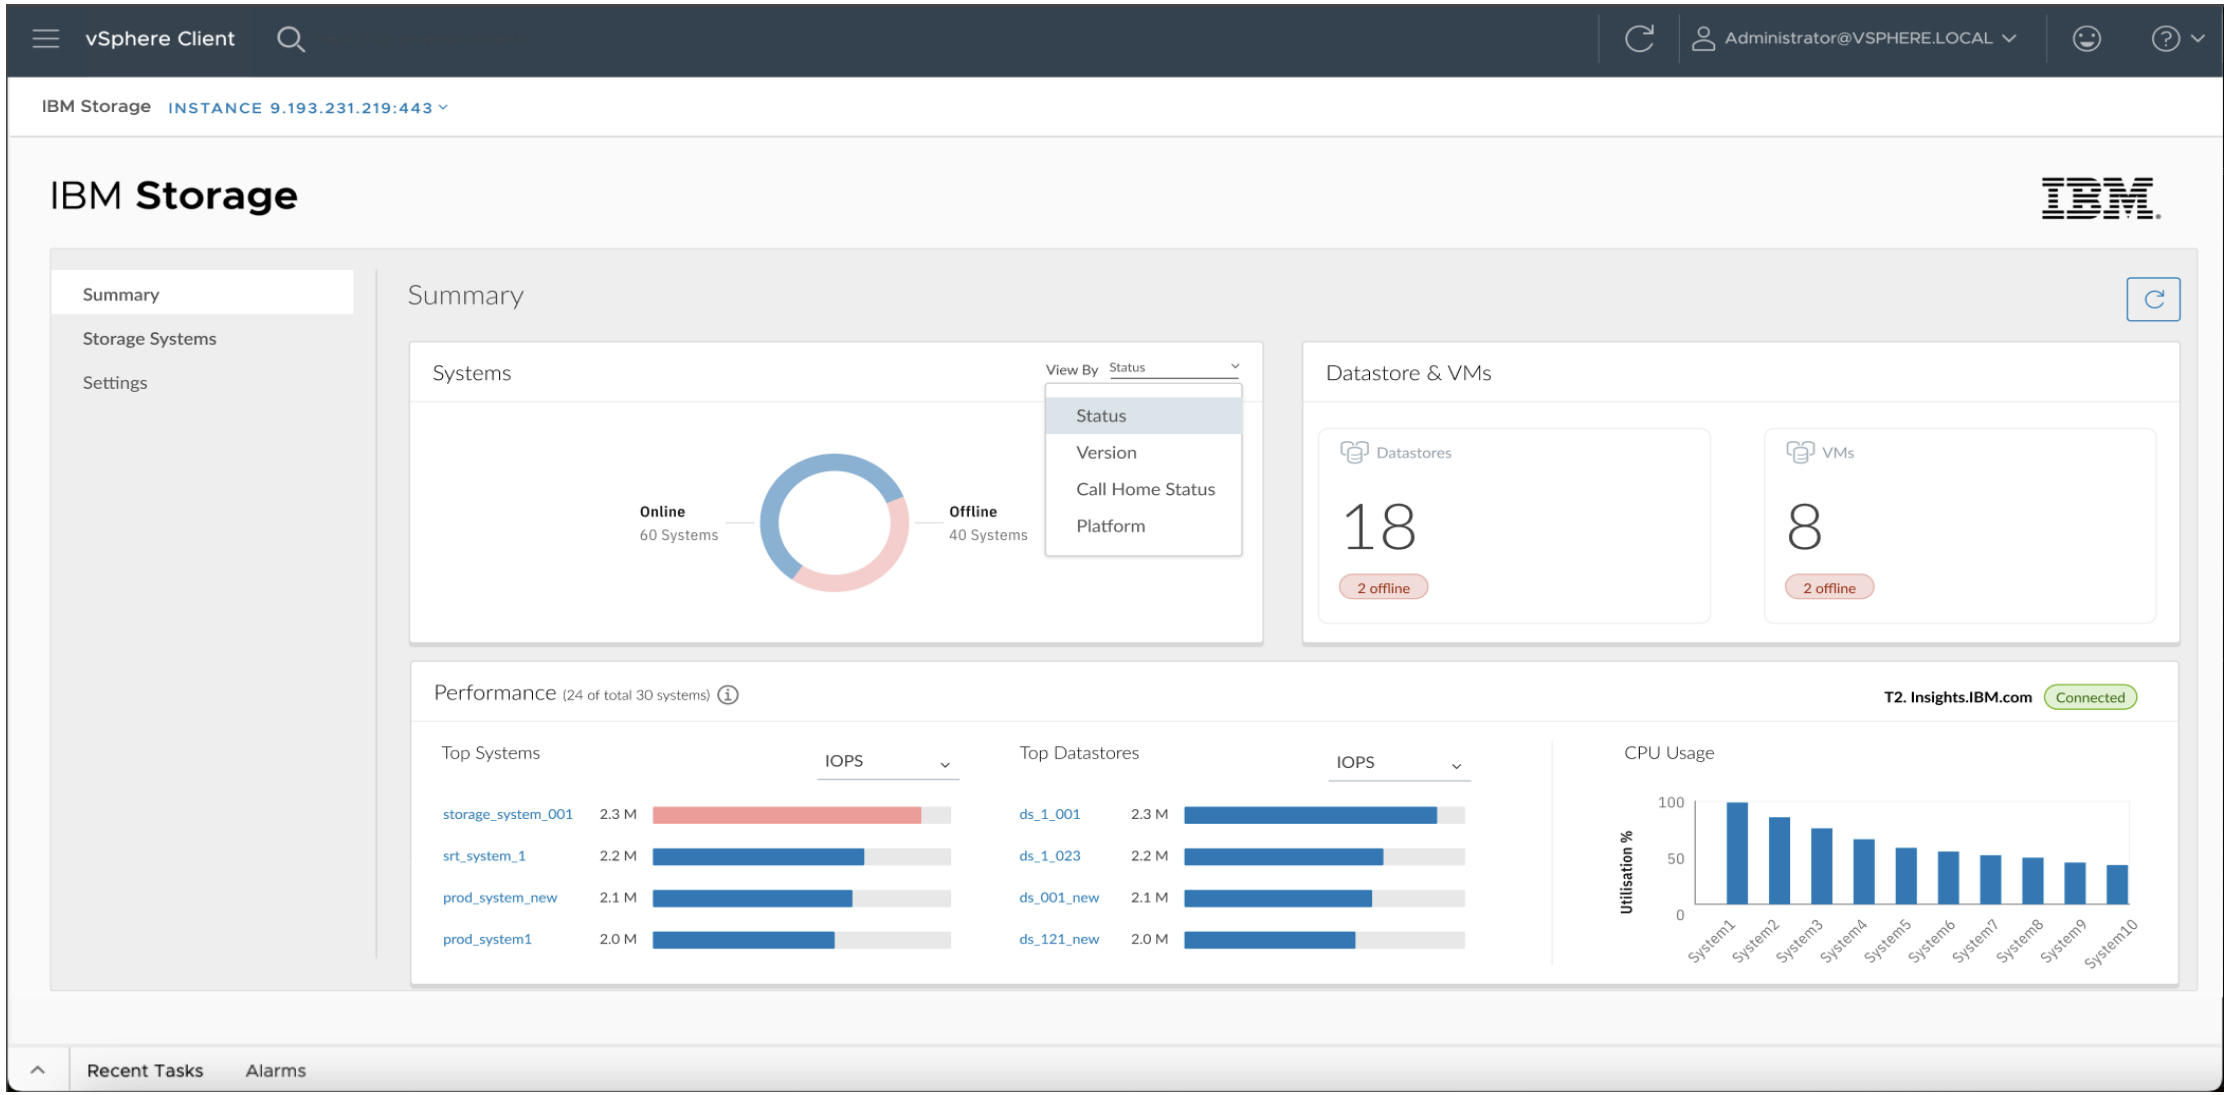

Performance Monitoring Dashboard

System Level Statistics

- View system level performance statistics in Storage System's Overview tab
- View current and historical performance data having both read and write flavor (Latency, IOPS, Bandwidth) in chart/table form.

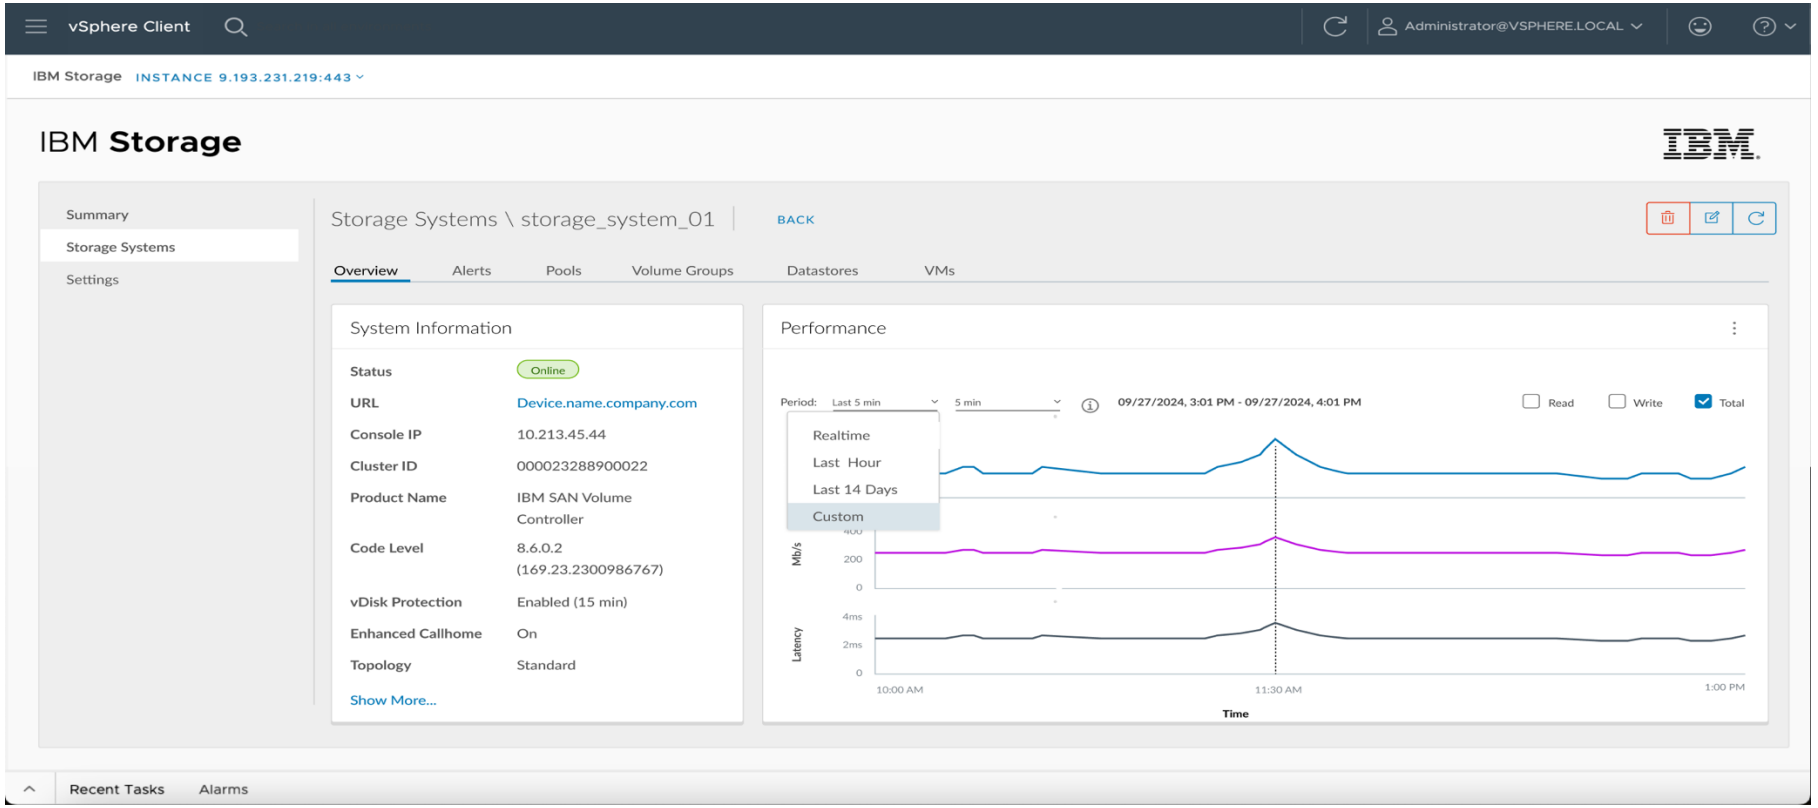

System Level Statistics

Volume Level Statistics

- View volume level performance statistics in Monitor tab of datastore
- View current and historical performance data having both read and write flavor (Latency, IOPS, Bandwidth) in chart/table form.

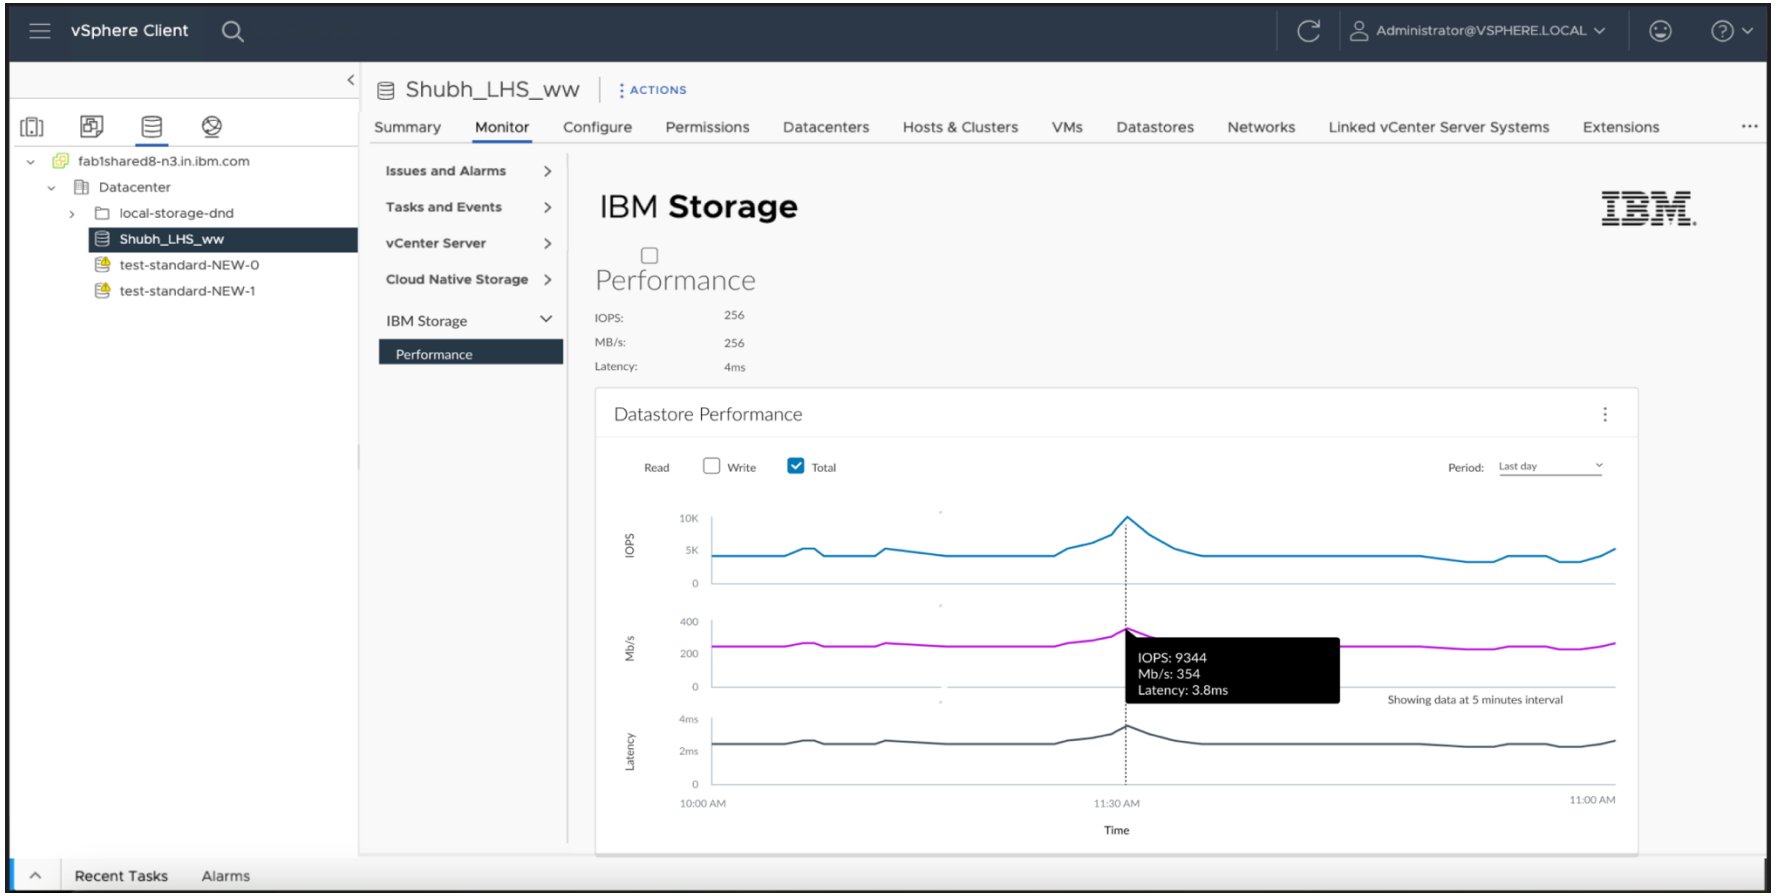

Volume Level Statistics

What's New in vSphere Plugin 2.0(continued...)

iSCSI host support

- Provision to add and edit iSCSI host through vSphere plugin
- Provision to create/reconcile iSCSI sessions
- Host Cluster Management support now comes with both FC and iSCSI type of hosts

Ransomware alerts integration through SI

Ransomware alert will be highlighted in following locations in vSphere plugin, if detected:

- Overall summary on dashboard
- Storage systems
- Datastore summary

List of alerts (apart from ransomware) from FlashSystem

Add iSCSI Host Through vSphere Plugin

- Add host configured with iSCSI protocol on storage system from plugin
  IBM Storage → host connections which is part of host configure tab in vSphere
- Provide required input like name, CHAP details, portset etc.
- Automatic discovery of target and session creation

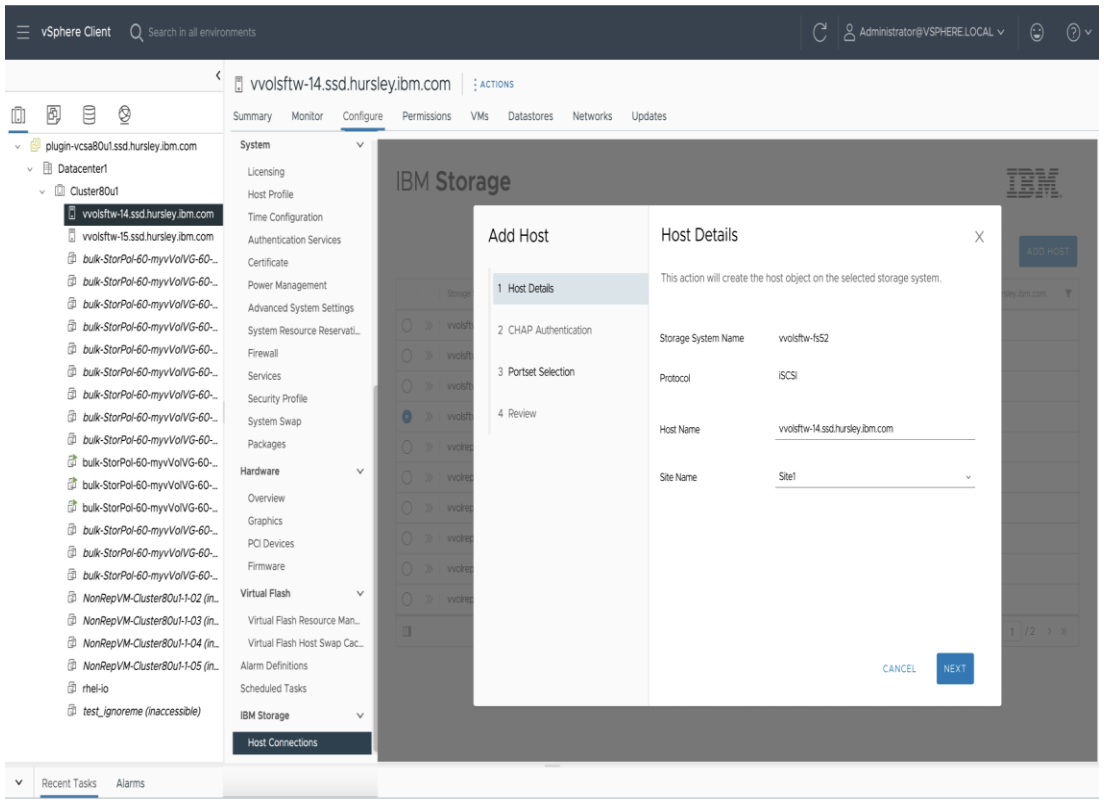

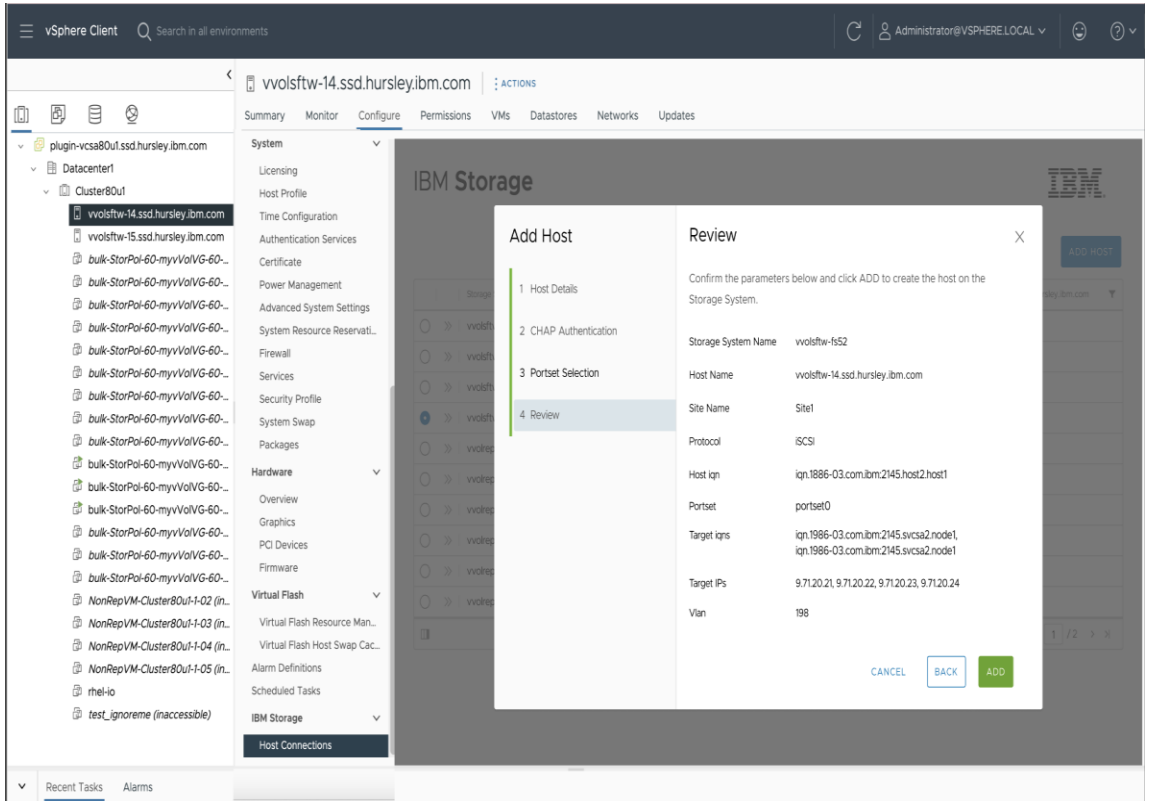

Add iSCSI host

Edit Host Through vSphere Plugin

- Edit host details like name, site, CHAP details, portset etc.

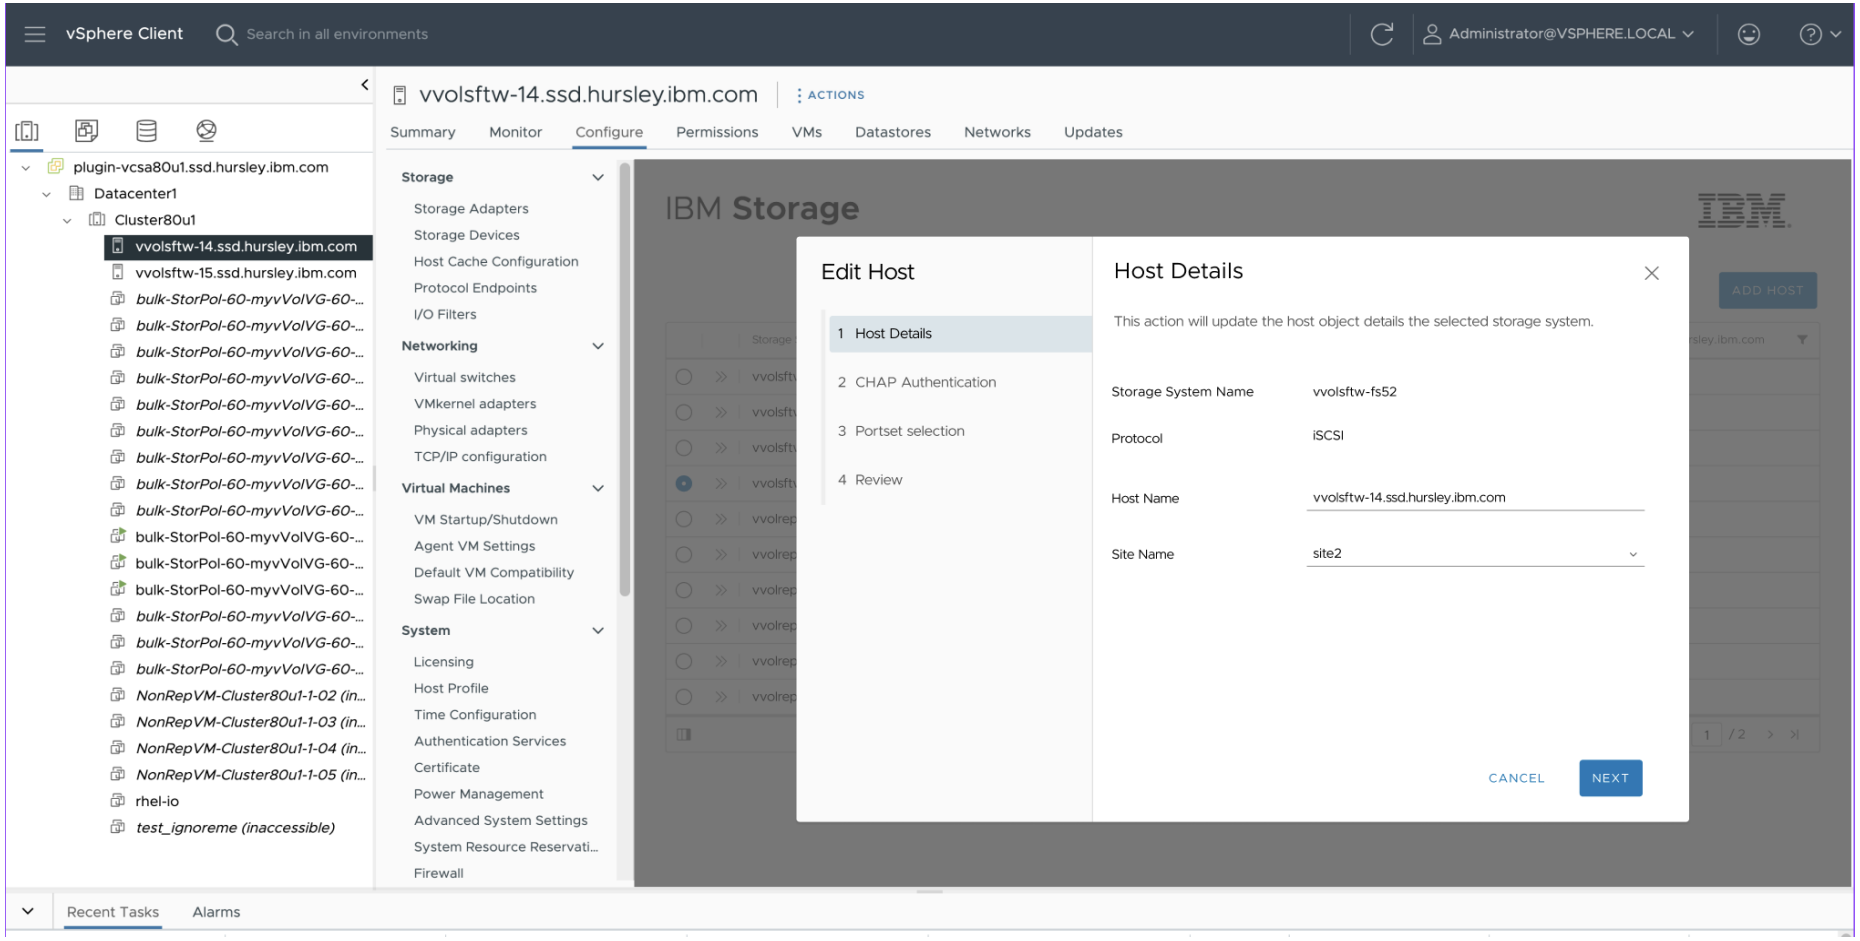

Edit host

Host Cluster Management for iSCSI Hosts

- Add multiple iSCSI hosts to the existing or a new cluster on storage system using vSphere plugin host cluster management in a single workflow

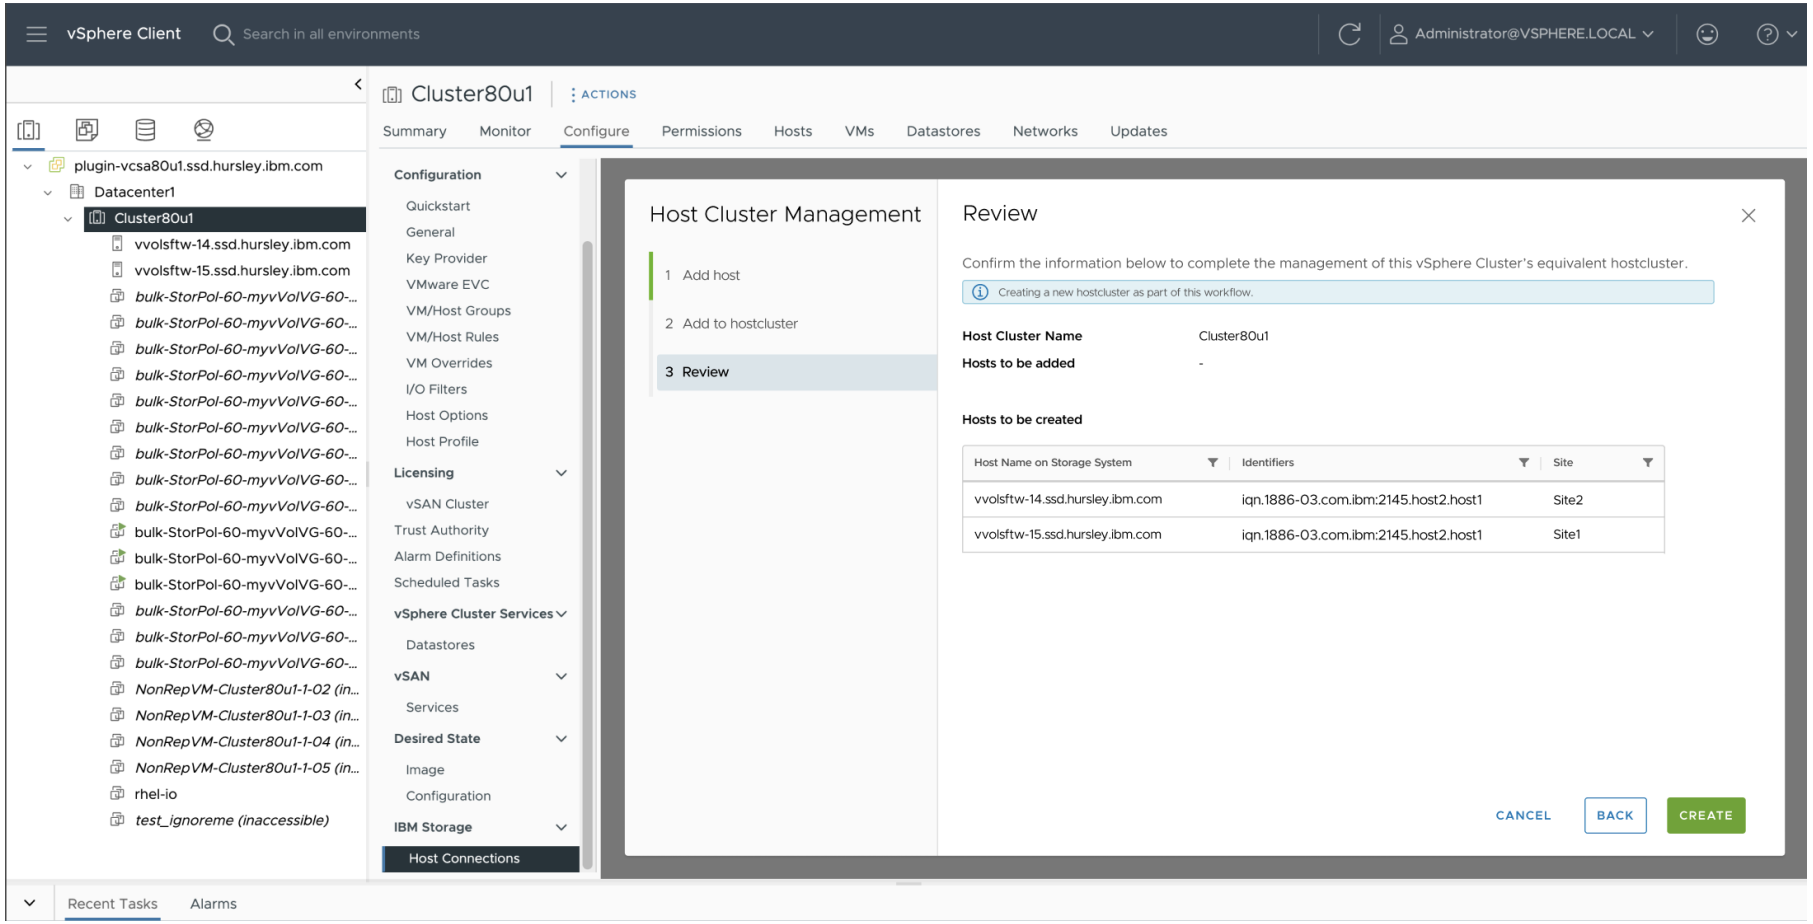

Host cluster management for iSCSI hosts

Ransomware Alerts Through SI

- Ransomware alert will be highlighted in important locations in vSphere plugin, if detected:

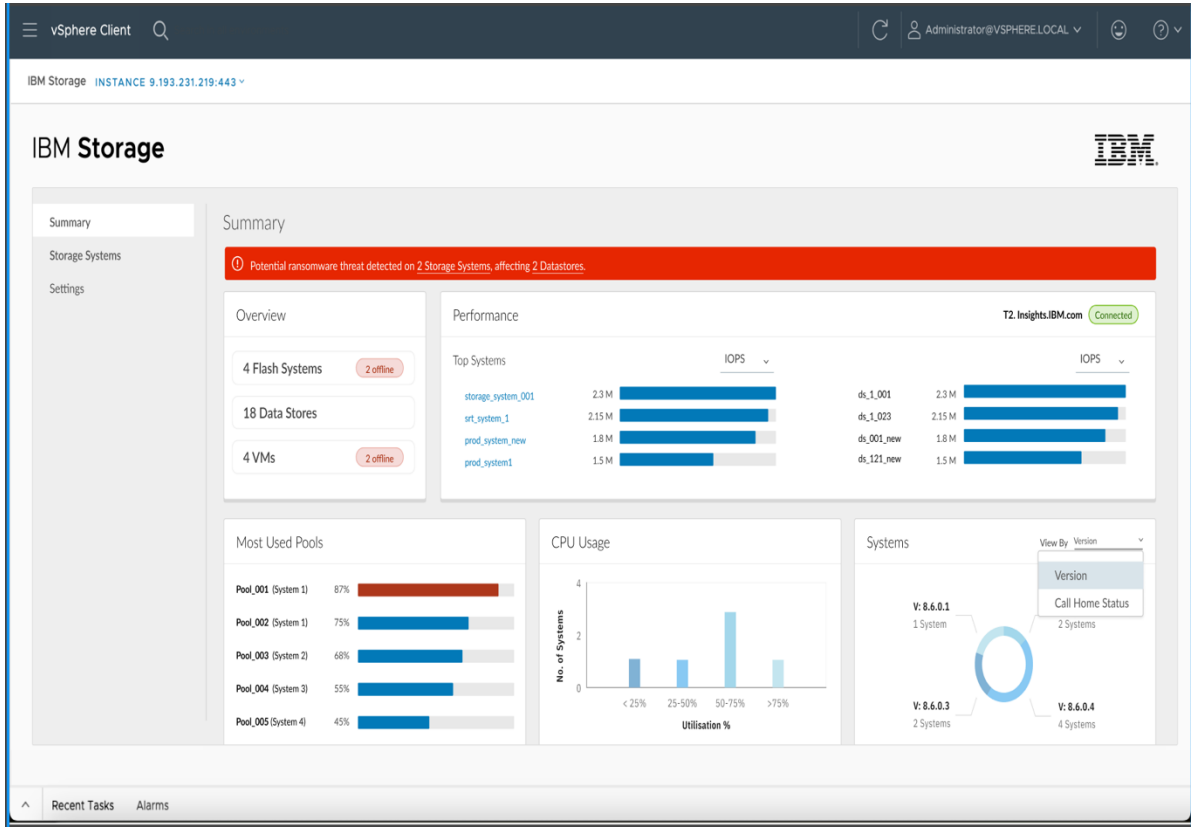

Ransomware alert on Overall Summary

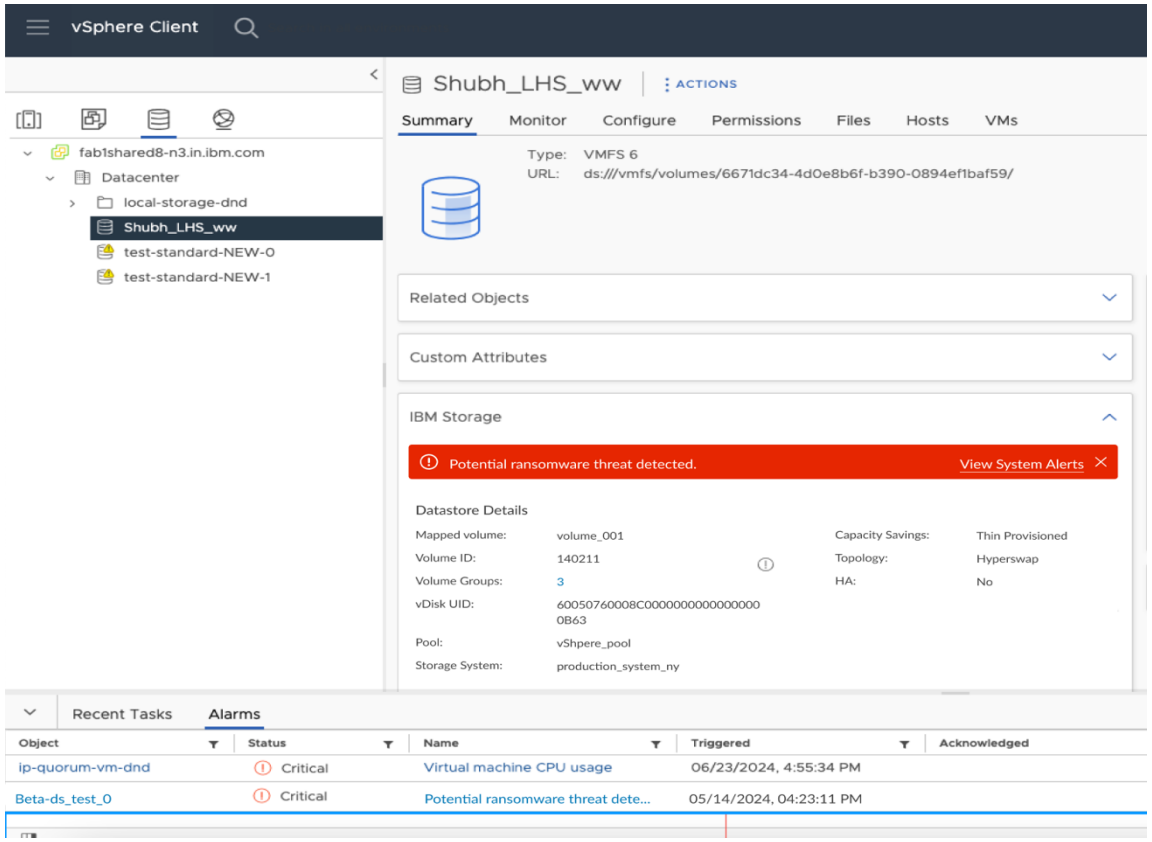

Ransomware alert on Datastore Summary

Alerts From Storage System

- View critical alerts from Storage System apart from ransomware

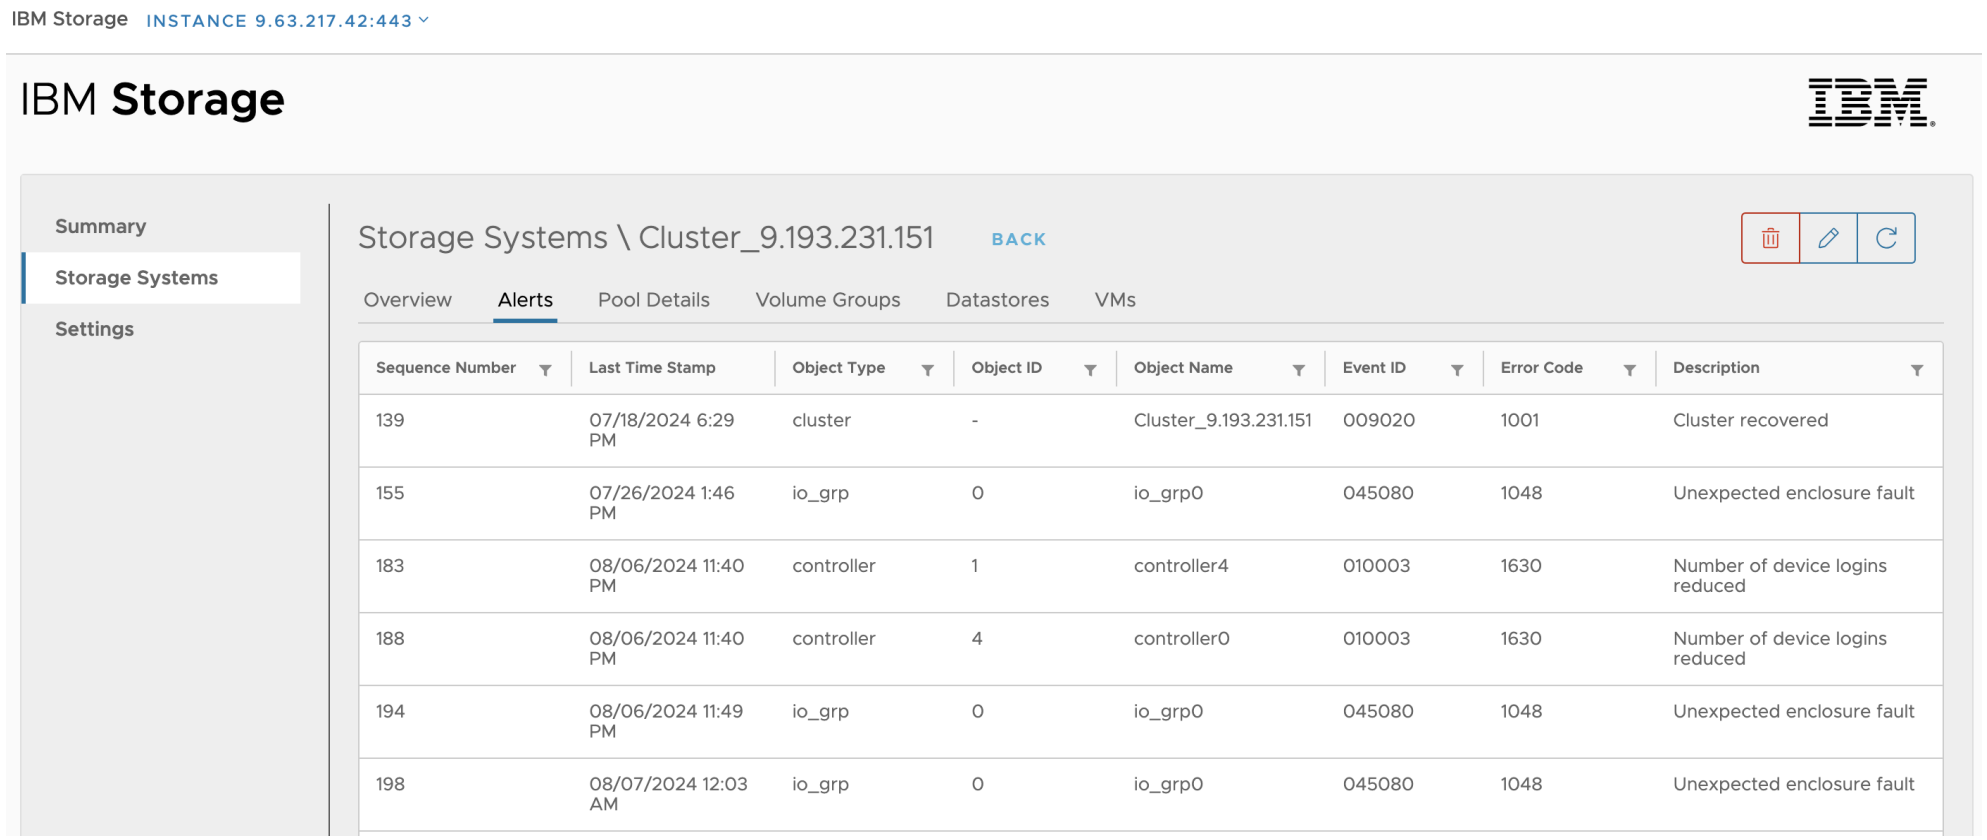

Sequence Number	Last Time Stamp	Object Type	Object ID	Object Name	Event ID	Error Code	Description
139	07/18/2024 6:29 PM	cluster	-	Cluster_9.193.231.151	009020	1001	Cluster recovered
155	07/26/2024 1:46 PM	io_grp	0	io_grp0	045080	1048	Unexpected enclosure fault
183	08/06/2024 11:40 PM	controller	1	controller4	010003	1630	Number of device logins reduced
188	08/06/2024 11:40 PM	controller	4	controller0	010003	1630	Number of device logins reduced
194	08/06/2024 11:49 PM	io_grp	0	io_grp0	045080	1048	Unexpected enclosure fault
198	08/07/2024 12:03 AM	io_grp	0	io_grp0	045080	1048	Unexpected enclosure fault

Alerts from Storage System

Plugin Upgrade

- Supporting upgrade from vSphere Plugin 1.1.x/1.2.0/1.3.0 to v2.0.0
- Seamless upgrade –
- Upgrade ensures continuous management of old datastores by new plugin instance.
- All registered storage systems and pools remains intact in vSphere plugin inventory

Restrictions And Limitations

- Only fiber channel (SCSI) and iSCSI connectivity between the ESXi hosts and the storage systems would be supported through plugin 2.0.0.
- Standard, Stretched and HyperSwap topologies are supported by the plugin.
- Plugin does not support multiple vCenters unless they are configured in linked mode
- Recommended to have user, hosts and pools either to be part of same ownership group or not part of any ownership group
- Partition based features like PBHA, 3site PBR + PBHA etc. are not yet supported

IBM Storage Virtualize

Plugin Updates

CSI v1.12
December

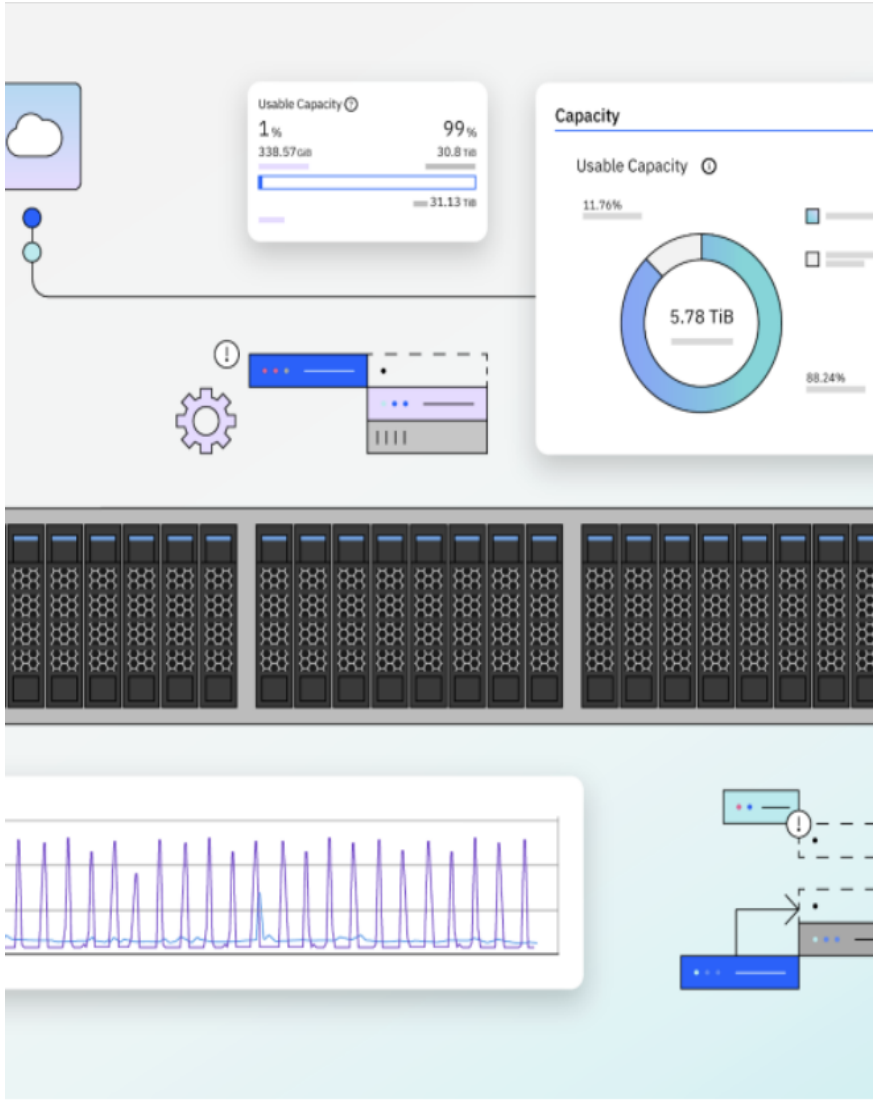

PVC Shared Access mode - RWX

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: demo-pvc-file-system
spec:
  volumeMode: Filesystem  # Optional. The default is Filesystem.
  accessModes:
  - ReadWriteMany
  resources:
    requests:
      storage: 1Gi
  storageClassName: demo-storageclass
```

- IBM added RWX support to PVC, which allows the creation of VMs with OCPv that are "live-migration" capable.
- Existing PVC can also be converted to RWX manually by changing the accessModes to ReadWriteMany or by using the example script.

Portset for HostDefiner

Field	Description
`prefix`	Adds a prefix to the hosts defined by the host definer. **Note:** The prefix length is bound by the limitation of the storage system. When defined, the length is a combination of both the prefix and node (server) hostname.
`connectivityType`	Selects the connectivity type for the host ports. Possible input values are: - `nvmeofc` for use with NVMe over Fibre Channel connectivity - `fc` for use with Fibre Channel over SCSI connectivity - `iscsi` for use with iSCSI connectivity By default, this field is blank and the host definer selects the first of available connectivity types on the node, according to the following hierarchy: NVMe, FC, iSCSI. **Note:** When left blank, the connectivity type will update along with any changes within the host ports, according to the set hierarchy. If the value is set and there are host port changes, connectivity needs to be manually updated. For more information, see Changing node connectivity.
`allowDelete`	Defines whether the host definer is allowed to delete host definitions on the storage system. Input values are `true` or `false`. The default value is `true`.
`dynamicNodeLabeling`	Defines whether the nodes that run the CSI node pod are dynamically labeled or if the user must create the `hostdefiner.block.csi.ibm.com/manage-node=true` label on each relevant node. This label tells the host definer which nodes to manage their host definition on the storage side. Input values are `true` or `false`. The default value is `false`, where the user must manually create this label on every node to be managed by the host definer for dynamic host definition on the storage.
portSet	Specifies port set for new port definitions (ports already defined on SVC are not modified).

- With this feature enabled, all hosts dynamically created by HostDefiner will be set to the provided portset.
- There is no need to change the portset for existing hosts.
- See official docs for more information on host definer.

CSI Support Information

- General documentation:
- https://www.ibm.com/docs/en/stg-block-csi-driver/1.11.4

81

- Log collection instructions:
- https://www.ibm.com/docs/en/stg-block-csi-driver/1.11.4?topic=troubleshooting-log-status-collection

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 with code 5151 0447 or

Direct link https://www.menti.com/alhsf3bgvxu6
Or

QR Code