IBM

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar: Storage Virtualize 8.6.2 Update

Byron Grossnickle

IBM Advanced Technology Group - Storage Virtualize SME

byrongro@us.ibm.com

IBM

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.

2024 Upcoming Webinars – Register Here!

Exploring the IBM DS8 Safeguarded Copy Capacity Study – April 4th, 2024

Leveraging Your Data Lake with IBM Storage Fusion Data Cataloging
Part 1: April 9th, 2024
Part 2: April 16th, 2024

Important Links to bookmark:

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

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

Offerings

Client Technical Workshops

- **IBM DS8900F Advanced Functions:** May 8-9 in Chicago, IL
- **IBM Cyber Resiliency with IBM Storage Defender:** June 5-6 in Tucson, AZ
- IBM Storage Fusion
- IBM Storage Scale & Storage Scale System
- **IBM FlashSystem Deep Dive & Advanced Functions:** May 22-23 in Atlanta, GA
- IBM Storage Ceph
- IBM FlashSystem and Storage Virtualize
- IBM Cloud Object Storage

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.**

Save the date

Storage @ IBM TechXchange Conference 2024

October 21-24, 2024
Mandalay Bay | Las Vegas
#IBMTechXchange

Key Learnings

- Practical how-to advice
- Patterns and best practices
- Success stories, IBM PoV, proven techniques

Featured Products

IBM Storage Defender
IBM Storage Fusion
IBM Storage Scale + IBM Storage Ceph
IBM Tape + IBM SAN
IBM Storage FlashSystem + IBM Storage DS8000

Collaborate. Learn. Play.

Community

IBM Champions
User Groups
Tech Peers
Business Partners

Sandbox

Network
Learn
Collaborate
Play

Roadmaps

Go deep with people in the know and set the stage for where IBM is going in the future

Breakout Sessions

Trends and Directions
User Groups
Product Deep Dives
Meet the Expert
Professional Development
Show the Code
Birds of a Feather
Academic/Research

Accelerate your Career

Labs (Instructor-Led, Self-paced)
IBM Certification Testing
Earn up to 25 hours in CPE credits

https://www.ibm.com/community/ibm-techxchange-conference/

Game On! →

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 1708 6924 or

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

Or

QR Code

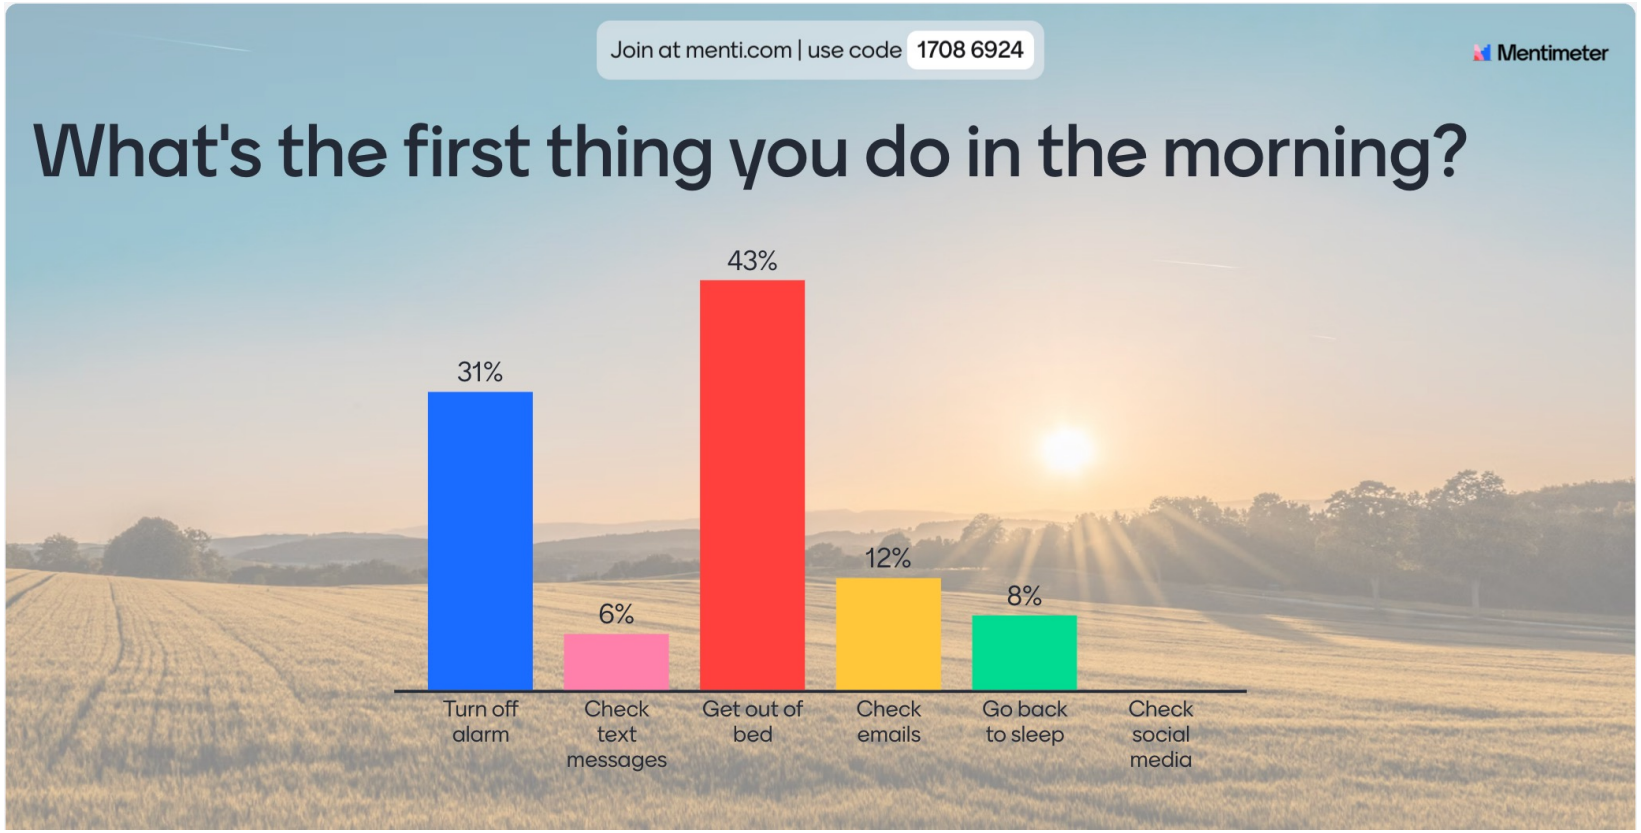

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar: Storage Virtualize 8.6.2 Update

Byron Grossnickle

IBM Advanced Technology Group - Storage Virtualize SME

byrongro@us.ibm.com

Meet the Speakers

Byron Grossnickle is an IBM Storage Technical Specialist concentrating on Storage Virtualize software. This include FlashSystem, SVC, and Storage Virtualize for Public Cloud. Byron has been with IBM 18 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.

Agenda

- Release Schedule
- NVMe/TCP Performance Enhancements
- High Speed Ethernet (iSER) Replication
- Performance – Lee and Yijie
- VMware Messaging
- vVol 2.0 Replication
- VMware Plugin 1.2
- Restore In Place/Single Volume Clone
- Misc
  - Dual Voltage PS for FS5200
  - Encryption Recovery Key
  - SNMP Security Improvements
  - 4 I/O Group Support for PBR
  - 4 PB Replication Capacity for SV3
  - DRAID Smart Rebuild with Extent Awareness
  - FIPS-140
  - Scalability Increases
  - GUI Enhancements
  - Statistics Enhancements/Environmental Stats

Release Schedule

- RFA Announce – 8.6.2 – December 5, 2023
- eGA – 8.6.2 – December 21, 2023

8.6.2.1 is a Non-LTS Release. This means that it will get no patches or updates until the next Non-LTS or LTS release

Storage Virtualize was previously known as Spectrum Virtualize

Ethernet and NVMe/TCP Enhancements

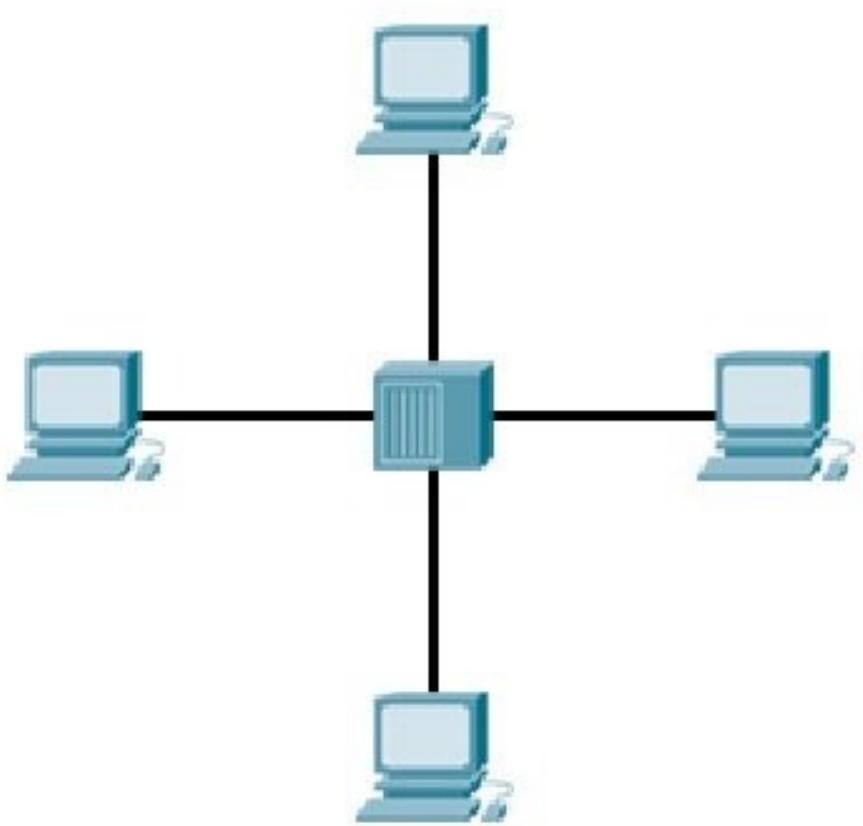

Jumbo Frames on RoCE Adapters

- In 8.6.0.0 jumbo frames were disabled on RoCE (Mellanox) adapters due to the driver crashing
  - In 8.6.2 jumbo frames have been re-enabled giving significant performance advantages to iSCSI/iSER
  - This fix is backported to 8.6.0.2+
- NVMe/TCP and NVMe/RDMA are still limited to a 1500 MTU. Hopefully this limitation will be lifted in the future.

NVMe/TCP Enhancements

- lssystemstats and lsnodecanisterstats will show NVMe/TCP and NVMe/RDMA IO statistics, in MB/s and IOPs, for all nodes as well as per node canister for later
- The above statistics are available in GUI as well

Host Attach Changes

- mkhost/chhost/mkip CLIs will now throw error message when NVMe/TCP and NVMe/RDMA host attach is getting used with portset having Ethernet ports with MTU other than 1500

High Speed Ethernet Replication (iSER)

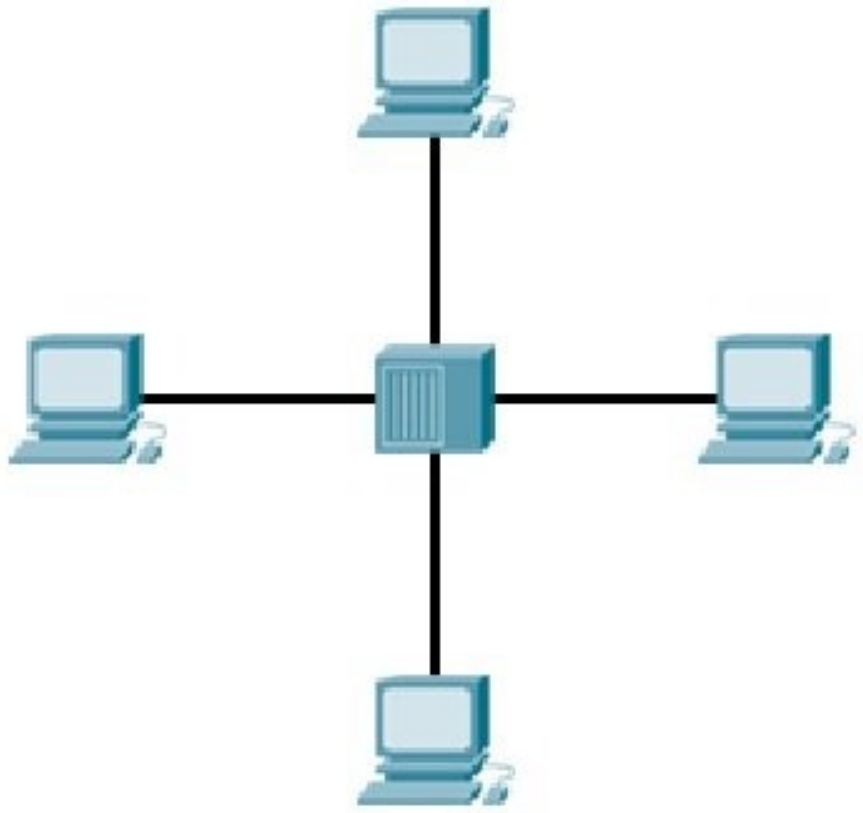

Short Distance/High Speed Ethernet Link

- This was designed for synchronous replication
  - 1ms rtt or less
  - Zero packet drop is preferred on the network
  - iSER protocol used
- Short distance partnership using RDMA allows to connect two systems and establish an IP partnership. This partnership uses RDMA capable Ethernet ports to accelerate the traffic between Storage Virtualize systems
- iWarp cards only

- Use Cases
  - MetroMirror
  - Volume Mobility
  - Potentially PB-HA in the future

Admin Model

1. Create highspeedreplication portsets
2. Associate IP addresses with highspeedreplication portsets
3. Create partnership using mkippartnership command in same way as you would for native IP partnership except use highspeed replication portsets for link1 and link2
4. There is no change in lspartnership view. Users may look at link1, link2 field and match them into lsportset view to find out if the partnership uses TCP or RDMA. GUI already displays this information in partnership list.
5. Each node displays its connectivity under “sainfo lsnodeipconnectivity” view.

```
IBM_FlashSystem:FS9500_Cluster:superuser> sainfo lsnodeipconnectivity
status          local_port_id local_vlan local_rdma_type local_ip_addr remote_port_id remote_vlan remote_rdma_type remote_ip_addr remote_wwnn      remote_panel_name cluster_id       error_data
Connected:iWARP 4             30         iWARP           17.17.17.41   6              30          iWARP           17.17.17.21    500507681000110C 01-1              0000020420002218
Connected:iWARP 4             30         iWARP           17.17.17.41   6              30          iWARP           17.17.17.22    5005076810001143 01-2              0000020420002218
Connected:iWARP 4             30         iWARP           17.17.17.41   6              30          iWARP           17.17.17.23    500507681000AA02 02-1              0000020420002218
Connected:iWARP 4             30         iWARP           17.17.17.41   6              30          iWARP           17.17.17.24    50050768100032CA 02-2              0000020420002218
Connected:iWARP 6             30         iWARP           17.17.17.61   6              30          iWARP           17.17.17.21    500507681000110C 01-1              0000020420002218
Connected:iWARP 6             30         iWARP           17.17.17.61   6              30          iWARP           17.17.17.22    5005076810001143 01-2              0000020420002218
Connected:iWARP 6             30         iWARP           17.17.17.61   6              30          iWARP           17.17.17.23    500507681000AA02 02-1              0000020420002218
Connected:iWARP 6             30         iWARP           17.17.17.61   6              30          iWARP           17.17.17.24    50050768100032CA 02-2              0000020420002218
```

Notes and Limitations

1. Ethernet HyperSwap and Short distance partnership using RDMA are mutually exclusive on system
2. Short distance partnership using RDMA does not support Encryption of data in flight and Compression
3. Short distance partnership using RDMA is only supported using iWARP
4. This feature requires users to open network TCP ports 3260 and 21456 in their network firewall for management and data connection between systems
5. Most of native IP rep DMPs (2021/2022/2023) are used by short distance partnership as well. DMP 2020 is only for native IP replication (uses TCP).

Performance Summary

- HSE replication demonstrates equivalent or better performance when compared to FC replication. It offers user an alternative mirroring connectivity (other than IP replication) when high replication performance is in demand and short distance is allowed.
  - Note: PBHA will be supported for HSE replication in the near future, but not yet available in v8.6.2 release
- MTU9000 enablement improves iscsi performance by up to 2x, and the performance improvement is more obvious in FS9500 than in FS7300.

vVol Replication

vVol Replication Agenda

Overview

Planning

Implementing

Monitoring

Managing

Each site requires:

- A storage system
- A vCenter Server
- One or more ESXi hosts

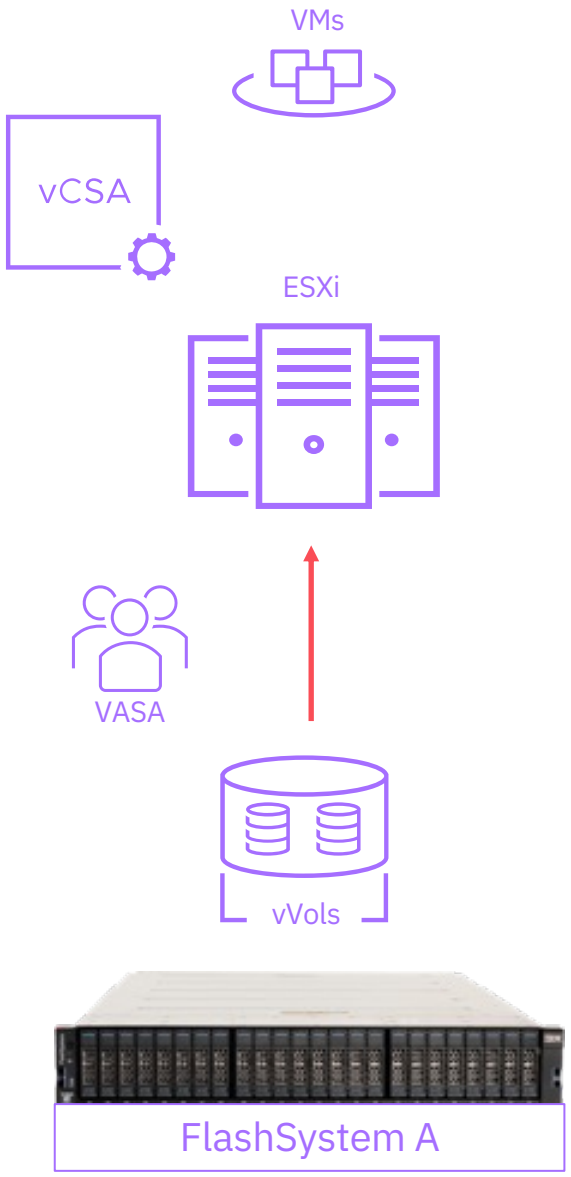

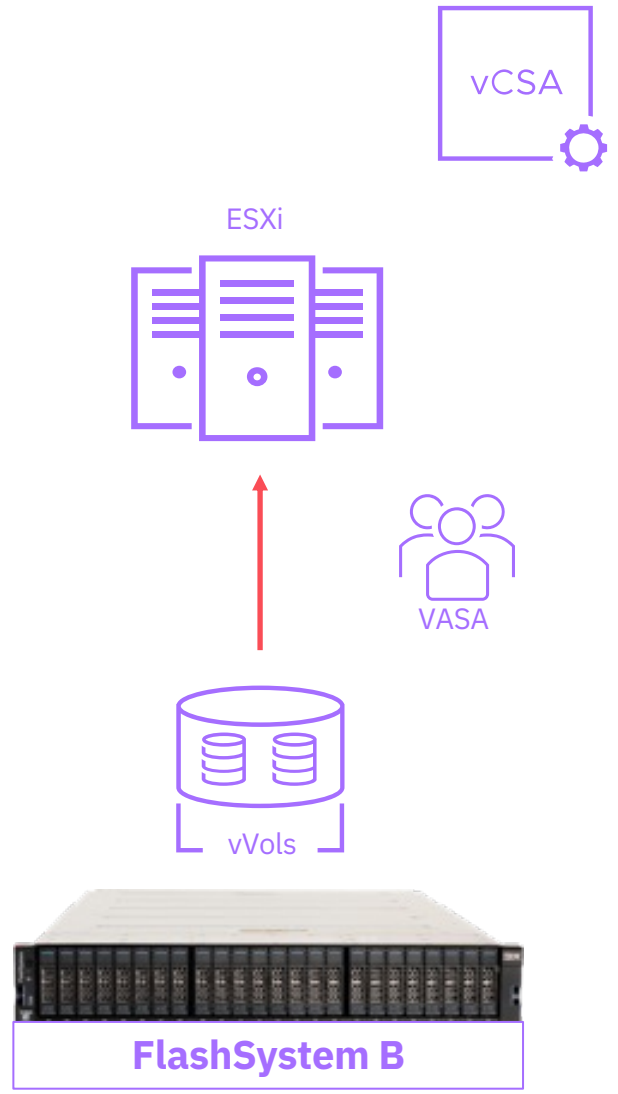

Configure a 2-site partnership between systems using async policy-based replication

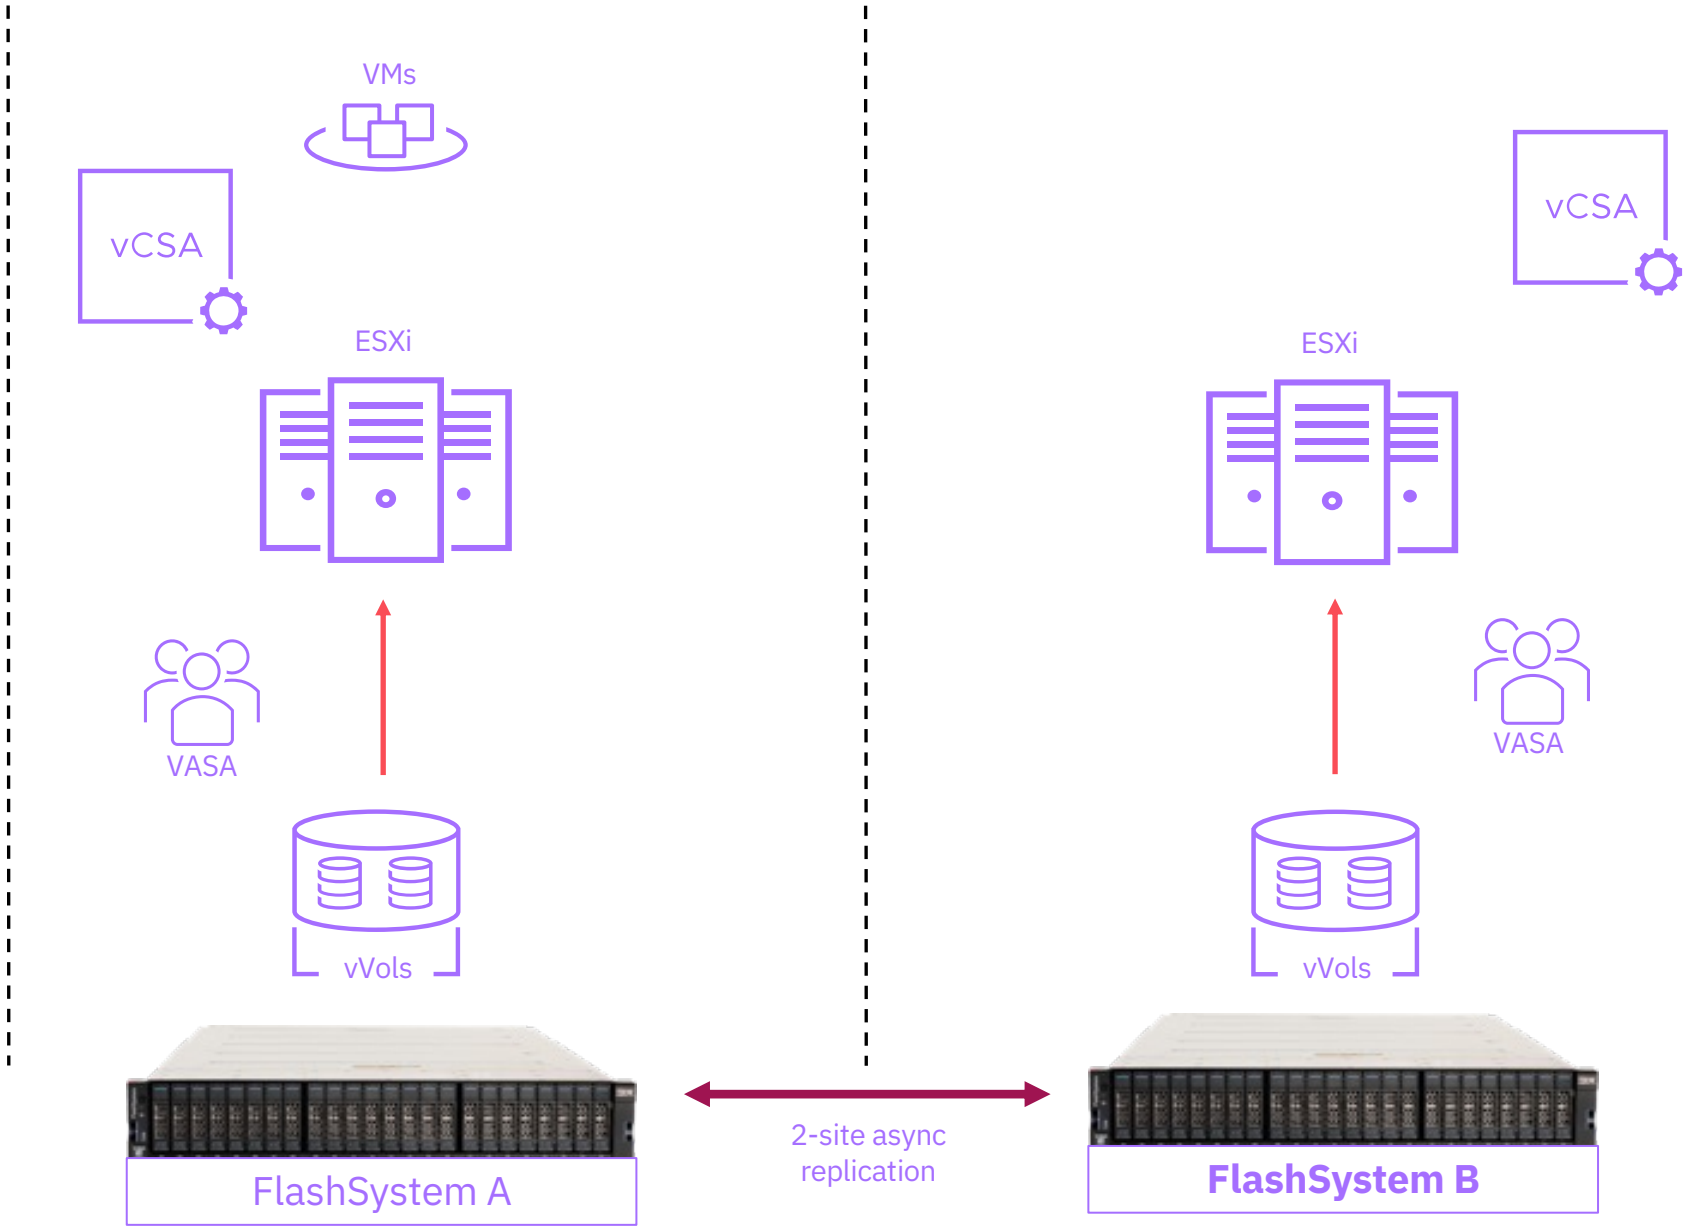

Enable vVol replication to:

- Link ownership groups
- Link child pools
- Create a volume group with a replication policy

Additional volume groups and child pools can be configured later

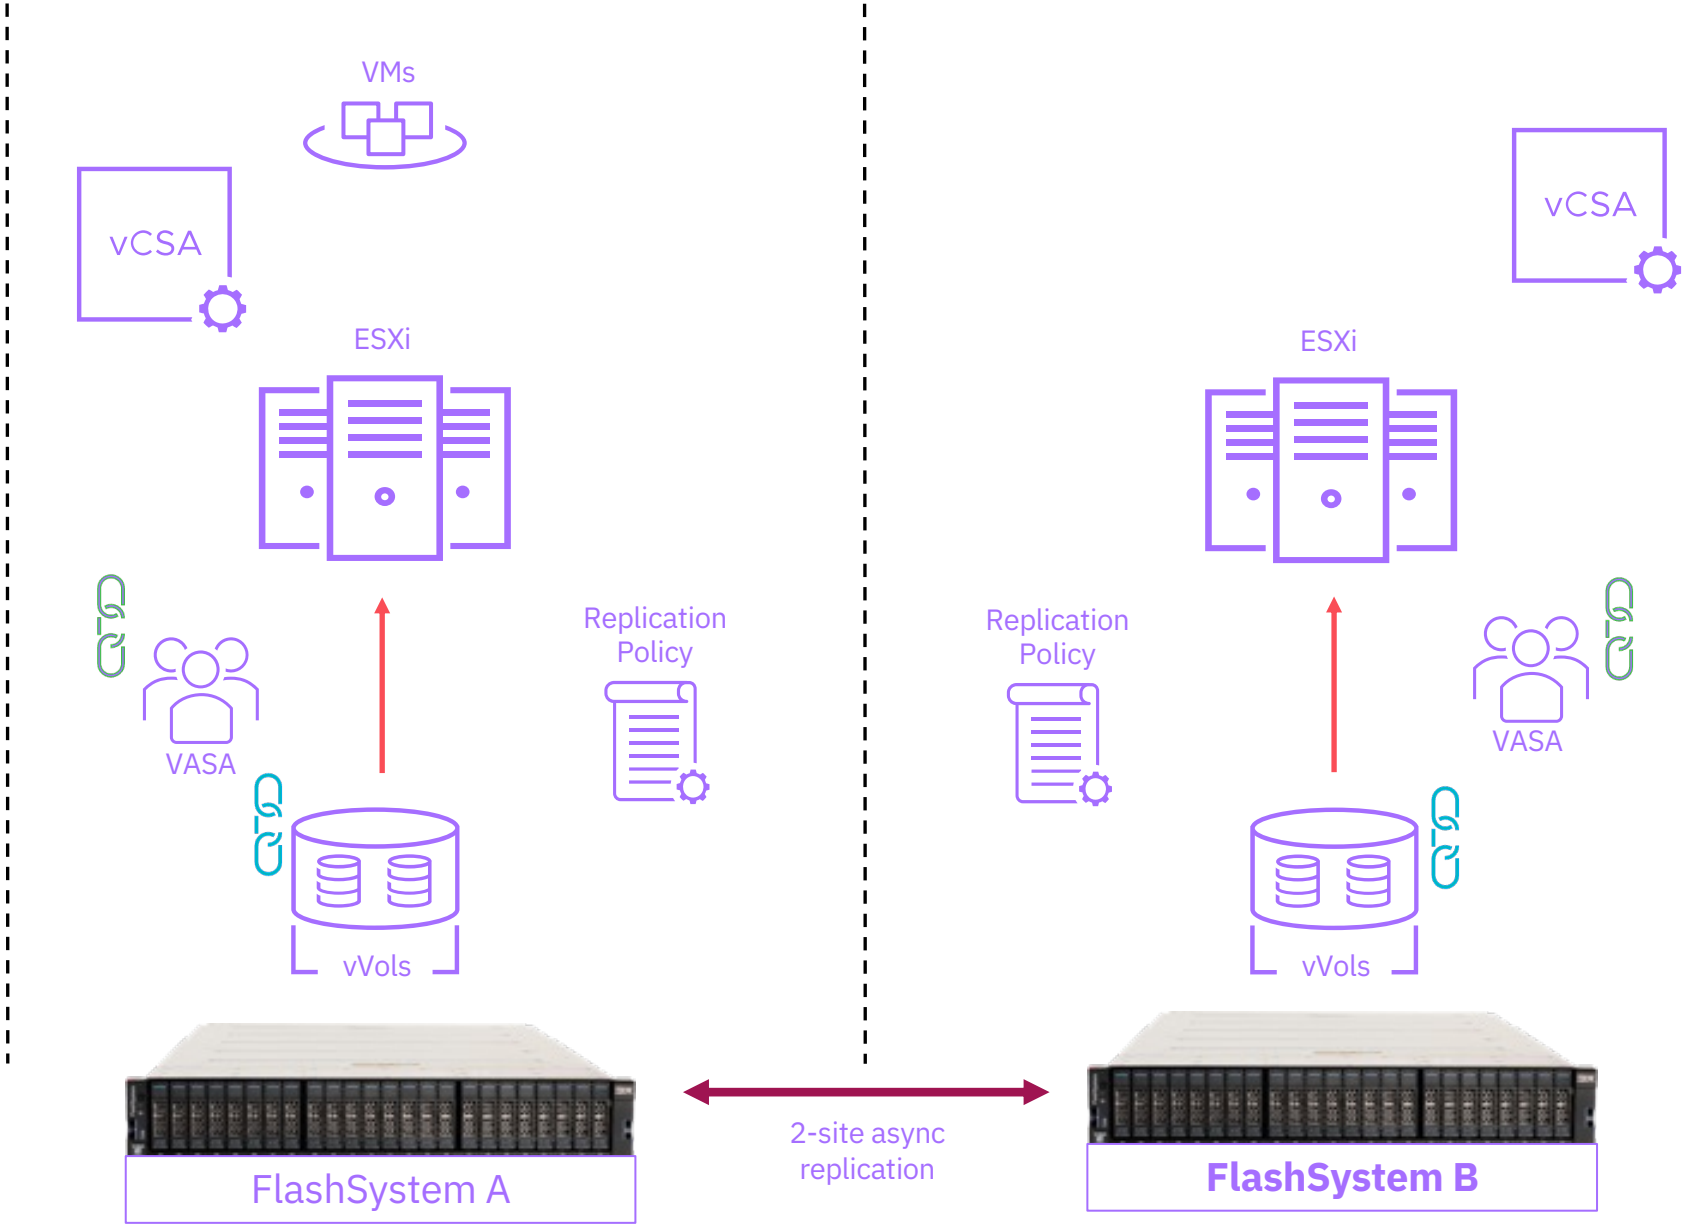

Each vVol volume group with a replication policy is a Replication Group for VMware

A VM Storage Policy can be created in vSphere that specifies the replication settings:

- Remote system
- RPO

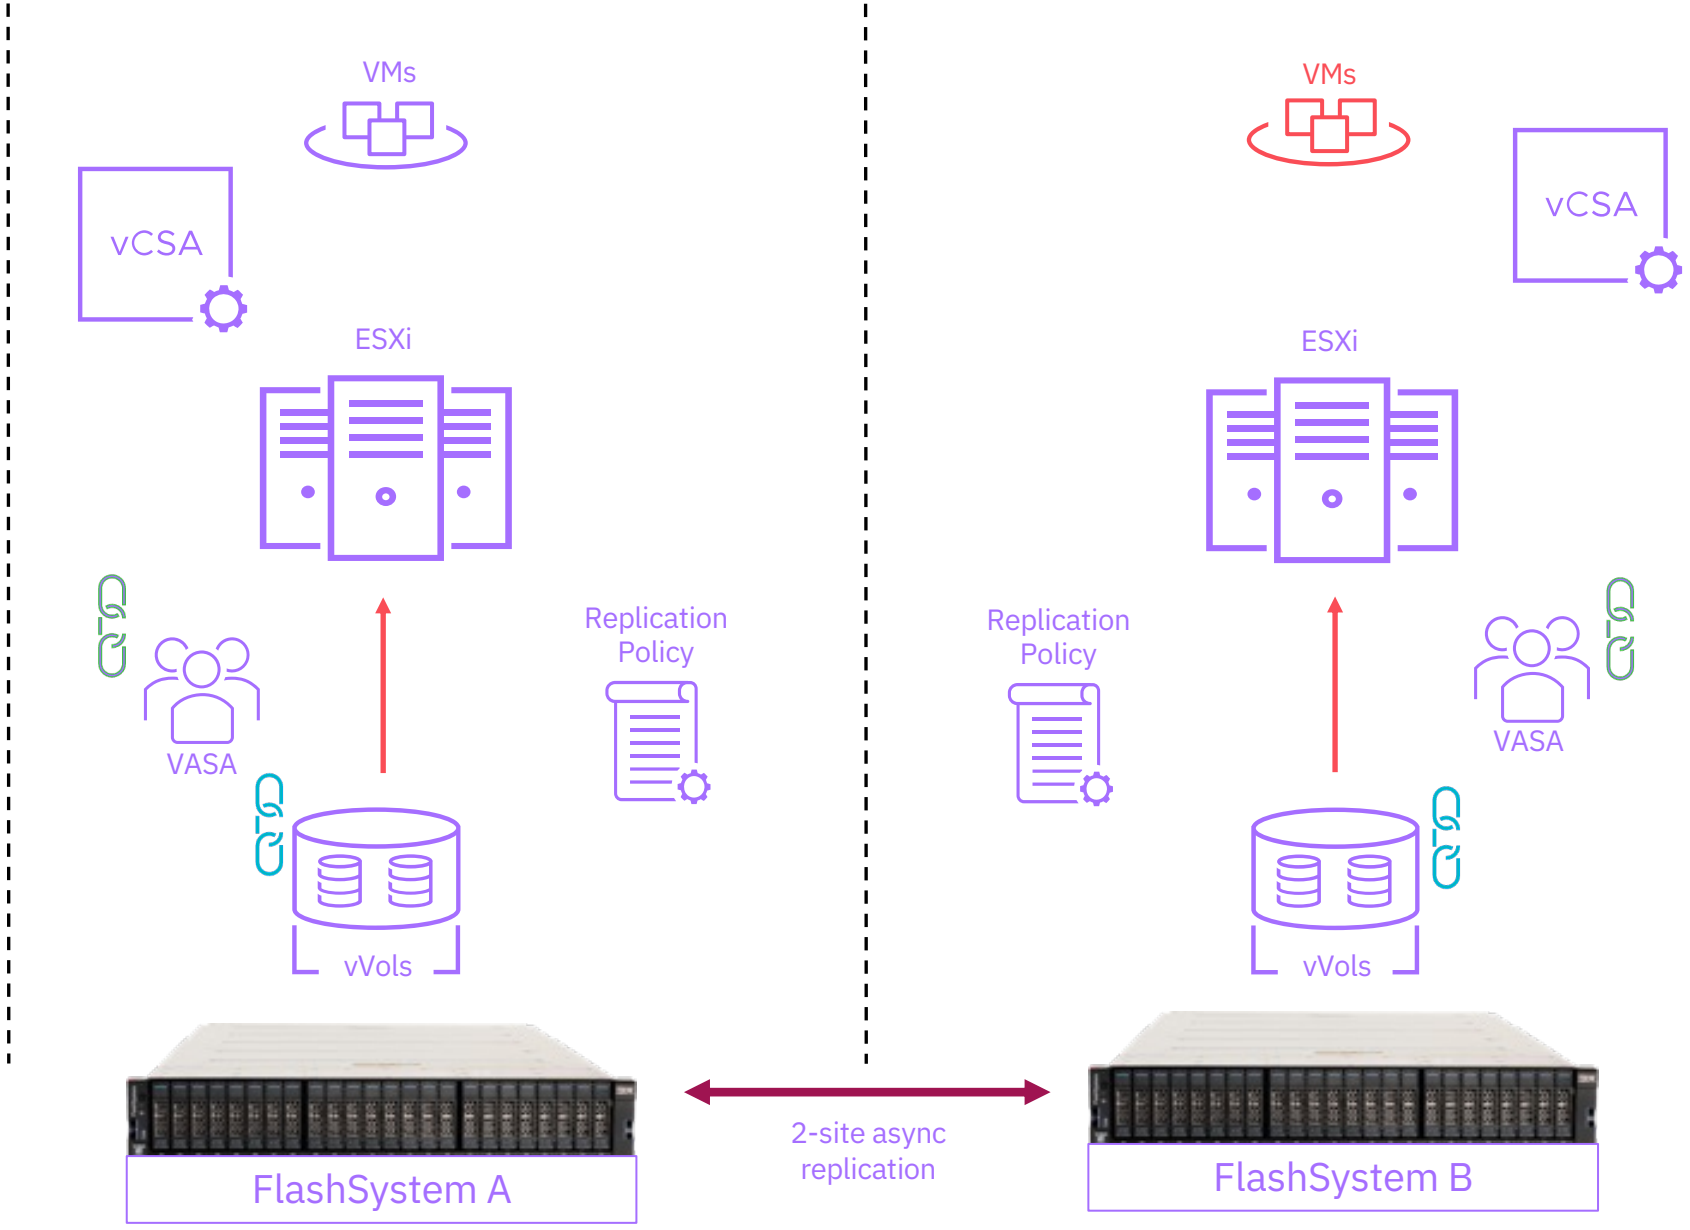

Replication is managed in the VMware environment using Site Recovery Manager or PowerCLI

Virtual Machines can be failed-over to the recovery site

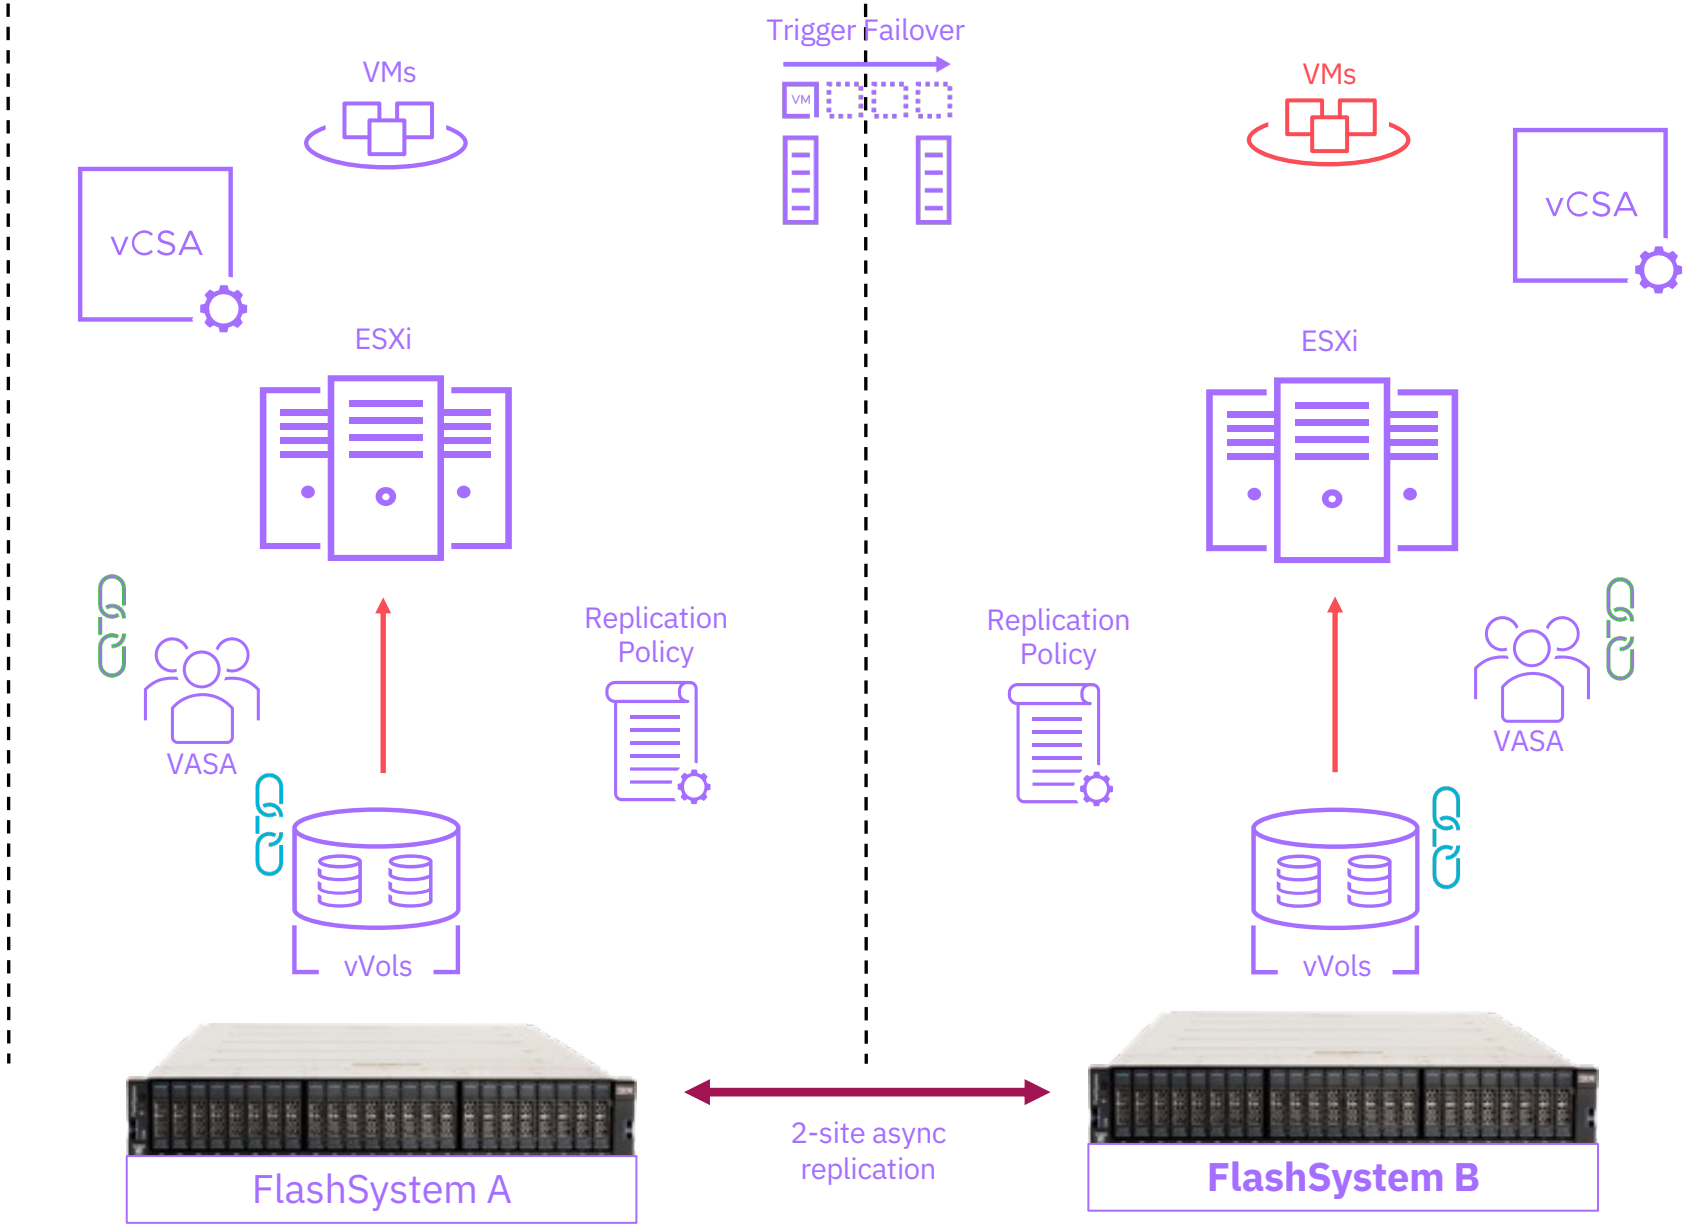

Replication is managed in the VMware environment using Site Recovery Manager or PowerCLI

Virtual Machines can be failed-over to the recovery site

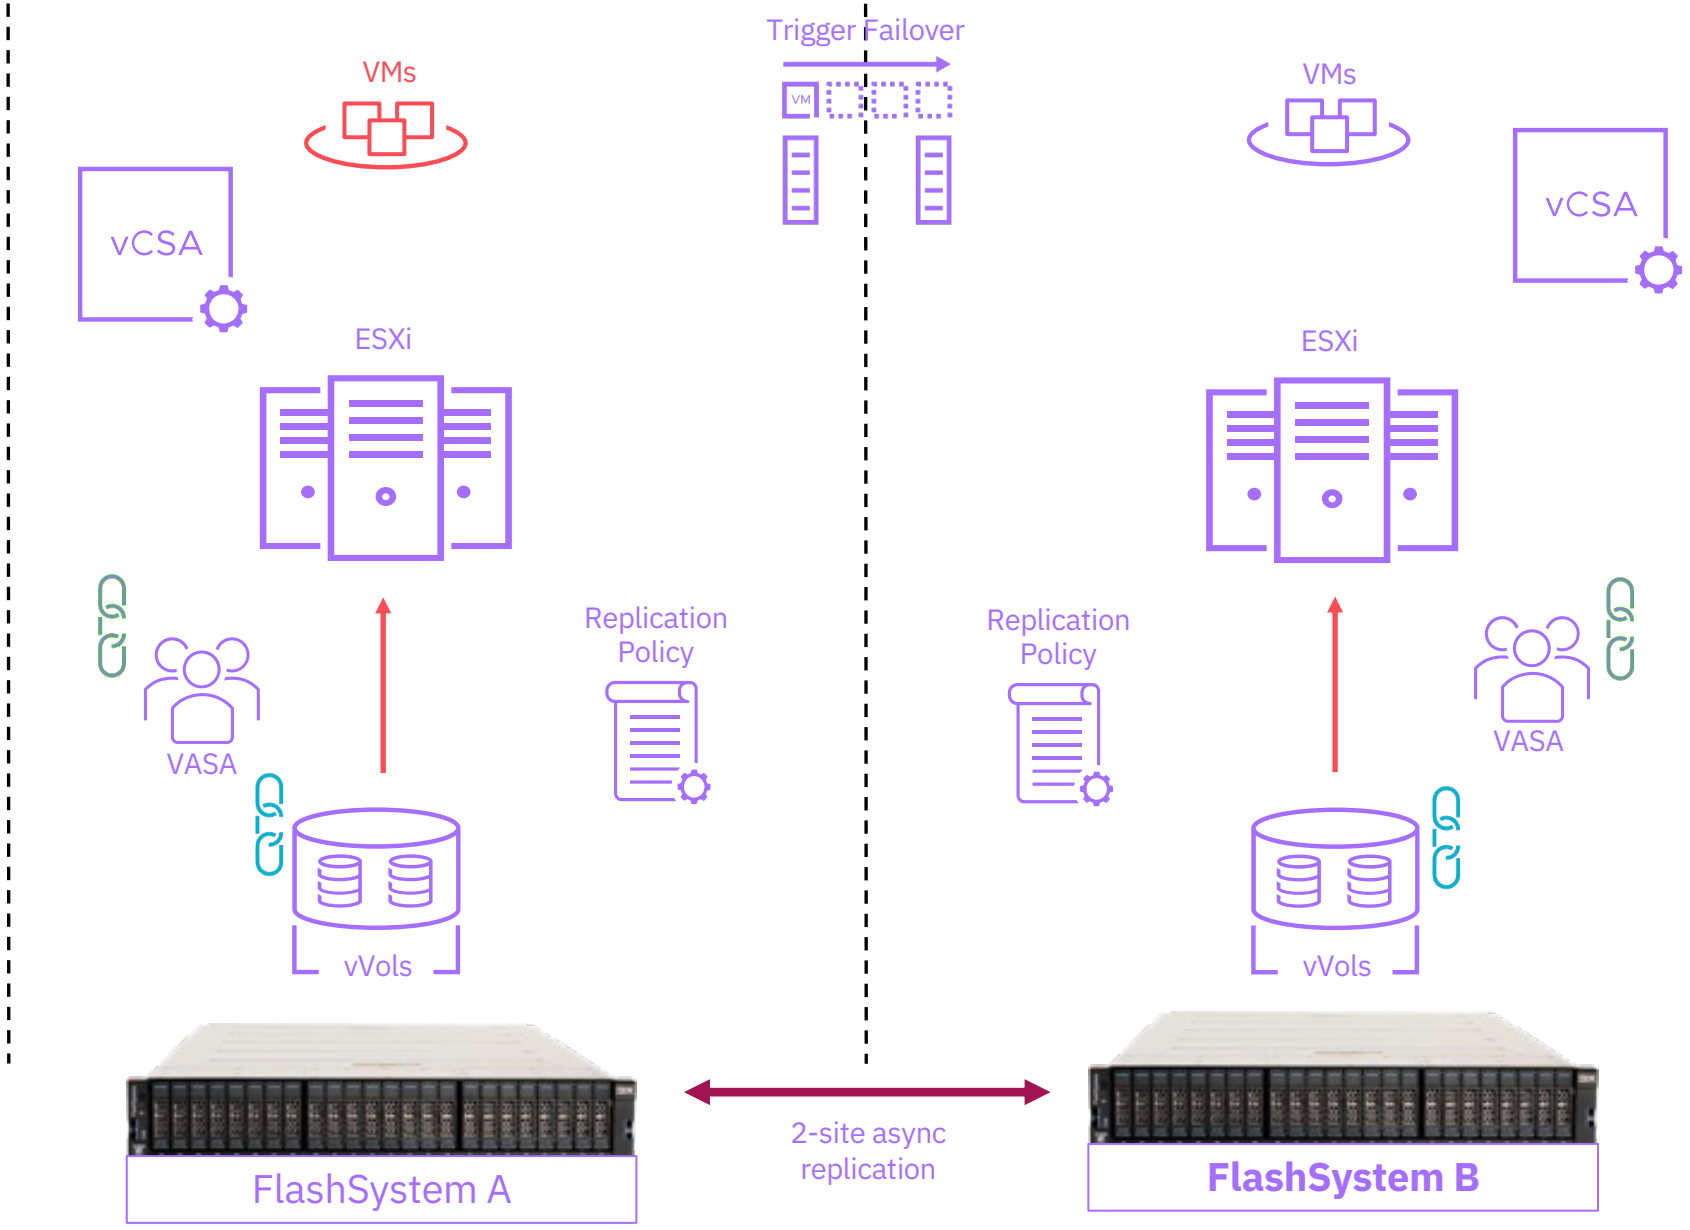

Replication can be reversed to reprotect VMs after a failover

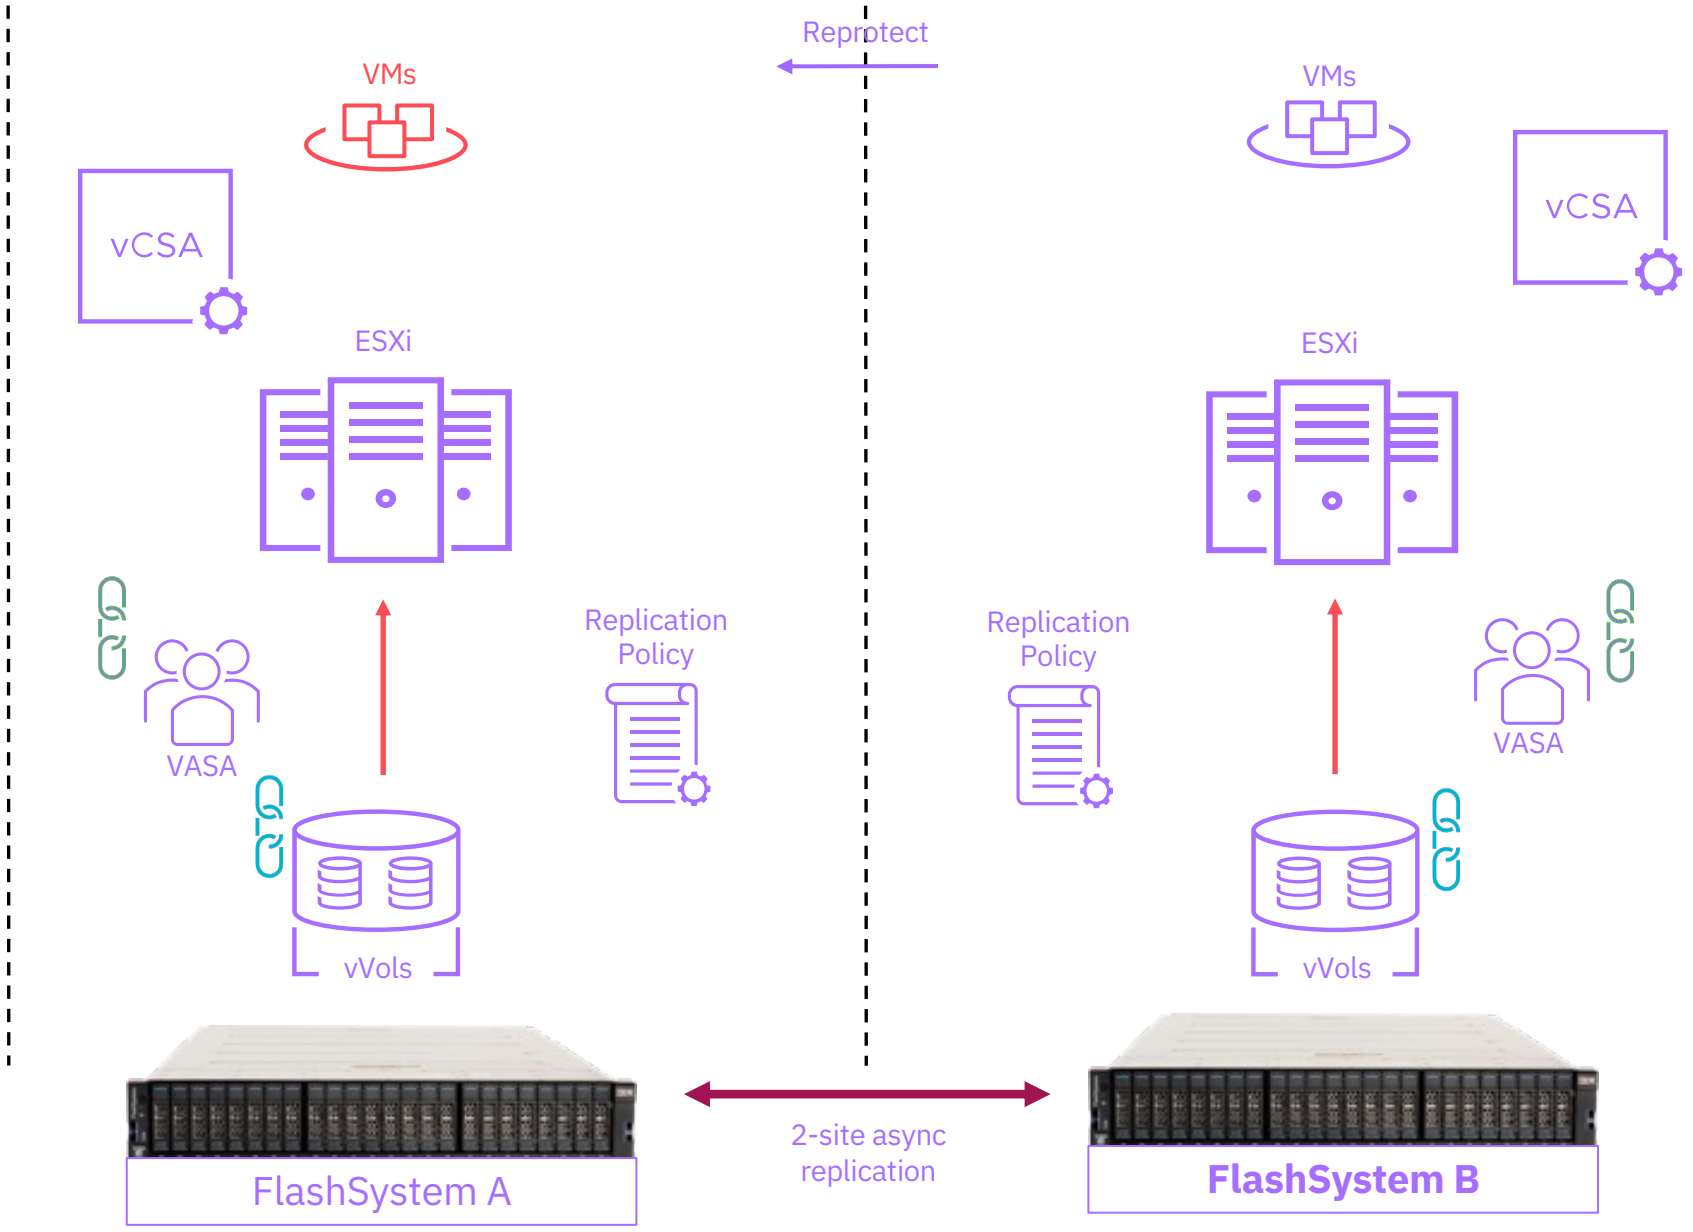

vVol Replication Agenda

Overview
Planning
Implementing
Monitoring
Managing

Planning for vVols replication

System limits:

- 1 partnership per system using vVol replication.
  - Additional partnerships can be configured, but only one can be used for vVol replication.
- Single I/O group systems
- Standard topology systems
- At least one standard storage pool is required. Data reduction pools are not supported.
- Each system must support vVols and asynchronous policy-based replication.
- The standard configuration guidelines for policy-based replication apply.
- A vCenter Server is required in each location.
- Use vCenter Server and ESXi versions that support vVols replication.
- Failover and disaster recovery operations are managed through vCenter Server integrations such as:
  - VMware Site Recovery Manager (SRM)
  - PowerCLI

Supported platforms for vVol Replication

Model	vVol replication
SAN Volume Controller	Supported
FlashSystem 9500	Supported
FlashSystem 9200	Supported
FlashSystem 91x0	Supported
FlashSystem 7300	Supported
FlashSystem 7200	Supported
FlashSystem 5200	Supported*
FlashSystem 50x5	Not supported
Storage Virtualize for Public Cloud	Not supported**

* Requires minimum of 128 GB of memory per canister
** Needs 8.7 update to be supported

vVol Replication Agenda

Overview
Planning
Implementing
Monitoring
Managing

Implementing

Enable vVol replication using the GUI

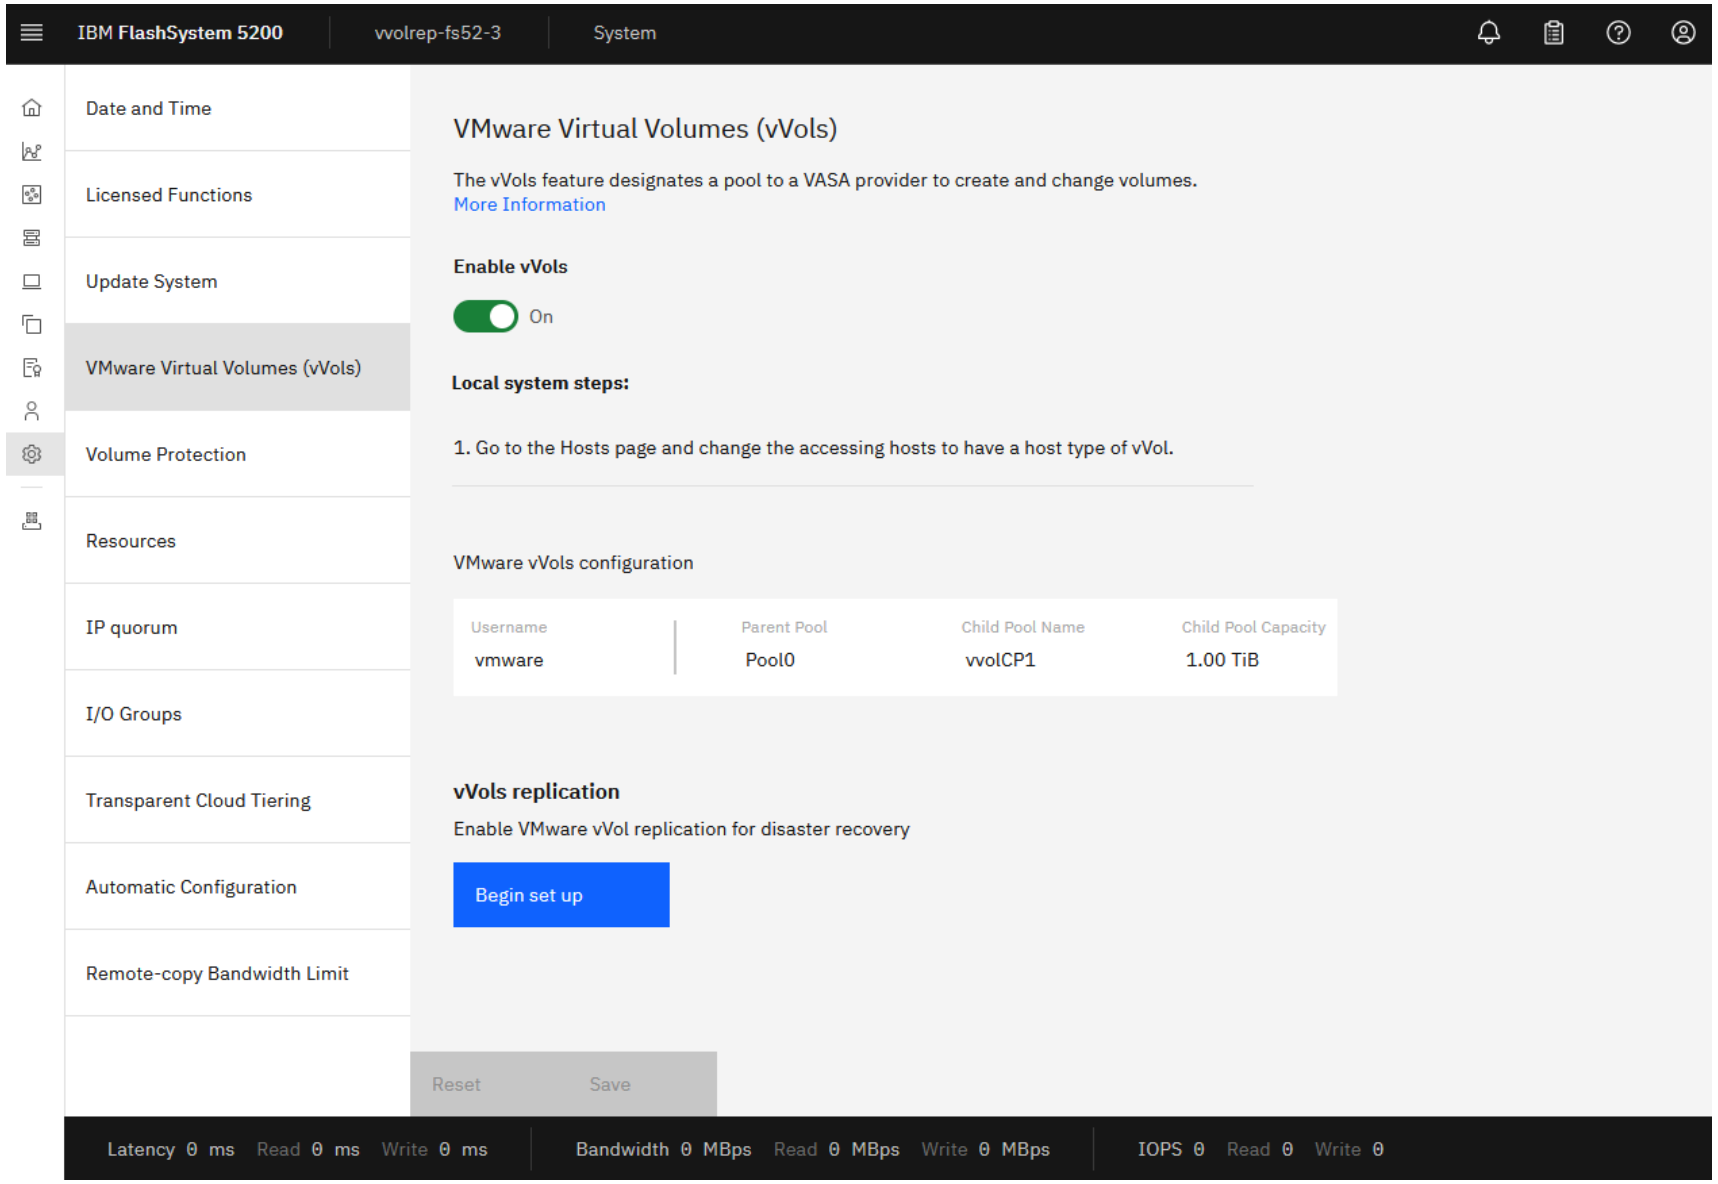

Implementing

Enable vVol replication using the GUI

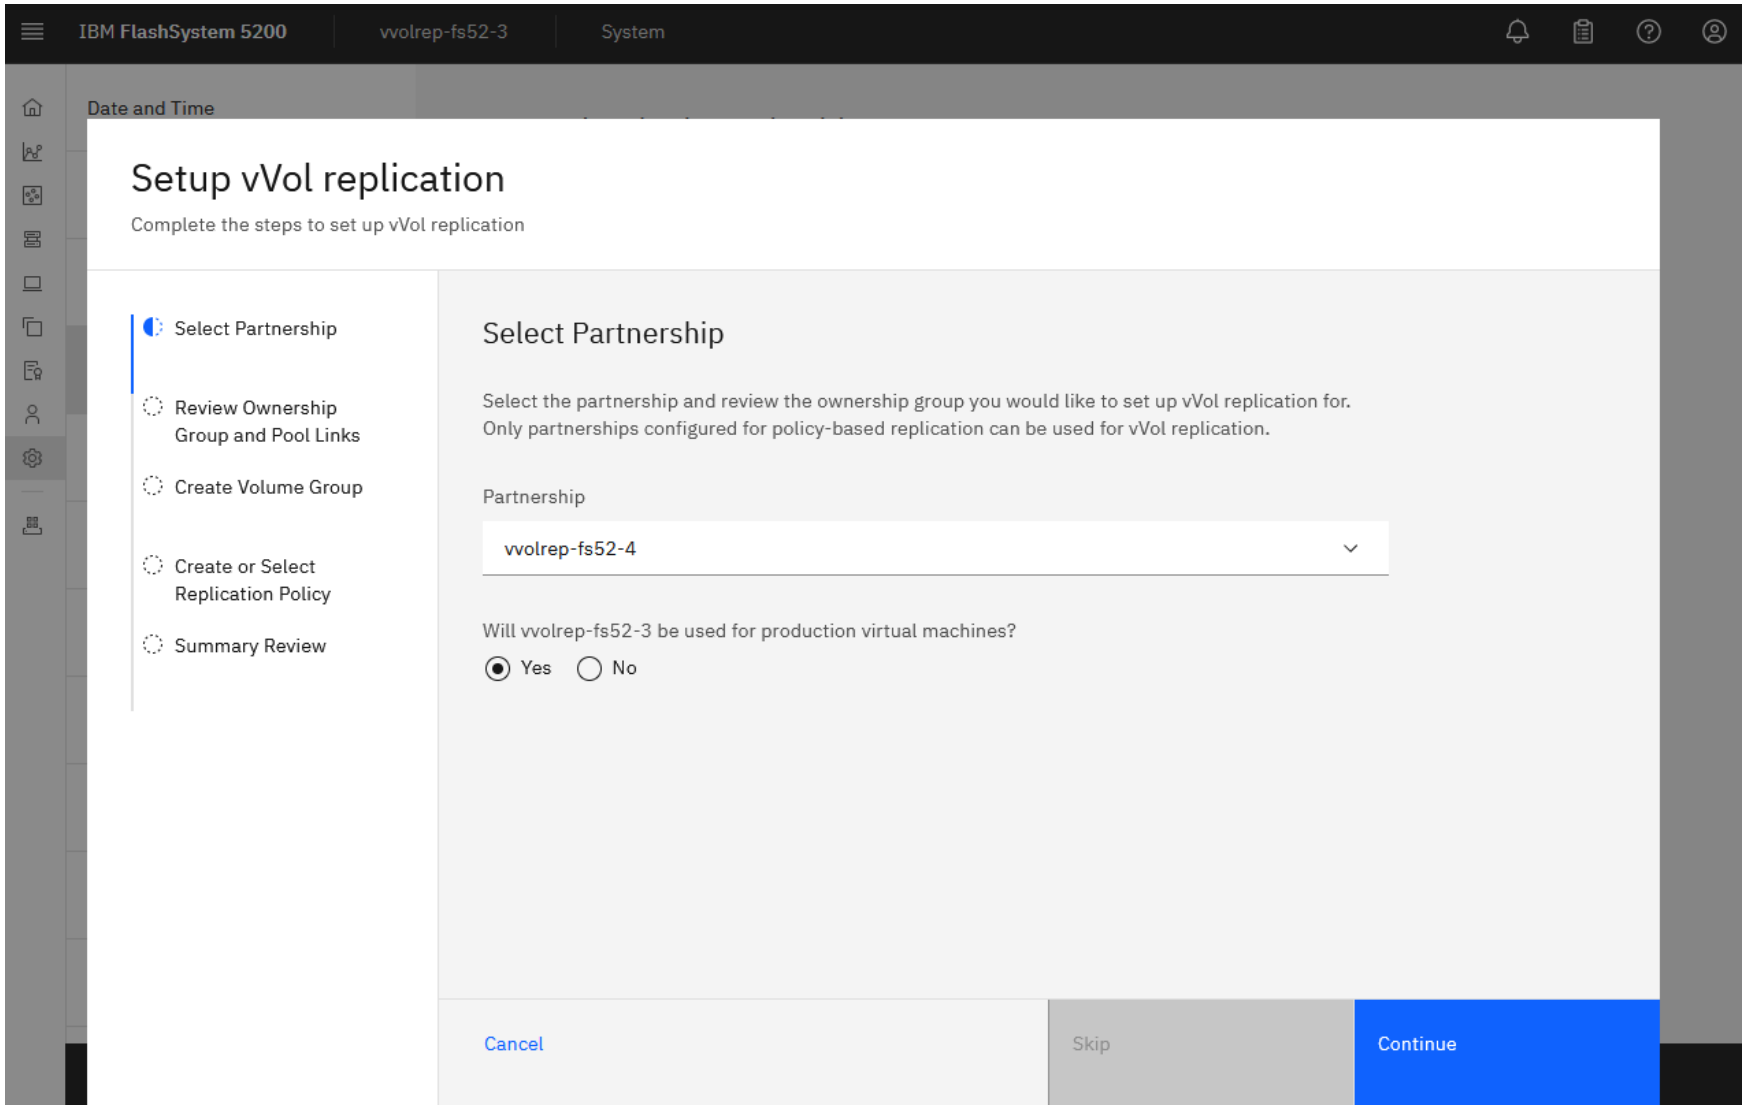

Implementing

Enable vVol replication using the GUI

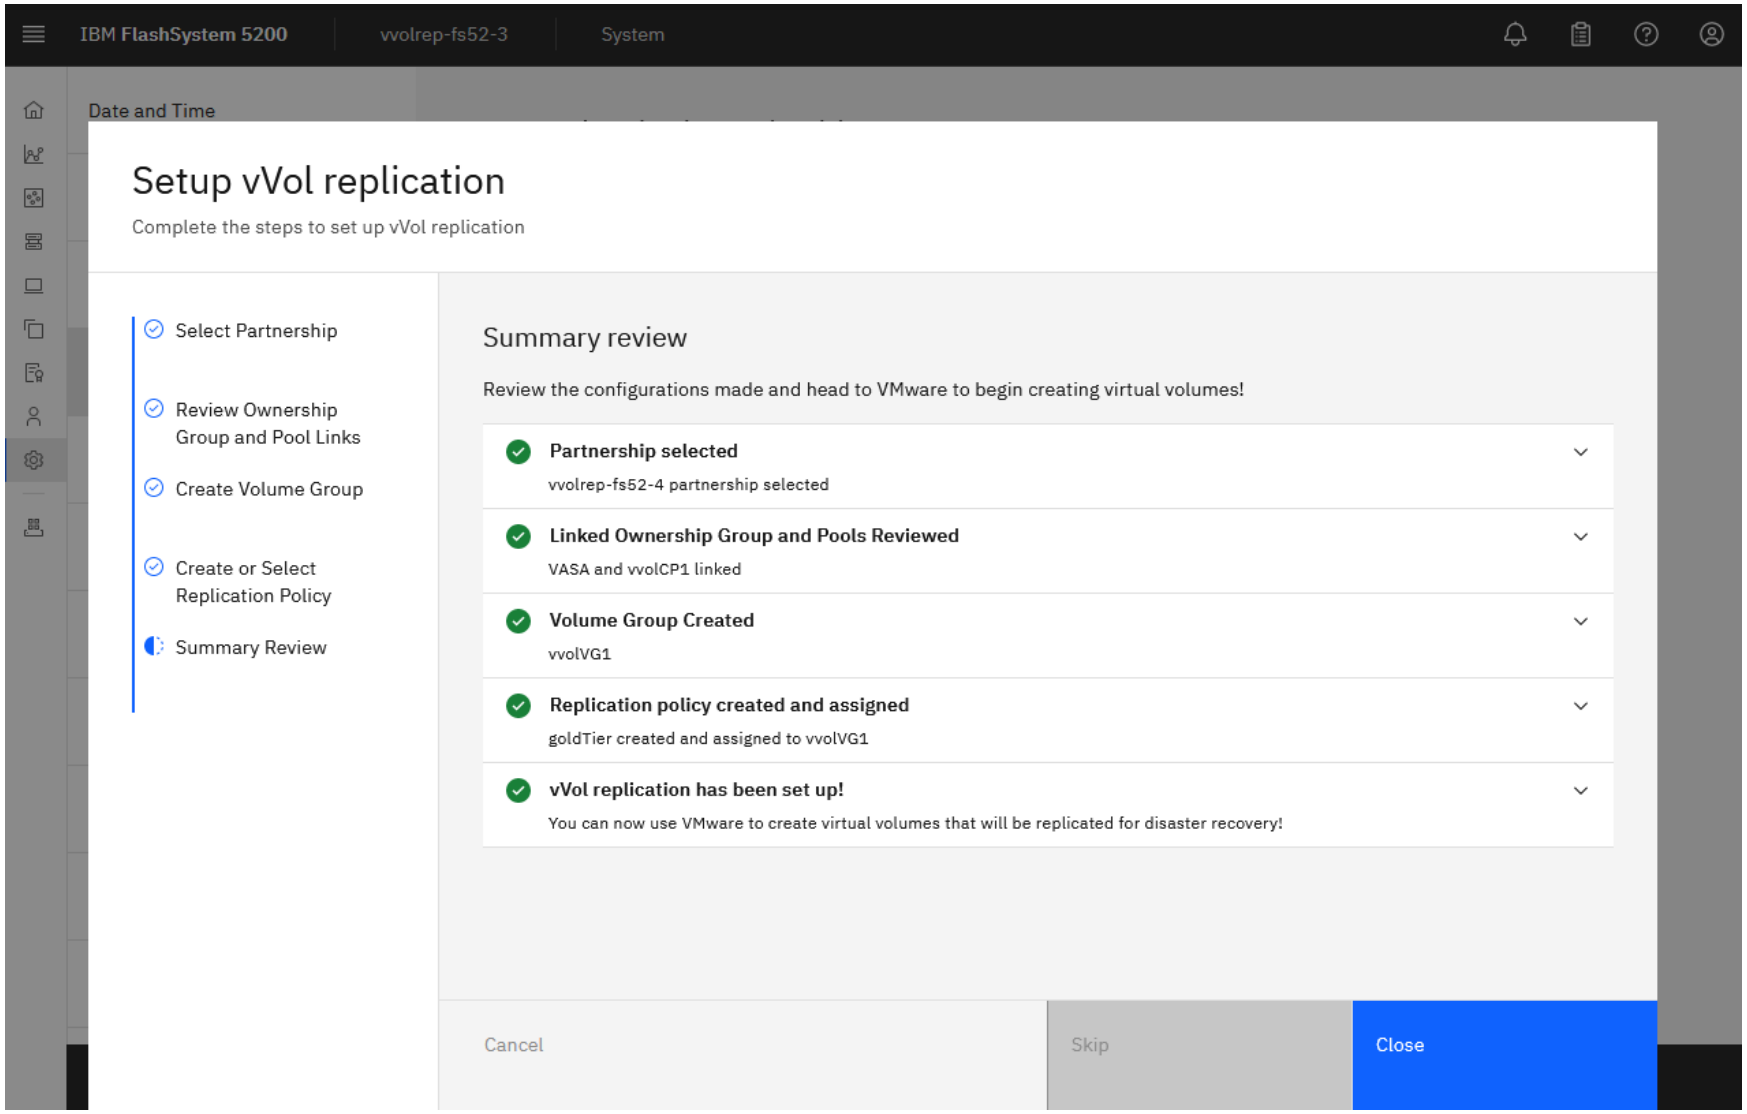

Implementing

Enable vVol replication using the GUI

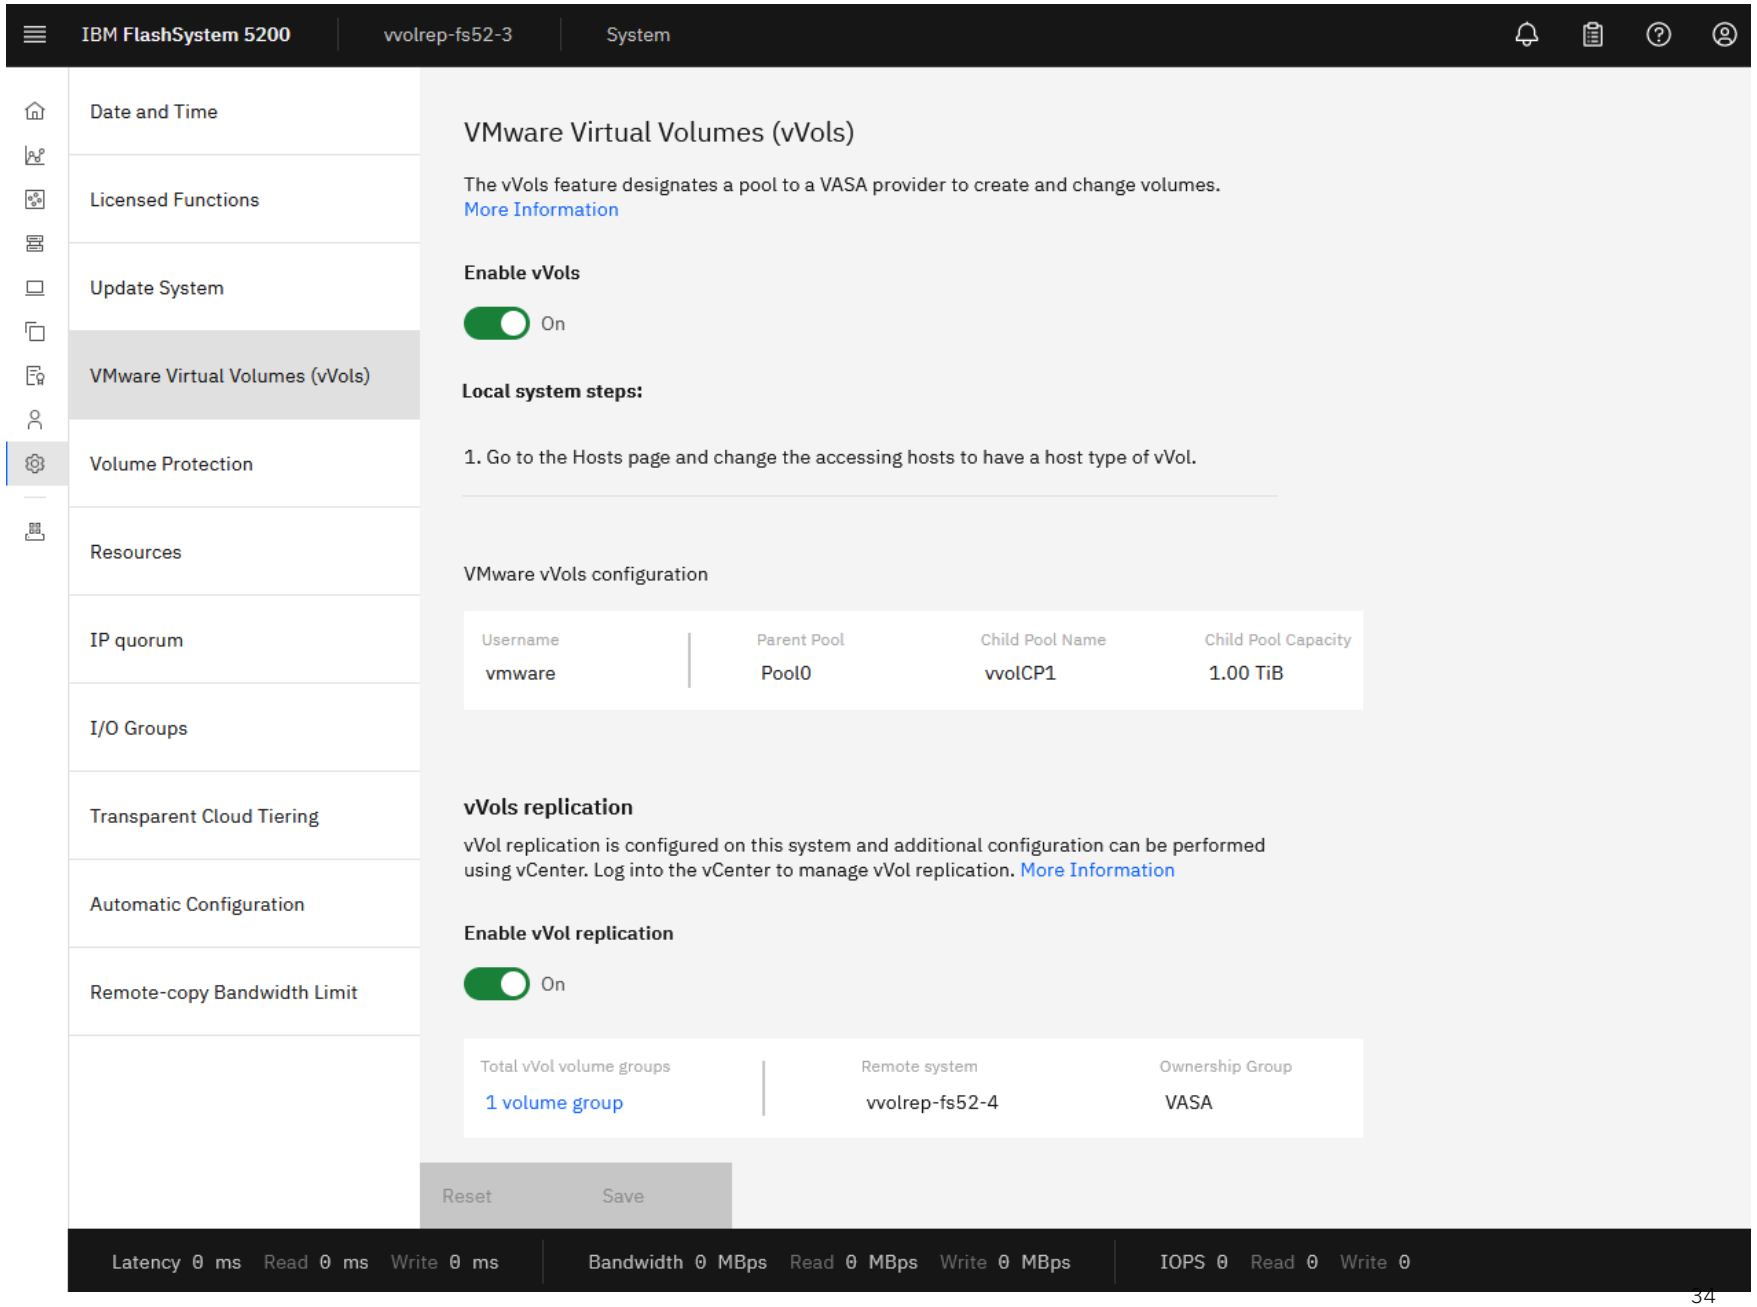

Implementing

Additional **Replication Groups** can be configured for use by VMware by creating volume groups with a replication policy

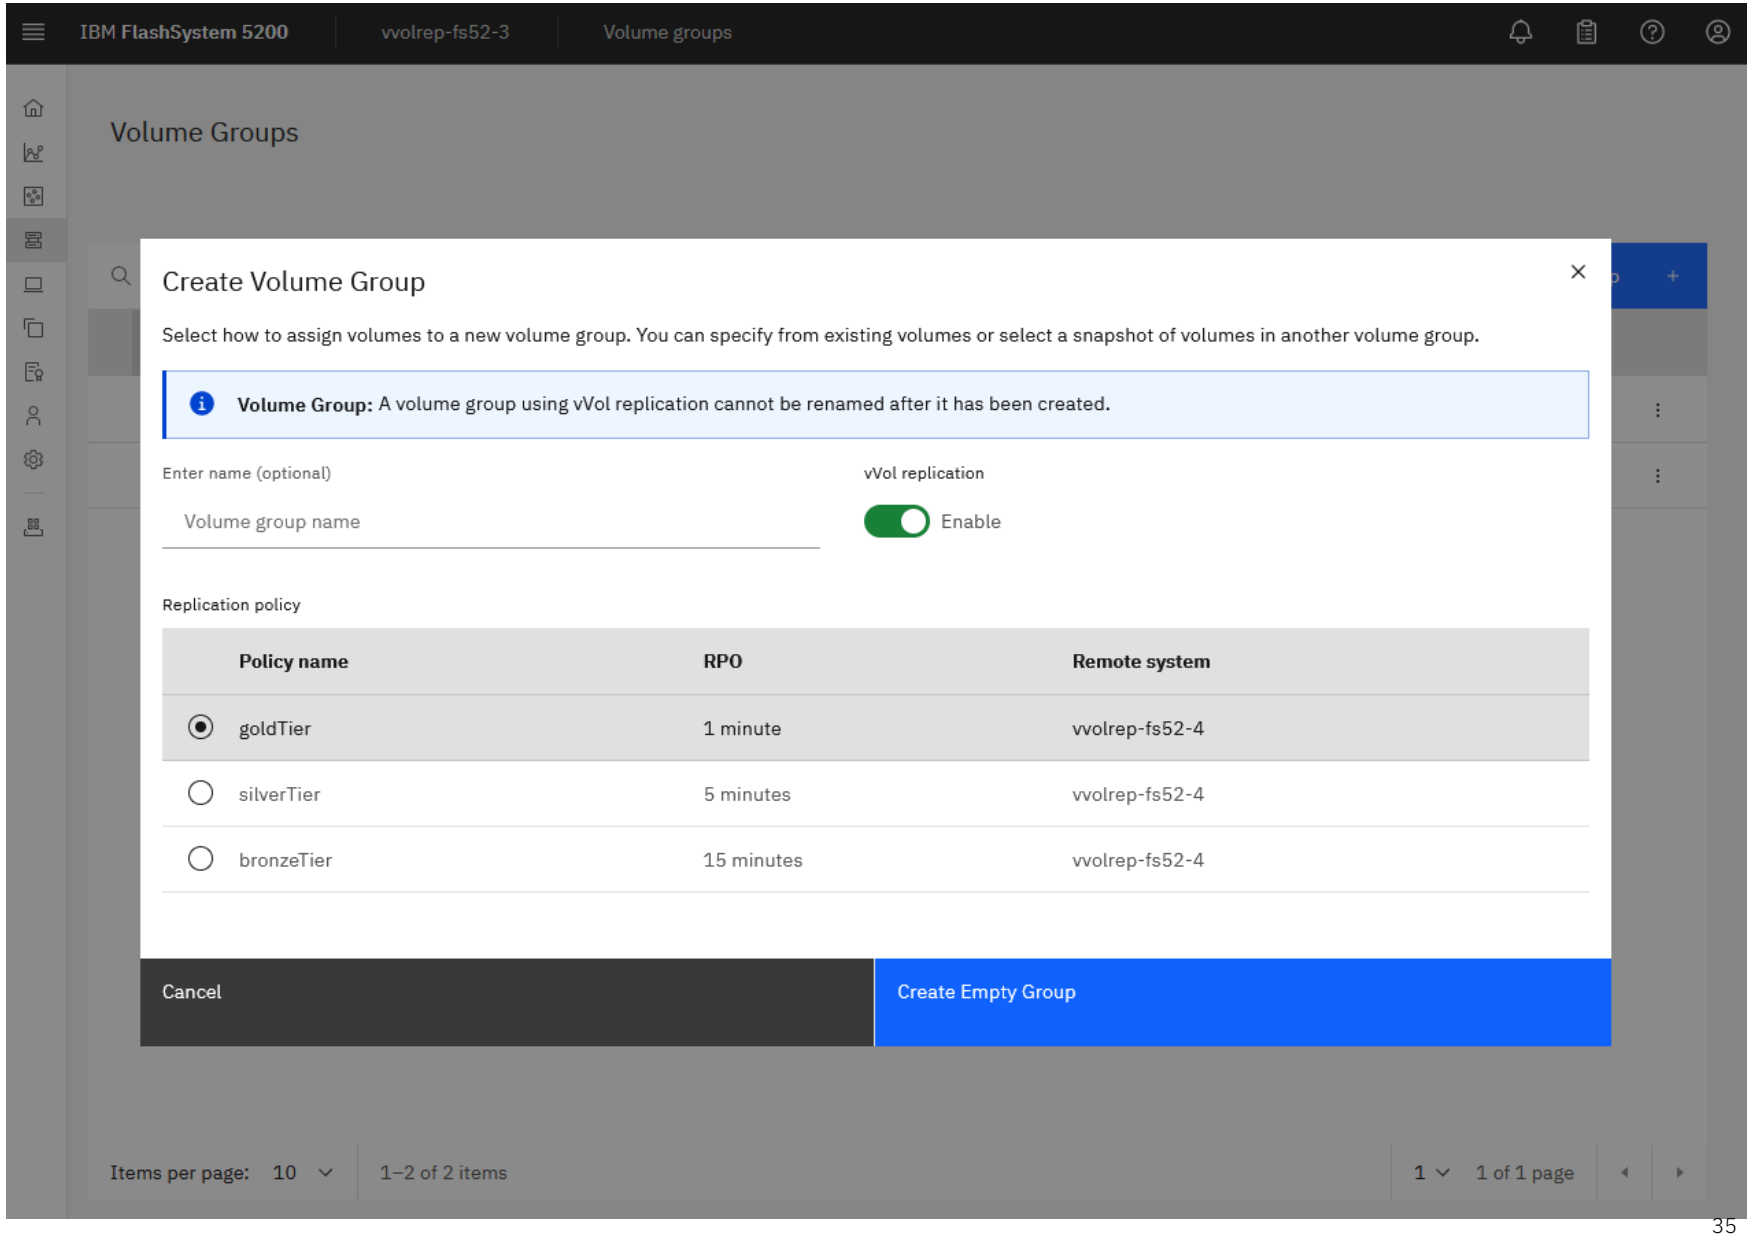

Implementing

Create VM **Storage Policies** in vSphere that specify the replication requirements

The remote system (fault domain) and RPO choices here reflect the capabilities of the vVol volume groups configured on the storage system

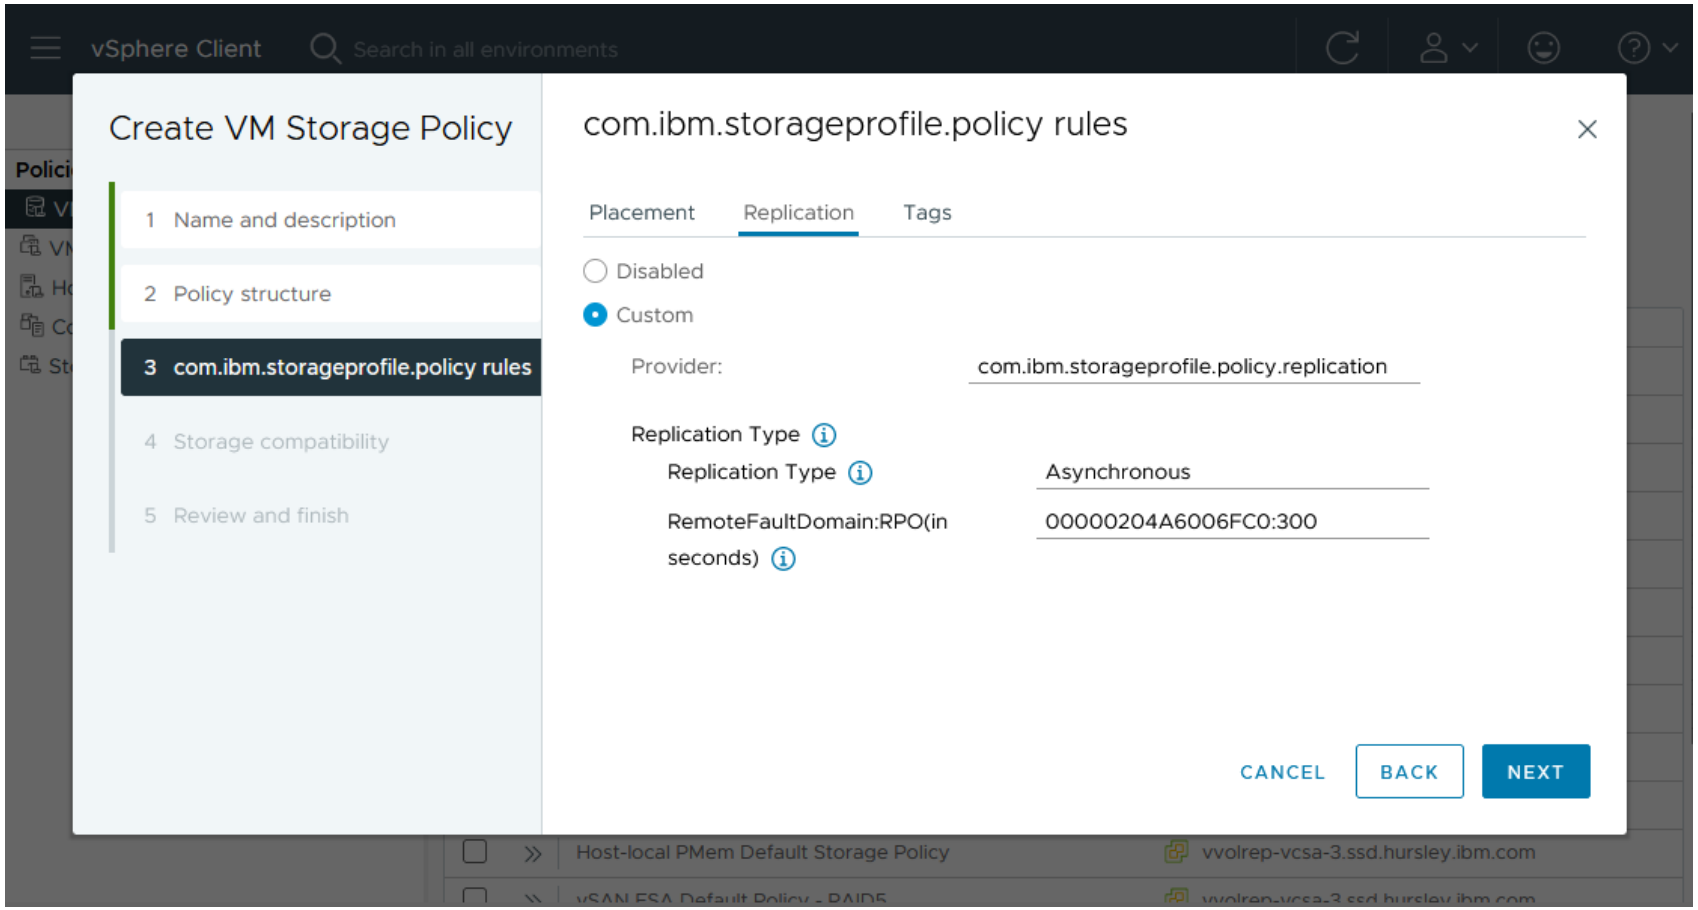

Implementing

Create new Virtual Machines in vSphere, selecting a Storage Policy that includes replication and a compatible vVol data store

The **Replication Group** choices here reflect the choice of compatible volume groups configured on the storage system

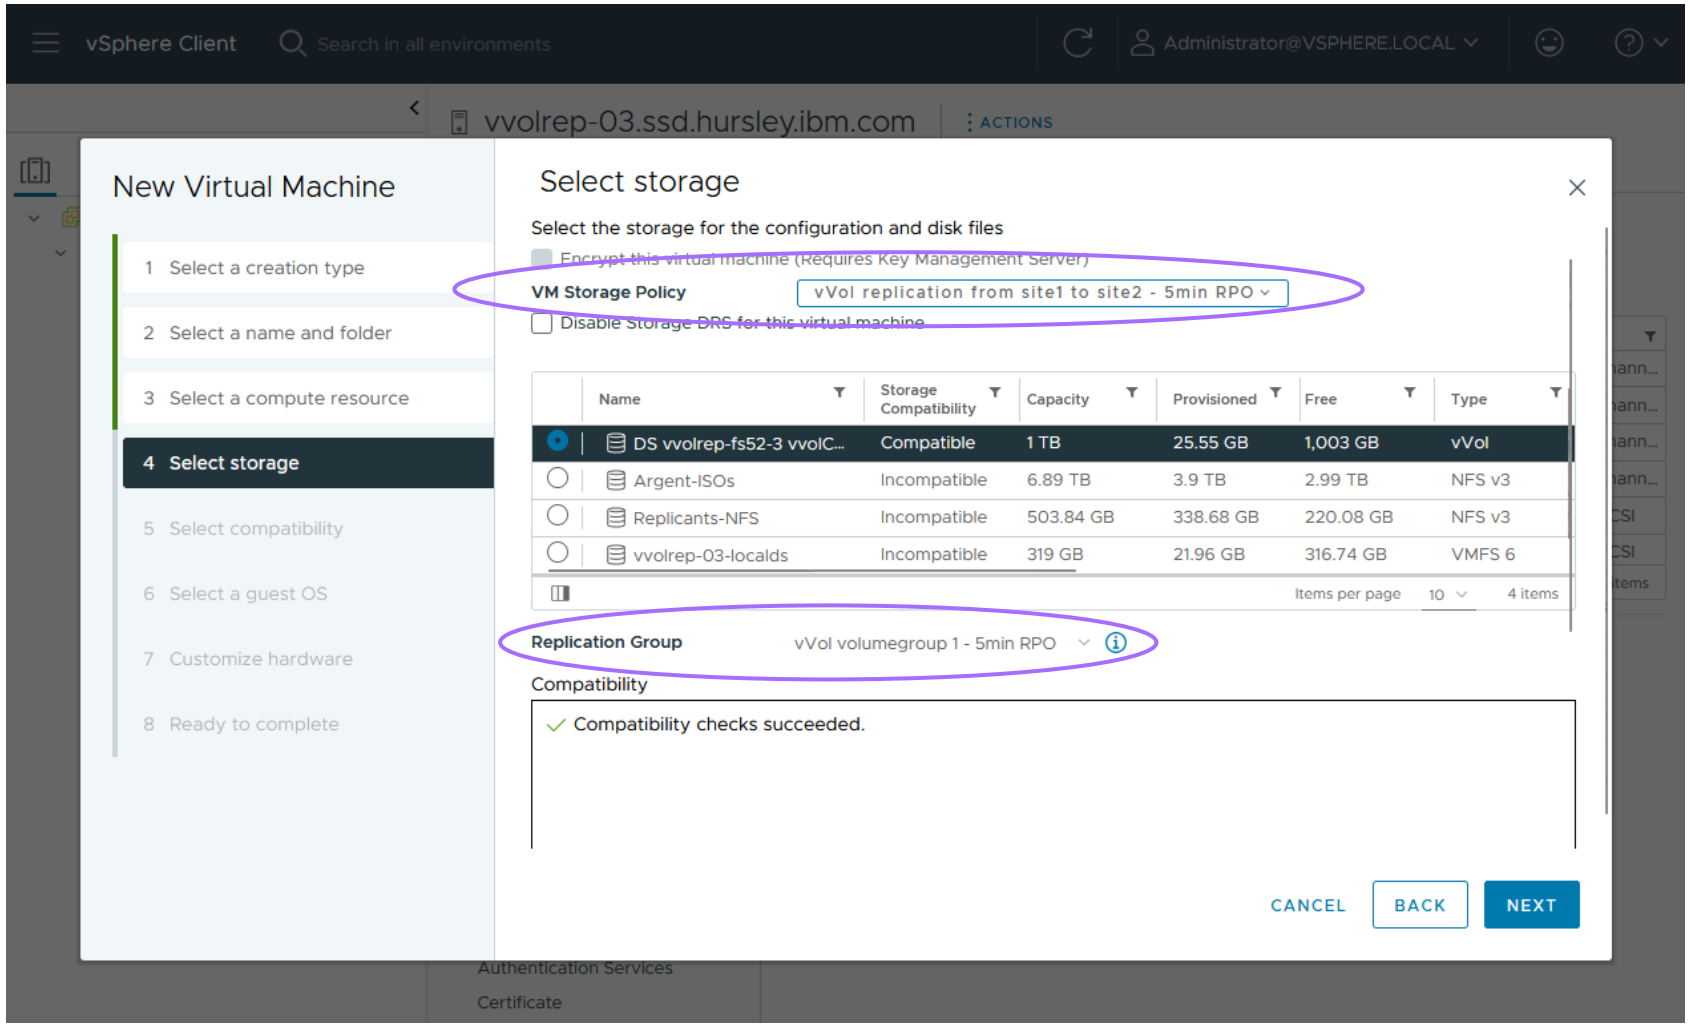

vVol Replication Agenda

Overview
Planning
Implementing
Monitoring
Managing

Monitoring

Replication status can be seen under

Volumes > Volume Groups

High-level summary of all volume groups

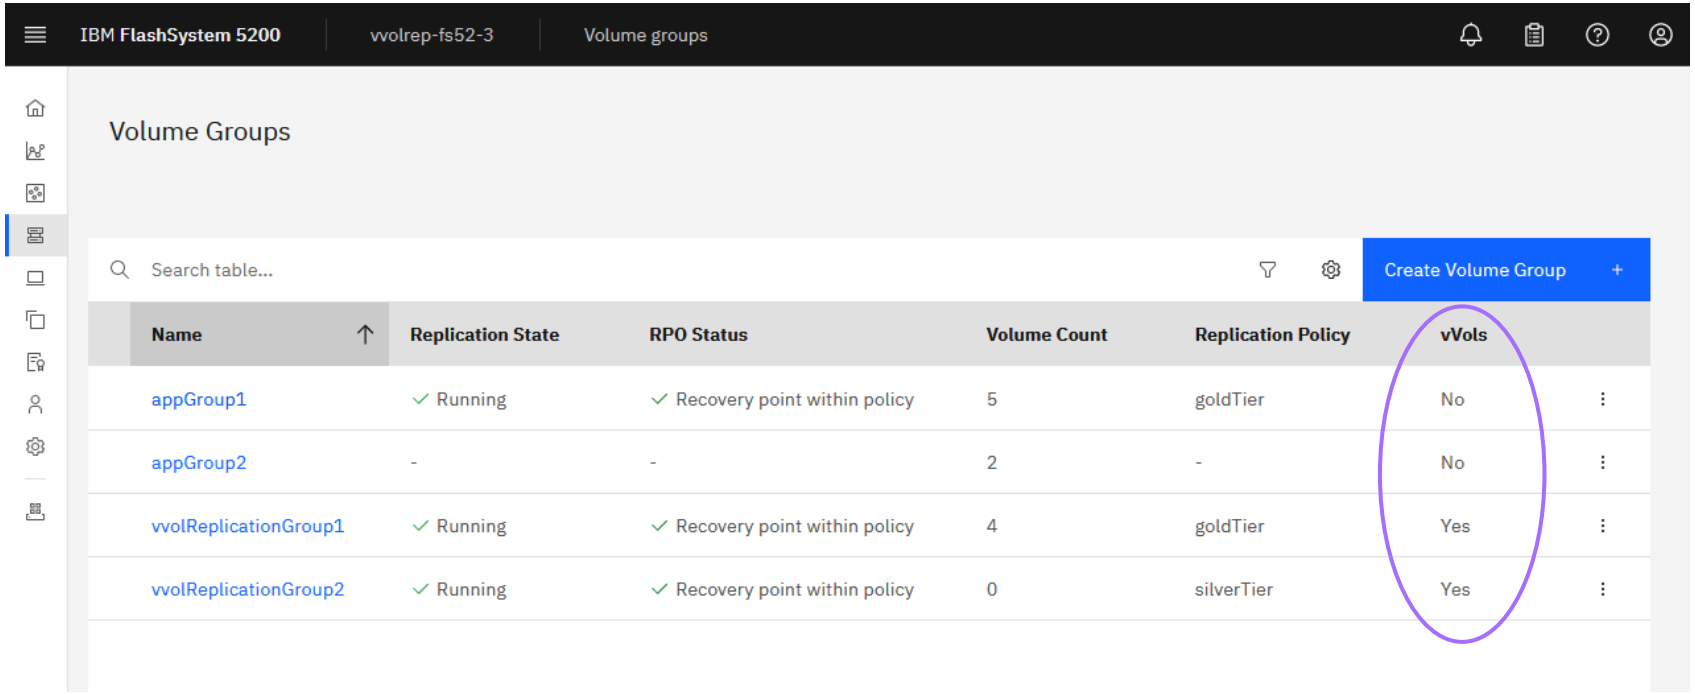

Monitoring

Replication status can be seen under

Volumes > Volume Groups

Table can be searched, sorted and filtered to easily identify groups of interest

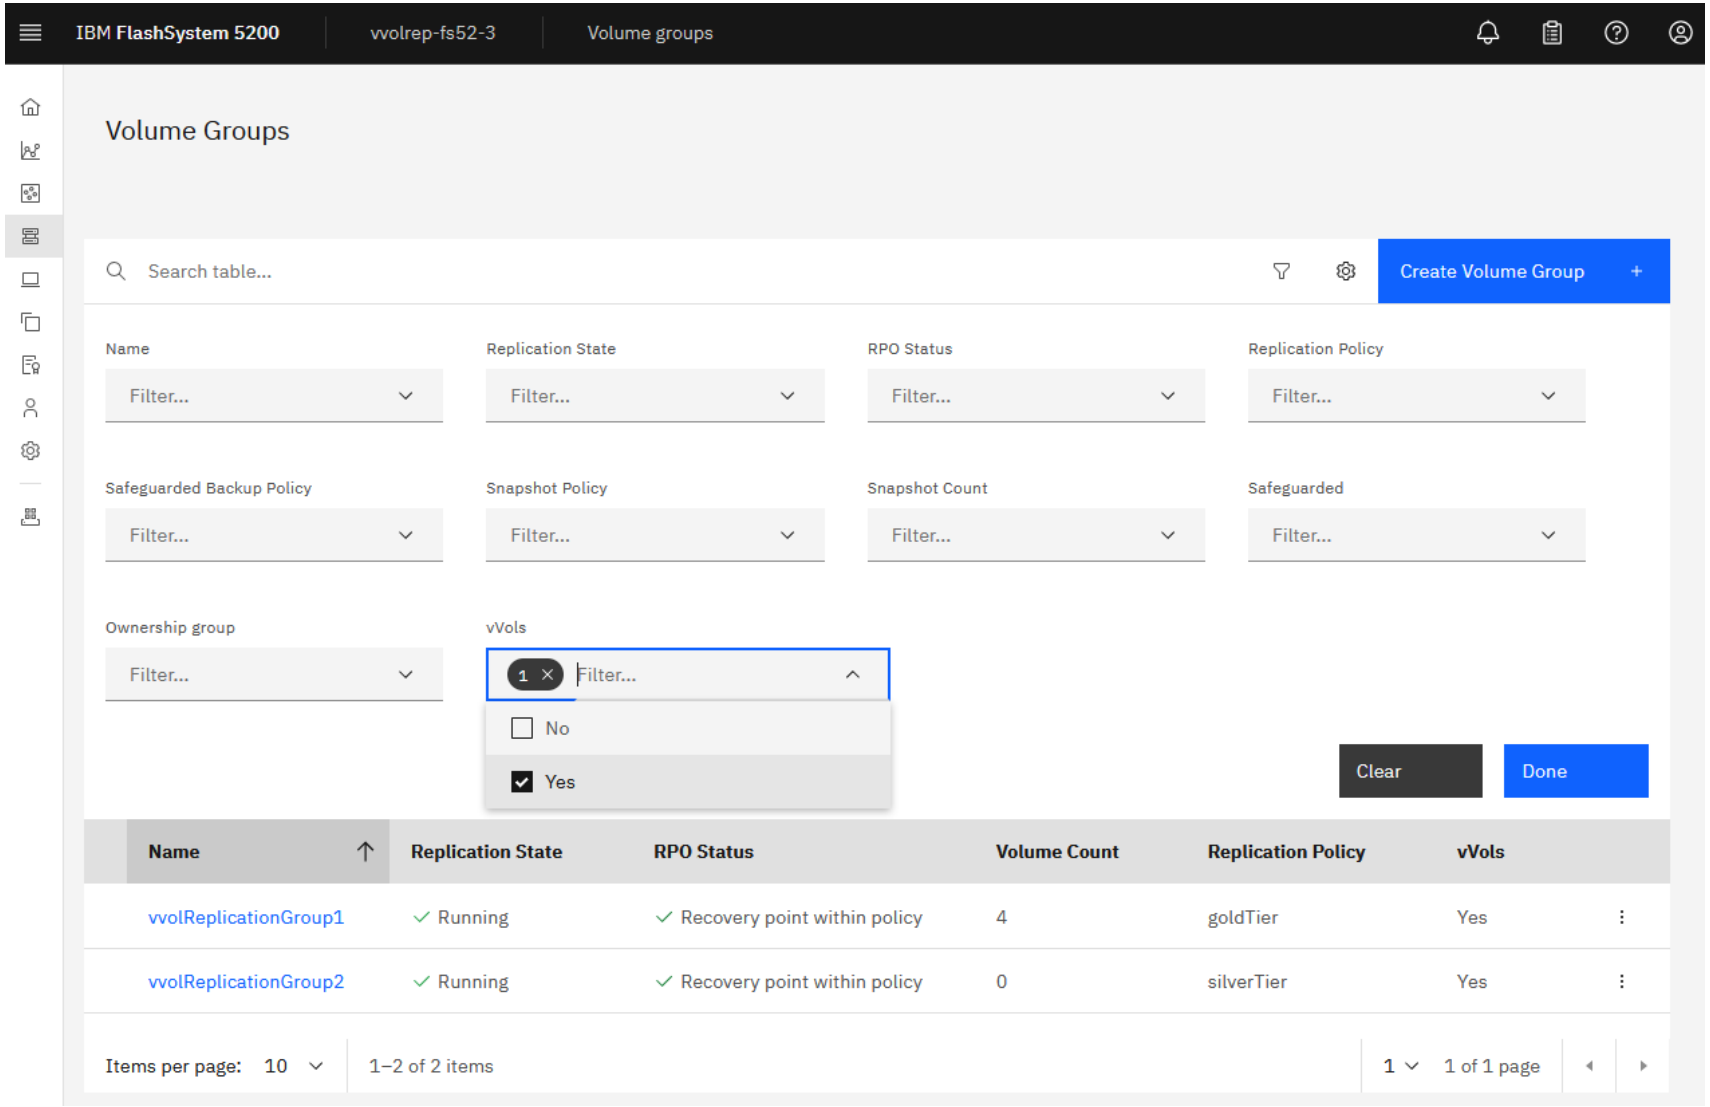

Monitoring

Replication status can be seen under

Volumes > Volume Groups

Selecting a volume group and selecting the **Policies** tab:

- displays detailed replication status for the group

- provides links for the actions for replication that can be performed on the group

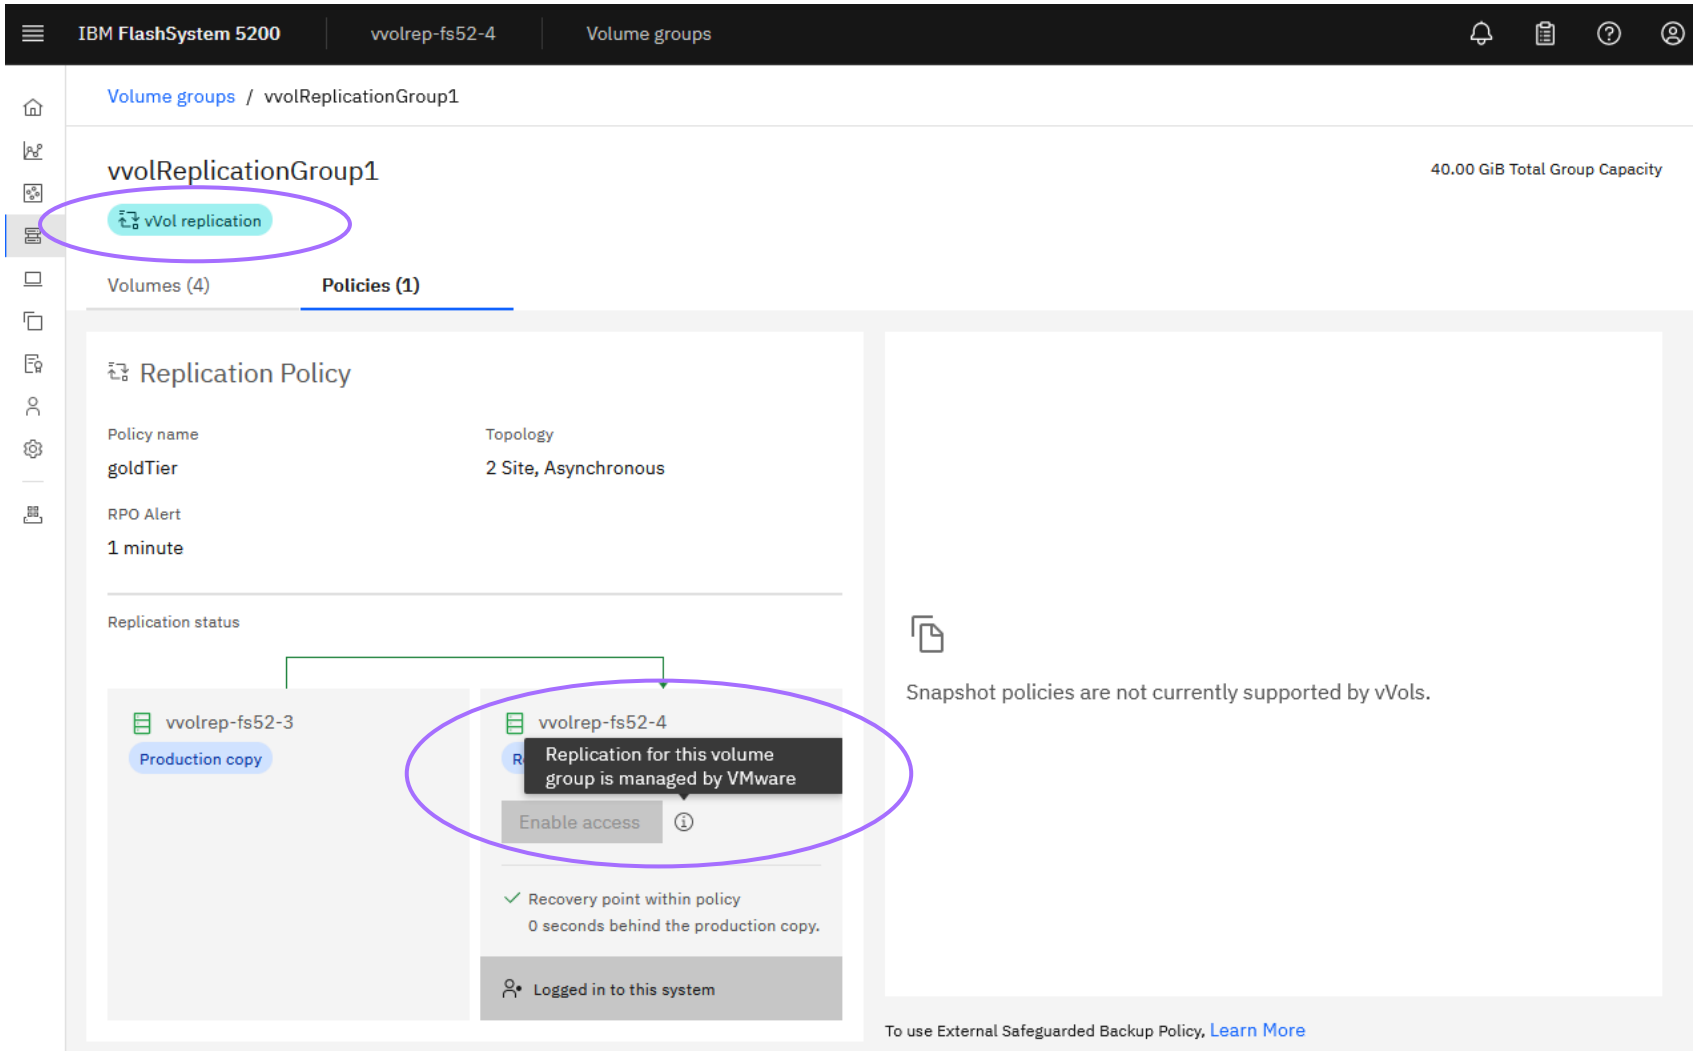

Monitoring

Storage Policy compliance can be seen in vSphere

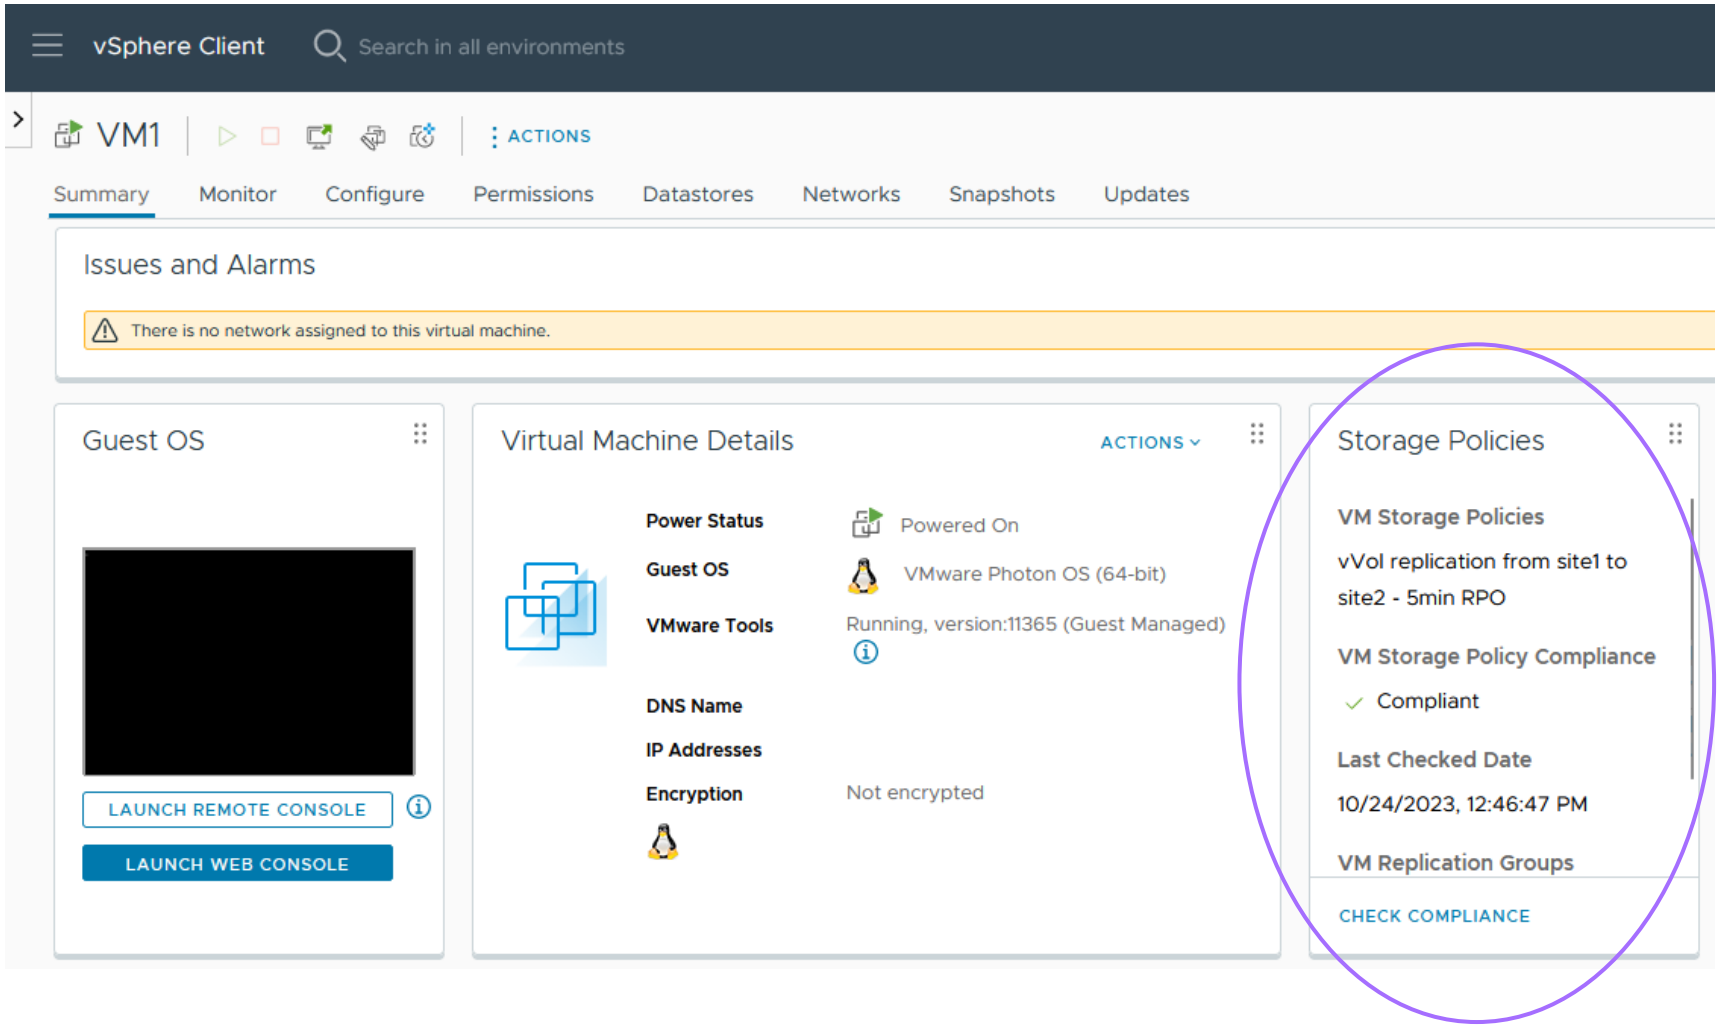

vVol Replication Agenda

Overview
Planning
Implementing
Monitoring
Managing

VMware can perform three types of failovers:

- Planned failover

- Unplanned Failover

- Test Failover

Movement of a VM from one site to another for a planned migration.
The systems at both sites are required to be accessible throughout the failover process.
Failover of a VM without data loss (RPO=0).
Once a planned failover is complete, replication can optionally be reversed so that the failed-over VM can be reprotected.

VMware can perform three types of failovers:

- Planned failover

- Unplanned Failover

- Test Failover

Movement of a VM in response to a failure in the production site.
Only the recovery system is required to be accessible.
If the original production site recovers after the failover, replication can optionally be reversed so that the failed-over VM can be reprotected.
Alternatively, a new storage policy with alternative replications settings can be configured.

VMware can perform three types of failovers:

- Planned failover

- Unplanned Failover

- Test Failover

Allows the recovery copy of a VM to be brought up for testing, without taking down the production VM.
Test failover enables temporary access to the recovery environment to allow a failover plan to be verified before an actual disaster or planned migration.
When the test failover is stopped, any changes made to the recovery copy during the test are discarded.

Example planned failover using VMware Site Recovery Manager

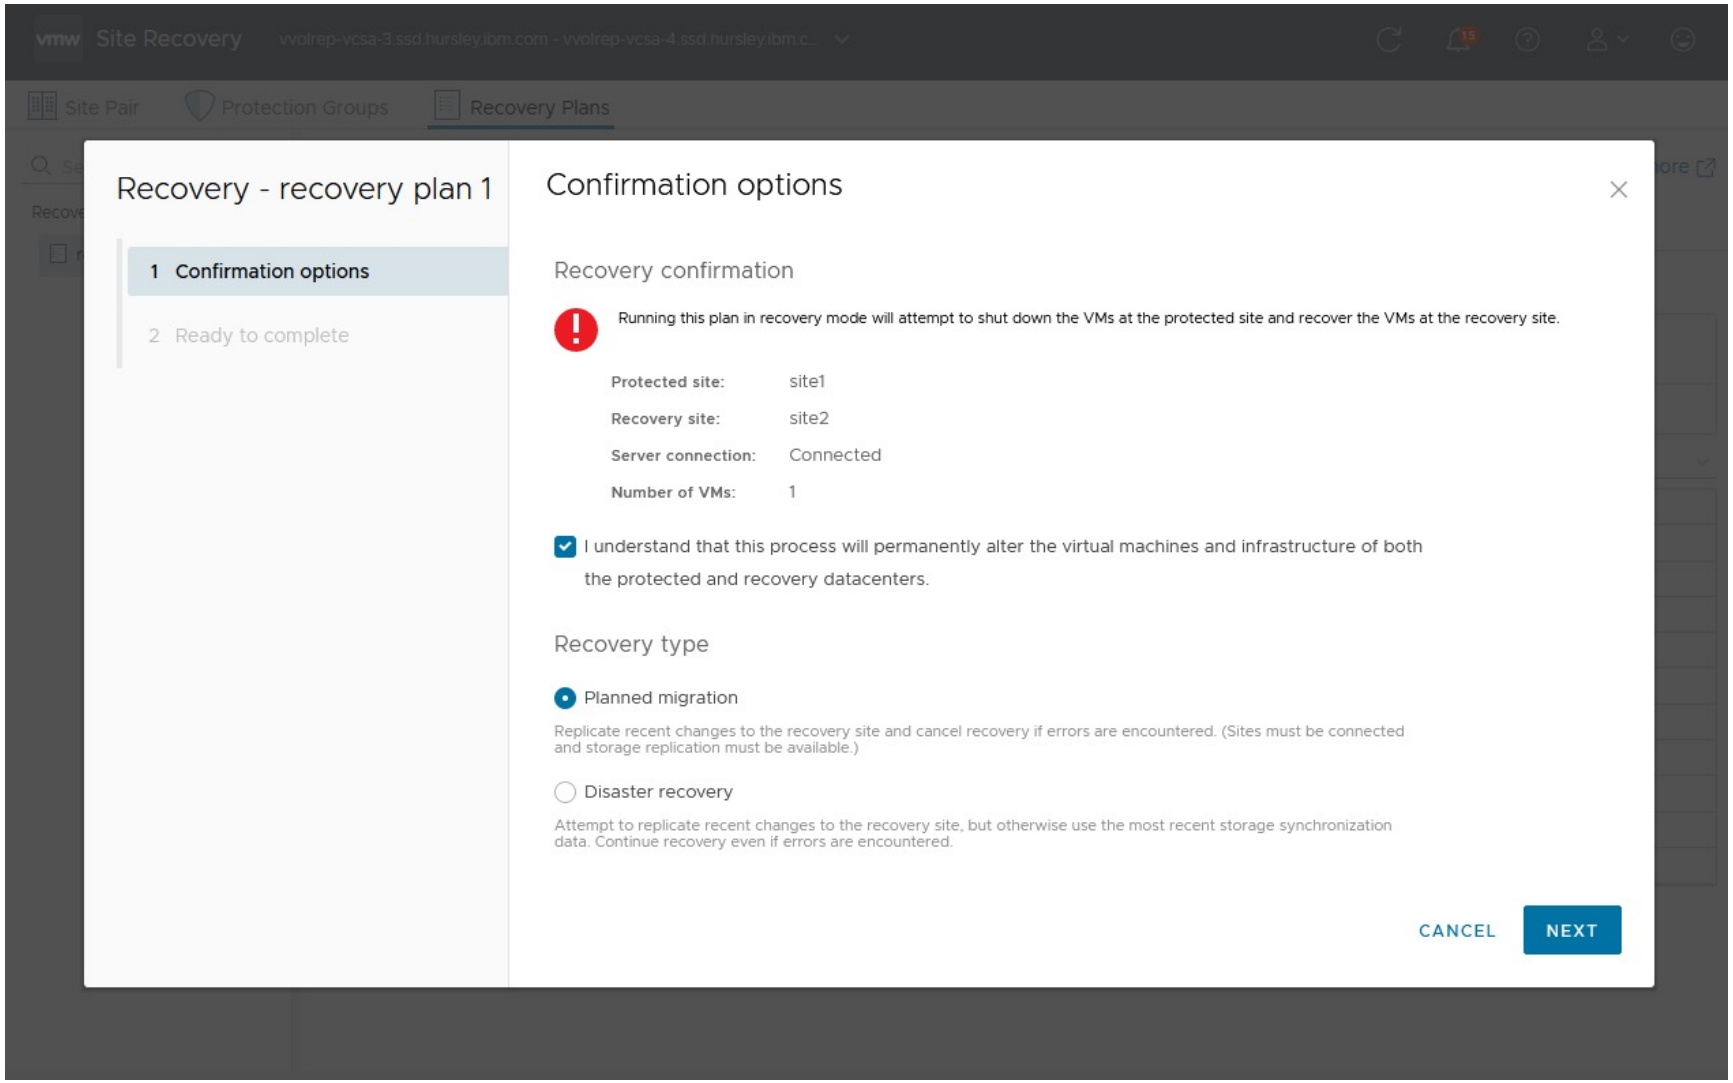

Example planned failover using VMware Site Recovery Manager

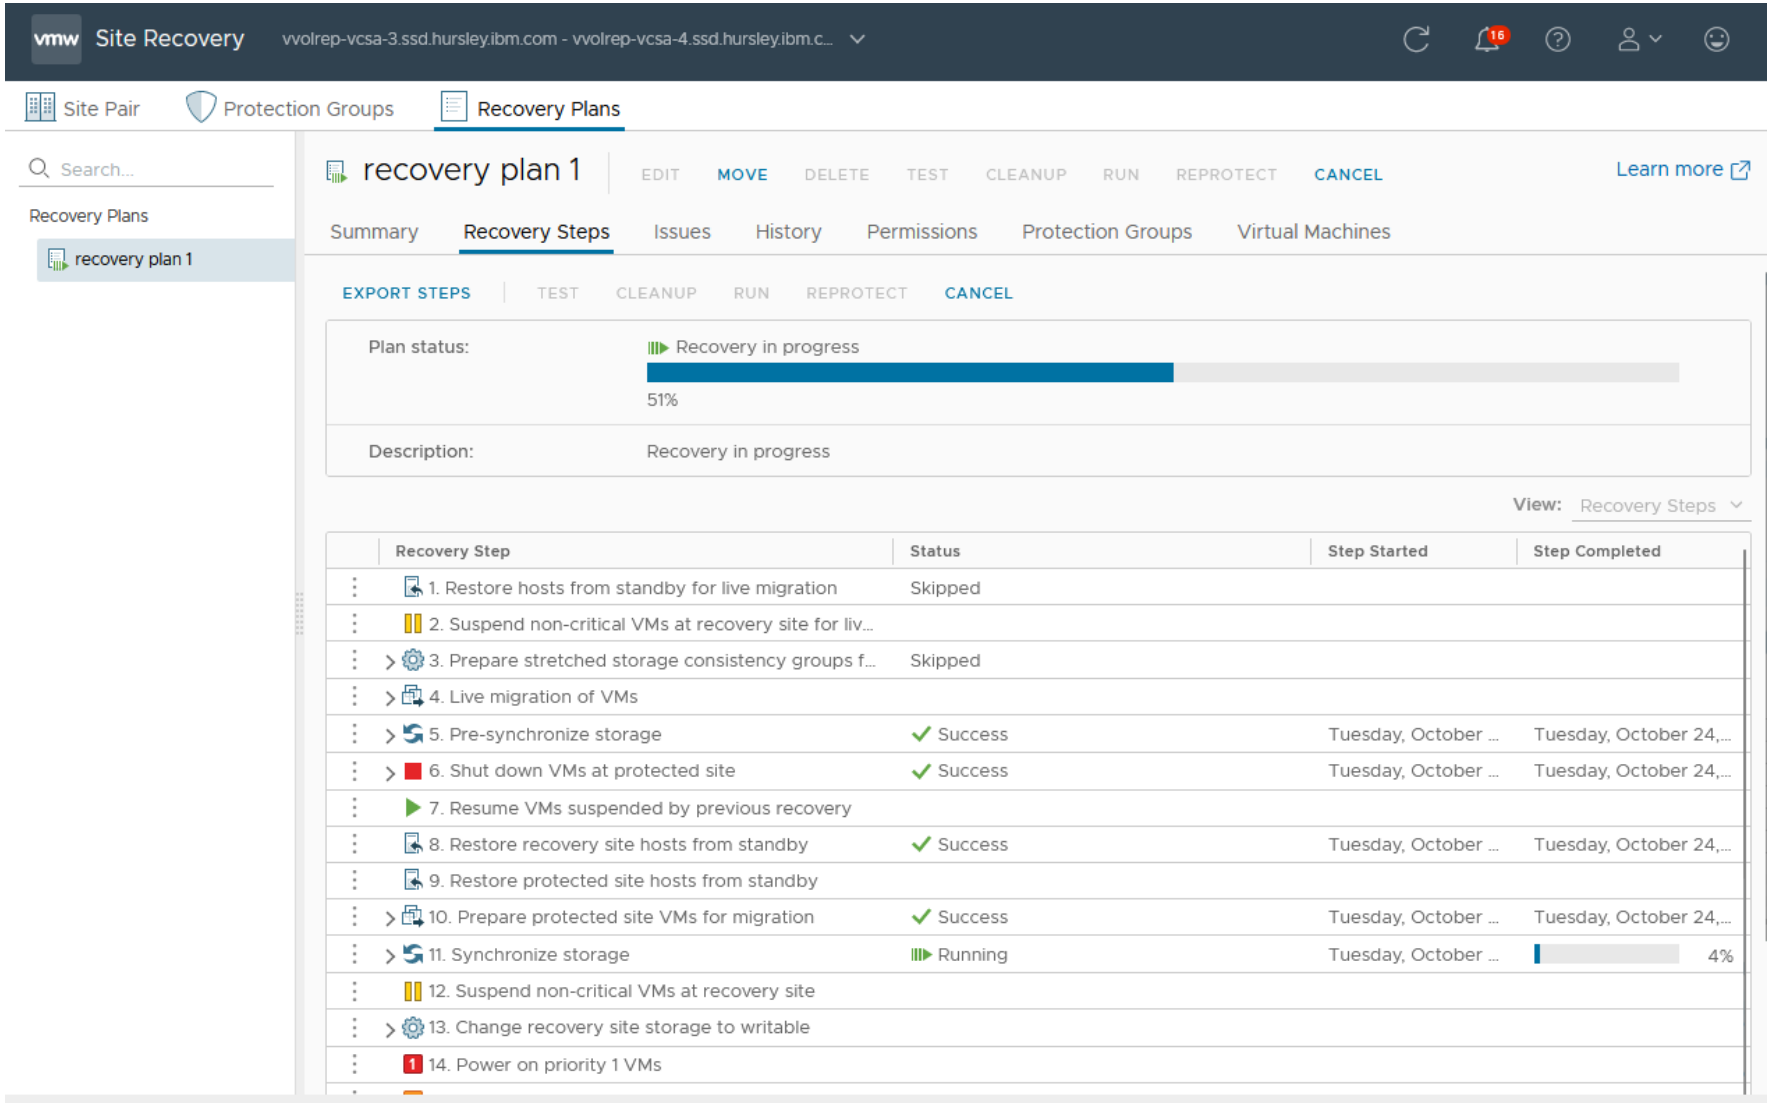

Example using PowerCLI

Preparing for a Planned Failover

```
PS /Users/warren> Get-Cluster Cluster1 | Get-VM -server vvolrep-vcsa-1 | Format-Table -AutoSize

Name                                PowerState Num CPUs MemoryGB
----                                ---------- -------- --------
bulk-TestSPBMPolicy-myvVolVG-60-1-1-02 PoweredOn  1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-05 PoweredOn  1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-01 PoweredOn  1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-03 PoweredOn  1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-04 PoweredOn  1        2.000
vvolrep-11-1                        PoweredOn  2        8.000
vvolrep-12-1                        PoweredOff 2        8.000

PS /Users/warren> Get-Cluster Cluster2 | Get-VM -server vvolrep-vcsa-2 | Format-Table -AutoSize

Name                                PowerState Num CPUs MemoryGB
----                                ---------- -------- --------
vvolrep-21-1                        PoweredOff 2        6.000
vvolrep-22-2                        PoweredOff 2        8.000

PS /Users/warren> $ReplicationGroups = Get-SpbmReplicationGroup

PS /Users/warren> $ReplicationGroups | Format-Table -AutoSize

Name                                     ReplicationState
----                                     ----------------
myvVolVG-60-1                            Source
rfc4122.d150a8dc-fb02-44e7-a6cb-acbf9c4620f3 Target


PS /Users/warren> $ReplicationGroupSource =  Get-SpbmReplicationGroup | Where-Object {$_.State -eq "Source"}

PS /Users/warren> $ReplicationGroupPair = Get-SpbmReplicationPair -Source $ReplicationGroupSource

PS /Users/warren> $ReplicationGroupTarget = $ReplicationGroupPair.Target

PS /Users/warren> Sync-SpbmReplicationGroup -PointInTimeReplicaName "Checkpoint 1" -ReplicationGroup $ReplicationGroupTarget

Name                    CreationTime              ReplicationGroup
----                    ------------              ----------------
rfc4122.d150a8dc-fb02-44… 25/10/2023 12:42:06       rfc4122.d150a8dc-fb02-44…


PS /Users/warren> Start-SpbmReplicationPrepareFailover $ReplicationGroupSource
```

"bulk" VMs exist in vCSA-1

No "bulk" VMs exist in vCSA-2

Synchronize IO checkpoint

Prepare Failover

Example using PowerCLI

Triggering Planned Failover

```
PS /Users/warren> $VMs = Start-SpbmReplicationFailover -ReplicationGroup $ReplicationGroupTarget

Confirm
Are you sure you want to perform this action?
Performing the operation "Starting failover on" on target "Replication group 'rfc4122.d150a8dc-fb02-44e7-a6cb-acbf9c4620f3'".
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help (default is "Y"): a

PS /Users/warren> $VMs
[vVolCP1-vvolrep-fs52-2] rfc4122.07e8f9be-4fe5-4083-80ad-ca29abe87eed/bulk-TestSPBMPolicy-myvVolVG-60-1-1-01.vmx
[vVolCP1-vvolrep-fs52-2] rfc4122.46dae322-a2b4-4ab5-a4c1-3789df3487c2/bulk-TestSPBMPolicy-myvVolVG-60-1-1-04.vmx
[vVolCP1-vvolrep-fs52-2] rfc4122.63d85174-9aa0-4c20-baef-50289c5404a2/bulk-TestSPBMPolicy-myvVolVG-60-1-1-02.vmx
[vVolCP1-vvolrep-fs52-2] rfc4122.bda8fe2e-26ee-4c09-b738-3f4e4e790a87/bulk-TestSPBMPolicy-myvVolVG-60-1-1-03.vmx
[vVolCP1-vvolrep-fs52-2] rfc4122.cf984cb7-9b01-4038-8252-4f476a006c55/bulk-TestSPBMPolicy-myvVolVG-60-1-1-05.vmx

PS /Users/warren> $VMsToRegister = @()
PS /Users/warren> foreach ($VM in $VMs){
>>     $VMsToRegister += New-VM -VMFilePath $VM -ResourcePool Cluster2
>> }

PS /Users/warren> Get-Cluster Cluster2 | Get-VM | Format-Table -AutoSize


Name                                 PowerState Num CPUs MemoryGB
----                                 ---------- -------- --------
bulk-TestSPBMPolicy-myvVolVG-60-1-1-05 PoweredOff 1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-03 PoweredOff 1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-02 PoweredOff 1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-04 PoweredOff 1        2.000
bulk-TestSPBMPolicy-myvVolVG-60-1-1-01 PoweredOff 1        2.000
vvolrep-21                           PoweredOff 2        6.000
vvolrep-22                           PoweredOff 2        8.000


PS /Users/warren> Start-SpbmReplicationReverse -ReplicationGroup $ReplicationGroupTarget

Name                 ReplicationState
----                 ----------------
myvVolVG-60-1        Source
```

Initiate Planned Failover

Register VMs into vCSA-2

“bulk” VMs now in vCSA-2

Initiate reprotection / reverse replication

VMware Plugin 1.2

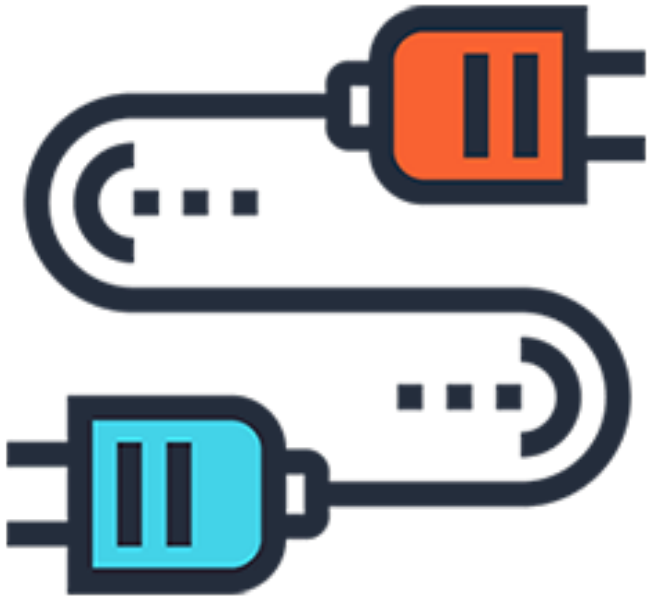

IBM Storage Virtualize plugin for vSphere

Automates vSphere and storage tasks together All from the vCenter UI

- Based on VMware remote plugin architecture
- Supports all Storage Virtualize platforms:
  FlashSystem (8.4.2+), SAN Volume Controller and Storwize products)
- Manage multiple storage systems from within vSphere
- Simplifies Storage Provisioning & Management
- Provides storage context for vSphere objects
- OVA Download from IBM Fix Central

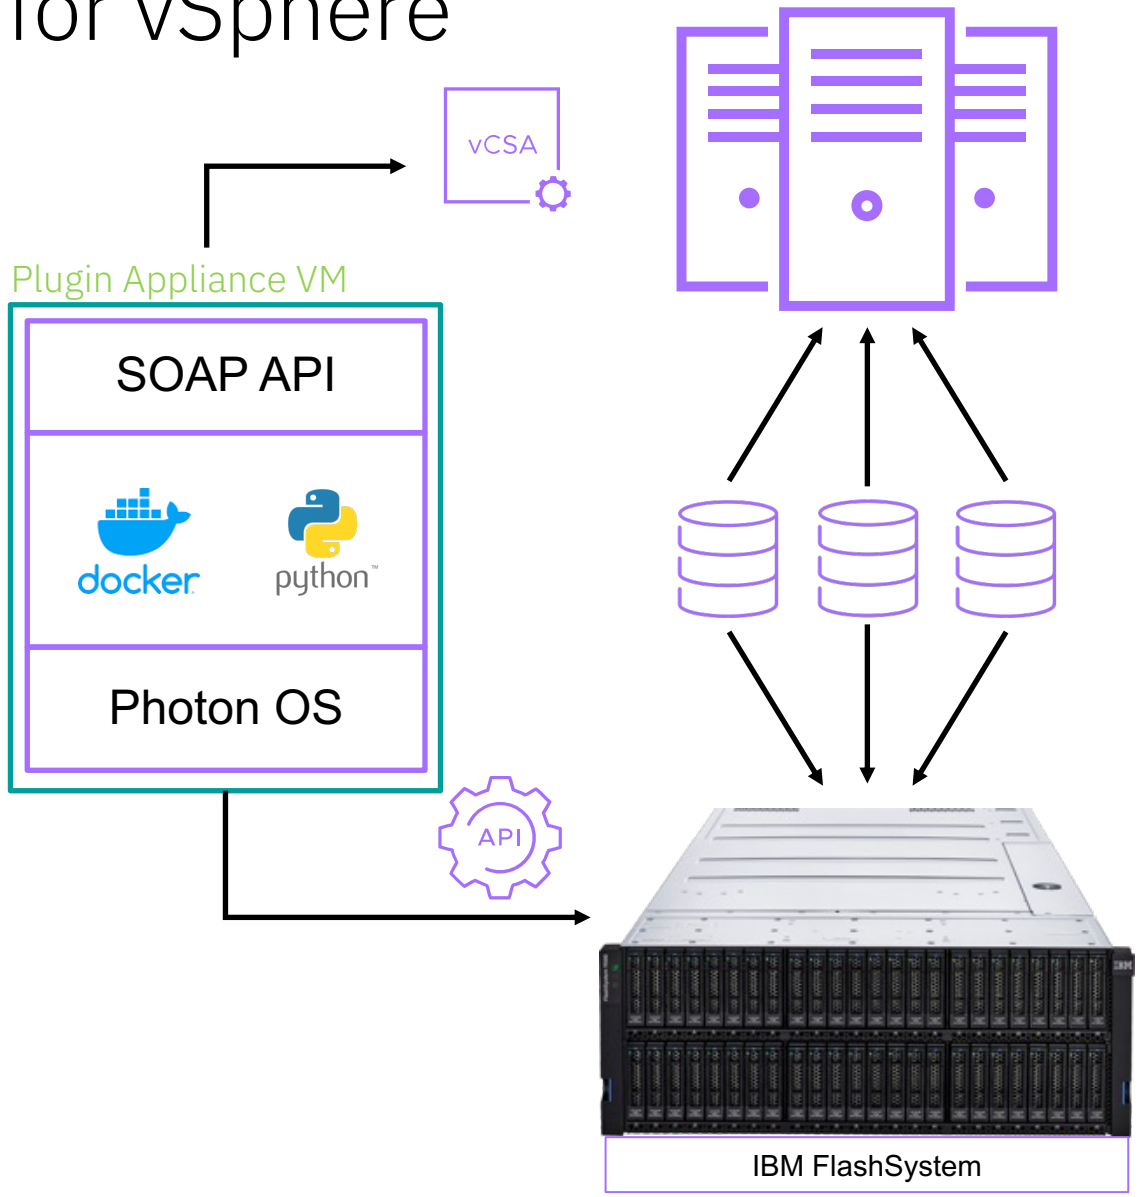

Current vSphere Plugin Capabilities

- **Create Multiple Datastores**
  - Automates creation of volumes, mapping to hosts and creation of VMFS datastores

- **Expand Datastore**
  - Combines volume expansion and increase datastore capacity

- **Delete Datastore**
  - Combines unmount datastore, delete volume and delete datastore
  - Checks for existing VMs registered and Snapshots

- **Snapshot**
  - Volume Group Snapshot support
  - Creates volume group, adds volume to VG and initiates a snapshot

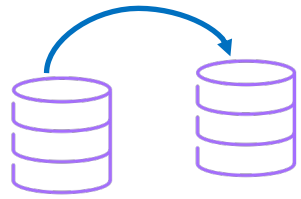

1.1.1 Updates

- Hyperlink to Storage Insights (if enabled)
- RBAC on the VMWare Cluster
- Manually Rescan all objects managed by the plugin

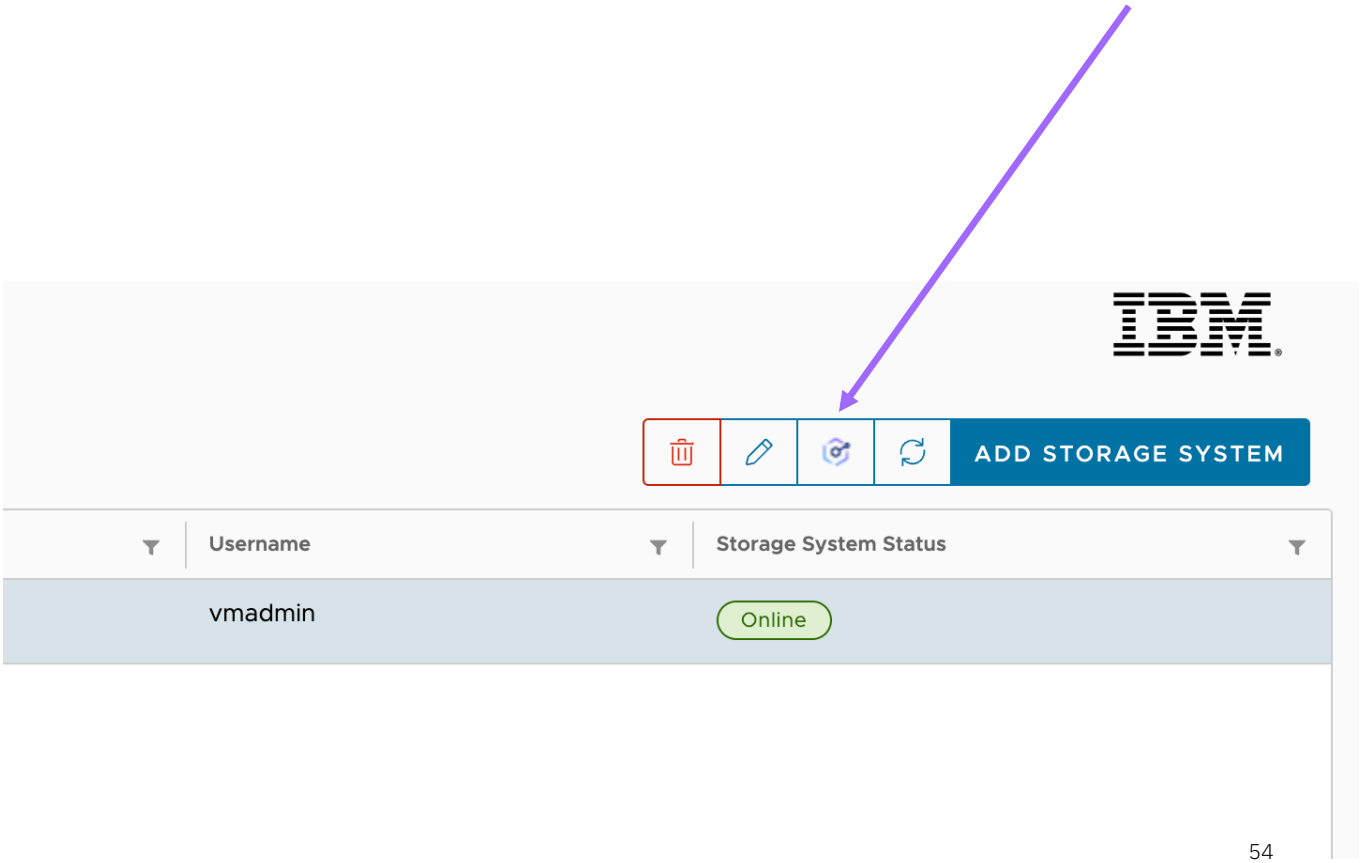

IBM Storage Virtualize plugin for vSphere v1.2

Volume Groups
- Snapshot Policies
- Replication Policies

Snapshot Management

Multiple-Pools

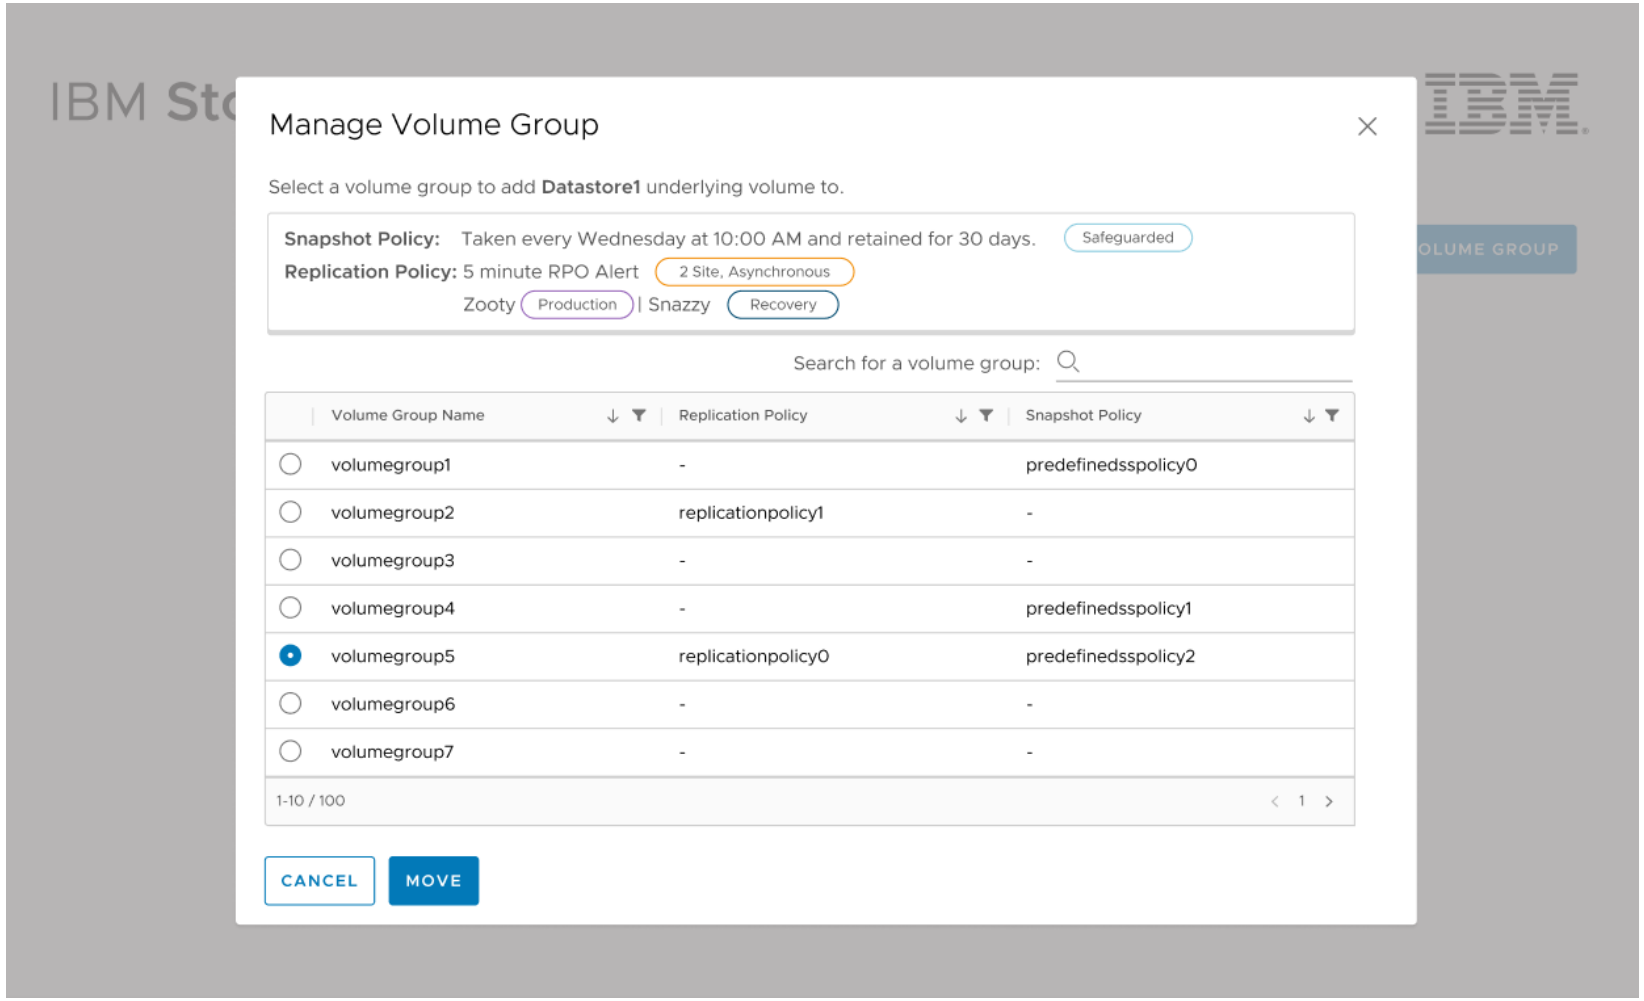

Plugin Dashboard Updates

- The dashboard presents all the registered IBM storage systems and specific information for each of them.

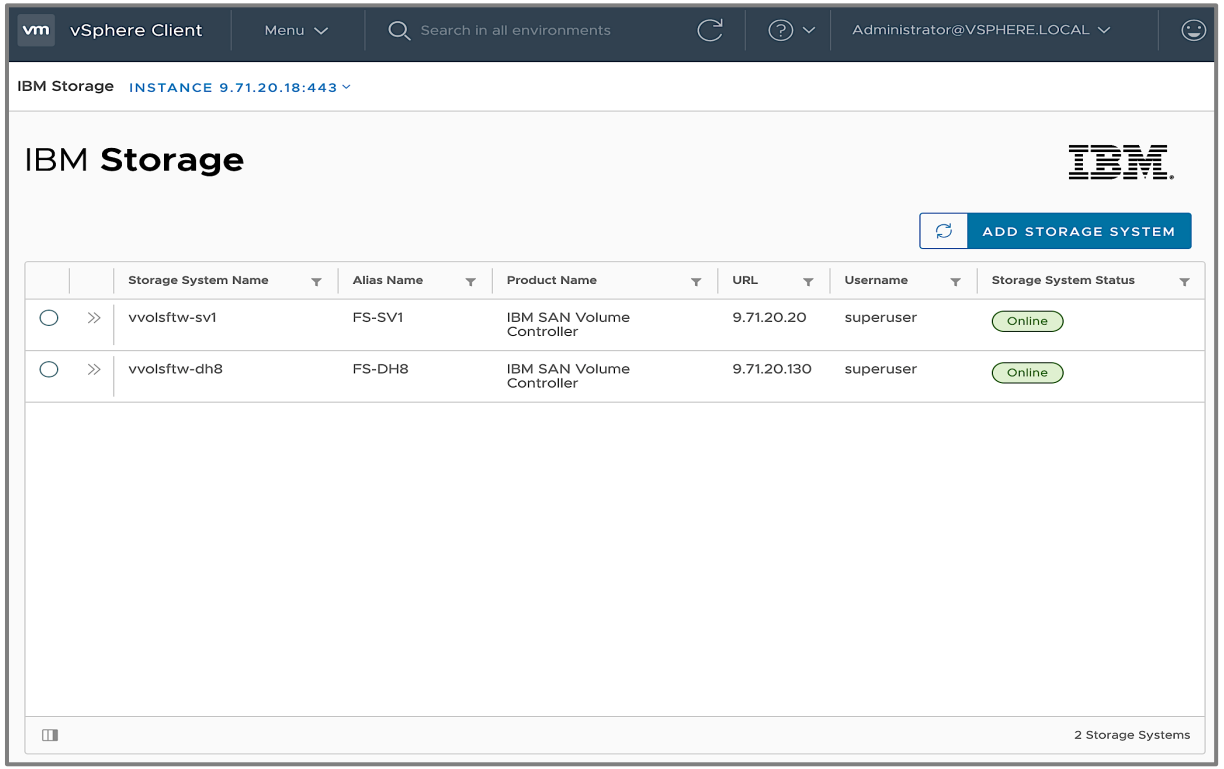

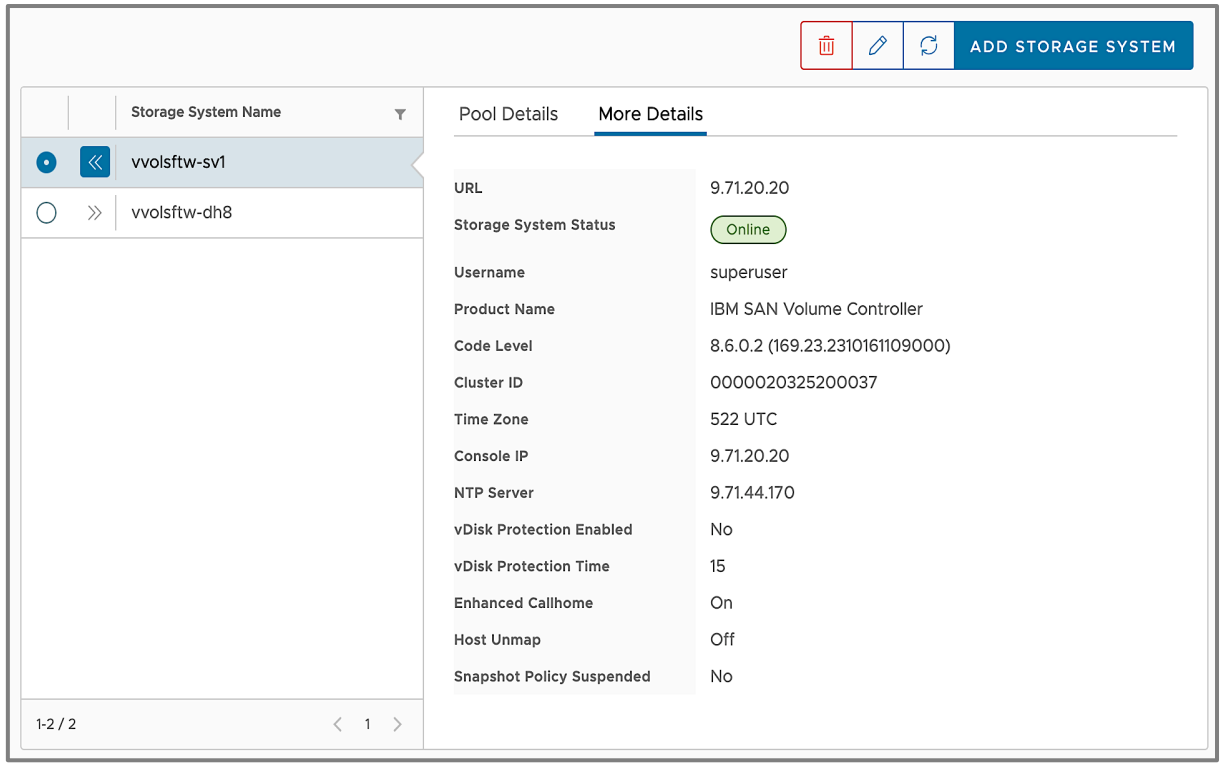

Support for Multiple Pools

- When adding a new storage system, you can now select multiple pools for management under the plugin

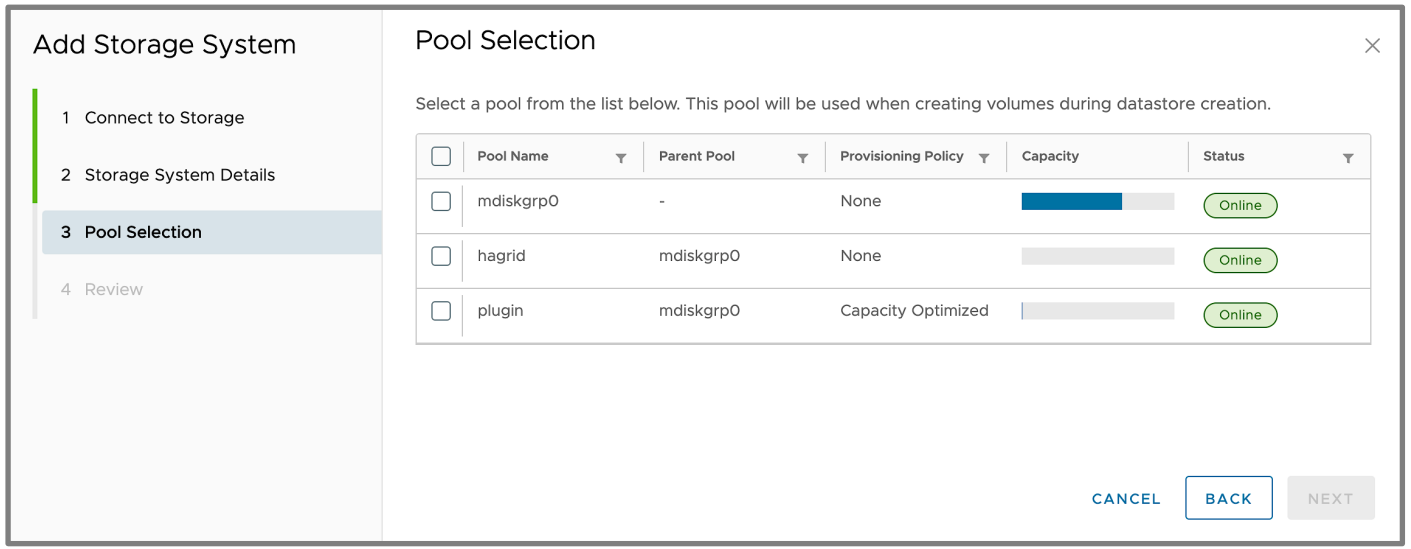

Create VMFS Datastore with Volume Groups

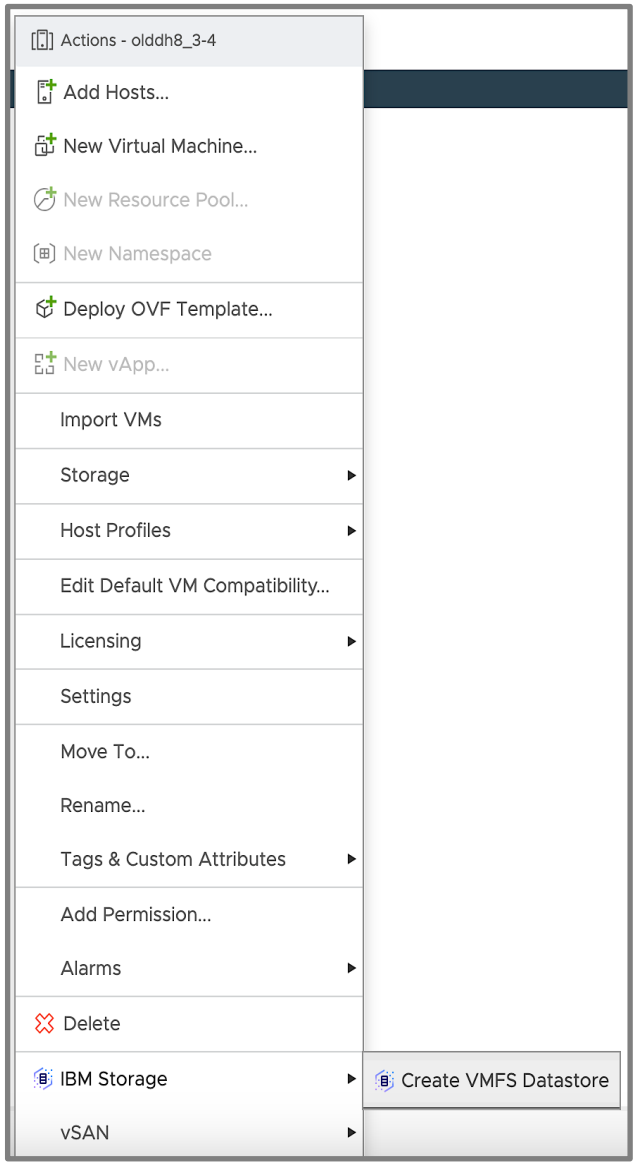

- During VMFS datastore creation, the user can now select a volume group for the datastore to reside in
- Support for showing snapshot and replication policies of the volume group
  - Snapshot policy will show schedule details
  - Replication policy will show locations and RPO

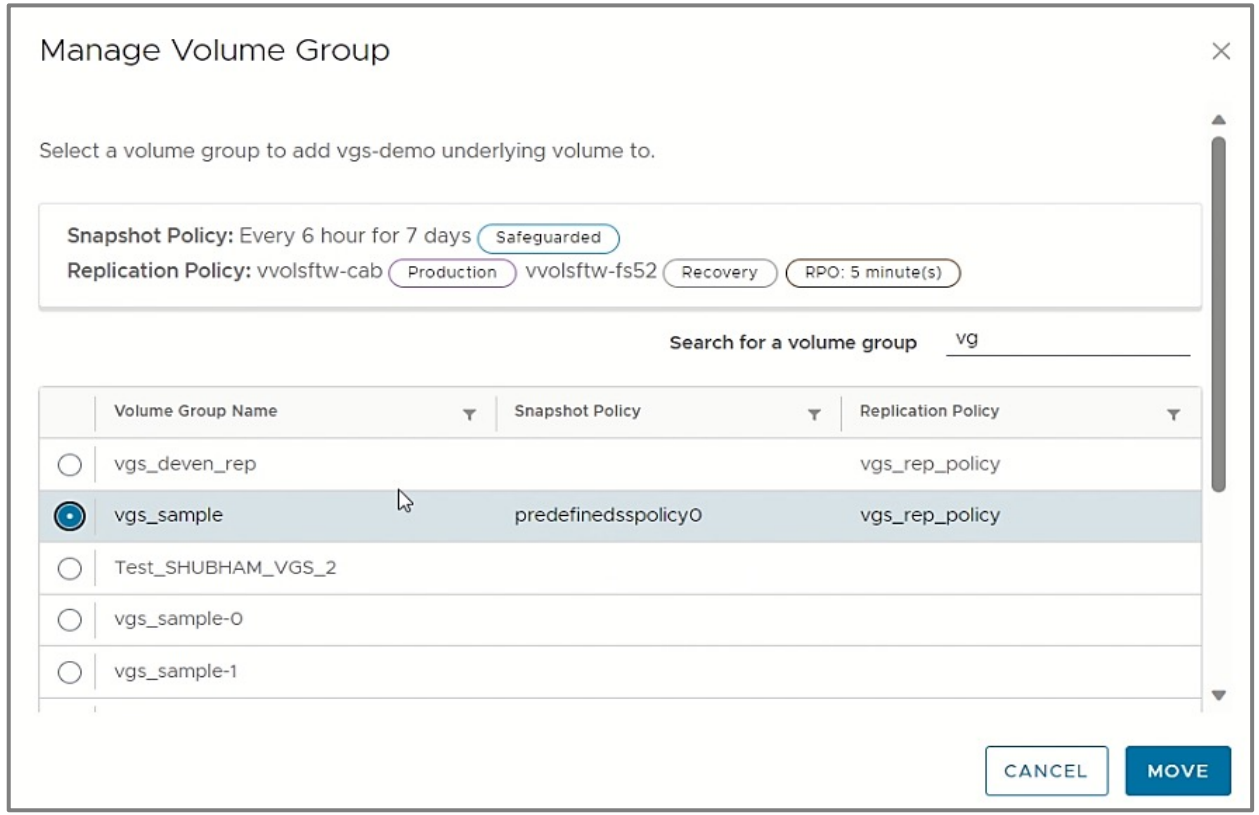

Datastore Summary

- VMware Admin can see **“Datastore Details”** in summary section
- VMware Admin can also see **“Volume Group Details”** in summary section

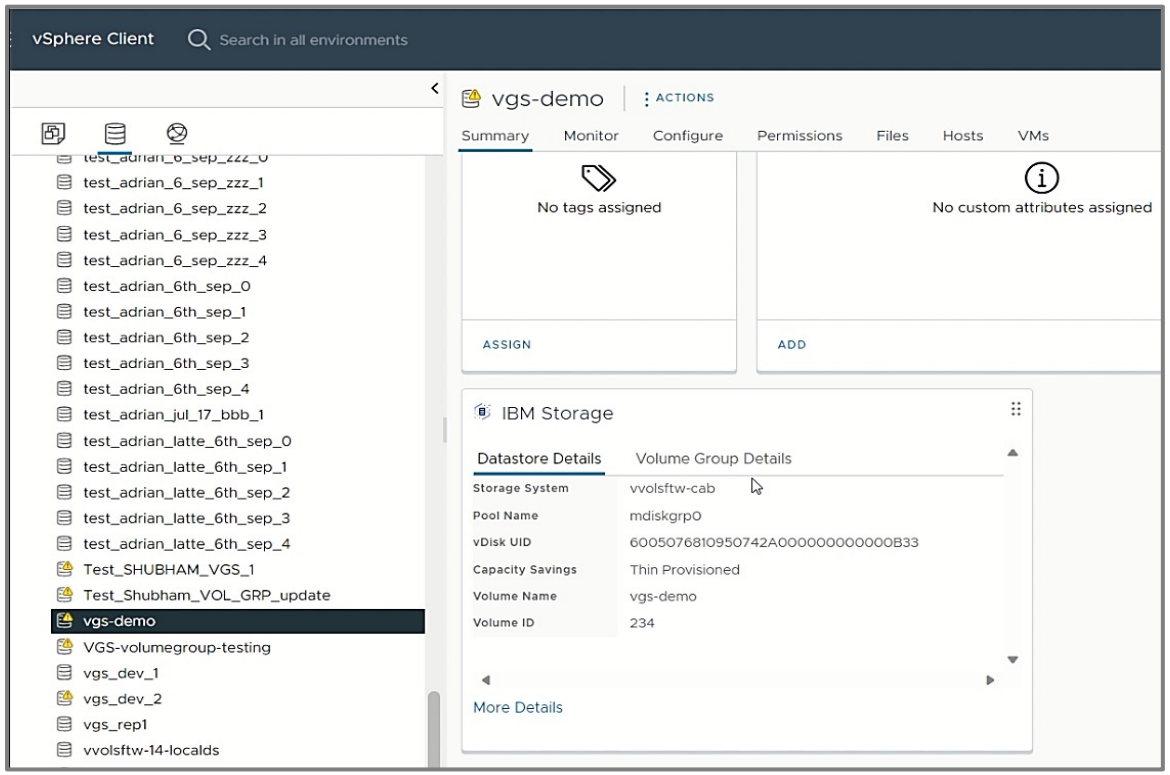

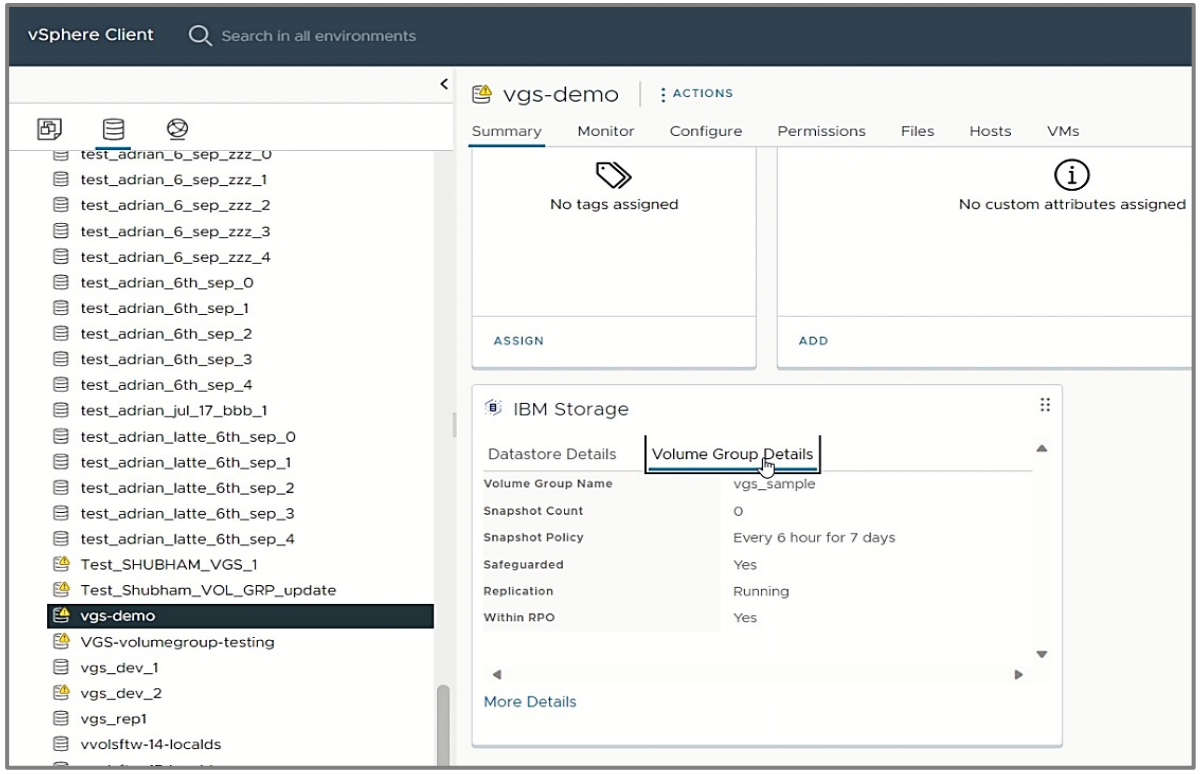

Datastore Snapshots

- Provision to **take a snapshot** of volumegroup associated to datastore.
- List of snapshots taken on a volumegroup are visible in this section with its attributes like safeguarded, creation and expiry time.

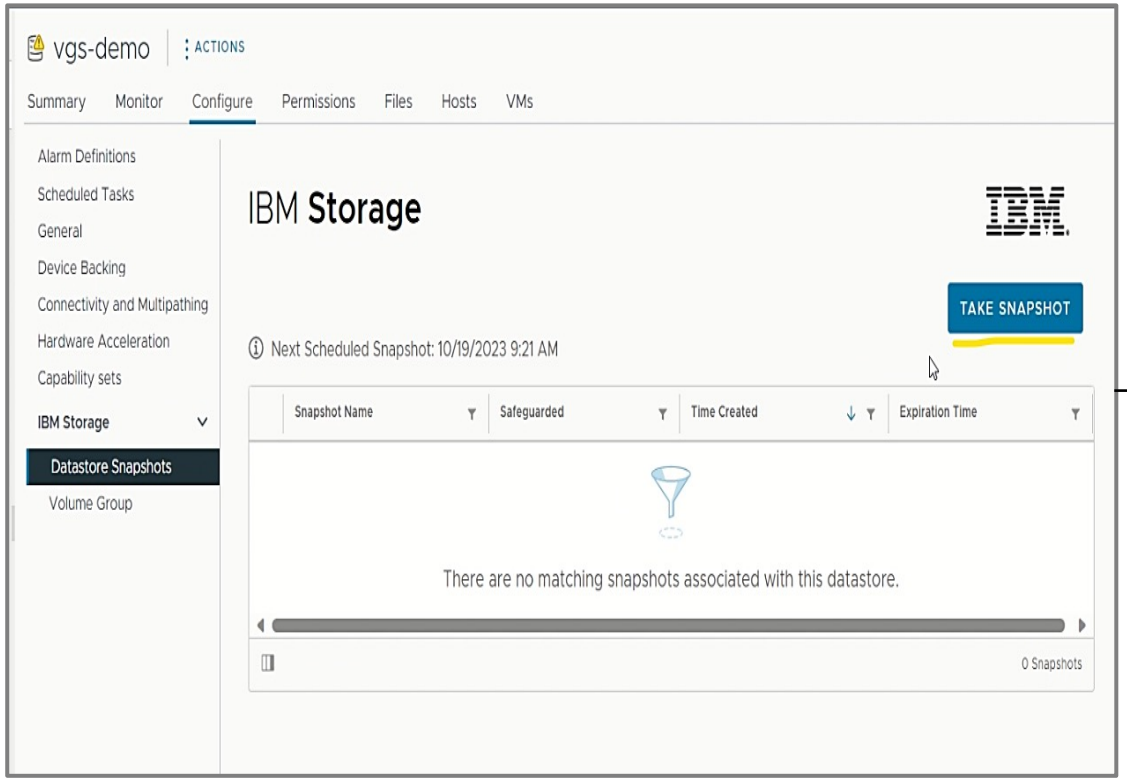

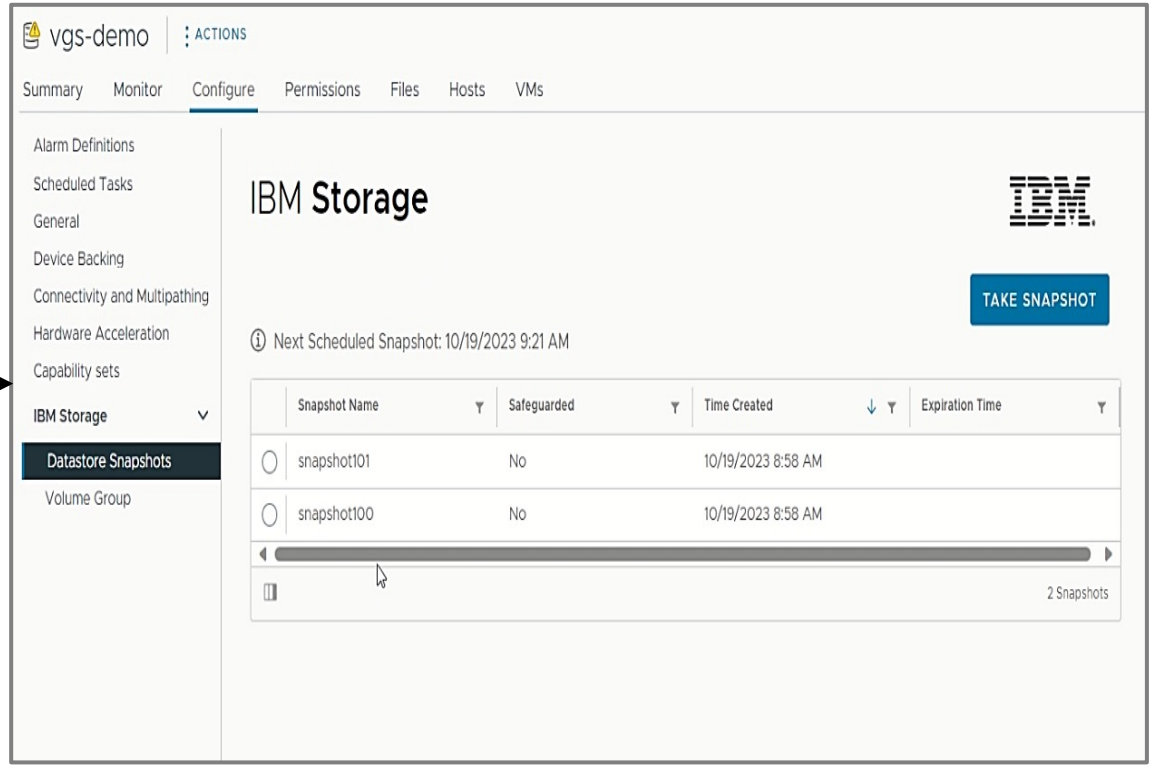

Delete a Snapshot

- Provision to **delete a snapshot** of volumegroup associated to datastore.

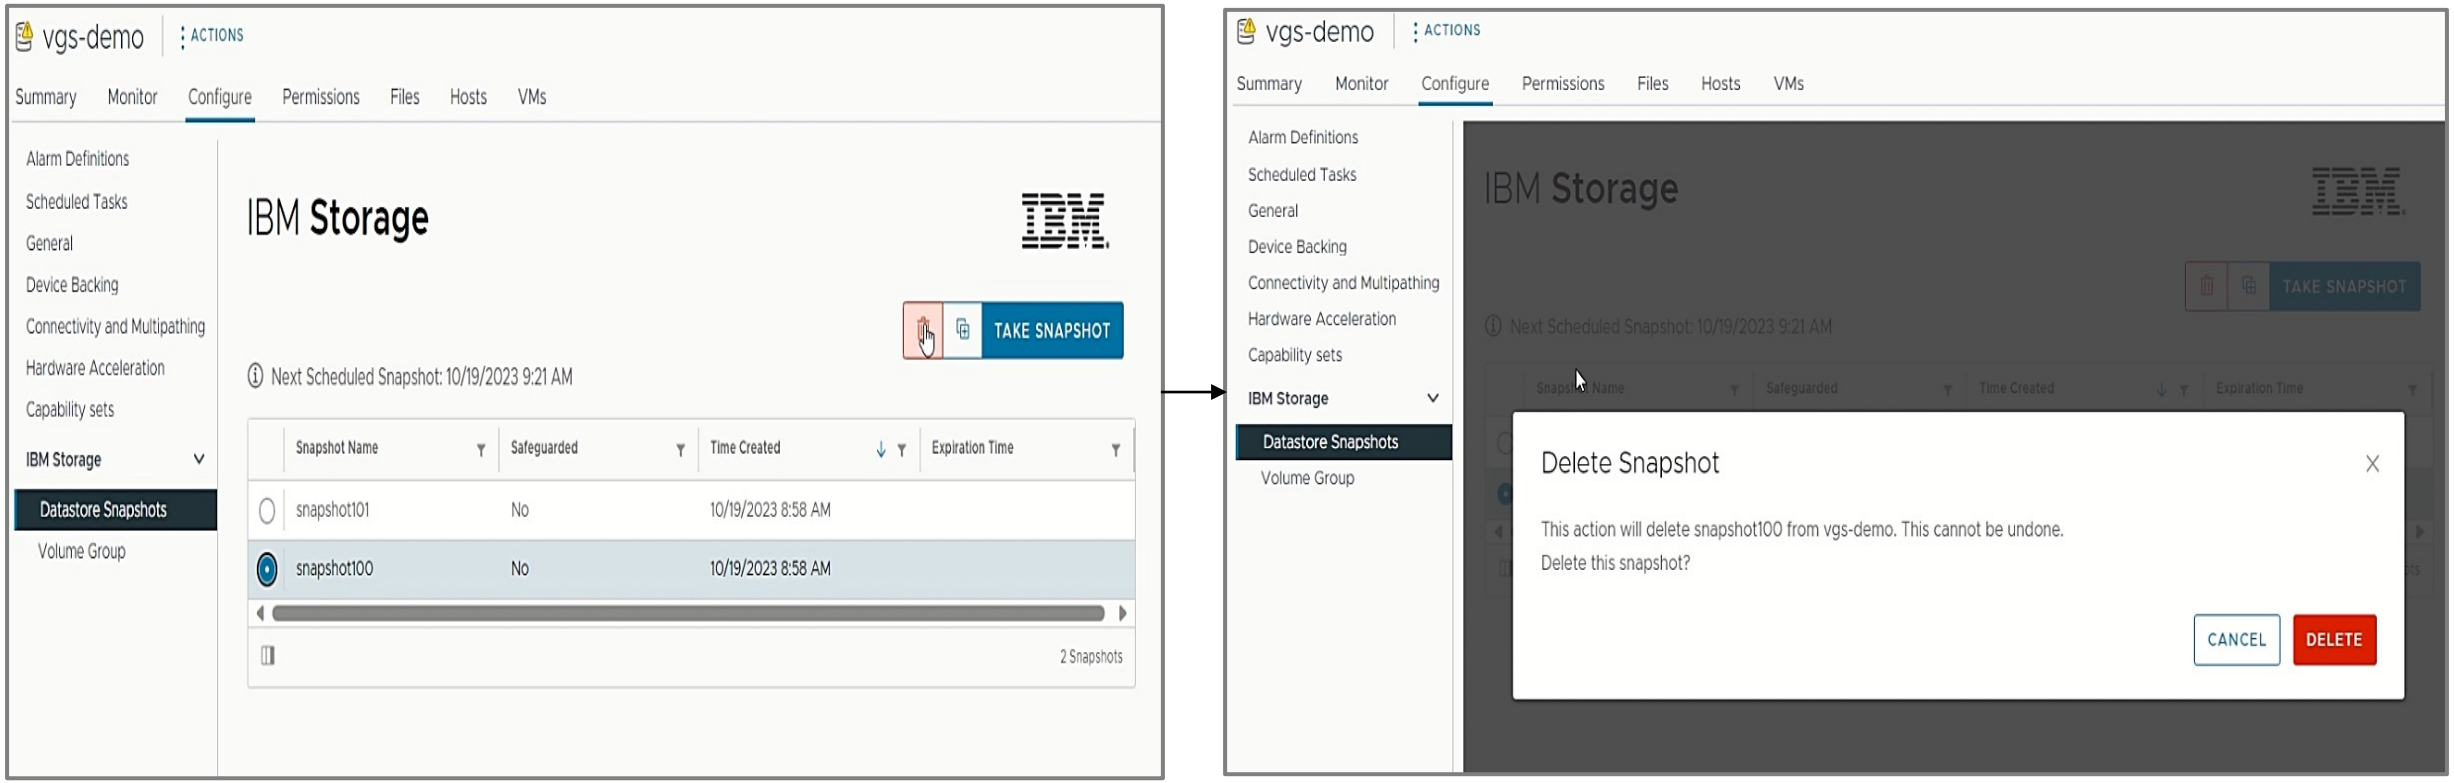

Copy Snapshot to a New Datastore

- Provision to perform copy of snapshot which creates a new datastore using thin clone or clone.

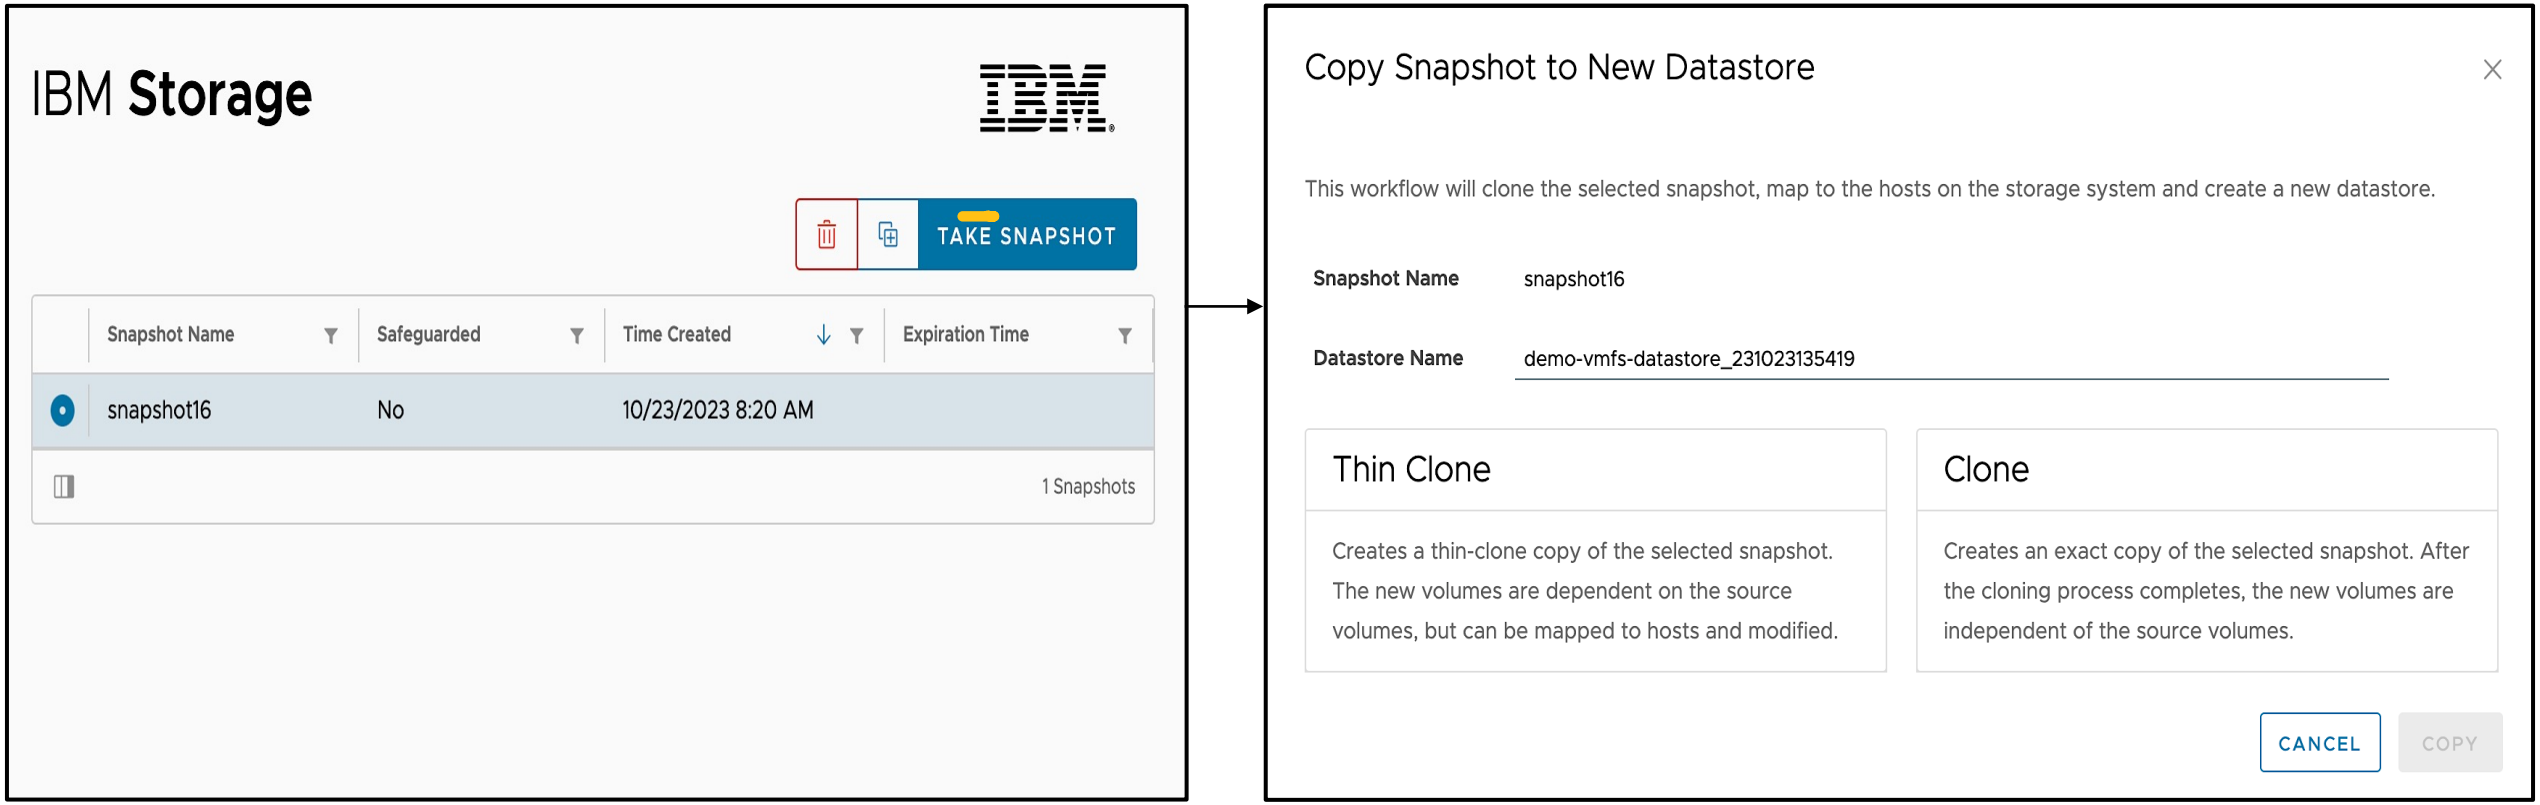

Volume Group Details of a Datastore

- VMware Admin can see **“Snapshot Policy”** details of Volume Group, if assigned
- VMware Admin can see **“Replication Policy”** details of Volume Group, if assigned

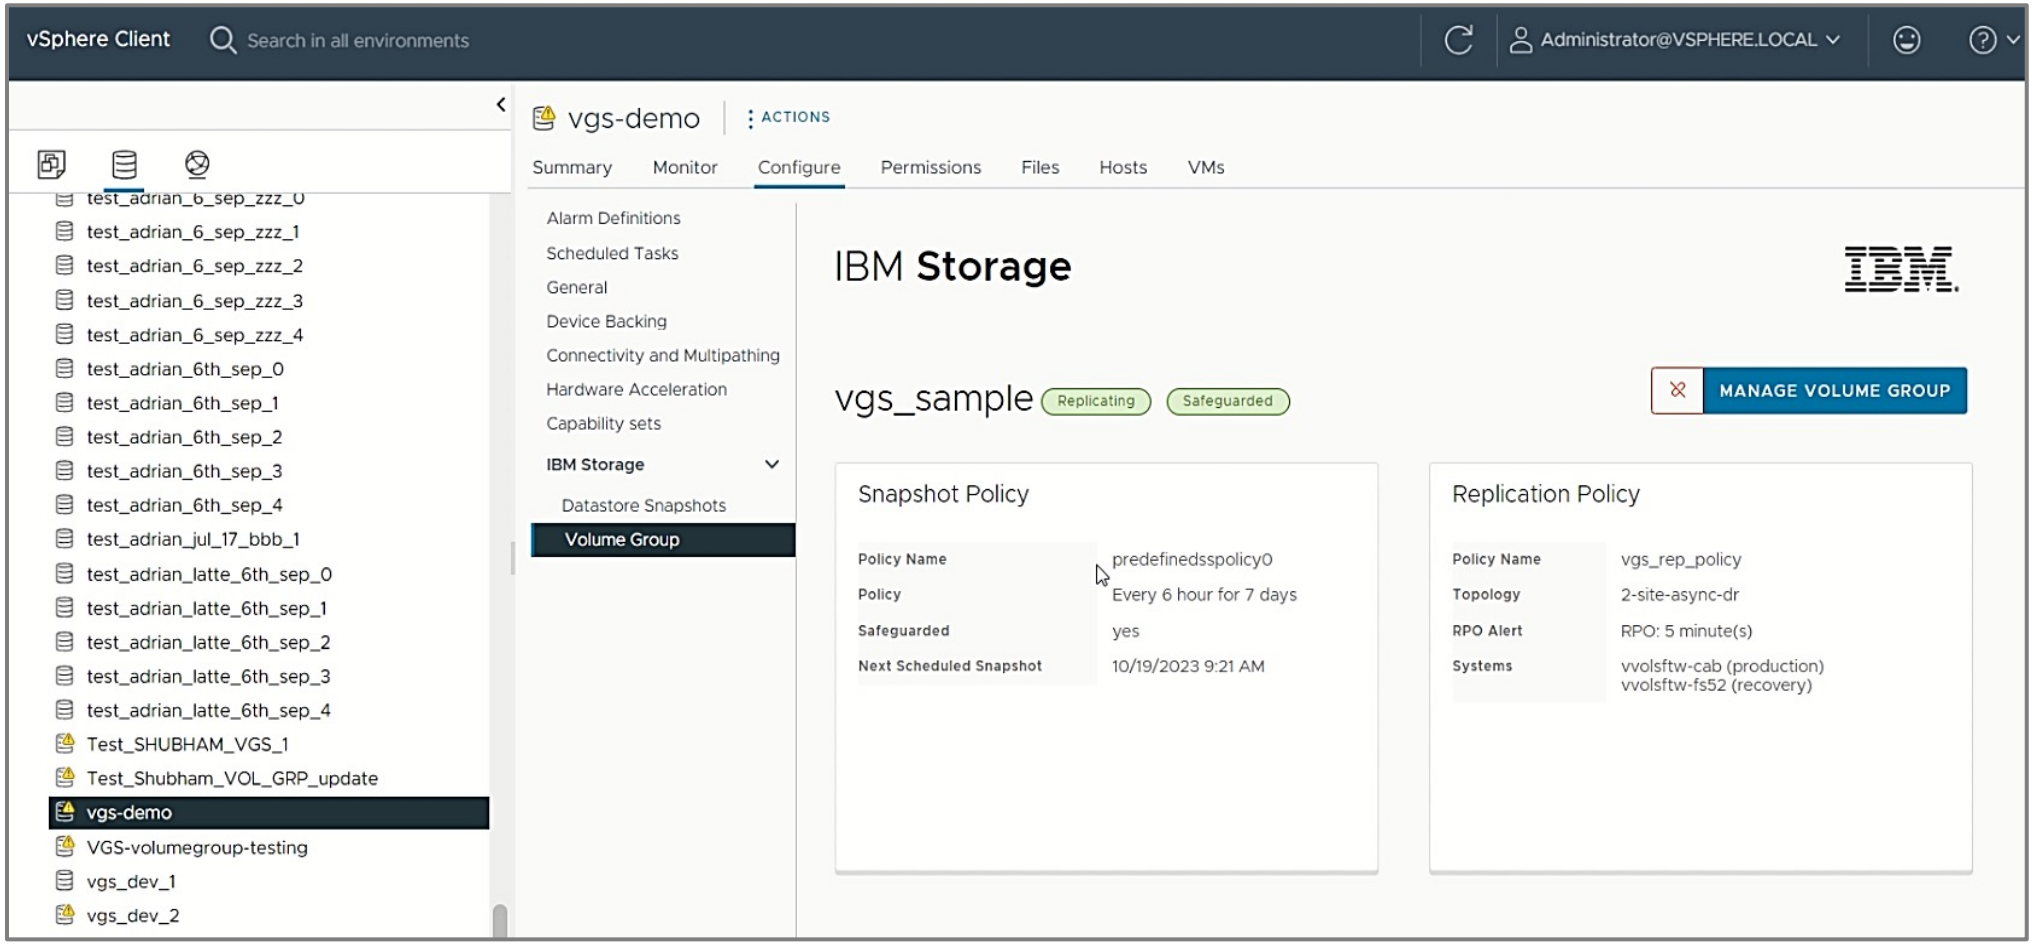

Manage Volume Group for a Datastore

- Provision to move datastore from one Volume Group to other on a given storage system, based on policy requirements

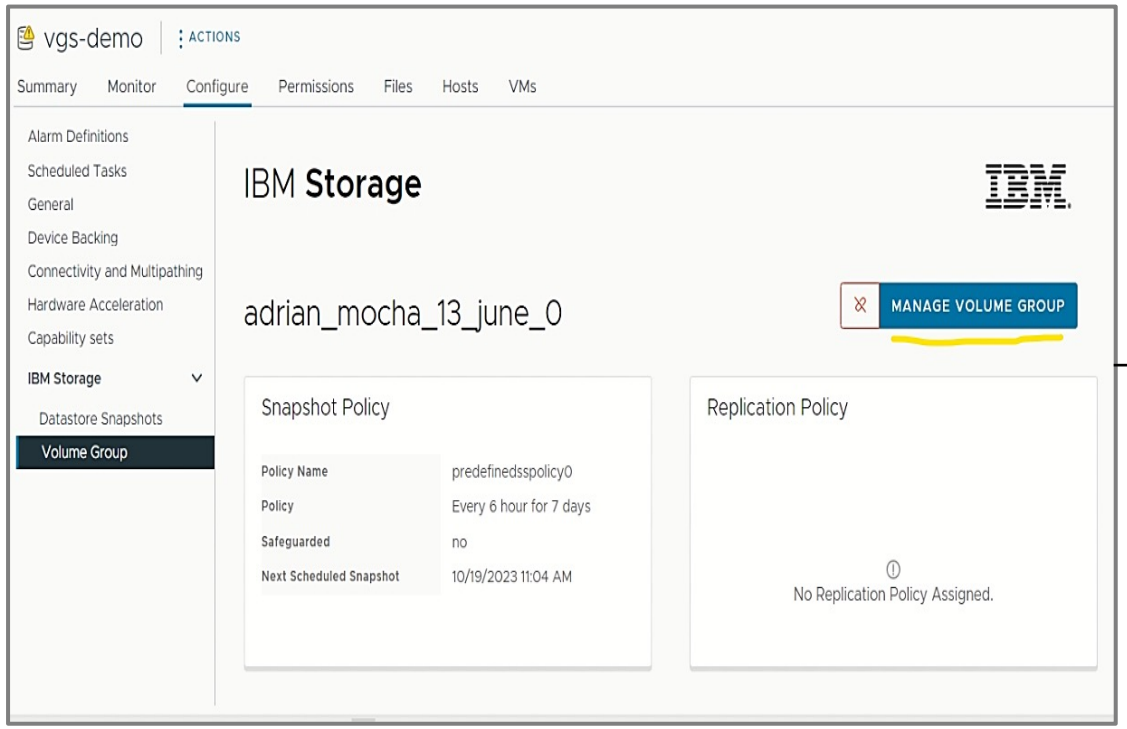

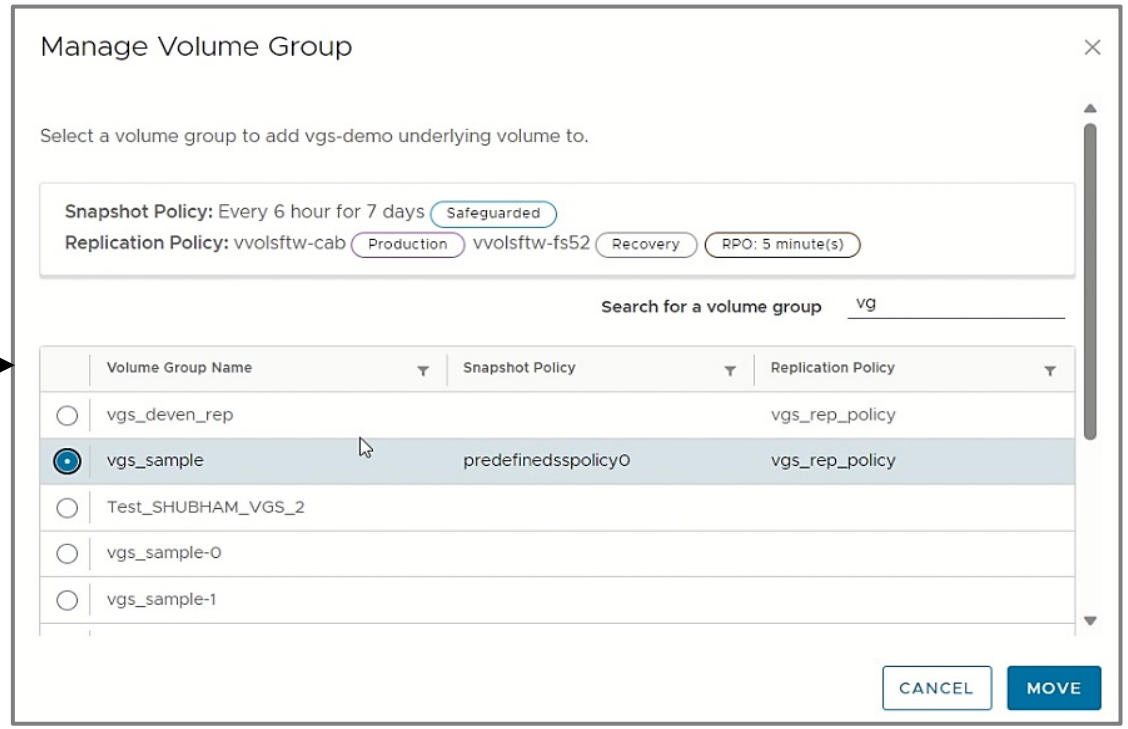

Remove Volume Group for a Datastore

- Provision to remove a datastore from a Volume Group

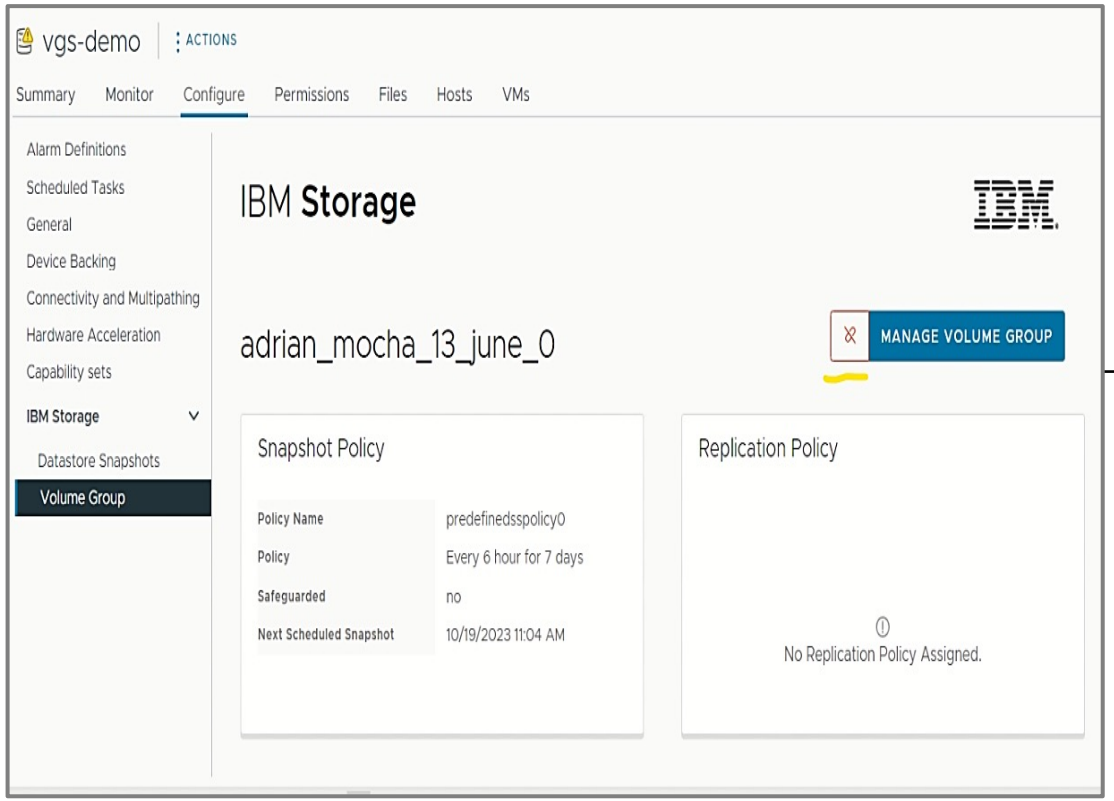

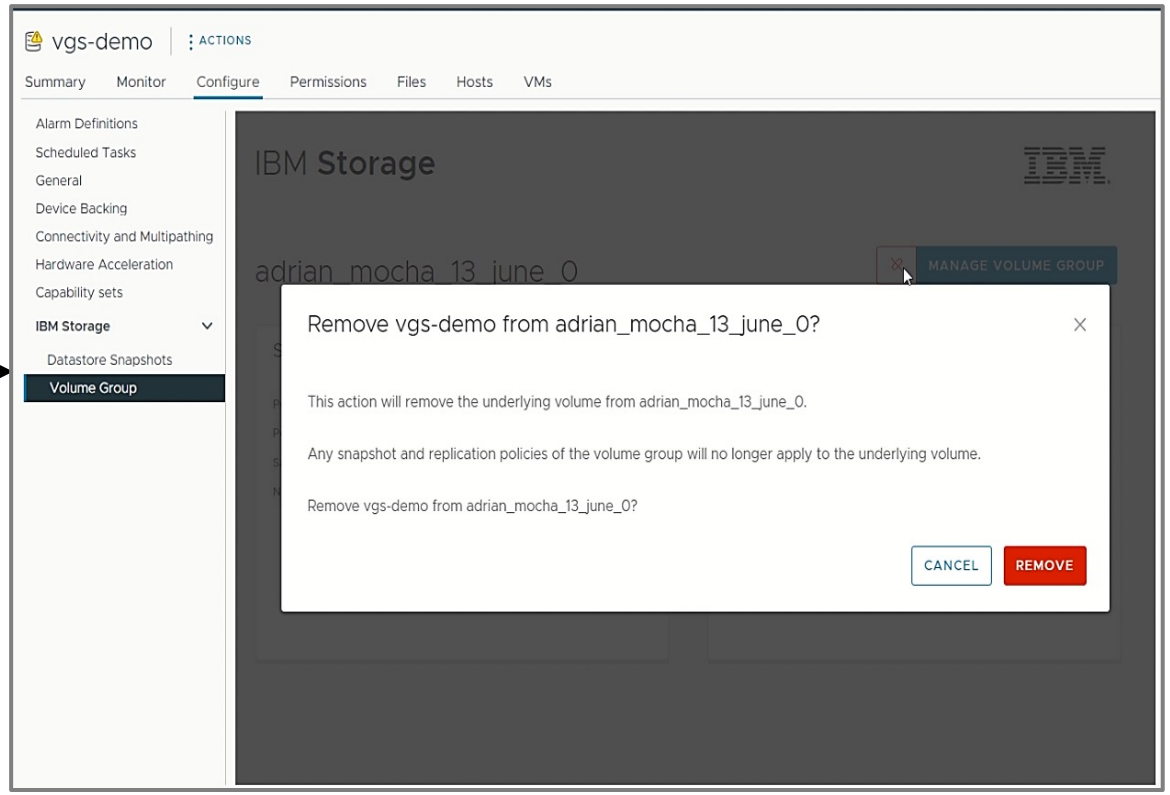

Upgrade Plugin From v1.1

Database upgrade

- We will support database migration so that existing data is not lost when upgrading from from v1.1/1.1.1 to v1.2.0.

- Seamless upgrade - It would ensure continuous management of old datastores by new plugin instance.

Restore In Place/Single Volume Clone

Snapshot Restore

- In 8.5.1 we added the ability for a user to take a snapshot of a volume group. If they wanted to access the data from one of those snapshots they had to create a volume group from the snapshot and set up the required host mappings.
- 8.6.2 adds full support for the restoration of a volume group from one of its snapshots thereby removing the need to create volume groups and set up host mappings.
- This functionality is provided by the addition of a new command:
  - restorefromsnapshot
- The management GUI supports snapshot restores from the Volume Group panel.

Currently Snapshot Restore **not** supported for:

- TCT enabled volumes
- Volumes involved with policy-based replication and policy-based HA
- Metro Mirror, Global Mirror and HyperSwap
- 3 site configurations

Prerequisites

- A volume group can only be restored from a snapshot taken from that volume group
- The composition of the volume group at the time of the restore from snapshot must be the same as when the snapshot was taken (no extra volumes in the volume group)
- If a volume has been deleted from the volume group after a snapshot was taken, restoring from that snapshot will result in the volume being changed back to the active state and being re-added to the volume group

Restore From Snapshot

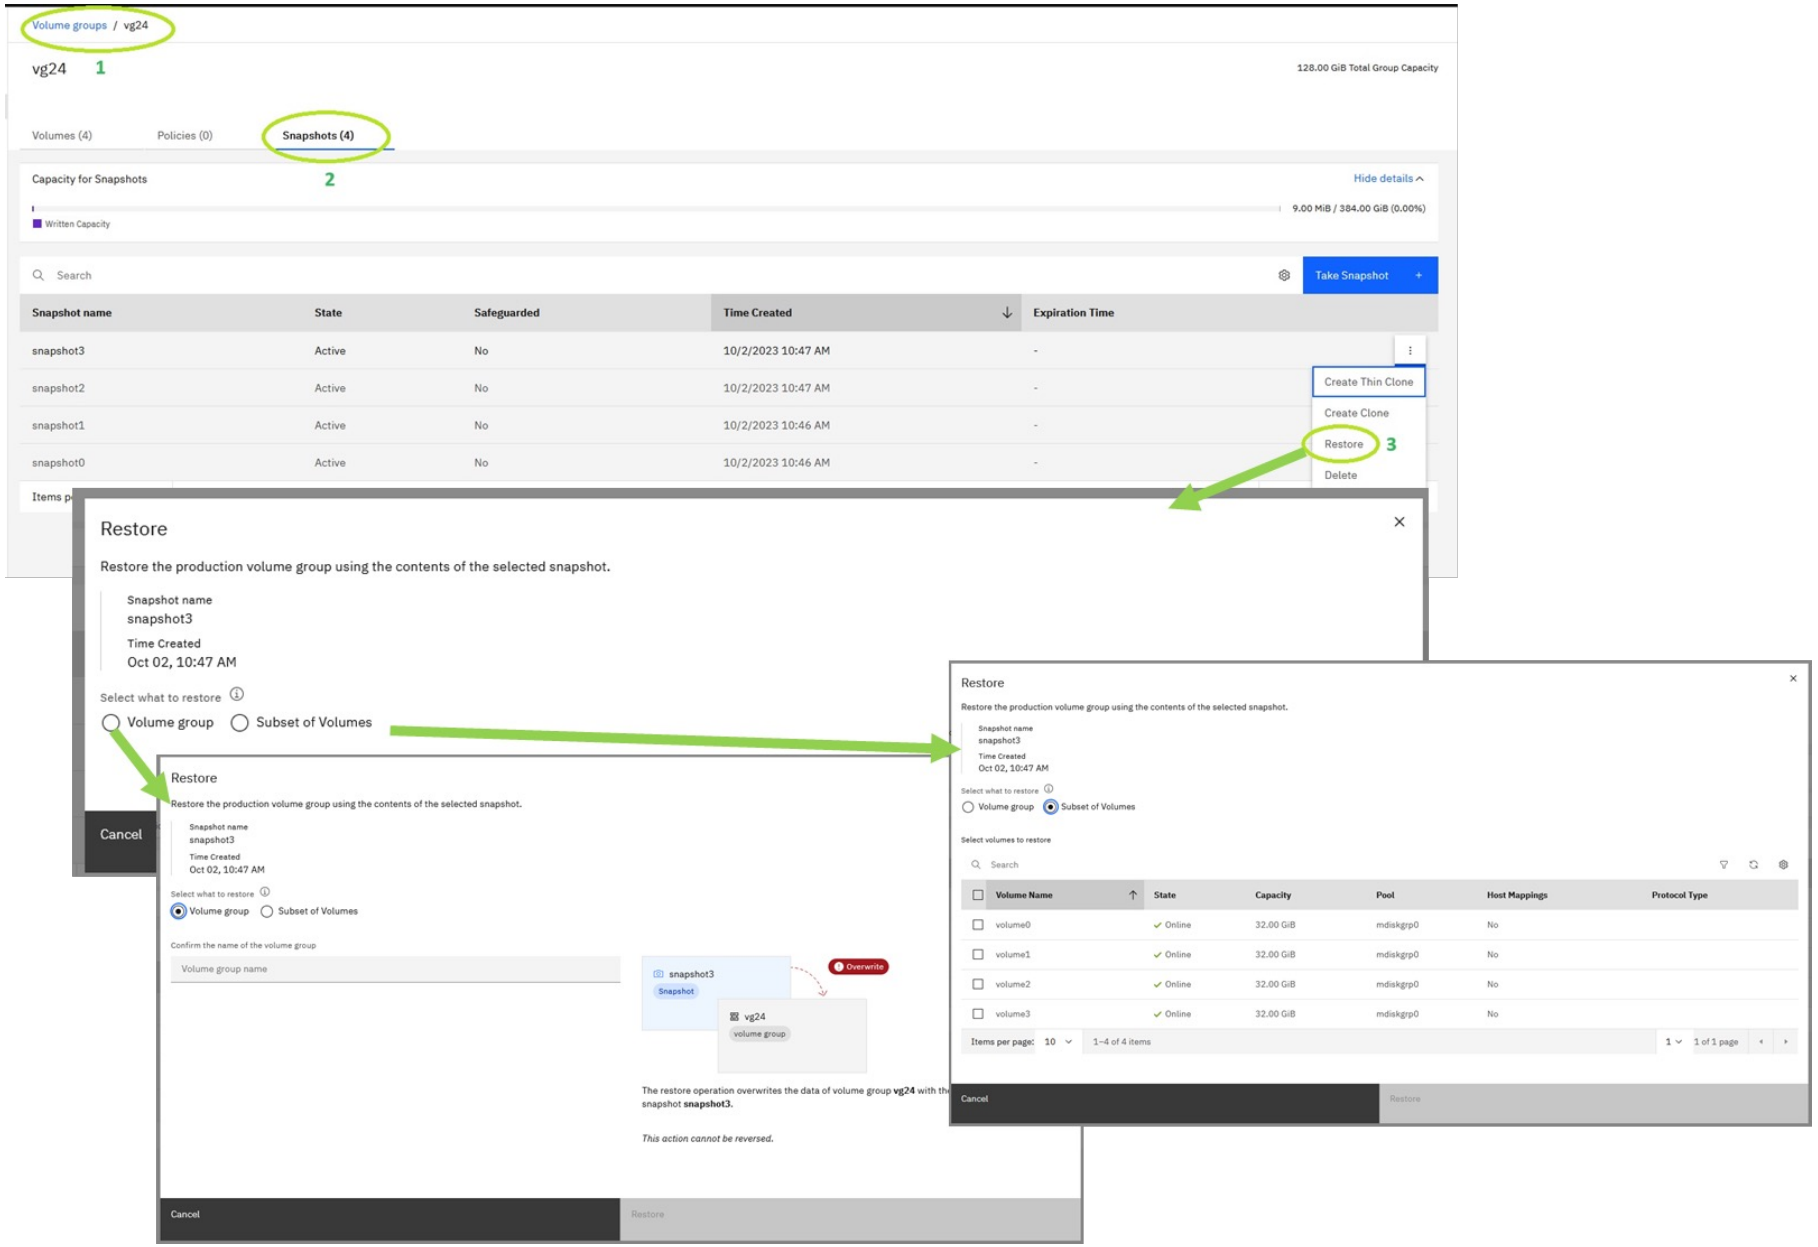

Monitoring Restores

- lsvolumegrouppopulation – shows progress at a volume group level. Some additional fields have been added:
  - data_to_move - the amount of data yet to be restored. As the restore progresses the value in this field will decrease. An empty field indicates no more data to restore.
  - restore_start_time - The time that the latest restore operation was started. Blank if no restore operation is ongoing
  - restore_snapshot_name - The name of the snapshot used to populate the volume group. Blank if no restore operation is ongoing.
  - restore_estimated_completion_time - An estimate for when the background restore of the volumes in the volume group will complete. This estimate will be calculated using the used_capacity (used_capacity_before_reduction for DRP) of the volumes and the actual background copy rate being achieved to estimate how long each volume is going to take and then choosing the longest time remaining.
- lsvolumepopulation – shows progress at a volume level.
  - A new field, data_to_move, has been added. The amount of data yet to be restored. As the restore progresses the value in this field will decrease. An empty field indicates no more data to restore.
- lsvolumegroup – a new field, restore_in_progress, has been added. This is yes / no field.
- lsvdisk – a new field, restore_in_progress, has been added. This is a yes / no field.

We have added a **volume_size_mismatch** field to **lsvolumesnapshot**. If this field is 'yes' it indicates that the volume has been expanded since the snapshot was taken and so a restore using that snapshot will not proceed until the size has been restored to what it was when the snapshot was taken.

Thin Clone Refresh

- The user creates a thin-clone volume group and populates it from a snapshot of the source volume group. These volumes can be mapped to a host and accessed. The user wants to examine a different snapshot without having to map a new set of volumes to the host. They also want to be able to discard changes made to the volumes by host IO activity without having to map a new set of volumes to the host.
- In addition, there is the same use-case as the volume group refresh discussed above but for a vector of volumes instead of a whole volume group. This is primarily aimed at automation software rather than a manual operation because of the difficulties in managing arbitrary lists of volumes.

- A new command, refreshfromsnapshot, is available to refresh a thin-clone volume group or a subset of thin-clone volumes from a snapshot.
- The thin-clone volumes being refreshed can’t have any snapshots.
- The functionality to refresh a subset of thin-clone volumes is only available via the CLI

Thin Clone Refresh

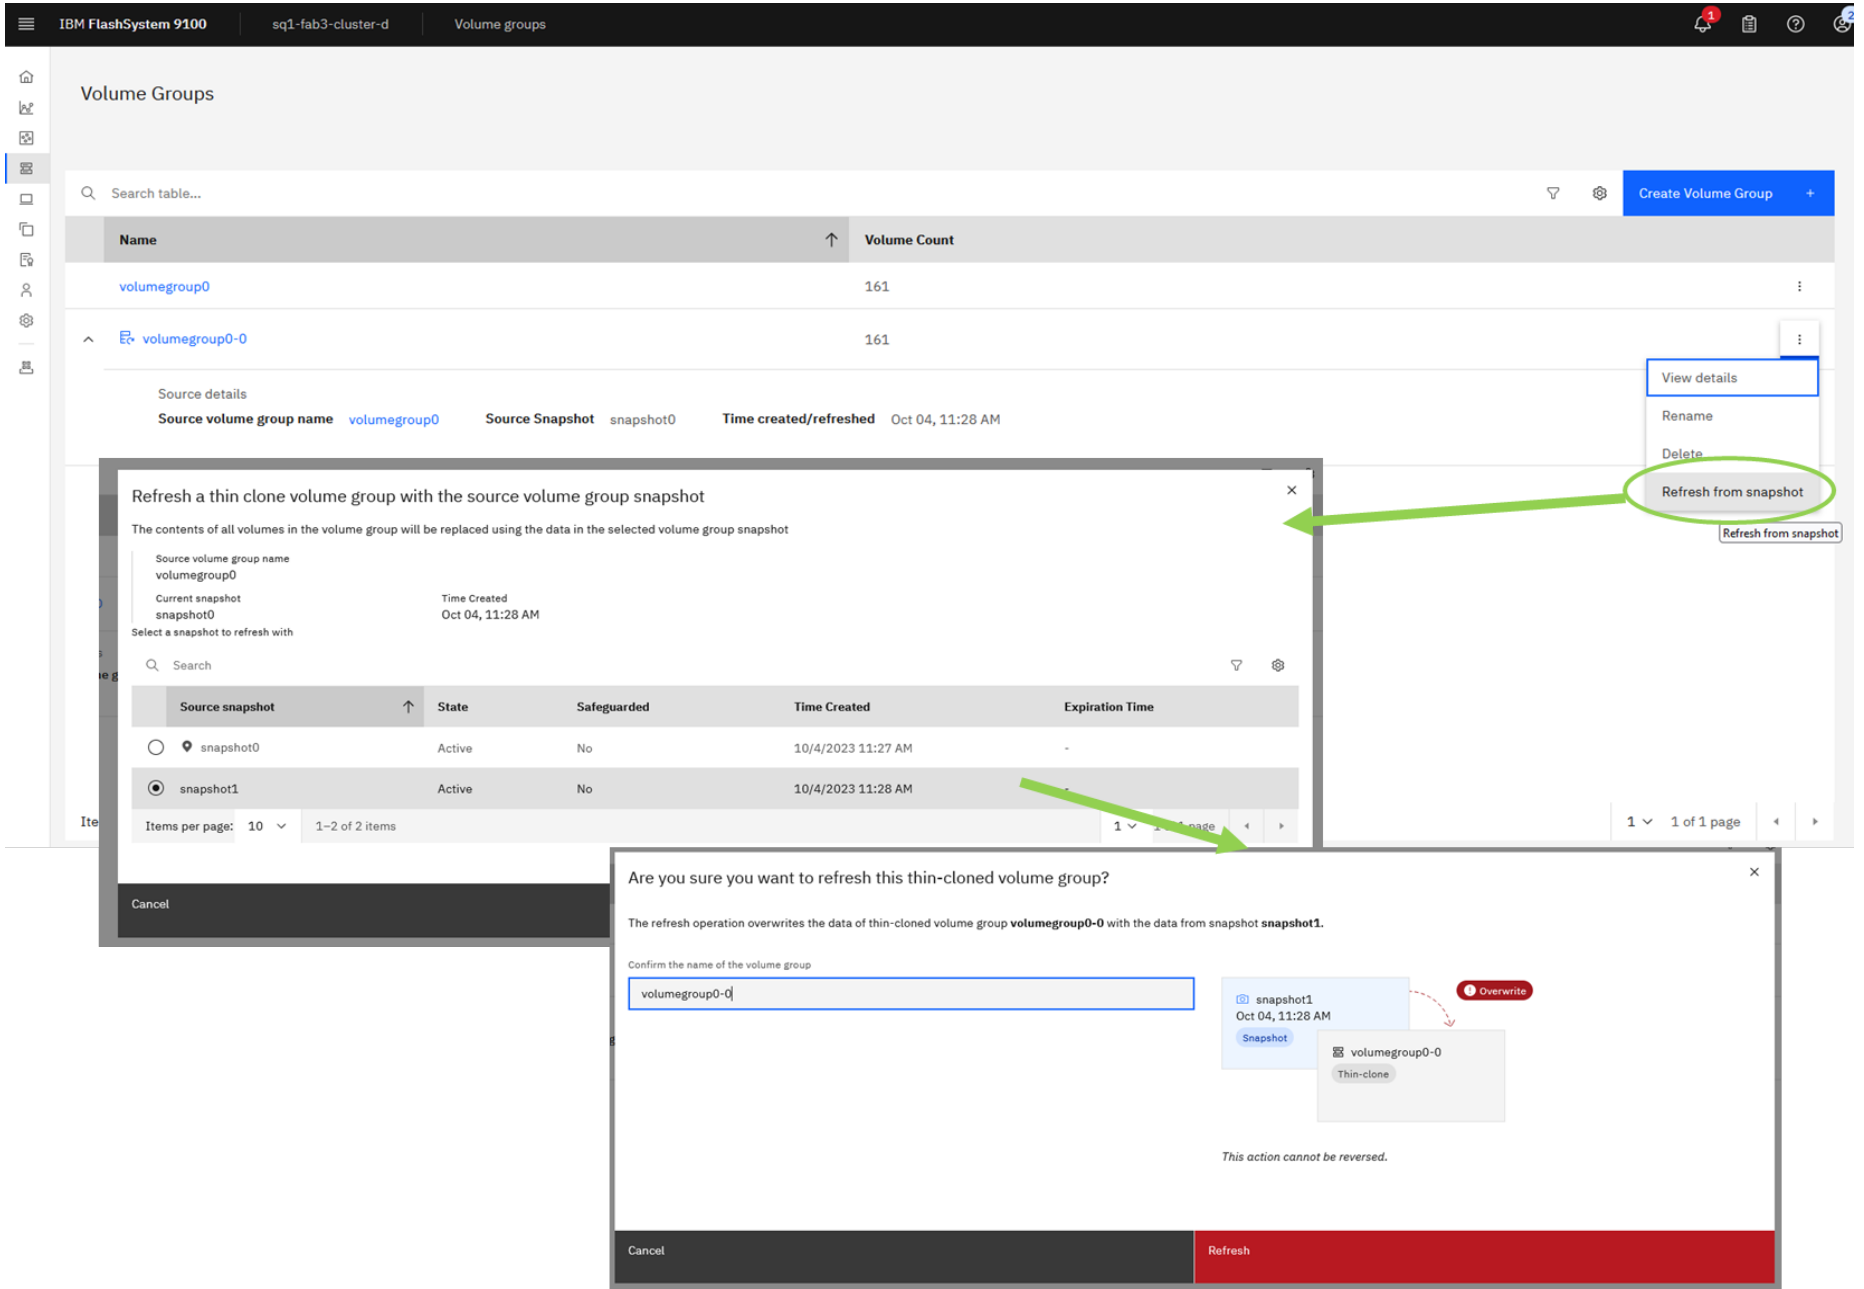

New Volume Creation Options

- **mkvolume** has been extended to allow a clone or thin-clone of a volume to be created from a volume snapshot of the required source volume.
- This would create a new volume that is based on the specified source volume and pre-populated with the contents of the specified snapshot.
- The pool attribute must be specified.
- It is possible to create a mirrored volume by specifying a colon separated list of pools. (SVC only)

Misc Snapshot Changes

- When calling out individual volumes in the **addsnapshot** command, the command can now accept up to 512 individual volumes instead of the previous 128

- The **rmvolumegroup** command now has an **–evictvolumes** flag to remove the volumes when the volume group is deleted
  - Prior a **chvdisk –removevolumegroup** needed to be issued for each volume prior to deleting the volume group
  - The volume group cannot have:
    - Replication policy
    - Snapshot policy
    - Storage Partition
    - Snapshots

Misc Snapshot Changes - Continued

Maintain Source Volume ID

- If a volume is created from a snapshot, the source volume ID and source volume name will be listed on the lsvdisk command

Clean up Snapshots

When using automation there is a potential that certain snapshots could be orphaned on the system

- The addsnapshot command has been given new parameters:
  - -retentiondays
  - -retention minutes

Using these parameter will make sure snapshots are cleaned up after a certain length of time

Misc. Updates

Auto-Ranging Power Supply for FS5200

Starting on October 10, 2023, all new FS5200s will ship with an auto-ranging (100-240 volt) power supply

- FS5200 Auto-sensing Power Supply
- Feature Code **AHPG**
- This fulfills a long time ask from small customers and BPs

Storage Virtualize encryption methods

Primary encryption methods (existing)

Encryption with USB flash drives

IBM Security Guardium Key Lifecycle Manager

THALES

CipherTrust Manager

Backup encryption method (NEW)

Encryption Recovery Key

– The Encryption Recovery Key is a secondary method and not intended to be used as a primary key method

Motivation for encryption recovery key

Some clients want a backup encryption key method to compliment key server encryption, but don't want or can't use encryption with USB flash drives.

Encryption Recovery Key Features

No Additional Hardware

The encryption recovery key is presented to the user as an ASCII text string

Rapid Data Access

The encryption recovery key will unlock encrypted storage as quickly as the key can be entered

No need to have an admin on site to fetch USB drives from the safe, the encryption recovery key can be entered remotely

Peace of Mind from Hardware Failure

Worried about a USB drive batch failure or key server outage?

The encryption recovery key offers an additional layer of protection

Use Your Existing License

An encryption recovery key can be enabled on systems with existing encryption without any additional cost or license requirements

Simply upgrade, enable and go

Encryption Admin Security

- Current key must be supplied when rekeying
- Keys for existing types must be supplied when enabling new type
- Disabling encryption must be carried out using the technician port

SNMP Security Improvements

- Strengthen security for SNMP by delivering the following improvements:
  - More secure algorithms for authentication and privacy
  - Encrypt authentication and privacy passphrases at rest and ensure that they are redacted in all logs
  - Automatically encrypt existing passphrases on upgrade to 8.6.2+
  - Support SNMP over TLS and server certificates
- Improve usability by providing a command to test that SNMP servers are correctly configured

New SNMP Configuration Options

- Two security models for SNMPv3 servers:
  - User-based Security Model (USM) provides application layer security using authentication and privacy passphrases
  - Transport Security Model (TSM) provides transport layer security using TLS and certificates
- Both USM and TSM can be enabled for the same SNMP server
- New test command **testsnmpserver** can be used to verify that the configuration works

- Supported algorithms (new in bold):
  - Authentication: MD5, SHA-1, **SHA-224, SHA-256, SHA-384, SHA-512**
  - Privacy: DES, AES-128, **AES-196, AES-256**
- MD5, SHA-1 and DES are insecure and not recommended, but supported for backwards compatibility
- Algorithms used for TLS are determined by the system-wide security setting (sslprotocol)
- Certificates are configured using trust stores, which now have an option to enable use for SNMP

Limitations

- Max passphrase length reduced from 255 characters to 64 characters
- TLS not supported over IPv6

PBR Enhancements

- Four I/O groups for async policy-based replication
  - Each I/O group can replicate to a maximum of six I/O groups in partnered systems (see illustration)
  - High availability permits, and will continue to permit, only one I/O group to exist in each system
  - Remember that the partnership bandwidth applies to each I/O group connection

DR Recovery Test Function

- Starting a recovery test makes the DR volumes accessible while replication continues
- Can promote the test image to independent if a real disaster occurs
  - Recommended to take a snapshot when starting the test so a ‘true’ recovery point is preserved
- Reduced resynchronization once the test stops
- CLI and REST API only
  - The GUI will show that a DR test is active, but doesn’t allow starting/stopping of a test
  - Initiated from the target side

Increased volumes and capacity

Model	HA/replicated volume count	HA/replicated capacity (per I/O group)
FlashSystem 50x5	Not supported	Not supported
FlashSystem 5200*	7932	1024 TiB
FlashSystem 7200	7932	2048 TiB
FlashSystem 7300	15864 → 16050	2048 TiB
FlashSystem 91x0/9200	7932	2048 TiB
FlashSystem 9500	15864 → 32500	4096 TiB
SAN Volume Controller (SA2/SV2)	7932	2048 TiB
SAN Volume Controller (SV3)	7932	2048 TiB → 4096 TiB
SV Public Cloud*/**	7932	1024 TiB

* Requires minimum of 64GiB of memory per node. ** Async only, only on certain clouds

DRAID Smart Rebuild

- Rebuild, Rebuild-In-Place, Copyback and Scrub will only perform on parts of an array that have extents allocated (extents are allocated for things like vdisks/cdisks, FC snapshots, etc).
  - For example, this would mean a 3PB array with 10% extents allocated that would normally take 30-40 hours to rebuild would now only take 3-4 hours.
- FCM drives only (All Generations)
- DRAID6, DRAID5 and DRAID1 arrays only.
- Progress will be calculated according to the allocated capacity of the array when the rebuild, RIP, copyback or scrub was triggered, not the full capacity.
- Upgrade from a release that does not support extent awareness tasks shall not affect any already running tasks and they shall either complete regularly or be restarted.
- Any DRAID array that has supporting member drives shall become extent aware after the CCU. New tasks will then use the extent awareness feature.

Rebuild Performance Comparison with Extent Awareness On vs. Off

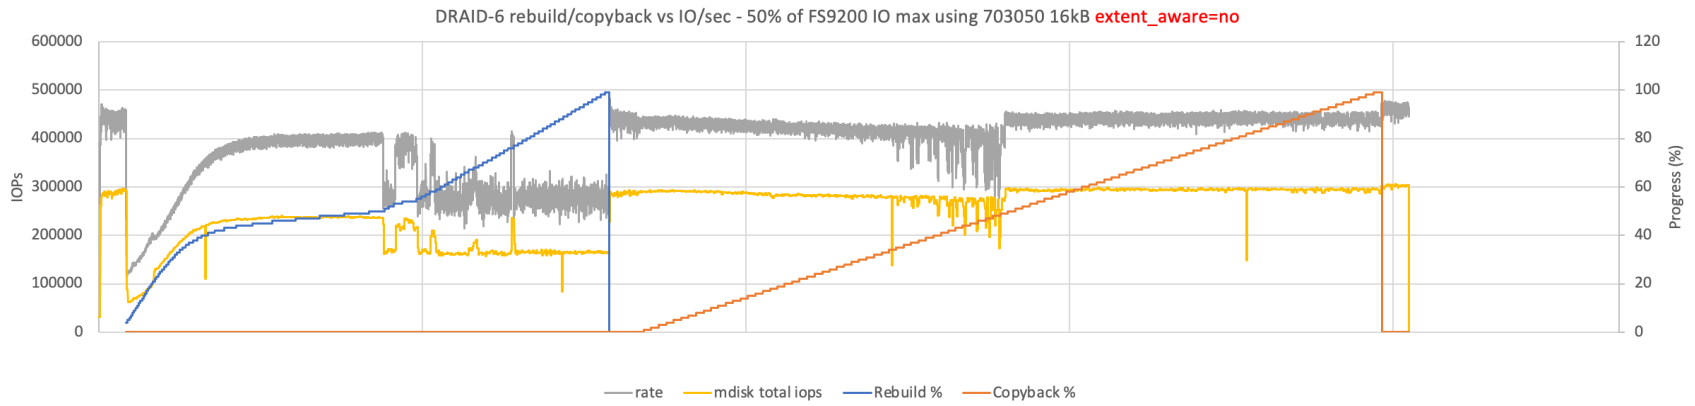

extent_awareness = no
Rebuild time: 9hr 45min
Copyback time: 14hr 59min

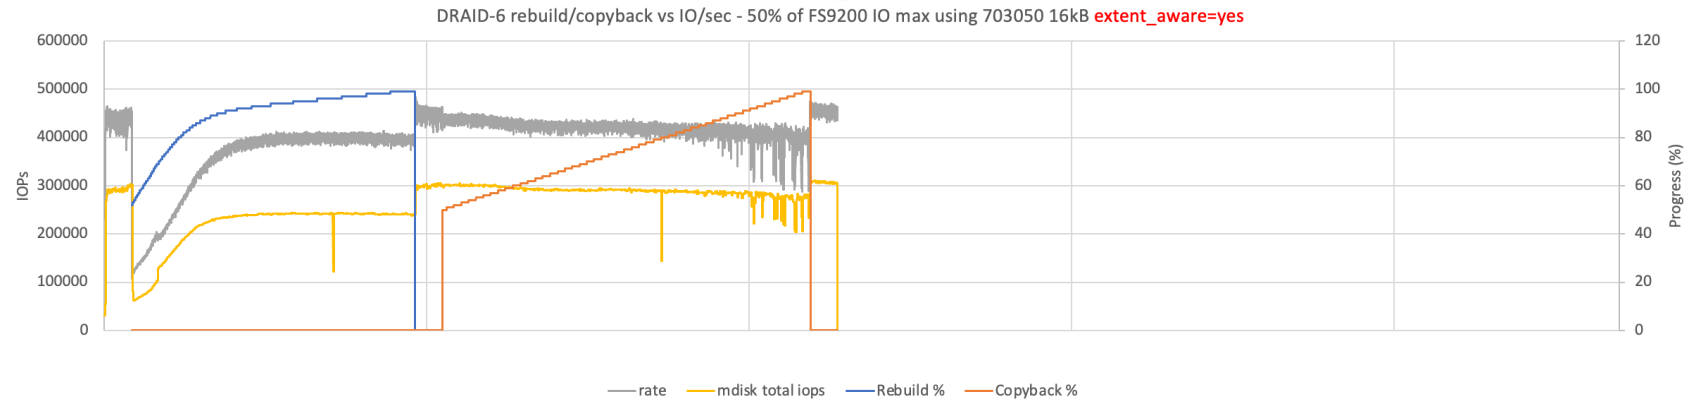

extent_awareness = yes
Rebuild time: 5hr 57min
Copyback time: 7hr 24min

- Test condition: FS9200 w/24x 4.8TB FCMs, 50% physical capacity used
- Note: Extent awareness also lowers host IO impact and the IO impact duration. Host IO is recovered to expected performance sooner.

FIPS - Federal Information Processing Standards

- FCM 1 drives have been certified to FIPS 140-2 level 1
  - FCM1 - NIST Certification Statement
- FCM 2 drives have been certified to FIPS 140-2 level 2
  - FCM2 - NIST Certification Statement
- FCM 3 undergoing FIPS 140-3 certification
- FCM 4 certification is planned

- Due to the time taken (12+ months), only one FIPS firmware per FCM generation is likely to undergo compliance testing
  - Bugs found in FIPS certified versions will likely be fixed with a FIPS non-certified (but compliant) version
- Customers can choose between:
  - FIPS certified
  - FIPS non-certified (compliant)
  - Non-FIPS (validated)
- By default, drives from the factory ship with non-FIPS (validated) firmware
- The firmware on Fix Central is non-FIPs

– Requests for FIPS certified code can be directed to: jorgeesc@mx1.ibm.com

FIPS

Pre 8.6.2

- Data at rest was AES-256 encrypted
- Security PINs were sent to the drive in cleartext over the PCI bus
- A bad actor with physical access could listen to that traffic for the PIN and use that to unlock the drive outside of the Flashsystem enclosure

8.6.2+

- Data at rest is AES-256 encrypted (just like FCM 1 & 2 before)
- Security PINs are sent to the drive in encrypted form over the PCI bus with SKP (Secure Key Passing)
- FCM3: RSA public key cryptography
- FCM4: RSA and CRYSTALS Kyber cryptograph
  - SKP data is encrypted twice, once by each cypher
- CRYSTALS Kyber is a Quantum Safe Cryptography (QSC) algorithm
  - AKA Post-Quantum Cryptography (PQC)

8.6.2+ and SKP (Secure Key Passing)

- SKP is used on all FCM 4 drives, whether in FIPS mode or not
- FCM 4 is a FRU for FCM 3 drives
- From 8.6.2 onwards, if SKP is available it will be enabled
  - Brings additional security benefits, including PQC, to all customers
  - Ensures the SKP code is well tested and reliable for FIPS customers
  - Non FIPS arrays can have mixed SKP support

FIPS Mode

- FIPS mode ensures that everything in the unit is FIPS certified/compliant
- Moving into or out of FIPS mode is data destructive!
  - Do it BEFORE putting data on the box
  - FIPS firmware levels will not be available on fix central – only via support
  - FIPS spec requires that the key used on a FIPS drive is NOT shared with any other drive
  - An existing array cannot be upgraded into FIPS mode
  - Data would need to be copied off, FIPS upgraded, then copied back
  - Could be accomplished with swing storage, but painful

CLI Changes

- lsdrive and lsarray
  - fips_enabled [yes|no] field in the detailed view
  - For lsdrive – the drive is running a FIPS firmware
  - For lsarray – all members are running FIPS firmware
- lsdriveclass
  - fips [yes|no] in the normal and detailed view

FIPS Feature Matrix

Drive	AES-256 Encrypted Data	FIPS Tested	SKP with FIPS	SKP without FIPS	SKP with RSA	SKP with Crystals Kyber
FCM 1	✓	✓	×	×	×	×
FCM 2	✓	✓	×	×	×	×
FCM 3	✓	In Progress	✓	3.2 PTF onwards	✓	×
FCM 4	✓	Planned	✓	✓	✓	✓

Scalability Increases

Limit	5200	7300	9500	SV3
Host Mappable Volumes	8,192	16,050	32,500	15,864
Snapshots	15,863	32,099	64,999*	15,863
Snapshot Capacity	4PiB	20PiB	40PiB	40PiB**
Policy Based Rep Async volumes	7,932	16,050	32,500	7,932

* Assuming 1 volume on the system
*** 40PB assuming no volume mirroring. Bitmap space is shared between snapshots and volume mirroring.

GUI Enhancements

- Duration field added to the audit log
  - Tracks how long a process takes
- GUI support for security patch installation
- Wattage consumption added to dashboard
- New statistics intervals added to performance screens

Proof Keycode Exchange (PKE) Support

- Improves security between FlashSystem and supported SSO/MFA providers
- The security admin must enable PKCE support in their SSO or MFA application definition
  - FlashSystem storage admin configures connection to SSO / MFA provider, but no setup needed for PKCE
- Allows the SSO or MFA provider to verify that the FlashSystem that initiated the authorization flow is the same FlashSystem that completes the flow.
- Based on a unique token called a code_challenge that is compared to a code_verifier token to ensure integrity of the authorization flow

Statistics Improvements

- Add power usage and temperature for SVC nodes (lsnodestats)
- Increase available history with minute, hour and daily averages
- Save history to file to avoid loss on node failover or T2
- Add power and temperature to the XML performance stats (Nn_stats)
- Include lssystemstats -days in the daily inventory

Cascading Statistics

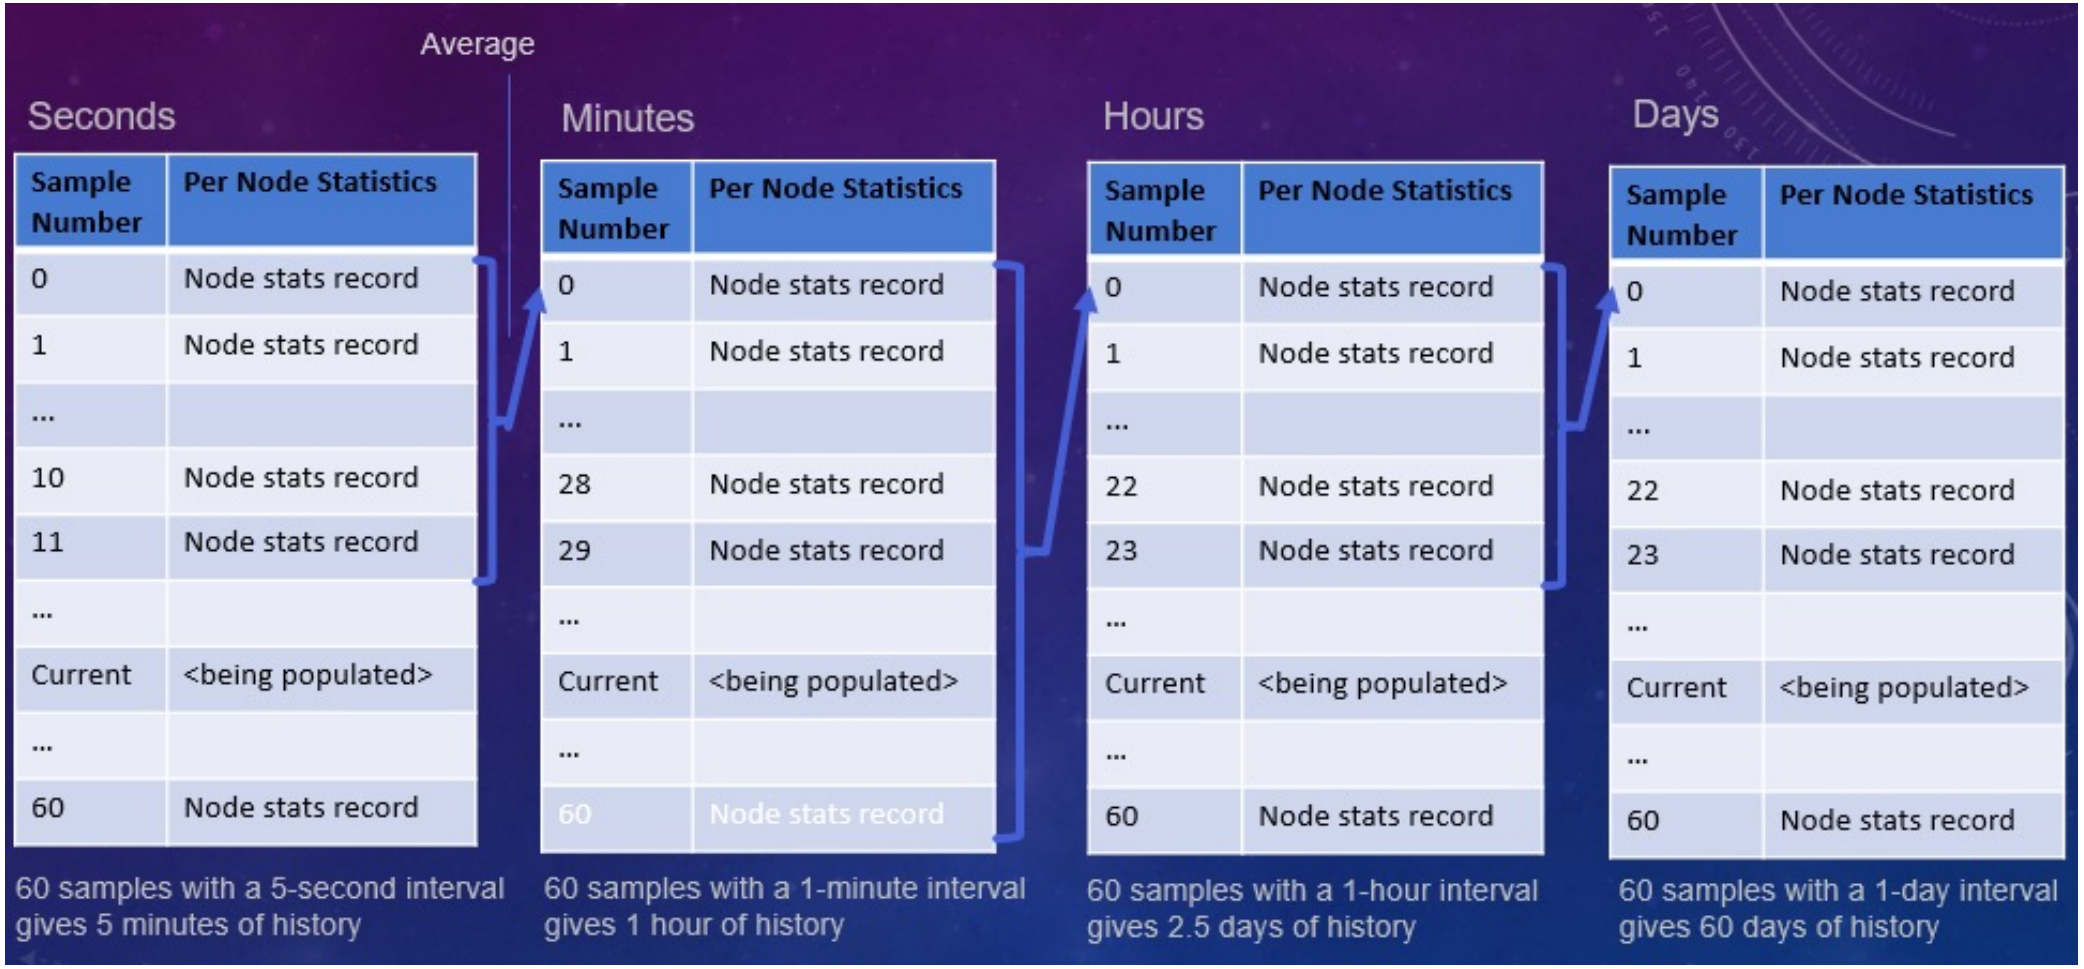

Sample Number	Per Node Statistics
0	Node stats record
1	Node stats record
...	
10	Node stats record
11	Node stats record
...	
Current	<being populated>
...	
60	Node stats record

Sample Number	Per Node Statistics
0	Node stats record
1	Node stats record
...	
28	Node stats record
29	Node stats record
...	
Current	<being populated>
...	
60	Node stats record

Sample Number	Per Node Statistics
0	Node stats record
1	Node stats record
...	
22	Node stats record
23	Node stats record
...	
Current	<being populated>
...	
60	Node stats record

Sample Number	Per Node Statistics
0	Node stats record
1	Node stats record
...	
22	Node stats record
23	Node stats record
...	
Current	<being populated>
...	
60	Node stats record

Stats CLI Changes

- New parameters: –minutes / -hours / -days for existing CLI lsnode(canister)stats, lsenclosurestats and lssystemstats
  - Can be used with or without –history
  - Inventory includes lssystemstats –days:

stat_name	*stat_current*	*stat_peak stat*	*peak_time*
compression_cpu_pc	*0*	*0*	*231023124821*
cpu_pc	*1*	*1*	*231023124821*
…			
power_w	*1162*	*1162*	*231023124821*
temp_c	*23*	*23*	*231023124821*
temp_f	*73*	*73*	*231023124821*

Integration Updates

- IBM Sentinel now supports multifactor authentication
- PowerHA SystemMirror 7.2.8 that now integrates SGC2 into PowerHA
  - Announced Oct10 and GA’s Dec 15

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 1708 6924 or

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

Or

QR Code