HPE GreenLake

Technical white paper

Deployment considerations for HPE MSA Gen6 with VMware vSphere 7

Contents

Executive summary

This technical paper discusses deployment considerations and best practices for the sixth-generation HPE MSA 1060/2060/2062 with VMware vSphere® 7. This document explores the unique capabilities and integration of HPE MSA Storage with vSphere. In addition, the document covers the recommended settings and configurations needed to optimize HPE MSA Storage for deployment with vSphere. When deployed together, HPE MSA Storage with vSphere provides small and midsize businesses with the benefits of increased consolidation savings by increasing virtual machine density, lowering storage costs, and realizing time savings from simplified storage management and provisioning.

This document outlines recommendations for configurations, software settings, and design architectures to get best results for HPE MSA storage systems with vSphere. The document outlines the features and recommended configurations for the HPE MSA 1060/2060/2062 product line.

This document is not a user guide and does not list all features or explain how to configure them. For detailed information regarding the features of an HPE MSA sixth-generation storage system, consult the links on the last page of this document.

Intended audience

This document is intended for IT administrators, vSphere administrators, and solution architects planning a server virtualization deployment with sixth-generation HPE MSA storage systems. This and other documents pertaining to virtualization with HPE Storage and VMware® software are available at hpe.com/us/en/alliance/vmware.html and hpe.com/storage.

VMware vSphere/VMware ESXi™ administrators planning to set up hosts with HPE MSA Storage should have a working knowledge of storage area network (SAN) concepts.

HPE MSA Gen6 product overview

The HPE MSA Gen6 storage system features an active/active architecture that provides both flexibility and resiliency to failure. It ships in a rack-mountable 2U form factor that contains:

- Disk drive bays (either 24 x SFF[1] or 12 x LFF[1, 2])
- Two hot-swappable power supplies units, each with integrated cooling fans
- Two hot-swappable controller units
- A passive midplane to which all components are connected
- Optional lockable bezel

HPE MSA array enclosures contain Fibre Channel, iSCSI, or SAS controller modules and optional expansion disk enclosures that house additional disk drives. Expansion disk enclosures include I/O modules in place of controller modules that provide SAS connectivity between disk drives and controller units.

[1] Small form factor (SFF) drives are 2.5-inches; large form factor (LFF) drives are 3.5-inches
[2] HPE MSA 2060, 2062, and expansion enclosures only

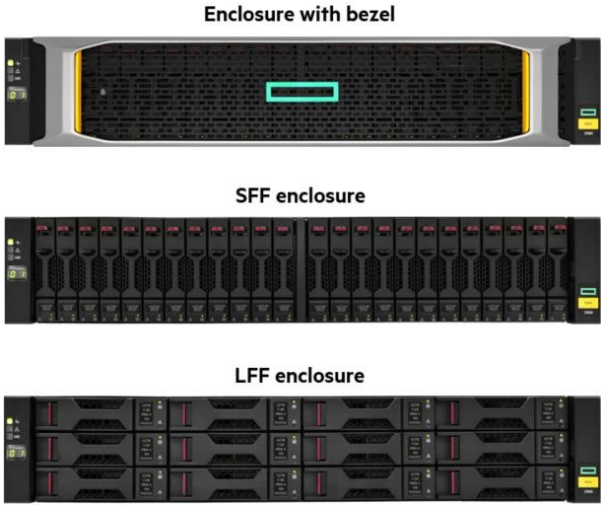

Figure 1. HPE MSA enclosures

HPE MSA 1060 storage systems have a total of four host ports compared to the eight of the HPE MSA 2060 and 2062 arrays. However, HPE MSA 1060 SAS models support an optional fan-out cable that doubles the host port count by reducing the number of SAS lanes per port from four to two. Nevertheless, the fan-out cable provides increased scalability without compromising the performance of the array. HPE advises using fan-out cables even if they are not initially needed to avoid interruptions when connecting additional hosts.

Figure 2. Rear view of an HPE MSA 1060 array enclosure

Figure 3. Rear view of an HPE MSA 2060/2062 array enclosure

The HPE MSA 1060 array enclosures ship in SFF only but support a mix of up to three LFF and SFF expansion disk enclosures. This configuration allows for a total of 96 SFF drives per array or 36 LFF drives and 24 SFF drives. HPE MSA 2060 and 2062 arrays support up to nine expansion disk enclosures and up to 120 LFF drives or 240 SFF drives.

Important
Sixth-generation HPE MSA storage systems do not support disk enclosures or drives from previous generations of HPE MSA.

Figure 4. HPE MSA array enclosure naming

HPE MSA Storage features and concepts

Growing storage needs for virtualized servers require greater levels of storage performance and functionality at a lower cost of ownership. HPE MSA Gen6 storage systems are positioned to provide excellent value for SMB customers that need increased performance to support initiatives such as consolidation and virtualization.

HPE MSA Storage is built for virtualization environments and delivers key advantages for vSphere:

- High-performance I/O throughput to meet the peak virtualization demands of cloning or migrating multiple VMs
- A SAN attached storage array with cost-effective scale-up options
- Virtualized storage technology that provides non-disruptive scalability to VMs
- Integration with key vSphere Storage APIs, including vSphere Storage APIs for Array Integration (VAAI) and VMware Site Recovery Manager™ (SRM)
- Easy-to-use web-based HPE Storage Management Utility (SMU)
- Integration with HPE Storage Integration Pack for VMware vCenter®
- Support for 64 TB vSphere Virtual Machine File System (VMFS) datastores

Architecting the HPE MSA with VMware

Because storage plays a critical role in the success of a vSphere deployment, proper configuration of the HPE MSA Gen6 system is of great importance. This section covers recommended HPE MSA configuration settings that help you get the best experience with your VMware environment.

HPE MSA disk group considerations

MSA-DP+ is a RAID type that offers superior performance, availability, flexibility, and rebuild times when compared to other RAID types, especially those that employ parity-based RAID schemes.

Best practice: Use MSA-DP+ disk groups for both the standard and archive tiers (HDDs) in the HPE MSA Gen6 array. MSA-DP+ is only recommended for the performance tier when configured as an all-flash pool.

Naming hosts and host groups

Create meaningful host and host group names for ease of use when mapping HPE MSA volumes for use with vSphere.

Best practice: Create friendly hosts and host group names as seen in the following two-node cluster example.

Two-node cluster example: Host group name: “ESXi7_cluster_1”

Host name: “ESXi7_host_1”

Initiator 1 nickname: “ESXi7_host_1_port_0”

Initiator 2 nickname: “ESXi7_host_1_port_1”

Host name: “ESXi7_host_2”

Initiator 1 nickname: “ESXi7_host_2_port_0”

Initiator 2 nickname: “ESXi7_host_2_port_1”

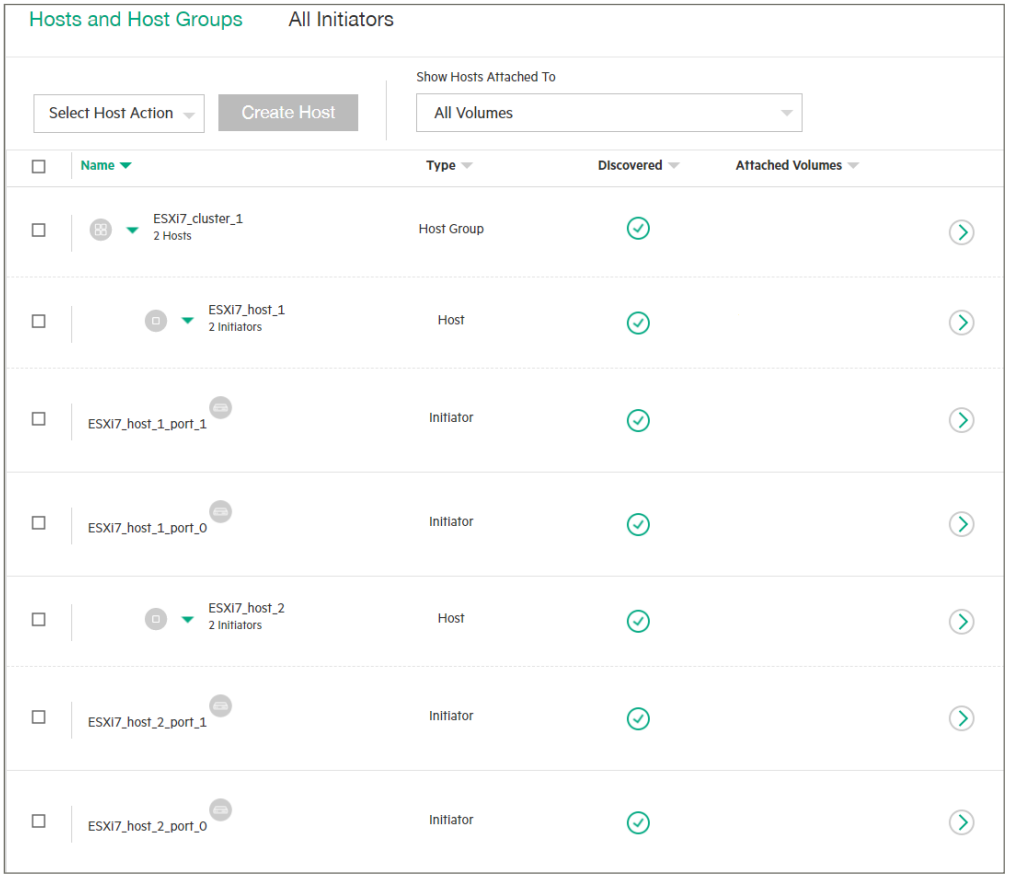

Figure 5. HPE MSA hosts and host groups screen

HPE MSA considerations and best practices for vSphere

The following section highlights recommended practices for setup and configuration of the HPE MSA Gen6 arrays best suited for virtualization environments running vSphere.

HPE MSA storage configuration

Configuring HPE MSA arrays correctly is important in vSphere. The following section highlights some typical HPE MSA and vSphere scenarios and the recommended configuration settings.

Scenario 1: You require datastore volumes that do not exceed the single pool capacity of the MSA storage system.

Best practice: Create datastore volumes using one storage pool.

Detail: Create a single pool using a standard tier created with MSA-DP+. With this configuration, all storage presented has the same performance and capabilities.

This configuration also allows for incremental expansion of the existing MSA-DP+ disk group or expansion using a performance or archive tier.

Single controller HPE MSA performance and addressable capacity scales beyond the requirements of most customers. Additionally, single controller configurations ensure the minimum impact to performance in the unlikely event of controller unavailability during a peak in I/O demand.

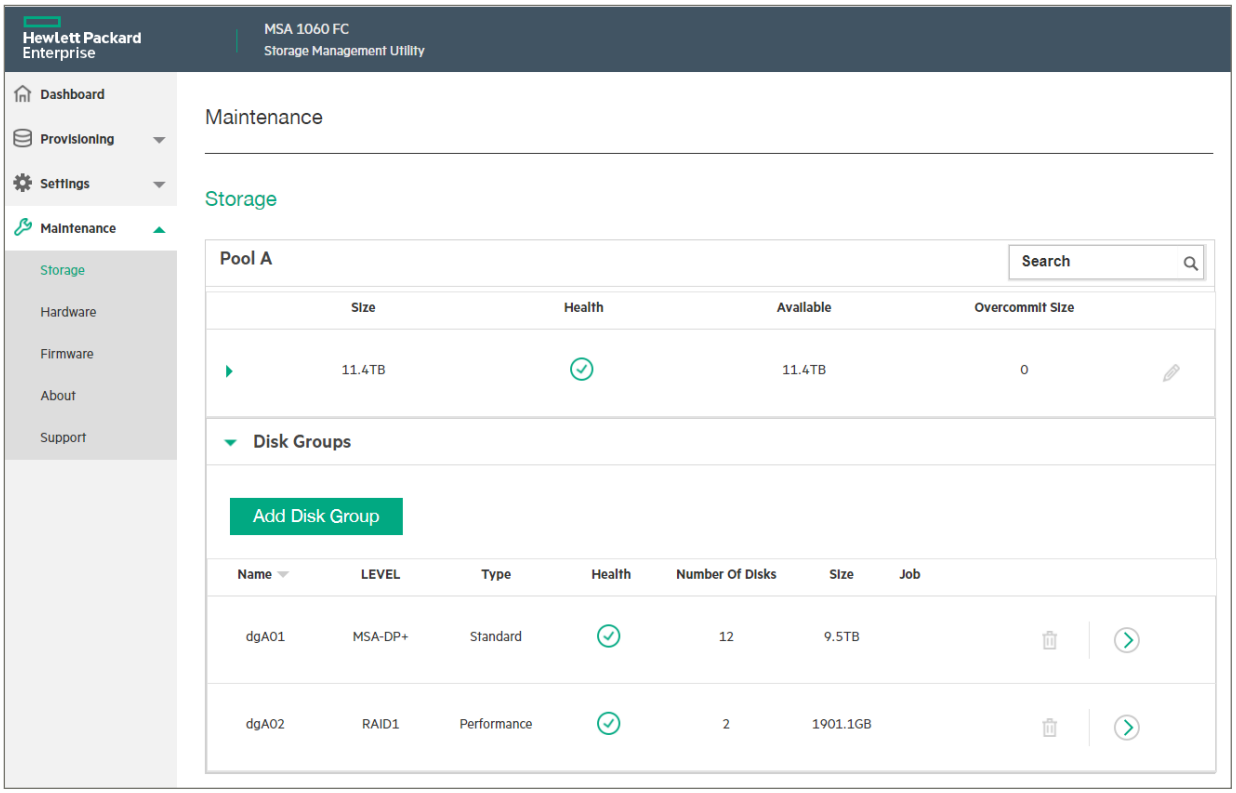

Figure 6. Single pool with standard and performance tiers

Scenario 2: You require datastore volumes and exceed the performance and capacity requirements of a single pool.

Best practice: Create volumes on both storage pools.

Detail: Using multiple storage pools can increase the capacity limitations and adding the second pool increases overall performance because both controllers are used. However, to ensure that performance is equal for all volumes, configure each pool using the same disk group configuration and distribute volumes and application loads evenly between pools.

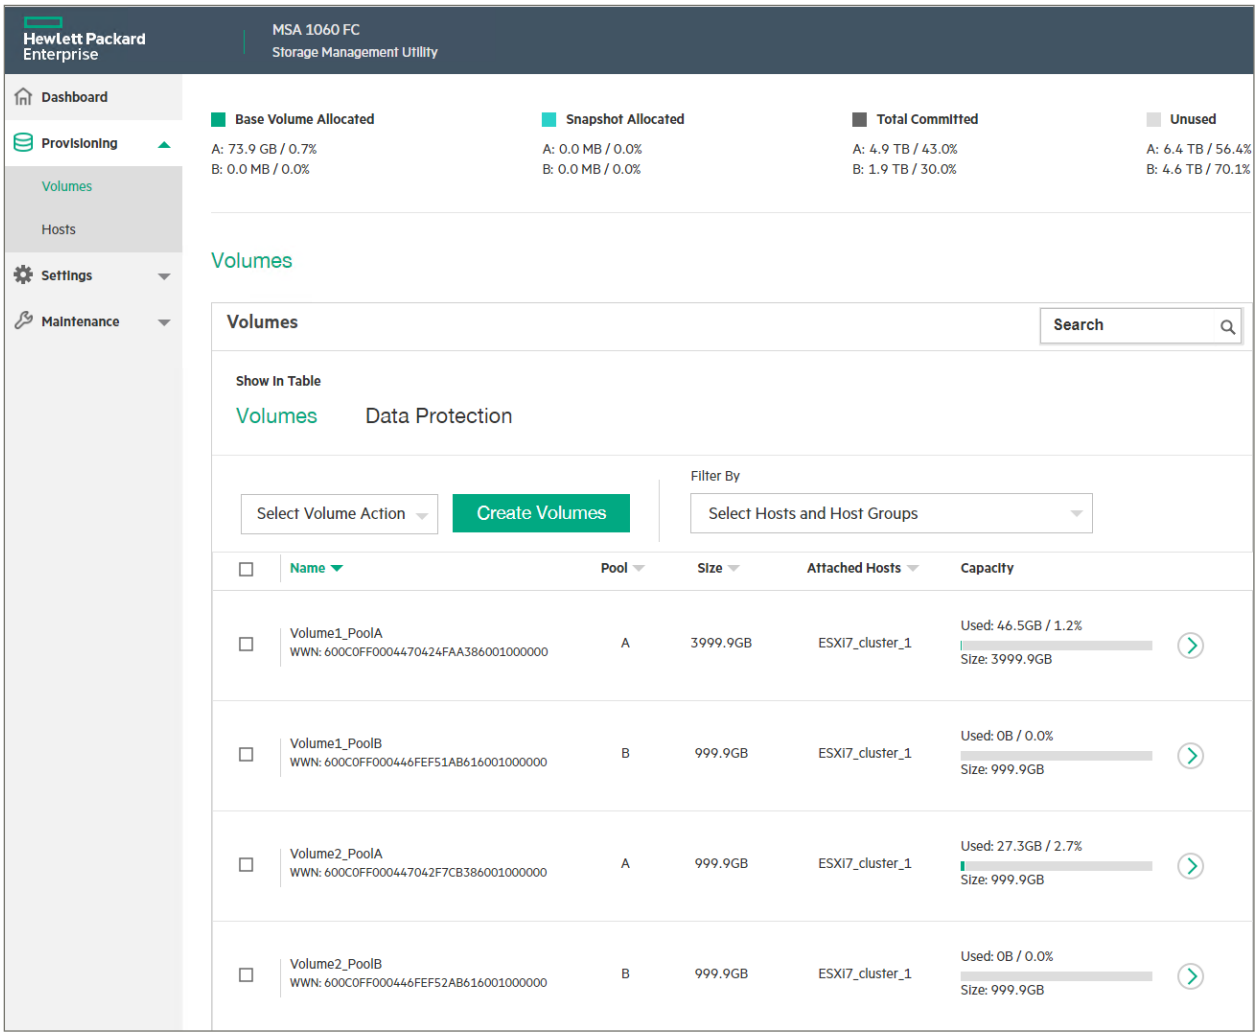

Figure 7. Dual pools with volumes mapped to a host group

Scenario 3: You require virtual machines or storage with different performance needs.

Best practice: Create multiple pools to segregate volumes on specific drive types.

Detail: Create an all-flash pool for the high-performing VMs and volumes. Next, create a pool of spinning drives and place the lower performing VMs on the volumes of this pool.

Creating multiple volumes configured in this manner allows for volumes dedicated for highest performance on one pool and volumes for spinning drive performance on another pool.

Note

Configuring both HDDs and SSDs as capacity disk groups within the same system requires an HPE MSA Advanced Data Services (ADS) license. Using an HPE MSA 2062 storage system would be a good starting point for this scenario because it includes the ADS license and SSD capacity.

Boot from SAN

When booting ESXi hosts from SAN, each HPE MSA volume used for booting the host should be mapped to only one ESXi host. Each ESXi host must have access only to its own boot HPE MSA volume and not of others.

Mapping a boot volume to multiple hosts should be used only when recovering from host hardware failure or device changes. In this instance, map the boot volume to the replaced unique device IDs such as World-Wide Port Names (WWPNs), IP addresses, or IQNs introduced when replacing an ESXi host or host bus adapter (HBA). After the new hardware devices have been mapped, the old volume mapping should be removed.

For more information regarding vSphere installation and boot from SAN configurations, refer to: VMware ESXi Installation and Setup.

Missing LUN Response setting

ESXi will not search for other volumes to which it may have access if it does not discover LUN 0 or if LUN IDs are not contiguous. The Missing LUN Response parameter handles these situations by enabling the host drivers to continue probing for LUNs until they reach the LUN to which they have access. This parameter controls the SCSI sense data returned for volumes that are not accessible because they do not exist or have been hidden through volume mapping. The Illegal Request parameter sends a reply that there is a LUN but that the request is illegal.

Best practice: Keep the default Missing LUN Response value of "Illegal Request" for HPE MSA Gen6 arrays when using vSphere.

Use the following command from the HPE MSA Gen6 storage system to reset it to the factory default:

```
# set advanced-settings missing-lun-response illegal
```

Refer to the HPE MSA 1060/2060/2062 CLI Reference Guide for further guidance on how to apply this setting.

ESXi handling of SCSI Queue Full and Busy messages

ESXi uses an algorithm to respond to Queue Full and Busy SCSI messages from storage systems. HPE MSA Gen6 arrays do not require these parameters to be enabled. These settings are disabled by default (value of 0) on HPE MSA arrays.

Use the following command from the ESXi host to verify the Queue Full Sample Size and Queue Full Threshold values of an HPE MSA array:

```
# esxcli storage core device list -d naaID

       Queue Full Sample Size: 0

       Queue Full Threshold: 0
```

If these values are not disabled (value of 0), use this command to set these values to 0:

```
# esxcli storage core device set -d naaID --queue-full-threshold 0 --queue-full-sample-size 0
```

Best practice: Keep the default values for QFullSampleSize and QFullThreshold.

Multipath considerations

ESXi software supports multipathing to maintain a redundant connection between a vSphere host and storage. To take advantage of this feature, the ESXi host requires multiple paths to the HPE MSA volumes. Multiple ports or multiple adapters on the ESXi hosts are recommended for improved performance and availability.

vSphere supports active/active multipathing to maintain a redundant connection between the ESXi host and the HPE MSA array. HPE MSA arrays and vSphere have been tested with three path policies:

- Fixed
- Most Recently Used (MRU)
- Round Robin

By default, ESXi systems use only one path from the host to a given volume at any time. This is defined by the path selection policy named "MRU path." If the path actively being used by the ESXi system fails, the server selects another of the available paths. Path failover is the detection of a failed path by the built-in ESXi multipathing mechanism, which switches to another path by coordinating MPIO software, VMware Native Multipathing (NMP), and the HPE MSA firmware.

The default storage array type for the HPE MSA system is VMW_SATP_ALUA. By default, the path selection policy is set to use the Most Recently Used (VMW_PSP_MRU) path. This selects the first working path discovered during boot up. If the path becomes unavailable, it moves to another path.

Best practice: Change the VMware preferred selection path (PSP) policy to **Round Robin** on all ESXi hosts for best performance and load balancing on the HPE MSA array.

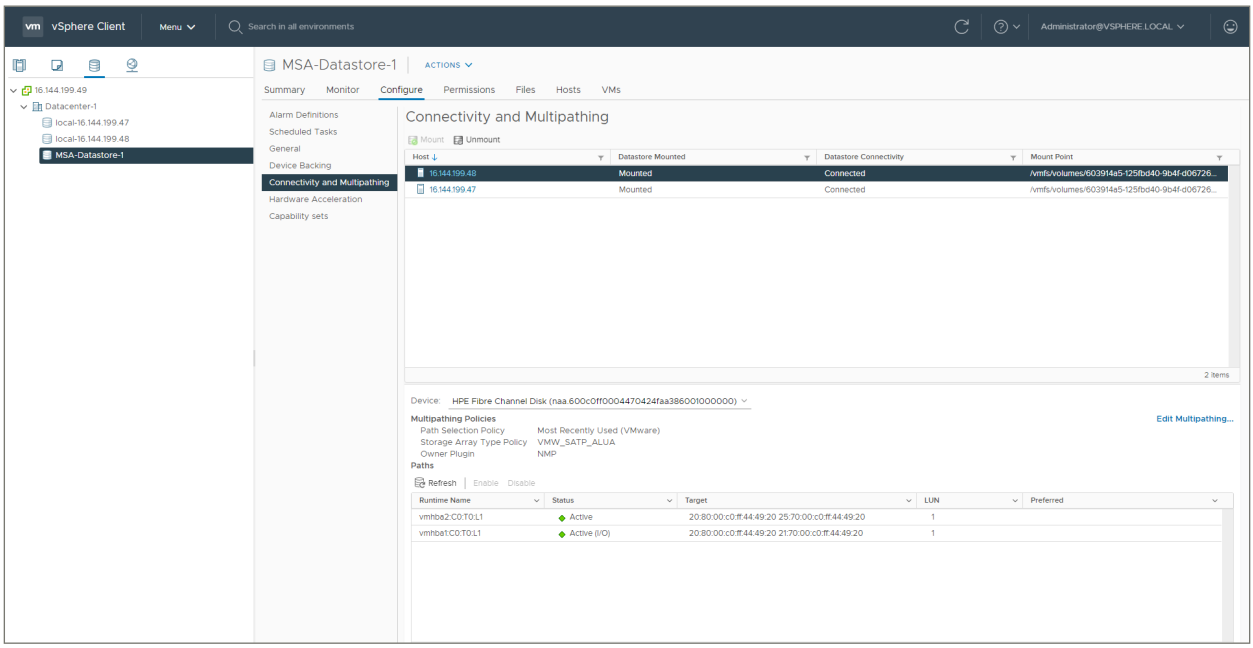

Figure 8. vCenter volume multipathing screen

The Round Robin PSP balances the load across all active storage paths. The storage path will be chosen and used until a specific quantity of data has been transferred. The PSP chooses the next path in the list after that quantity is reached. The quantity at which a path change is triggered is known as the "limit."

Round Robin PSP supports two types of limits:

- **IOPS limit:** The default IOPS limit is 1000. This means a new path will be used after 1000 I/O operations are reached.
- **Bytes limit:** The default bytes limit is 10,485,760. This means a new path will be used after the bytes limit is reached.

In summary, Round Robin will attempt to rebalance paths after every 1000 I/O operations or 10,485,760 bytes.

Best practice: Keep the default IOPS and bytes limit. HPE MSA storage arrays do not see a benefit by changing these limits.

HPE MSA Gen6 direct attach support

The HPE MSA Gen6 array can be beneficial if you chose to deploy a direct attach environment using vSphere. For example, the HPE MSA Gen6 SAS array is an ideal deployment for a small VMware cluster configuration because of its simplified configuration. The HPE MSA SAS direct connect VMware cluster has all of the performance of a Fibre Channel configuration, but at a lower cost.

Note

HPE MSA Gen6 arrays currently do not support direct attached configurations with vSphere 7 via iSCSI. Network switches are compulsory when using vSphere 7 and iSCSI with HPE MSA Gen6 arrays.

VMware does not currently provide a testing methodology for iSCSI configurations as it does with Fibre Channel and SAS. Therefore, HPE MSA Gen6 arrays do not support iSCSI direct attach.

Best practice: Consult HPE Storage Single Point of Connectivity Knowledge (SPOCK) for supported connectivity options.

HPE MSA considerations and best practices for vCenter

The following section highlights recommended configuration practices of the HPE MSA Gen6 arrays from within VMware vCenter.

Datastore clusters

One of the features that simplify the management of storage in vCenter is the datastore cluster. Creating a datastore cluster for the HPE MSA volumes allows for vCenter automations to help balance the storage load. For example, when a VM is created, the load is balanced across the volumes of the datastore cluster.

Another benefit of the datastore cluster is that when the datastore is created, the vCenter software shows only those volumes that are visible to all machines associated with the datastore cluster. This provides a simple way of verifying that mappings on the HPE MSA array have been configured properly.

Note

In the vCenter administration tools, use the datastore cluster configuration to help distribute VMFS creation for VMs being added to the data center.

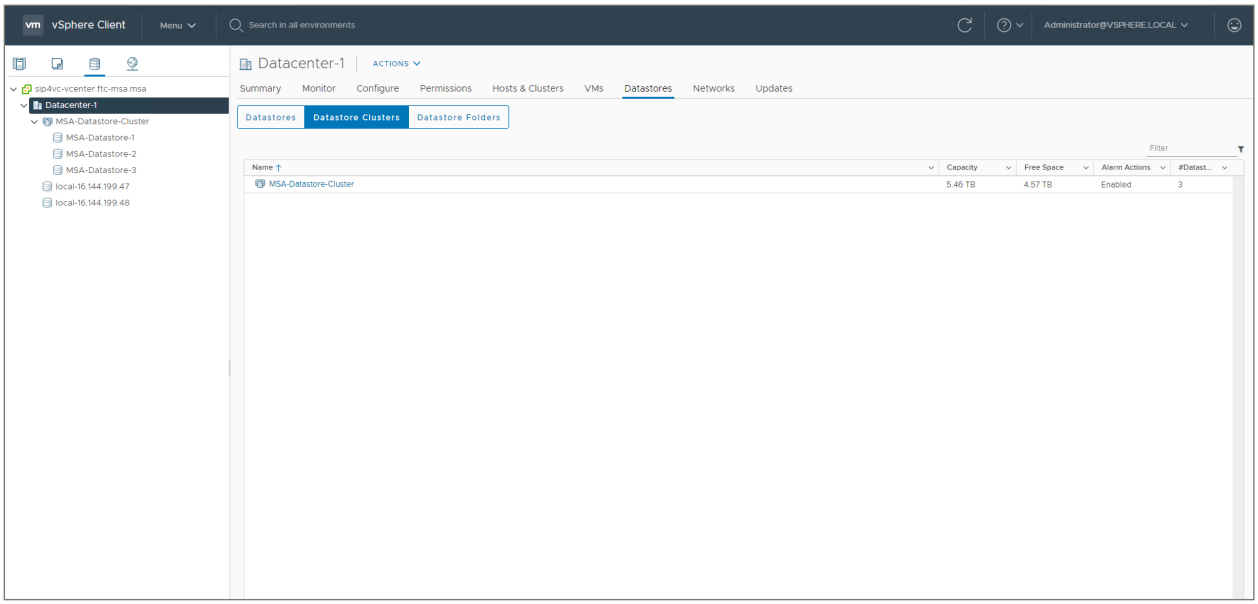

Figure 9. VMware vCenter Datastore Clusters screen

Storage Distributed Resource Scheduler

vSphere datastore clusters provide the Storage Distributed Resource Scheduler (DRS) feature with I/O load balancing. The intent of this feature is to balance and distribute VM storage needs across traditional physical disk-based LUNs.

Because the HPE MSA controllers and HPE MSA volumes provide the benefits of Storage DRS and I/O balancing dynamically, Storage DRS and I/O load balancing are not needed with HPE MSA volumes if all hosts in the cluster share the same HPE MSA volume mappings.

Note

In the vCenter administration tools, do not enable Storage DRS and I/O load balancing on data storage clusters that are solely based on HPE MSA volumes.

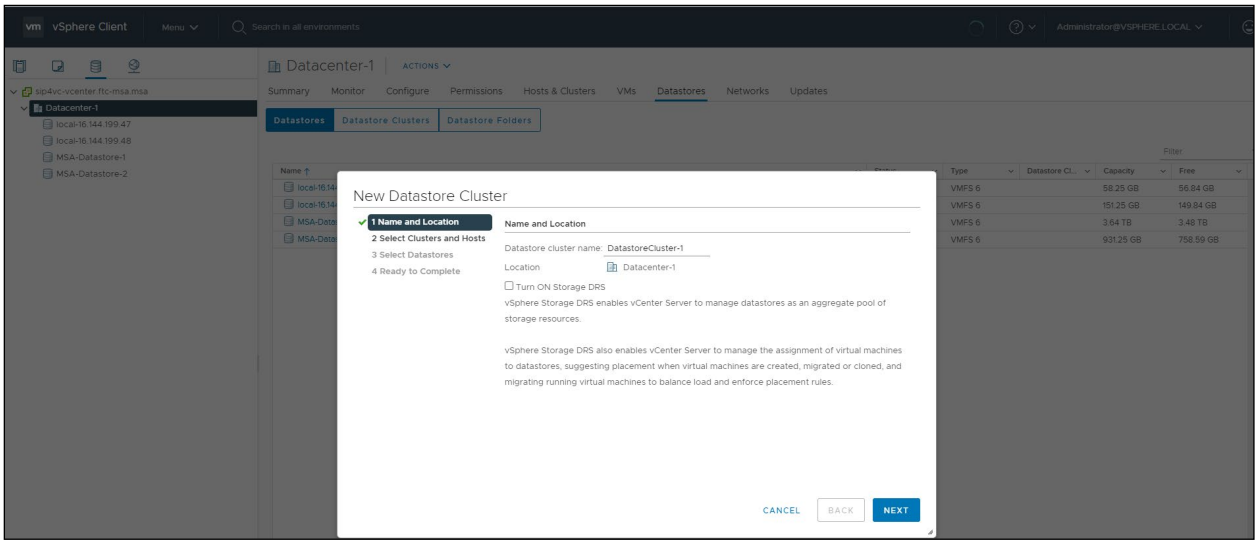

Figure 10. VMware vCenter Datastore Cluster Creation screen

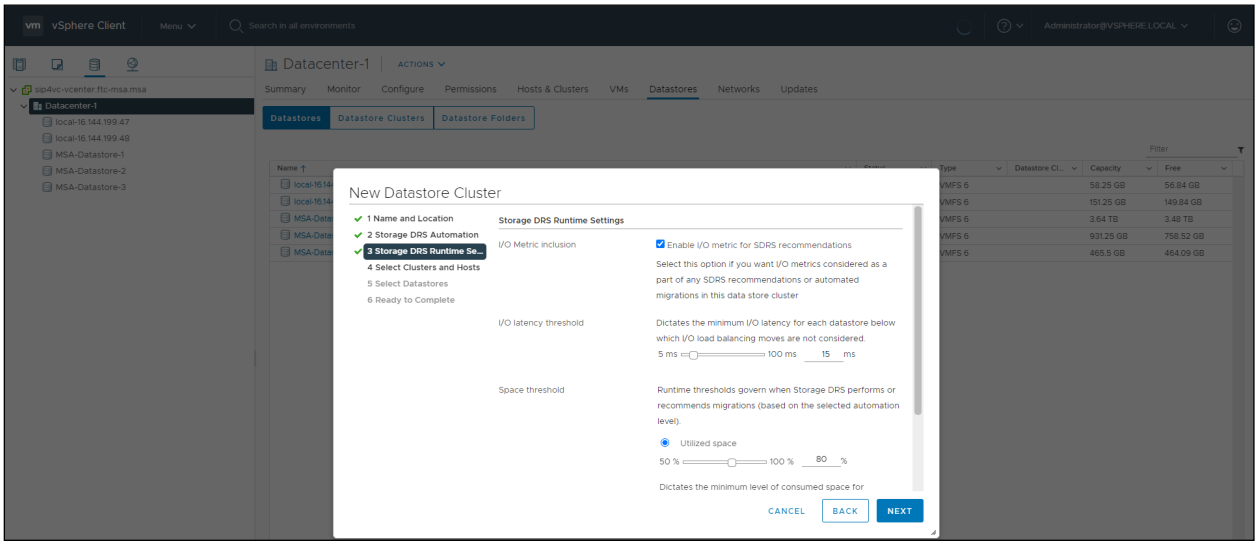

Figure 11. VMware vCenter Storage DRS Configuration screen

VMFS and raw device mapping

Virtual machines can access data using two methods: VMFS (.vmdk files in a VMFS file system) and raw device mapping (RDM). The only difference between them is that RDM contains a mapping file inside VMFS that behaves like a proxy to the raw storage device. The virtual machine can access and use the storage device directly. RDMs contain metadata for managing and redirecting disk access to the physical device.

Sample use cases for RDMs include:

- SAN snapshot or other layered applications run in the virtual machine. RDM enables backup offloading systems by using features inherent to the SAN storage array.
- A Microsoft Cluster Server (MSCS) running as a guest that spans physical hosts, such as virtual-to-virtual clusters and physical-to-virtual clusters. In this scenario, cluster data and quorum disks are configured as RDMs rather than as virtual disks on a shared VMFS.

There are two compatibility modes for RDMs:

- **Virtual mode:** The RDM acts like a virtual disk file and can use snapshots.
- **Physical compatibility mode:** The RDM offers direct access to the SCSI device for those applications that require lower-level control.

Important
Due to design limitations of SAS HBAs, RDM is not supported with SAS models of HPE MSA arrays.

Because HPE Gen6 MSA arrays offer the ability to map volumes to virtual machines directly via iSCSI, physical mode RDMs are not required.

Best practice: Use Virtual Machine Disks (VMDKs) when creating virtual machines except for the use cases listed for RDM.

VMware provides three VMDK provisioning types on VMFS datastores:

- Thin Provision
- Thick Provision Eager Zeroed
- Thick Provision Lazy Zeroed

Table 1. Provisioning types for datastores

VMDK provisioning type	VMware behavior	HPE MSA Gen6 behavior	HPE MSA Gen6 recommendation
Thin	The VMDK is allocated and zeroed on first write.	Unless pool over-provisioning is disabled, capacity usage on the HPE MSA will eventually be thin, regardless of the chosen VMDK provisioning type. With the default pool setting of over-provisioned enabled, zeroed pages within the pool are automatically reclaimed over time, ranging from hours to days. Therefore, eager zeroing an area of the VMFS file system will not necessarily lead to the measurable performance gain seen in other storage arrays. It could, however, cost valuable short-term performance and may also negatively affect the tiering engine of the HPE MSA. **Note:** HPE does not recommend disabling the pool over-provisioning if using array-based snapshots with HPE MSA Gen6 storage systems. In such a scenario, pool capacity usage would be greatly amplified and difficult to manage. **Note:** Although VMFS 6 introduced automatic space reclamation when using thin VMDK files, the 1 MB file system granularity is incompatible with the 4 MB page file size of the HPE MSA. It is therefore recommended that a periodic manual UNMAP be run from the ESXi command line to ensure unused capacity is returned to the pool.	Yes
Thick Eager Zeroed	All of the VMDK space is allocated and zeroed out at the time of creation.		No
Thick Lazy Zeroed	All of the VMDK space is allocated at the time of creation, but each block is zeroed only on first write.		No

Use the following command to UNMAP the MSA datastore:

```
# esxcli storage vmfs unmap -l MyDatastore
```

Best practice: Use the Thin VMDK provisioning type.

Note
Refer to the HPE MSA Gen6 Virtual Storage Technical Reference Guide for tiering information.

HPE MSA Storage Replication Adapter and vCenter Site Recovery Manager

The HPE MSA Storage Replication Adapter (SRA) integrates with VMware vCenter Site Recovery Manager (SRM) to manage replication features on HPE MSA storage systems. Combine SRA with the Remote Snap replication to automate failover, failback, and data migration between geographically separated sites.

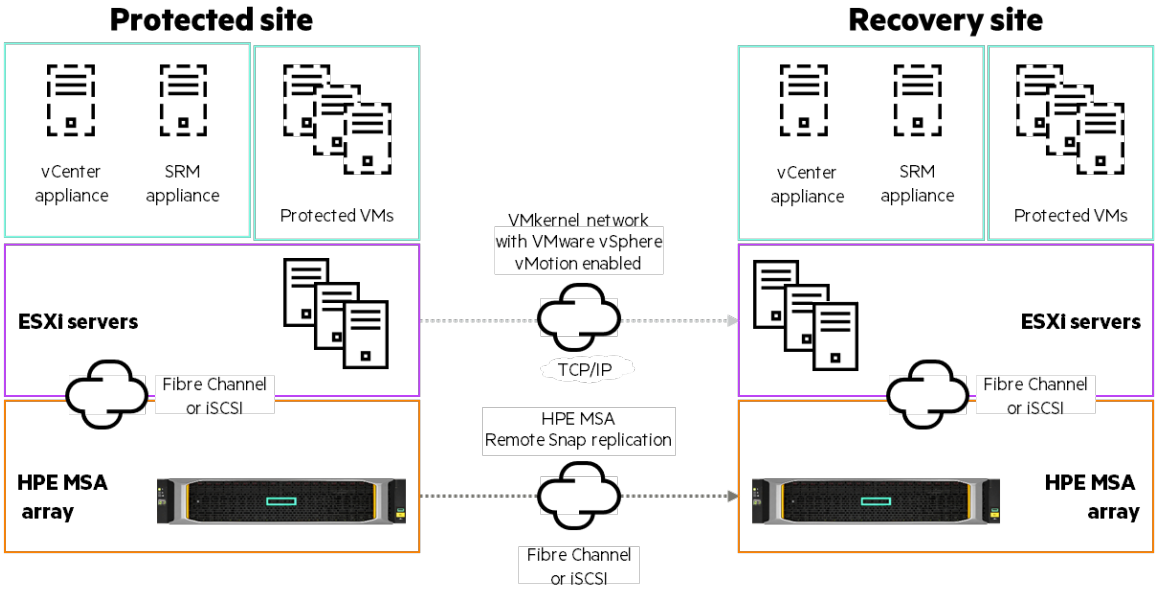

Figure 12. Array-based replication of VMs with HPE MSA and SRM

Best practice: Use SRA/SRM for greater availability.

For SRA configuration and troubleshooting information, consult the HPE MSA 1060/2060/2062 Storage Replication Adapter User Guide.

Near synchronous replication with Zerto, a Hewlett Packard Enterprise Company

You can leverage Zerto Virtual Replication to replicate applications and data from one HPE MSA array to another. Popular use cases include departmental HPE MSA storage replicated to enterprise storage, enterprise storage replicated into an HPE MSA array, or protection of HPE MSA workloads in the cloud.

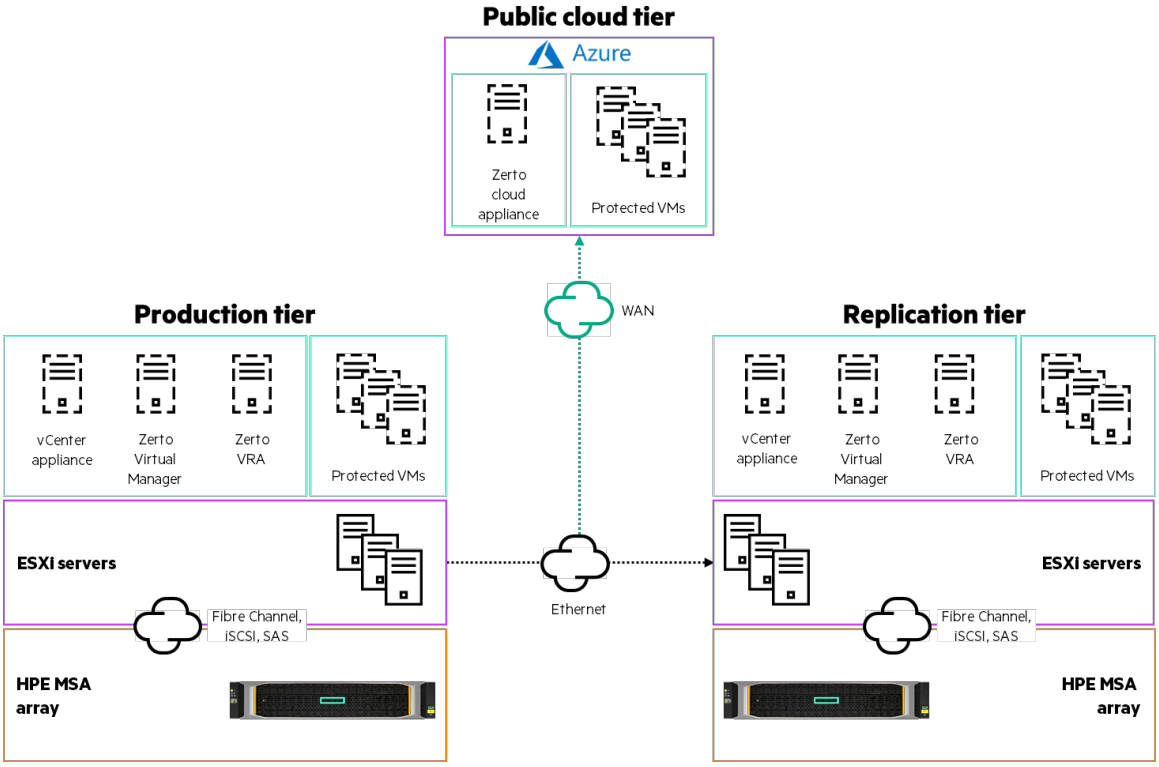

Figure 13. Three-tier replication using HPE MSA storage, Zerto, a Hewlett Packard Enterprise Company, and Windows Azure

For more information on the HPE MSA and Zerto solution, consult this technical paper: hpe.com/psnow/doc/a00053770enw?from=app&section=search&isFutureVersion=true

HPE Storage Integration Pack for VMware vCenter

HPE Storage Integration Pack for VMware vCenter (SIP4VC) is a component within the HPE OneView for vCenter plug-in. It enables vSphere administrators to quickly obtain context-aware information and manage supported HPE storage devices such as the HPE MSA in their vSphere environment directly from within vCenter. This plug-in operates independently of the core HPE OneView product and does not require a license to use. By providing a clear relationship across VMs, datastores, and storage, the VMware administrator's productivity increases, as does the ability to ensure quality of service. Roles for administrators can be defined on an individual basis, providing the ability to apply specific permissions for both view and control functions.

HPE Storage Integration Pack for VMware vCenter supports mixed array environments including HPE MSA Storage and other HPE Storage systems including HPE Primera Storage and HPE 3PAR Storage.

When deployed with an HPE MSA array, HPE Storage Integration Pack for VMware vCenter:

- Enables active management for HPE MSA arrays:
  - Create/expand/delete a datastore
  - Create a virtual machine from a template
- Monitors the health and status of the HPE MSA array
- Displays LUN/volume connections from VMs and vSphere servers to the arrays and provides the location and attributes of the HPE MSA array within the SAN
- Identifies which storage features are available to allow administrators to match the features available on the HPE MSA array to their requirements
- Provides a cluster-level view of the storage

HPE Storage Integration Pack for VMware vCenter is downloadable from My HPE Software center: myenterpriselicense.hpe.com/cwp-ui/free-software/HPESIPVC.

Best practice: Implement HPE Storage Integration Pack for VMware vCenter for ease of management and monitoring of storage from within vCenter.

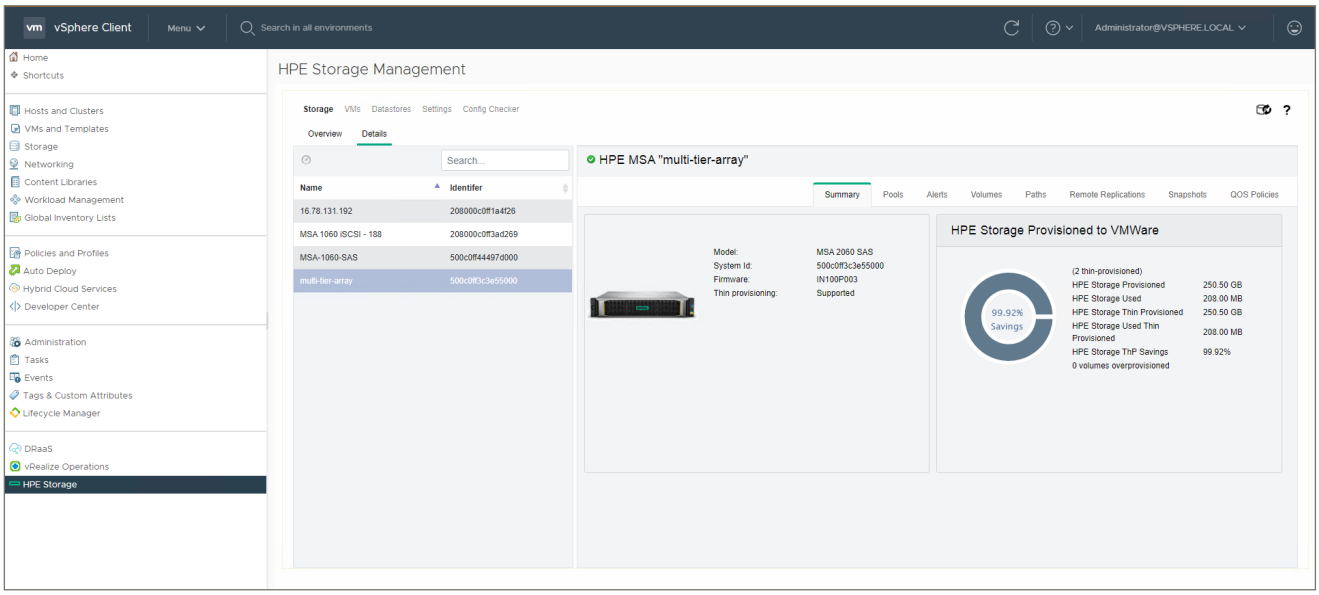

Figure 14. HPE Storage Integration Pack for VMware vCenter screen

Summary

The HPE MSA Gen6 system provides exceptional value for any business. The rich feature set makes the array attractive as an entry-level storage solution with the capacity to scale to meet business needs of a growing virtualized server environment. The ease of storage administration and integration with vCenter 7 makes the HPE MSA Gen6 a desirable storage array. HPE MSA Fibre Channel, iSCSI, and SAS connectivity options with multipath policies make the HPE MSA system highly resilient and ready to support test and development environments to full-production vSphere solutions. All of these features combined with the HPE MSA array-based snapshots and replication services create a versatile storage array for vSphere environments. By following these best practices and configuration recommendations, you can ensure that the HPE MSA array will become an integral part of any vSphere solution.

Appendix A

vSphere SAN troubleshooting

The following section lists some common troubleshooting techniques for use in a vSphere SAN environment.

Volume mapping

If the vSphere host does not see the volume, initiate a rescan of the storage adapters in the vSphere host. Then check zone configurations, fabric switches, and fiber cables for any damaged or nonfunctional components. With the correct drivers installed, provision the LUNs for the vSphere hosts in the cluster or data center. To verify storage connectivity, use the SMU option that shows discovery of the host port by the vSphere host.

Identifying HBA WWPNs on ESXi hosts

When creating hosts and host groups on the HPE MSA array, you need to identify the WWPNs associated with the initiators. You can find the WWPNs in vCenter by selecting the host configuration for storage adapters as seen in the following screenshot.

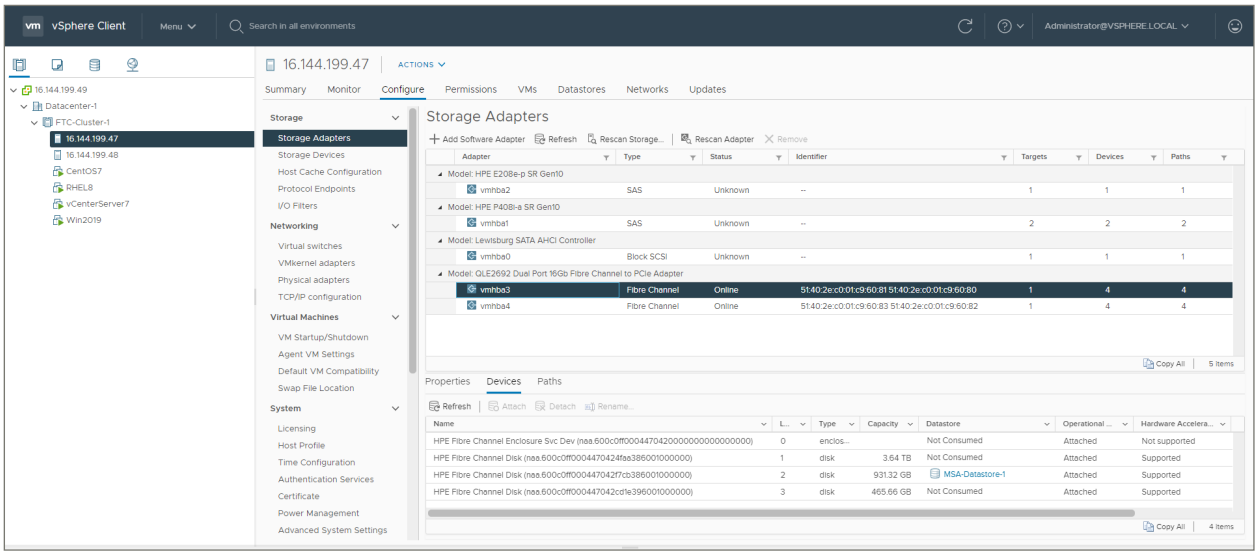

Figure 15. ESXi Host Fibre Channel Adapter Configuration screen

You can also find the WWPNs in the ESXi console by issuing the following commands:

```
# esxcli storage core adapter list
       HBA Name  Driver      Link State  UID
       --------  ----------  ----------  ------------------------------------
       Capabilities         Description
       -------------------  -----------
       vmhba3    qlnativefc  link-up     fc.51402ec001c96081:51402ec001c96080  Second Level Lun ID
       (0000:af:00.0) QLogic Corp QLE2692 Dual Port 16Gb Fibre Channel to PCIe Adapter
       vmhba4    qlnativefc  link-up     fc.51402ec001c96083:51402ec001c96082  Second Level Lun ID
       (0000:af:00.1) QLogic Corp QLE2692 Dual Port 16Gb Fibre Channel to PCIe Adapter
```

Or

```
# esxcli storage san fc list
       Adapter: vmhba3
       Port ID: 0000EF
       Node Name: 51:40:2e:c0:01:c9:60:81
       Port Name: 51:40:2e:c0:01:c9:60:80
       Speed: 16 Gbps
       Port Type: PTP
       Port State: ONLINE
       Model Description: HPE SN1100Q 16Gb 2p FC HBA
       Hardware Version: BK3210407-20  E
       OptionROM Version: 3.64
       Firmware Version: 9.03.00 (d0d5)
       Driver Name: qlnativefc
       DriverVersion: 4.1.9.0
```

```
Adapter: vmhba4
Port ID: 0000EF
Node Name: 51:40:2e:c0:01:c9:60:83
Port Name: 51:40:2e:c0:01:c9:60:82
Speed: 16 Gbps
Port Type: PTP
Port State: ONLINE
Model Description: HPE SN1100Q 16Gb 2p FC HBA
Hardware Version: BK3210407-20  E
OptionROM Version: 3.64
Firmware Version: 9.03.00 (d0d5)
Driver Name: qlnativefc
DriverVersion: 4.1.9.0
```

Boot from SAN
If booting from SAN, Agreed the following steps on the ESXi host to identify the boot virtual volume (LUN).

To obtain the host bootbank and altbootbank UUIDs, log in to the ESXi host with an SSH client and run the command:

```
# ls -l /
```

To obtain the disk ID, run the command:

```
# vmkfstools -P path_of_the_LUN
```

To check the storage device properties, run these commands:

```
# esxcli storage nmp device list -d naaID
# esxcli storage core device list -d naaID
```

Appendix B

VAAI integration

The vSphere Storage APIs are a set of technologies and interfaces that enable vSphere to leverage storage resources to deliver improved efficiency, control, and ease of use. The vSphere Storage APIs for Array Integration (VAAI) is one of these technologies. The VAAI initiative introduces APIs to improve performance, resource usage, and scalability by leveraging more efficient storage array-based operations.

Primitives are specific functions used with VAAI that serve as integration points to storage arrays. When supported by the array, primitives in VAAI allow the hypervisor to communicate directly with storage arrays to offload storage functionality traditionally handled by the hypervisor. The VAAI commands are automatically sent to the storage array, bypassing the ESXi host.

Storage arrays can handle these functions more intelligently and efficiently because they are purpose-built to perform storage tasks and can complete the request much faster than the host could complete it.

HPE MSA Gen6 systems natively support the following VAAI primitives:

- **Hardware-Assisted Locking:** Also known as “Atomic Test & Set (ATS).” This primitive protects metadata for VMFS cluster file systems at the block level rather than at the volume level, reducing SCSI reservation contention between vSphere hosts by allowing simultaneous access to different parts of the vSphere datastore.
- **Copy Offload:** Also known as “XCOPY.” This primitive copies VMDKs, enabling full copies of data to be made within the storage array, reducing data reads/writes required by both the vSphere host and network infrastructure.
- **Block Zeroing:** This primitive allows the array to handle the process of zeroing disk blocks. Instead of requiring the host to wait for the operation to complete, the array signals that the operation has completed immediately, handling the process on its own without involving the vSphere host.

- **Space Reclamation:** Also known as "SCSI UNMAP." This primitive originates from a VMFS datastore or a VM guest operating system. The command assists thin provisioned storage arrays such as HPE MSA Gen6 arrays to reclaim unused space from VMFS datastores and VMs that might have been deleted or migrated.

Note
ESXi issues the UNMAP command at 1 MB granularity, but the HPE MSA array operates with a 4 MB page size. As a result, pages do not free automatically, but can be released when the UNMAP command is manually invoked using the following ESXi CLI command:
`# esxcli storage vmfs unmap -l MyDatastore`

To confirm the support of the four VAAI primitives on an HPE MSA Gen6 array, issue the following command:

```
# esxcli storage core device vaai status get -d naaID
       naaID
              VAAI Plugin Name:
              ATS Status: supported
              Clone Status: supported
              Zero Status: supported
              Delete Status: supported
```

XCOPY is listed as "Clone" and SCSI UNMAP is listed as "Delete" when running the command.

Remember to verify the version of the HPE MSA firmware installed on the array. Consult the VMware Compatibility Guide for detailed compatibility information regarding path failover and VAAI plug-in support.

VAAI benefits and use cases

VAAI helps reduce the storage bandwidth consumed by a vSphere host and improves data center scalability. Storage operations such as virtual machine provisioning, VMware vSphere® Storage vMotion®, virtual disks creation, and so on consume less CPU, memory, and network bandwidth when using a VAAI-compliant HPE MSA storage system.

The following use cases address the four VAAI primitives:

- **Use case 1: VM migration**—Using VMware vCenter Server®, a VMware administrator wants to migrate VMs between datastores using Storage vMotion. The vSphere host can take advantage of the VAAI XCOPY (Copy Offload) command to migrate VMs much faster by offloading the data transfer to the array, greatly decreasing the amount of server and network resources consumed. Using the VAAI feature set results in reduced VM deployment time and faster migration of VMs between clustered hosts.

  VM migration also uses SCSI UNMAP to reclaim space on the source datastore. Because the page size of the HPE MSA Gen6 array is 4 MB, the manual space reclamation command needs to be used periodically to reclaim and free the space.
- **Use case 2: Rapid VM deployment**—An administrator needing to provide 20 VMs for a training class can use vCenter to deploy the VMs using **Deploy from template**. With a VAAI-enabled array, the Deploy operation uses both XCOPY and Block Zero primitives to accelerate VMs creation.
- **Use case 3: Increased VM density**—The Hardware-Assisted Locking functionality mitigates the potential for SCSI reservation contention between vSphere clustered hosts, reducing I/O performance impact to those hosts. Because the chance for SCSI reservation contention is greatly reduced, you can increase the number of VMs for per server (also known as "VM density").

Appendix C

Storage I/O Control

Storage I/O Control (SIOC) is a VMware feature that enables the attenuation of I/O for each virtual disk you choose. SIOC enables you to define shares and IOPS limits at the level of each VMDK to ensure that critical VMs are treated as high priority, better performing, and configured individually. The SIOC feature in vSphere 7.x is disabled by default.

SIOC provides I/O prioritization of virtual machines running on a cluster of vSphere hosts that access shared storage. It extends the familiar constructs of shares and limits, which have existed for CPU and memory, to address storage use through a dynamic allocation of I/O queue slots across a cluster of vSphere hosts. When a certain latency threshold is exceeded for a given block-based storage device, SIOC balances the available queue slots across a collection of vSphere hosts; this aligns the importance of certain workloads with the distribution of available throughput. This balancing can reduce the I/O queue slots given to virtual machines that have a low number of shares, to provide more I/O queue slots to a virtual machine with a higher number of shares. SIOC reduces I/O activity for certain virtual machines, so that other virtual machines receive better I/O throughput and an improved service level.

SIOC and the HPE MSA array could be combined to provide a more performance-optimized storage solution. Enabling SIOC is a simple process. More important is an understanding of the virtual machine environment with regard to the I/O demand being placed on the array.

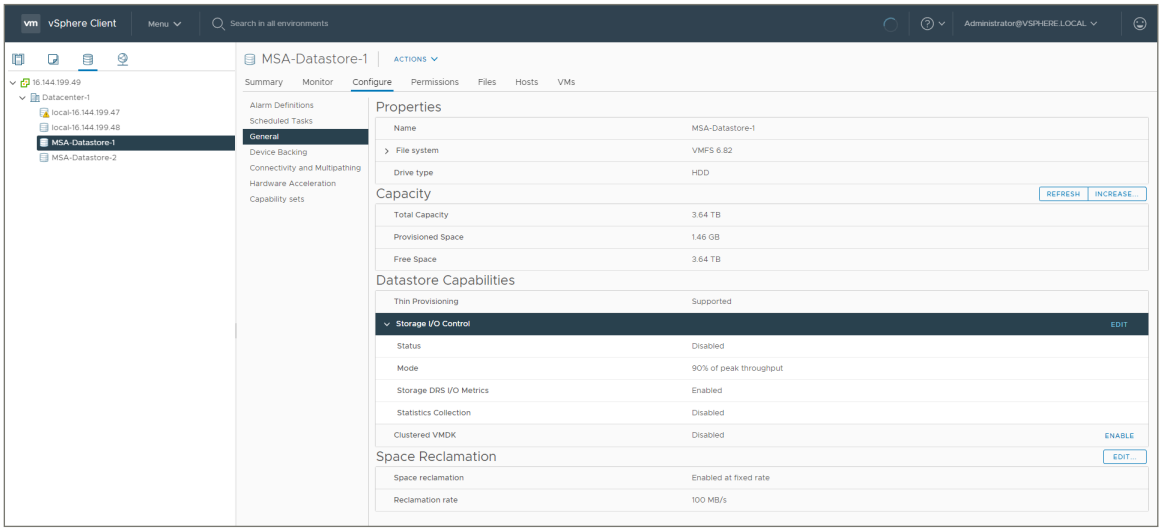

Figure 16. Storage I/O control screen

Resources

HPE MSA 1060 Storage QuickSpecs
hpe.com/support/MSA1060QuickSpecs

HPE MSA 2060 Storage QuickSpecs
hpe.com/support/MSA2060QuickSpecs

HPE MSA 2062 Storage QuickSpecs
hpe.com/support/MSA2062QuickSpecs

HPE MSA Gen6 Virtual Storage Technical Reference Guide
hpe.com/psnow/doc/a00103247enw?from=app&section=search&isFutureVersion=true

HPE MSA 1060/2060/2062 best practices
hpe.com/psnow/doc/a00105260enw?from=app&section=search&isFutureVersion=true

Learn more at

hpe.com/us/en/storage/msa-shared-storage.html

a00109648ENW, Rev. 2

Hewlett Packard Enterprise