



Hewlett Packard
Enterprise

HPE MSA Storage Configuration and Best Practices for VMware vSphere

Contents

Introduction.....	4
Intended audience.....	4
General product overview.....	5
HPE MSA 2052 Storage.....	5
Solution overview.....	6
HPE Storage Networking.....	6
HPE MSA Storage features and concepts.....	7
Storage Pools, Disk Groups, and Virtual Volumes.....	8
Supported RAID types.....	10
Virtual Volumes.....	12
Unified LUN Presentation (ULP).....	13
Data protection with HPE MSA Remote Snap.....	16
HPE MSA recommendations and common practices.....	16
HPE MSA 2040/2042/2050/2052 Storage installation.....	16
Expansion Cabling the HPE MSA Storage Array.....	16
Initial configuration.....	17
Disk Group RAID Type Considerations.....	18
Disk Sparing.....	18
Virtual Volume Affinity.....	18
Storage administration.....	18
HPE MSA Considerations and Best Practices for vSphere.....	23
Setup and installation.....	23
Storage cabling and network connectivity.....	23
Storage configuration.....	23
Tiered Storage.....	25
Boot from SAN.....	25
Balancing controller ownership.....	25
Changing LUN Response.....	25
Volume Mapping.....	26
Presenting Storage to Virtual Machines.....	28
Multipath Considerations for vSphere.....	29
HPE MSA Considerations and Best Practices for vCenter.....	31
Datastore Clusters.....	31
Distributed Resource Scheduler.....	32
Multipath and Datastore Clusters.....	34
Site Recovery Manager and the HPE MSA Storage.....	34
Configuring SRM and the Storage Replication Adapter.....	36

HPE OneView for VMware vCenter	37
Use cases for vSphere and HPE MSA Storage	38
Mixed Physical and Virtual Computing Environment	38
Multiple vSphere clusters on HPE MSA Storage	39
Distributed vSphere Environments	40
Data Protection with HPE MSA Remote Snap	41
Multi-Site Disaster Recovery	42
Summary and benefits	43
Appendix A	44
HPE MSA Firmware Update	44
HPE MSA Array Controller Firmware Update	44
Disk Drive Firmware Update	44
Appendix B	45
Array/Volume Replication and Remote Snap	45
Appendix C	45
Disk Background Scrubbing	45
Appendix D	46
vSphere SAN troubleshooting	46
Boot from SAN	48
Appendix E	48
VAAI integration	48
VAAI benefits and use cases	49
VMware Storage I/O Control	49
Multipath	50
Appendix F	51
Changing the LUN RESPONSE with SMU v2	51
Appendix G	51
Terminology	51
For more information	53
HPE resources	53
VMware resources	53
VMware Storage Solutions from HPE	53

Introduction

Storage is a critical component of the virtualization infrastructure. The storage platform must provide efficient capacity, deliver high performance, and scale easily. Storage solutions must be modular and scalable to meet the increasing demands of constant uptime and on demand purchasing models. Distributed and always available services demand no single point of failure in the infrastructure and replaceable components without interruption of service. The HPE MSA storage solution is an entry level product targeted to support the ever increasing storage demands of VMware vSphere®.

The latest generation of HPE MSA Storage is designed and built to exceed the economic and operational requirements of virtual data centers by providing the storage performance, scalability, availability, and simplified management needed by small and midsize businesses with growing storage needs. In this paper, we explore the unique capabilities and integration of HPE MSA Storage with VMware vSphere. In addition, the paper covers the best practice settings and configurations needed to optimize HPE MSA Storage for deployment with VMware vSphere. When deployed together, HPE MSA Storage with VMware vSphere provides small and midsize businesses the benefits of increased consolidation savings by increasing virtual machine density, lowering storage costs, and realizing time savings from simplified storage management and provisioning.

The outline of this document includes the features and recommended configurations for the HPE MSA 2040/2042/2050/2052 product line. The next section is [MSA Storage features and concepts](#) which outlines HPE's best practices for configuring the HPE MSA Storage Array. The next section is [MSA Considerations and Best Practices for vSphere](#). Although suggestions and considerations applicable to VMware vSphere are sprinkled throughout this document, this section contains the material most relevant to engineers and administrators managing the HPE MSA 2040/2042/2050/2052 in a virtualized vSphere environment. The next section [MSA Considerations and Best Practices for VMware® vCenter™](#) outlines specific items as they relate to the vCenter management environment. There are several Appendices included at the end of this document to provide information about less common activities and best practices around controller and drive firmware upgrades.

If the material in this document is helpful, please let us know by liking it on our website or provide feedback to hpe.com/contact/feedback.



This paper outlines HPE's recommendations for configurations, software settings, and design architectures to get best results for your HPE MSA storage with vSphere. These recommendations are highlighted throughout this document with the icon.

Intended audience

IT administrators, VMware vSphere administrators, and solution architects planning a server virtualization deployment with HPE MSA storage. This and other documents pertaining to virtualization with HPE and VMware® are available at hpe.com/storage/vmware and at hpe.com/storage.

VMware vSphere/VMware® ESXi™ administrators planning to setup hosts with the HPE MSA storage should have a working knowledge of storage area networks (SANs) concepts. Clustered, fault-tolerant, virtualization environments such as VMware vSphere rely heavily upon centrally-managed, scalable, SAN resources. The HPE MSA Storage system provides a versatile entry-level SAN solution for vSphere host clusters. The following sections provides an overview on the HPE MSA Storage system and the recommended configurations and best practices for its use with vSphere.

General product overview

Clustered, fault-tolerant, virtualization environments such as VMware vSphere rely heavily upon centrally-managed, scalable, SAN resources. The HPE MSA Storage system provides a versatile entry-level SAN solution for vSphere host clusters. The following section provides an overview on the HPE MSA Storage system.

HPE MSA 2052 Storage



Figure 1. HPE MSA 2052 Storage Array front view (SFF and LFF)

All HPE MSA Storage models offer a common set of features listed in Table 1

Table 1. Common features of HPE MSA models

Features	MSA family
Array	
Host protocols	8 Gb FC, 16 Gb FC, 1GbE iSCSI, 10GbE iSCSI, SAS
Maximum host ports	8 Host connect ports <ul style="list-style-type: none"> • 8 Fibre Channel (8 Gb or 16 Gb) • 8 iSCSI (1GbE or 10GbE) • Hybrid (4 FC and 4 iSCSI) • 8 SAS
Cache, per array	8 TB maximum read cache per array 16 GB data (read/write) cache + system memory per array
Maximum LUNs	512
RAID supported: Virtual mode	RAID 1, 5, 6, 10
Snapshots	64 snapshots standard/512 snapshots optional
Volume copy	Bundled
vSphere support	VAAI, VASA, SRM
Enclosures	
Rack/physical	2U (for controllers and expansions)
Expansion on drive enclosures	0–7 enclosures
LFF/SFF array/enclosures mixing	Supported
Maximum number of drives per array enclosure	24 SFF/12 LFF
Maximum number of drives per drive enclosure	24 SFF/12 LFF
Drive enclosure interface type	6 Gb SAS

For more in-depth information on MSA storage system models, visit the [HPE MSA Storage webpage](#)

Solution overview

Virtualization enables IT organizations to manage escalating demands on infrastructure and diminishing budgets. Server virtualization allows a business to run several server computing environments and operating systems on a single hardware platform. This approach results in a smaller physical footprint with reduced power and cooling costs.

Figure 2 illustrates a recommended configuration for a fault-tolerant HPE MSA 2042 Storage system supporting a VMware vSphere environment. In this configuration, two HPE 16 Gb FC switches are used in the storage fabric for redundancy. Two 16 Gb HBA FC adapters in each HPE DL server—deployed as a VMware vSphere host—provide redundant paths to the HPE MSA Storage. This configuration, with multiple paths to storage, improves I/O performance and eliminates a single point of failure between the vSphere hosts and the HPE MSA Storage. Direct Connection of host server to the MSA FC host ports is also supported.




Figure 2. Reference SAN Fabric Architecture

HPE Storage Networking

While this paper focuses on the best practices for deploying HPE MSA Storage for vSphere, it is important to ensure that the proper storage networking infrastructure exists to complement the server and storage requirements. HPE offers a full set of network solutions to complete your infrastructure.

A typical complement to an HPE MSA Storage and VMware vSphere deployment is the HPE SN6000B Fibre Channel Switch (Figure 3), which offers the features necessary to meet the demands of hyper-scale, private cloud storage environments by delivering 16 Gb/sec Fibre Channel technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the HPE SN6000B Switch is configurable in 24, 36, or 48 ports and supports 4, 8, 10, or 16 Gb/sec Fibre Channel speeds in an efficiently designed 1U package. It also provides a simplified deployment process and a point-and-click user interface, making it both powerful and easy to use.



Figure 3. HPE SN6000B Fibre Channel Switch



HPE recommended—Creating Domains and Zones when using the HPE SN6000B Fibre Channel Switch. vSphere prefers a single-initiator single-target zoning to restrict access and prevent misconfigurations.

HPE Switch highlights

- Ports-on-Demand capabilities for fast, easy, and cost-effective scaling from 24 to 48 ports in a 12-port increment.
- 16 Gb/sec performance with up to 24 ports in an energy-efficient, 1U form factor, providing maximum flexibility for diverse deployment and cooling strategies.
- Best-in-class port density and scalability for entry and mid-range SAN switches, along with hot-pluggable components and non-disruptive software upgrades.
- Exceptional price and performance value, exceeding comparable Ethernet storage-based alternatives.

HPE 16 Gb Host Bus Adapters (HBAs)

The HPE SN1200E 16 Gb single-port and dual-port Fibre Channel Host Bust Adapters are a step into the 16 Gbit/sec environment with greater performance and advanced, embedded support for virtualized environments. The HPE SN1200E 16 Gb Fibre Channel HBA purchased today is backward compatible with 8 Gb and 4 Gb storage networks and will protect future investments. When using storage intensive applications, such as backup, database transactions, virtualization, and rich media, the increased performance of the HPE SN1200E 16 Gb FC HBA enables more rapid storage and retrieval of critical information. Designed for environments with greater virtual machine density and bandwidth requirements, the HPE SN1200E 16 Gb Host Bus Adapters enable more applications and virtual machines to run on a single server and port, resulting in reduced cabling and higher return on IT investment.

Product highlights

- Enhances data center productivity by delivering twice the data throughput of 8 Gb FC HBAs
- Meets the needs of larger server virtualization deployments, cloud applications, and advanced server architectures
- Enables more applications and VMs to run on a single server and Fibre Channel port, resulting in reduced cabling and a higher return on IT investment
- Includes PCIe 3.0—Gen3 x8 lanes

For more information about HPE Storage Networking solutions and products, visit hpe.com/storage/storefabric

HPE MSA Storage features and concepts

Growing storage needs for virtualized servers now require greater levels of storage performance and functionality at a lower cost of ownership. The new HPE MSA 2050/2052 Storage arrays are positioned to provide an excellent value for SMB customers needing increased performance to support initiatives such as consolidation and virtualization.

The HPE MSA 2050/2052 delivers higher performance by leveraging the HPE MSA controller architecture, four 8/16 Gb Fibre Channel ports, and 8 GB of data cache per controller. The HPE MSA Storage also includes the ability to create Read Cache and Performance Tier using SSDs to further enhance performance of the HPE MSA Storage rich features.

The HPE MSA Storage is built for virtualization environments delivering the following key advantages for VMware vSphere:

- High performance I/O throughput to meet the peak virtualization demands of cloning or migrating multiple VMs.
- An affordable array storage product with cost-effective scale out options.
- Virtualized disk technology that provides non-disruptive scalability to VMs.
- Integration with key VMware vSphere Storage APIs, including VAAI, VASA, and SRM.
- Easy-to-use web-based storage administration and management software.
- Integration with HPE OneView for VMware vCenter.
- Support for very large VMFS datastores.
- Small footprint storage redundancy.

HPE MSA Storage delivers these advantages through a number of unique architectural features. Some of the virtualization benefits are detailed in the following sections.

Storage Pools, Disk Groups, and Virtual Volumes

The HPE MSA Storage Management Utility (SMU) has a new look and feel which support a new visual way of virtualizing volumes. This volume virtualization centers on the configuration of storage pools of disk groups where virtual volumes are created and managed. A general understanding of how virtual volumes works is essential to understand the best practices for the vSphere environment. For an in-depth discussion of how storage pools, disk groups, and virtual volumes support each other, see the [HPE MSA Best Practices Guide](#).

Storage Pools

Storage Pools in the HPE MSA Storage Array provide a way of virtualizing the physical disk array to support oversubscription, uninterrupted volume scaling, and storage tiering. Often referred to as a virtual pool, each storage pool is associated with one of the HPE MSA controllers. Each pool can support multiple disk groups to provide different tiers of storage—performant, standard, and archival types of storage. The HPE MSA storage also features a unique ability of adding a read cache disk group to a storage pool to provide caching of the most actively accessed data for the storage pool.

POOLS

Clear Filters

Show

10

 Showing 1-4 of 4 (1 selected)

Name	Health	Total Size	Class	Avail	Volumes	Disk Groups
A	<div><div></div></div> OK	8786.3GB	Virtual	8786.3GB	40	3
B	<div><div></div></div> OK	8986.2GB	Virtual	8986.2GB	25	3
LinRS_A01	<div><div></div></div> OK	3597.0GB	Linear	3097.1GB	10	1
LinR6_B01	<div><div></div></div> OK	3597.0GB	Linear	2997.0GB	10	1

Related Disk Groups

Clear Filters

Show

10

 Showing 1-3 of 3

Name	Health	Pool	RAID	Class	Disk Type	Size	Free	Current Job	Status	Disks
dgB01	<div><div></div></div> OK	B	RAID5	Virtual	SAS (Standard)	4793.9GB	4793.9GB		FTOL	9
dgB02	<div><div></div></div> OK	B	RAID1	Virtual	sSAS (Performance)	197.6GB	197.6GB		FTOL	2
dgB03	<div><div></div></div> OK	B	RAID6	Virtual	SAS MDL (Archive)	3994.6GB	3994.5GB		FTOL	6

Related Disks

Clear Filters

Show

10

 Showing 1-0 of 0

Location	Health	Description	Size	Usage	Disk Group	Status
----------	--------	-------------	------	-------	------------	--------

No data available in the table

Figure 4. HPE MSA SMU Storage Pool Screen

When configuring a storage pool on the HPE MSA Storage Array for vSphere, an administrator should consider the follow factors:

- The storage nature of the application being virtualized.
- The storage optimization objectives.
- Balancing of HPE MSA resources.
- Tiering of Storage resources.

Here are two examples to consider.

Example #1

A software application has an intensive read storage nature; however, the data is very static and requires very few updates to the data. Let's also say there is a large user base of the application and they expect sub-second screen refreshes. Knowing this information the user might consider configuring a Storage Pool with at least one Storage Group with multiple disks front ended by a read cache.

Example #2

A software application has a huge storage requirement that will expand on a daily basis and consists of historical records for the past 100 years. Let's also say this application has a large user base that expects the application to respond with sub-seconds refreshes. An appropriate storage pool for this application would include a performance disk group with a standard disk group as well as an archived tier disk group.

Disk Groups

Disk Groups in the HPE MSA Storage Array supports three types of disk groups—virtual, read cache, and traditional linear disk groups. Virtual disk-groups and Linear disk-groups are redundant array of independent disks (RAID) set of disks. Virtual disk-groups also have the advantage of being aggregated within a pool. The pool can then distribute the performance needs over multiple disk groups. Linear disk-groups will utilize the physical disks which create the RAID set only for performance, this is a good way to get known performance for a specific workload.

All disks in a disk group must be the same disk type (SAS, MDL SAS, or SSD). Each controller can have a maximum of 6 Virtual Disk groups + 32 Linear disk groups.

Disk Groups in the SMU default to RAID-6 but can be configured to support RAID 1, 3, 5, 6, 10, and 50, virtual disk-groups are limited to RAID 1, 5, 6, and 10. RAID levels NoRAID, 0, and 3 are not available in the Storage Management Utility (SMU). When creating a volume users can also specify what type of tier affinity the group should support. For more specifics on virtual disk group tier affinity types, see **“Drive Type and Capacity Considerations when using Tiering”** in the [HPE MSA Best Practices Guide](#).

Read Cache Disk Groups are created using SSD drives in the system. Read Cache is a non-fault tolerant set of 1 or 2 drives which contain a READ copy of recent accessed data on the spinning media tiers. Only one read cache disk group can be created for a pool. When designated as read cache, the SSDs in this disk group become an extension to the array controller's cache.

The screenshot displays the HPE MSA SMU Storage Pools screen. The interface includes a navigation sidebar on the left with icons for Home, System, Hosts, Pools, Volumes, Mapping, Replications, and Performance. The main content area is titled "POOLS" and shows a table of storage pools. Below the pools table, there are sections for "Related Disk Groups" and "Related Disks".

Name	Health	Size	Avail	Volumes	Disk Groups
A	OK	9.5TB	7573.4GB	2	3
B	OK	9.5TB	9.0TB	1	3

Name	Health	Pool	RAID	Disk Type	Size	Free	Current Job	Status	Disks
dgA01	OK	A	RAID10	SAS (Standard)	4793.9GB	3786.7GB		FTOL	16
dgA02	OK	A	RAID10	SAS (Standard)	4793.9GB	3786.7GB		FTOL	16
rcA1	OK	A	RAID0	sSAS (Read Cache)	797.1GB	797.1GB	VRSC (16%)	UP	2

Location	Health	Description	Size	Usage	Disk Group	Status
1:1	OK	SSD SAS	400.0GB	VIRTUAL POOL	rcA1	Up
1:24	OK	SSD SAS	400.0GB	VIRTUAL POOL	rcA1	Up

Figure 5. HPE MSA SMU Storage Pools Screen with Disk Groups



Best practice: Supporting large storage capacities requires advanced planning weighing the tradeoffs between performance and capacity. As a general guide:

- **To maximize capacity**—combine physical disks into a large storage pool of virtual disk groups and subdivide the pool into several virtual volumes.
- **To maximize performance**—create a storage pool with virtual disk groups in multiples of 2 data disks. (See “Power of 2 Method” in the [HPE MSA 2042 Best Practices Guide](#))
- **To maximize read performance**—create a storage pool with virtual disk groups and add a read cache disk group.
- **To maximize I/O performance**—create a storage pool with standard virtual disk groups and add a performance virtual disk group.

Note

Subdividing a disk group into virtual volumes with a capacity less than 2 TB was a requirement for VMware vSphere® 5.1; however, since ESXi 5.5 datastores can be as large as 62 TB.

RAID Considerations

This section contains considerations for RAID levels on the HPE MSA storage.

Table 2. Overview of supported RAID implementations for HPE MSA storage

RAID level	Cost	Performance	Protection level
RAID 0 Striping	N/A	Highest	No data protection available for linear disk groups through the CLI
RAID 1 Mirroring	High cost 2x drives	High	Protects against individual drive failure
RAID 3	1 drive	Good	Protects against individual drive failure
RAID 5 Block striping with striped parity drive	1 drive	Good	Protects against any individual drive failure; medium level of fault tolerance
RAID 6 Block striping with multiple striped parity	High cost 2x drives	Good	Protects against any two drive failures; medium level of fault tolerance
RAID 10 Mirrored striped array	High cost 2x drives	Good	Protects against any individual drive failure; medium level of fault tolerance

Supported RAID types

While all the RAID sets in this section have certain strengths and weaknesses, all are available for use in the vSphere environment.

RAID 0 (block level striping)

RAID 0 is also known as a non-stripe set or a non-striped volume. RAID 0 keeps blocks of data in sequence, one disk at a time, for all disks in a configuration. Because there is no overhead on a RAID 0 set, it is the fastest way to read and write data. Read and write speeds on the same disk are approximately equal. The main disadvantage of RAID 0 is that there is no parity, and consequently if a drive fails, all data is lost. The SMU does not support the creation of disk groups with RAID 0. Since RAID 0 is not recommended in production environments, these types of disk groups can only be created using the Command Line Interface.

RAID 1 (mirroring)

RAID 1 is a 1 to 1 mirroring of a physical disk. It protects against individual drive failure. Every write executed is written to both drives.

RAID 5+0 (RAID 50) is not supported with Virtual Disk-Groups

RAID 5 (distributed parity + striping) writes data across a set of hard disks, calculates the data parity, and writes that parity to one hard disk set. RAID 5 then writes the parity to a different disk in the set for every further stripe of data. Combining RAID 0 striping produces performance increases. However, RAID 5 has the disadvantage of increasing overall costs and lowering available capacity. In order to write to a RAID 5 environment, the affected blocks are first read, the changed data is entered, the new parity is calculated, and the block is then written. On systems with large RAID 5 sets, this means a write I/O is many times slower than a read I/O, which is undesirable in a server virtualization environment.

RAID 6 (multiple distributed parity, with striping)

RAID 6 is the default RAID level and the HPE recommend level of protection for HPE MSA Virtual Disk Groups. RAID 6 is identical to RAID 5 except for the addition of a second parity block. It does not have a performance penalty for read operations; however, it does have a performance penalty for write operations, due to the overhead incurred by parity calculations.

RAID 10 (mirroring + striping)

RAID 10 (also known as RAID 1+0) writes data in stripes across primary disks that have secondary disk mirrors. Server virtualization works well with this RAID level. Performance increases due to block-level striping, and the replication of volumes mirrored onto separate physical disks manage additional I/O requests.



Best practice: With large storage configurations, consider creating fewer disk groups containing many drives rather than creating many disk groups with few drives.

This practice not only utilizes storage more effectively but it also improves performance during writes to the disk group. For example, a single RAID 6 disk group of 12 SAS drives—2 will be for parity and 10 for Data. Compare the same number of drives in a four RAID 5 disk group. Each disk group will have a parity drive and two data drives. A difference of a 4 to 2 parity drive usage. Having four disk groups also impacts performance. (See Power of 2 Model).

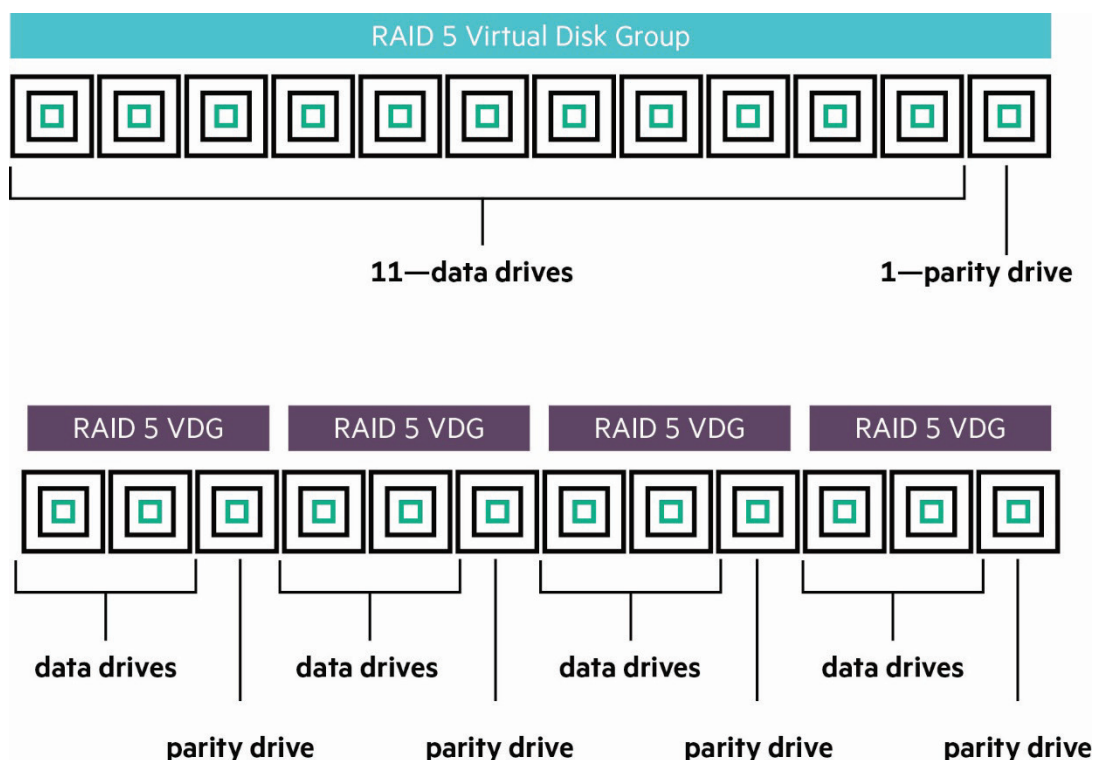


Figure 6. Virtual Disk Group RAID comparisons

Virtual Volumes

Virtual Volumes are the logical subdivision of the storage pool that are exposed through the controller ports to the vSphere host. They hide the complexities of the Storage Pool and Disk Group features (Striping, Read Cache, and Data Tiering) and map to the vSphere environment as VMFS datastores or Raw Disk Mappings (RDMs). The HPE MSA 2042 can have a maximum of 1024 Virtual Volumes per Storage Pool. Each Virtual Volume can be set with a Tier Affinity which maps to the Tier types created in the Virtual Disk Groups. By default Virtual Volumes are set to “No Affinity.” Hewlett Packard Enterprise recommends the default “No Affinity” option for most configurations. This setting attempts to balance the frequency of data access, disk cost, and disk availability by moving this volume’s data to the appropriate tier. To know more about Virtual Volume Affinity see the “**Mechanics of Volume Tier Affinity**” section of the [HPE MSA 2042 Best Practices Guide](#).



Best Practice: When possible use the “Power of 2 Model.”

When creating virtual disk groups, keep the number of data drives in the RAID Group to a multiple of 2. The HPE MSA writes data in 4 MB chunks (pages). Keeping the 4 MB chunk write balanced across a “Power of 2” number of data drives improves performance. For example, the optimal performance configuration for a virtual RAID 6 disk group would be 10 drives. Two drives for parity and 8 drives for data. Since the striping is across 8 drives at the same time, a single stripe is done for the 4 MB chunk (i.e., 8 x 512 KB = 4 MB).

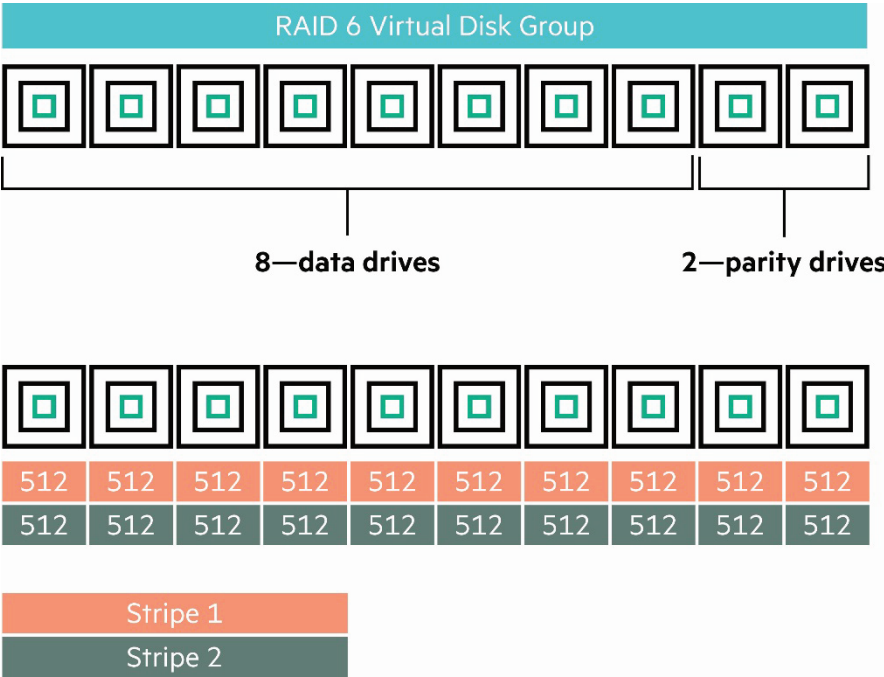


Figure 7. Virtual Disk Group RAID 6 Striping Model

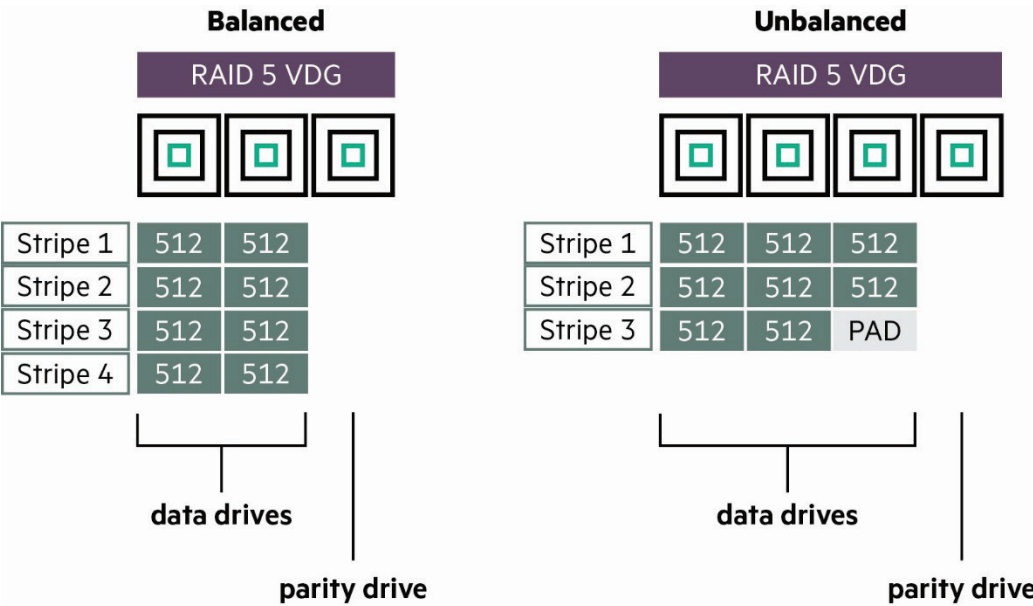


Figure 8. HPE MSA Balanced/Unbalanced Virtual Disk Group example




Figure 9. RAID 5 Unbalance Virtual Disk Group example

Unified LUN Presentation (ULP)

The HPE MSA Storage Array supports Unified LUN Presentation (ULP). In version 3 of the SMU, ULP is virtually hidden to the administrator through the use of Storage Pools and Virtual Volumes. Virtual Volumes are present through both controllers on the HPE MSA, eliminating the need to specify the interconnect paths between the specific storage controller and target hosts. Under the covers, ULP presents the same World-Wide Node Name (WWNN) for Fibre Channel or in the case of iSCSI the same IQN, for virtual Volumes on both controllers. Because pool ownership is associated to controllers in the HPE MSA, the SMU enforces ULP by preventing duplicate Virtual Volumes from being created within pools.

Because both controllers can access data for any of the specified virtual volumes in both Storage Pools, a preferred path will be identified following the Asymmetric Logical Unit Access (ALUA) specifications. This is part of the Reported Target Port Group (RTPG) protocol which identifies the virtual volume's pool and then identifies the optimal path to the Virtual Volume. The owning associated pool controller always performs the I/O to disk for these requests.

Virtual Volumes appear to the host as an active-active storage system where the host can choose any available path to access the Virtual Volume regardless of Storage Pool ownership.

Virtual Volumes uses the T10 Technical Committee of INCITS ALUA extensions, in SPC-3, to negotiate multiple paths to the ULP aware host systems. Unaware host systems see all paths as being equal.

The ability to identify and alter the LUN controller ownership is defined by the ALUA extensions in the SPC-3 standard. The HPE MSA Storage Array supports implicit ALUA modes. This means the array can assign and change the managing controller for a LUN; however, LUN ownership cannot be assigned to one particular HPE MSA controller.

Overview of ULP operation

ULP presents all Virtual Volumes to all host ports

- Removes the need for specifically identified controller interconnect paths.
- Presents the same WWNN for both array controllers.

ULP shares Virtual Volumes between controllers

- No duplicate of Virtual Volumes are allowed between array controllers.
- Either controller can use an unused Virtual Volume.

ULP recognizes "preferred" paths

- Preferred path indicates the owning controller of the Storage Pool of the Virtual Volume follows the ALUA specifications.
- "Report Target Port Groups" identify preferred paths.
- Performance is degraded by using non-preferred paths.

Processing Write I/O with ULP

Let's consider for a moment how ULP functions using the following Host to Virtual Volume example.

1. Write command to controller A for Virtual Volume A.
2. Controller A knows VVA is in pool B owned by Controller B.
3. Data is written to Controller A cache.
4. Controller A then broadcasts it to the VVA mirror on Controller B.
5. Controller A acknowledges I/O completion to the host.
6. Data is written to VVA by Controller B from Controller A's mirror.

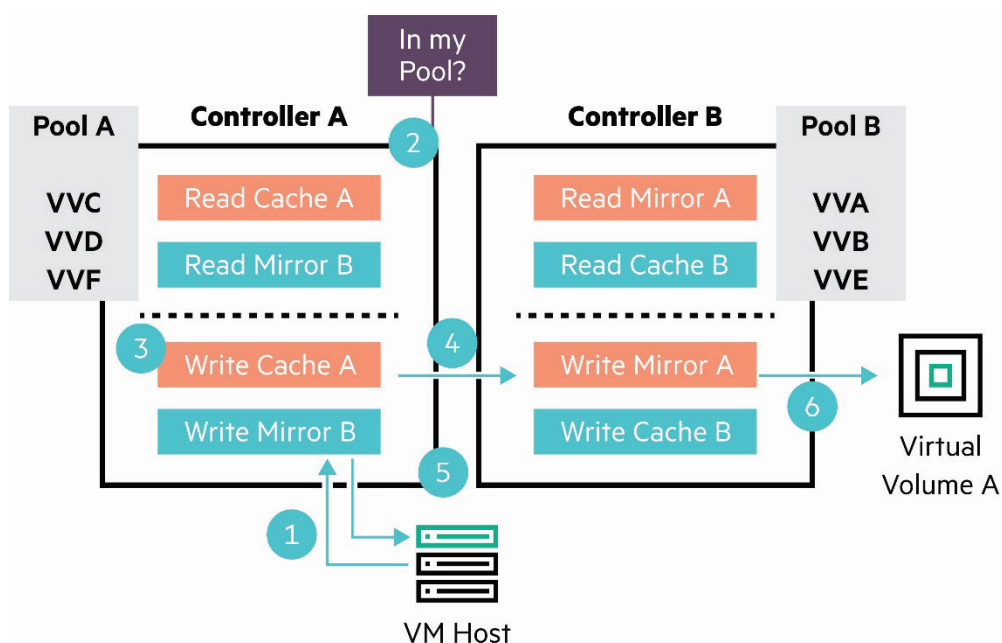


Figure 10. HPE MSA Unified LUN Presentation Model Process flow

Read I/O Processing with ULP

Read command to controller A for VVA owned by Controller B:

1. Controller A asks Controller B if VVA data is in Controller B Read cache.
2. If found, Controller B tells Controller A where in the Read cache it is and mirrors to Controller A.
3. Controller A sends data to the host from Controller B read mirror with I/O complete.
4. If not found in read cache, the request is sent to Controller B to retrieve data from disk.
5. Data from disk is placed in Controller B read cache and broadcasted to Controller A mirror.
6. The data is then sent to host by Controller A from Controller B read cache mirror with I/O complete.

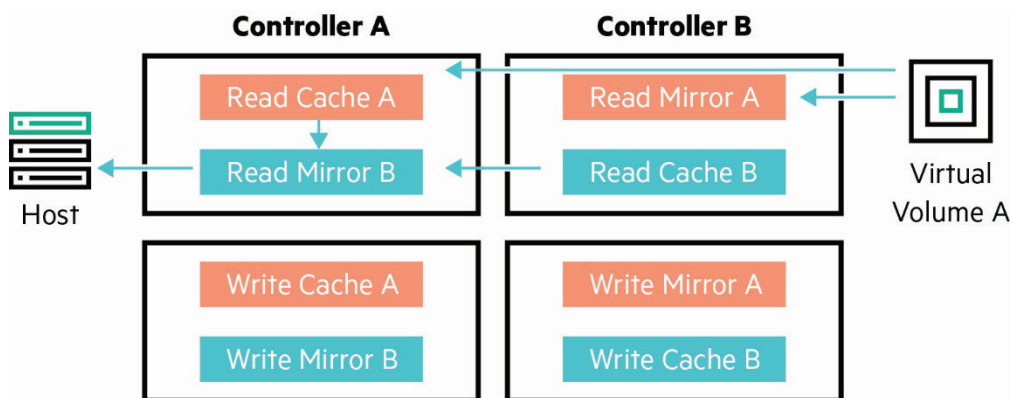


Figure 11. HPE MSA Unified LUN Presentation Model

Unified LUN Presentation (ULP) failover

If a controller unit fails on a dual controller HPE MSA Storage Array, Storage Pool ownership transfers from the failing controller to the secondary, or backup, controller in the array.

Upon failure the Storage Pool and all associated disk groups owned by the controller transfers to the secondary controller. Because of ULP only one World Wide Name Nodes (WWNN) presents itself to the surviving controller. Multipath software continues to provide I/O and the surviving controller reports all its paths as preferred.

The ability to identify and alter the Virtual Volume to controller ownership is defined by the ALUA extensions in the SPC-3 standard. The HPE MSA Storage Array supports implicit ALUA modes though virtualized through Storage Pools, Virtual Disk Groups, and Virtual Volumes.

Data protection with HPE MSA Remote Snap

Cost-effective data protection can be a challenge for today's small to midsize businesses. The HPE MSA Storage Array and its Remote Snap feature provides a real disaster-tolerant solution for a business needing data protection and redundancy.

Remote Snap is an array-base replication feature implemented between two HPE MSA storage systems. With a dual storage-array configuration, data is protected through a point-in-time data replication without breaking your budget.

For more specific information see the [HPE MSA Remote Snap Software](#) document.

HPE MSA recommendations and common practices**HPE MSA 2040/2042/2050/2052 Storage installation**

Because storage plays a critical role in the success of a VMware vSphere deployment, it is important to properly install and configure the HPE MSA Storage Array. The HPE MSA development team has set the default settings of the HPE MSA Storage Array for optimal performance for most storage deployments. This is generally true for virtualized environments as well. The following section highlights some key Installation and configuration concepts important for those managing and administering the HPE MSA Storage Array.

For specific recommendations for the HPE MSA Storage Array and the VMware vSphere virtualization environment see the section of this document entitled [HPE MSA Considerations and Best Practices for vSphere](#).

Other resources, such as HPE MSA Configuration Guides, provide detailed installation information. Refer to the following resources, which are available on the [HPE MSA webpage](#):

- The [HPE MSA 2040 User Guide](#)
- The [HPE MSA 2040 Cable Configuration Guide](#)
- The [HPE MSA 1040/2040 SMU Reference Guide](#)
- The [HPE MSA 2050 User Guide](#)
- [The HPE MSA 2050 SMU Reference Guide](#)

Expansion Cabling the HPE MSA Storage Array

HPE MSA Storage Array supports both fault-tolerant and straight-through SAS cabling to expansion enclosures. Fault-tolerant cabling allows any drive enclosure to fail or be removed, while maintaining access to other enclosures. When connecting multiple drive enclosures, use fault-tolerant cabling to achieve the highest level of fault tolerance.

Although the HPE MSA controllers support simultaneous connections to Fibre Channel and iSCSI connections, you should never configure both FC and iSCSI ports for the same volume to the same server. The host multipath solutions tend to get confused by trying to load balance across non-similar performance protocols. Although iSCSI supports out of sequence packets, Fibre Channel protocols cannot.

For guidance on cabling requirements, consult the [HPE MSA 2040 Cable Configuration Guide](#) or [HPE MSA 2050 Cable Configuration Guide](#).




Figure 12. HPE MSA Storage backplane with interfaces/ports

1. Power supplies
2. Host connection ports
3. CLI port (mini-USB)
4. Management Ethernet port
5. SAS Expansion port

A configuration with dual FC switches for host connectivity to the HPE MSA array is a best practice, offering excellent redundancy. For the HPE MSA Storage cabling requirements and options, consult the [HPE MSA 2040 User Guide](#) or the [HPE MSA 2050 User Guide](#).

FC Switches and SFP transceivers

Not all SFPs are equal! Although technology advances have tried to standardize the Small Form-Factor Pluggable (SFP) transceiver, not all SFP transceivers are universally adaptable to any devices. There are SFPs for Fibre Channel and for Ethernet switches and adapters and they come in a variety of speeds.

The use of the wrong SFPs transceivers can cause the storage on the HPE MSA Storage to not be visible. Consult the HPE MSA 2042 or HPE MSA 2052 QuickSpecs for more information on supported SFP transceivers for FC and iSCSI switches go to hpe.com/storage/san.

For more guidance regarding HPE MSA Storage cabling, consult the “HPE MSA 2040 Cable Configuration Guide” or “HPE MSA 2050 Cable Configuration Guide” on the [Hewlett Packard Enterprise Information Library](#).

Initial configuration

When setting up your HPE MSA Storage for the first time, consider the following general practices to ensure the best results for your installation. Specific best practices for the VMware vSphere environment can be found in the [HPE MSA Considerations and Best Practices for vSphere](#) section of this document.

Optimize for fastest throughput

When configuring your storage system for fastest throughput remember to:

- Configure the host ports for the fastest transmission speed available.
- Balance Disk Groups between two HPE MSA pools.
- Add a Read Cache Disk Group to each Storage Pool.
- To maximize sequential performance for a Storage Pool, create only one volume per pool. (Randomness in performance is created when multiple workloads to the same pool are being exercised).
- Map Virtual Volumes to dedicated host ports to ensure optimal performance.
- Distribute the load across as many drives as possible using the Power of 2 Model.

Highest fault tolerance optimization for HPE MSA

The following guidelines list the general best practices to follow when configuring your storage system for highest fault tolerance:

- Use dual controllers.
- Use two cable connections from each host.
- Provide multiple FC switches for SAN configuration (recommended for redundancy).
- Use Multipath I/O (MPIO) software.

Disk Group RAID Type Considerations

RAID 6 is recommended when using large capacity Midline (MDL) SAS drives in the Archive Tier. The added redundancy of RAID 6 will protect against data loss in the event of a second disk failure with large MDL SAS drives.

RAID 5 is commonly used for the Standard Tier where the disks are smaller and faster resulting in shorter rebuild times. RAID 5 is used in workloads that typically are both random and sequential in nature.

See the Best practices for SSDs section for RAID types used in the Performance Tier and Read Cache.

Disk Sparing

By default, the HPE MSA Storage enables Dynamic Sparing. This means any available drive not assigned to a Virtual Disk Group is available to dynamically replace a Virtual Disk Group drive that becomes degraded. When a disk fails in a virtual Disk Group, the HPE MSA software looks for a compatible global spare. If it finds one, the OS dynamically swaps it for the failing drive and automatically reconstructs the Virtual Disk Group. If no compatible disk is available, the array will send alerts of the failing disk. Replacing the failing disk with a compatible disk will trigger an automatic reconstruction.

During reconstruction of data, the effected Virtual Disk Group will be in a degraded or critical state until the parity or mirror data is completely written to the spare. Upon completion the Disk Group will return to a fault tolerant status.



Best practice: Have at least one compatible disk “spare” available for dynamic sparing for every fault tolerant Virtual Disk Group.

Virtual Volume Affinity

Hewlett Packard Enterprise recommends the default “No Affinity” option for most configurations. This setting attempts to balance the frequency of data access, disk cost, and disk availability by moving a volume’s data to the best disk tier for the data usage. Setting this affinity setting internally sets where data for the volume will be written first. For example, if this volume is for movies the users do not anticipate watching very often, then setting the volume affinity to “Archive” will mean the files you copy to this volume will be written to the archive tier of the Virtual Disk Group first. If a particular movie is watched frequently (daily) then it will be moved to a different tier to speed up access.

If the virtual volume uses mostly random or bursty low latency workloads such as Online Transaction Processing (OLTP), Virtual Desktop Infrastructure (VDI), or Virtualization environments, Hewlett Packard Enterprise recommends setting the preference to “Performance.” This setting keeps as much of this volume’s data in the Performance tier as long as possible.

If the virtual volume contains infrequently accessed workloads such as backup data or email archiving, Hewlett Packard Enterprise recommends setting the preference to “Archive.” This option will keep as much of this volume’s data in the Archive tier for as long a period as possible.



Best practice: Set the volume affinity setting to “Performance” for OLTP, VDI, and Virtualized environments.

Storage administration

After installation, the HPE MSA array is configurable and managed from a web browser interface called the Storage Management Utility (SMU). The SMU gives storage administrators a simple way to perform day-to-day administration tasks like: monitor the health of the system components, create Storage Pools, and manage Disk Groups and Virtual Volumes. The SMU allows for the creation of hosts which can be associated to host adapter WWNs and IQNs discovered as part of the SAN fabric. These hosts can then be mapped to Virtual Volumes to control access. In addition to the SMU, a command line utility is also available that does everything the GUI offers along with some more advanced features.

Storage Management Utility

Open a web browser and enter the IP address of the controller module's network port in the address field, then press Enter. The default user name is **manage** and password **!manage**. If the default user or password has been changed for security reasons, enter the secure login credentials instead of the defaults shown above.

Important

For detailed information on accessing and using SMU, see the "Getting started" section in the HPE MSA 1040/2040 SMU Reference Guide or HPE MSA 2050 SMU Reference Guide. The Getting Started section provides instructions for signing-in to SMU, introduces key concepts, addresses browser setup, and provides tips for using the main window and the help window.

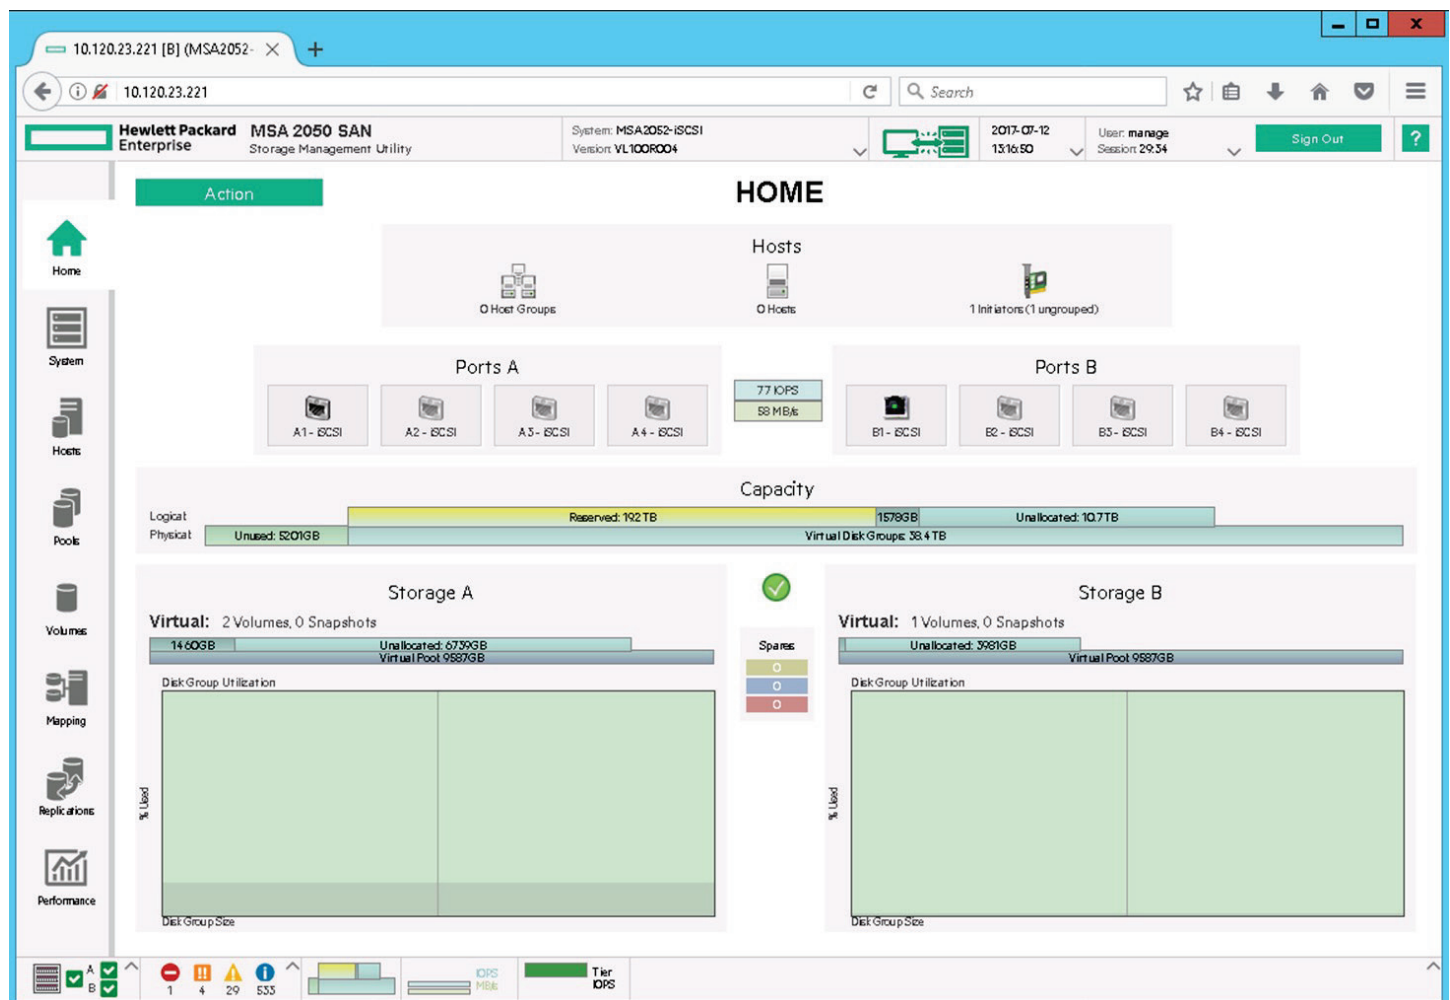


Figure 13. HPE MSA SMU interface

Configuring a new HPE MSA Storage system

Configuration Wizard

Configure basic settings for system operation.



Step 1 of 9: Introduction

This wizard helps you to initially configure the system, or to later change system settings. The wizard guides you through these steps:

- Configure system date and time settings.
- Change passwords for the default users.
- Configure each controller's network port.
- Enable or disable system-management services.
- Enter information to identify the system.
- Configure event notification.
- Configure host ports (if applicable).
- Confirm changes and apply them.

For each step you can view help by clicking the help icon . If you cancel the wizard at any point, no changes are made.

To continue, click **Next**



Figure 14. HPE MSA 2052 SMU Configuration Wizard

After signing into the SMU, use the Configuration Wizard to setup the HPE MSA for the first time. This will walk through the following setups:

1. Set the system date and time or connect to Network Time Protocol (NTP) Service.
2. Setup the password for the **manage** and **monitor** user accounts.
3. Change the Network configuration for the HPE MSA Controller(s).
4. Set the system-management services.
5. Add System contact information.
6. Configure SNMP, email, and log services for the HPE MSA.
7. Configure the FC or iSCSI port setting for the Controller(s).



Best practice: Configuring ALERTS is highly recommended as this is how the array will notify the admin of a fault.

Complete the initial installation and setup of the HPE MSA array before provisioning your vSphere hosts. Before provisioning Virtual Disk Groups and Virtual Volumes for vSphere see the section [HPE MSA Considerations and Best Practices for vSphere](#).

Using multiple enclosures

When using multiple enclosures, add Virtual Disk Groups to a Storage Pool by striping drives across the shelf enclosures using the Power of 2 Model. It is recommended that the user match the RAID level of other Virtual Disk Group(s) in the Storage Pool.

Storage pools and controller ownership

Storage pools are associated with either of the two controllers. By default, Pool A is assigned to Controller A and Pool B is assigned to Controller B. When creating Virtual Disk Groups, the SMU attempts to distribute disk group creation evenly across these two Storage Pools. Since both controllers are active, at least one Virtual Disk Group must be assigned to each pool to balance the controllers.

Naming hosts

The HPE MSA Storage version 3 of the SMU assist with the administration of your SAN environment by allowing the association of WWNs and IQNs with an initiator nickname, host, and host group. This provides a way to simplifying the mapping of WWNs to the Virtual Volumes on the array. Data center administrators may have a naming convention in use within the data center to keep track of servers and their components.

To create a host name association in the SMU, follow these steps:

1. Log in to the SMU.
2. Click on the Host Icon.
3. Select the WWN which corresponds to one of the adapters of the ESXi host's unique ID.

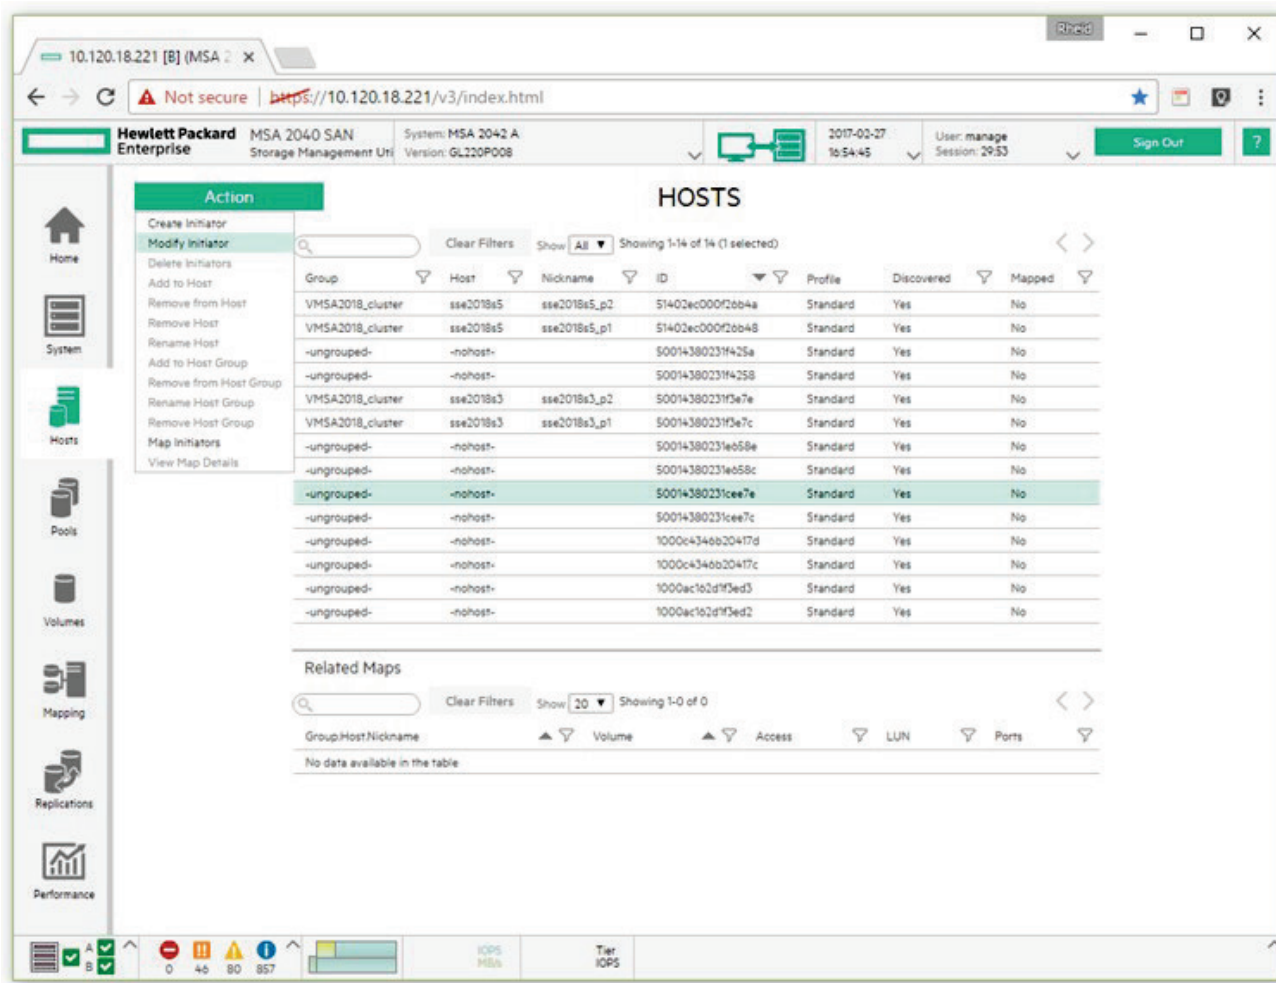


Figure 15. HPE MSA SMU Host Configuration Screen

4. Click the Action button and select the "Modify Initiator" option.
5. Type a name for the initiator. (Recommend using _p# to identify the port).

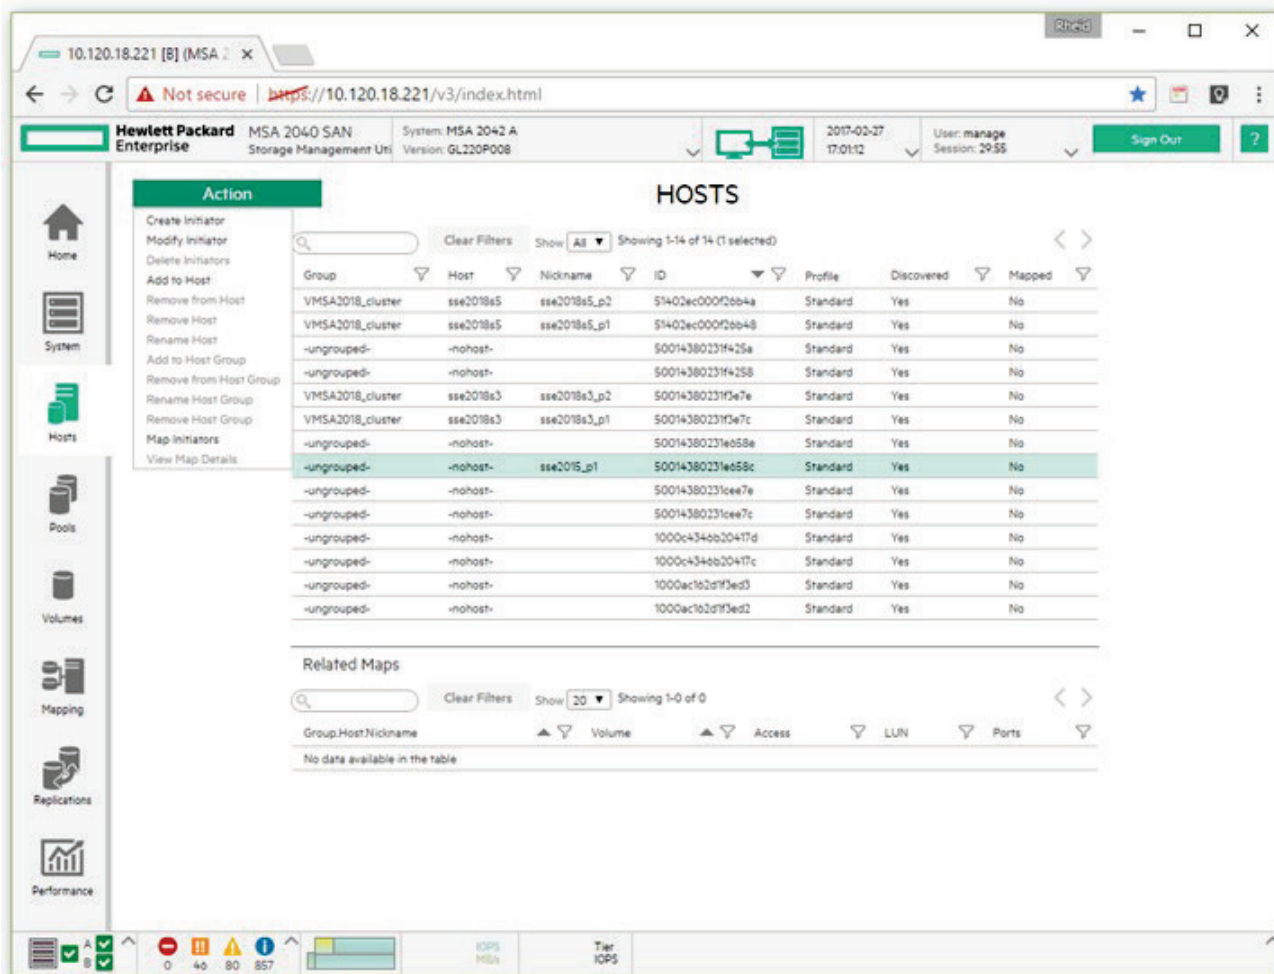


Figure 16. HPE MSA SMU Host Configuration with cluster

6. Save the change then click the "Action" button and select "Add to Host."
7. Type a name for the Host.
8. Save the change.

To create a Host Group at least 2 hosts must be selected from your list before the Add to Host Group will be a selectable option.

Once the Hosts and Host Groups have been created the volumes can be mapped not only to the individual initiator adapter's WWN; but, also to the Host and Host Groups to simplify administration.

In many cases, because the HPE MSA Storage is targeted to small businesses, the user may be the SAN, vSphere, and network administrator; it is particularly important to understand which hosts access which LUNs. This task is separate from working with zones at the Fibre Channel switch level.

HPE MSA Considerations and Best Practices for vSphere

Setup and installation

Storage plays a critical role in the success of VMware vSphere deployments. The following section highlights recommended practices for setup and configuration of the HPE MSA Storage Array best suited for virtualization environments running VMware vSphere.

Storage cabling and network connectivity

Production VMware vSphere implementations are based on clusters of servers requiring shared access to volumes on a storage array. For this type of configuration, HPE recommends a dual controller HPE MSA model supporting 16 Gb Fibre Channel host connections. Each controller should have a separate fibre connection to two Fibre Channel switches or switch fabrics to support redundancy and multi-path operations.



Best practice: For production environments use dual controller HPE MSA 2042 or HPE MSA 2052 Storage Array with redundant Fibre Channel connections to separate FC switches.

Storage configuration

The storage demands of a typical virtualized environment using VMware vSphere are not I/O intensive once a virtual machine is up and running. There are exceptions; however, generally the user wants the storage to be performant upon startup of VMs, generally fast for interaction, and then smart enough to off load snapshots to an archive tier. This is exactly what the HPE MSA can do.

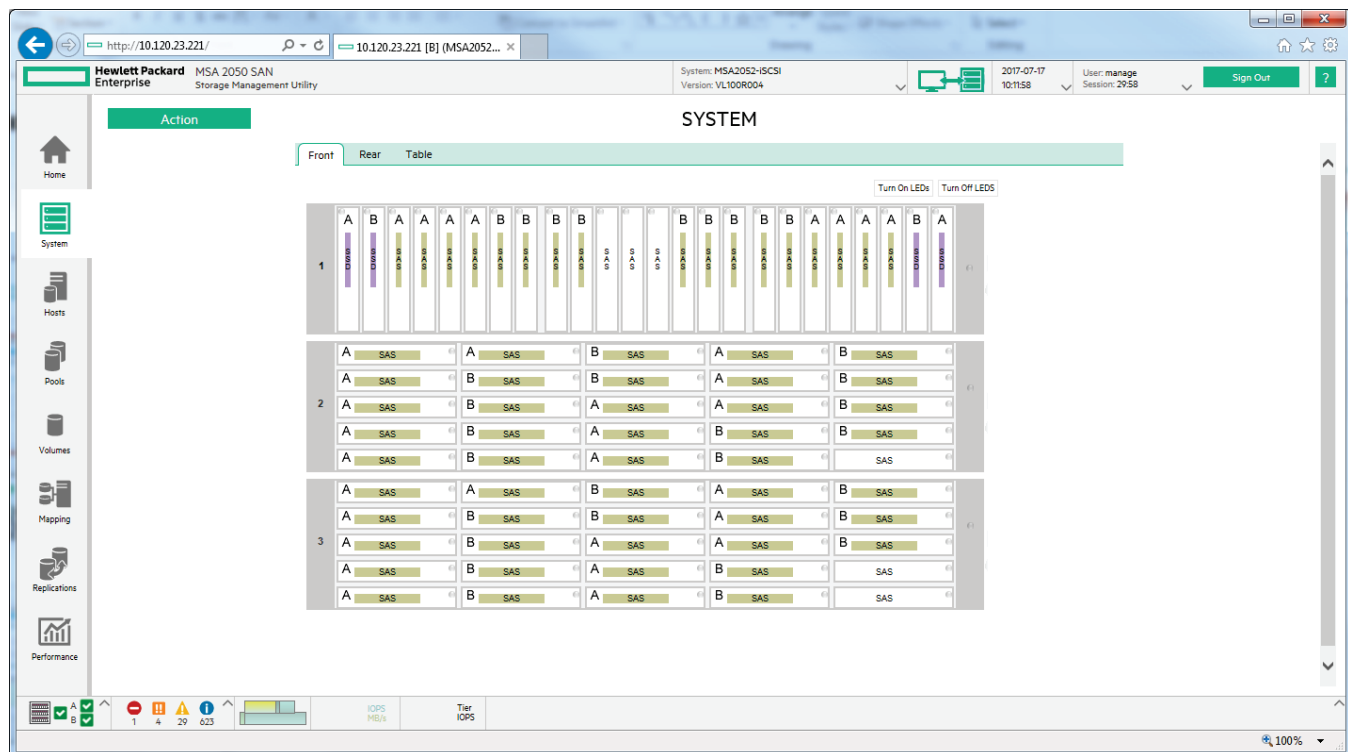


Figure 17. HPE MSA SMU system front view

For vSphere virtualized server environments, HPE recommends to setup and configure your HPE MSA Storage Array in the following manner.

1. For the base dual controller HPE MSA 2042 or HPE MSA 2052 enclosure, install the same capacity SAS drives in all but the first and last slots.
2. Install SSD drives in any slots of the enclosure.
3. Create two Storage Pools (A and B).
4. Create a performance RAID 6, Virtual Disk Group with 10 disk drives for each pool.
 - a. Pool A—dgA01 disks 1.2—1.11
 - b. Pool B—dgB01 disks 1.14—1.23
5. Create a Read Cache Disk Group for each pool.
 - a. Pool A—rcA1 disk 1.1
 - b. Pool B—rcB1 disks 1.24
6. Leave 1 disk per RAID Virtual Disk Group as a Dynamic Spare.
 - a. Disk 1.12 Spare
 - b. Disk 1.13 Spare
7. Create a single Virtual Volume for each pool.
 - a. Pool A—Vol0001
 - b. Pool B—Vol0002
8. Map both volumes to the hosts in your vSphere cluster.

The screenshot displays the HPE MSA SMU Storage Pool Screen. The interface includes a top navigation bar with the HPE logo, system information (MSA 2050 SAN, System: MSA2052-ISC1, Version: VL100R004), and user details (User: manage, Session: 29:43). The main content area is titled "POOLS" and shows a table of storage pools. Below this, there are sections for "Related Disk Groups" and "Related Disks".

Name	Health	Size	Avail	Volumes	Disk Groups
A	OK	9.5TB	7567.4GB	2	3
B	OK	9.5TB	9.0TB	1	3

Name	Health	Pool	RAID	Disk Type	Size	Free	Current Job	Status	Disks
dgA01	OK	A	RAID10	SAS (Standard)	4793.9GB	3783.7GB		FTOL	16
dgA02	OK	A	RAID10	SAS (Standard)	4793.9GB	3783.7GB		FTOL	16
rcA1	OK	A	RAID0	sSAS (Read Cache)	797.1GB	797.1GB		UP	2

Location	Health	Description	Size	Usage	Disk Group	Status
1.3	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
1.4	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
1.5	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
1.6	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
2.1	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
2.2	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
2.3	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up
2.4	OK	SAS	600.1GB	VIRTUAL POOL	dgA01	Up

Figure 18. HPE MSA SMU Storage Pool Screen

VMware vSphere® VMFS datastores on multiple Virtual Volumes on the same Virtual Disk Group can have an adverse effect on overall system performance of the vSphere cluster. Using the above recommended configuration of a single Virtual Volume for each Storage Pool with a Virtual Disk Group aligned with the Power of 2 Model will maximize performance and capacity.

Tiered Storage

The goal of Virtual Disk Group tiering of the HPE MSA Storage is to keep data in the fastest, highest tier as much as possible. Until the highest tier of storage is consumed, data will reside there until more active data needs to take its place. This forces the less used data to a lower tier. For this reason, adding an HPE MSA expansion of lower cost, larger capacity disks for an “Archive” tier when the performance Virtual Disk Group is not at capacity will appear unused.

Remember the affinity setting identifies where the data will be written first. Migration of the data will happen automatically.

Tiered Storage can benefit vSphere environments in a number of ways. Loading Virtual Machines is largely a read operation. Adding a Read Cache to your Storage Pool can boost the performance of loading virtual machines or creating VMs from templates that are used frequently.

If VM snapshots are created in the environment, having an “Archive” tier configured as part of the Storage Pool will provide automatic migrating of unused snapshots once your higher tier storage capacity is full.



Best practice: For production virtual environments create two Storage Pools that include: 1 RAID 6 Virtual Disk Group with 10 drives set to an affinity of “Performance.” Include in that Storage Pool a Read Cache Virtual Group with a single SSD. Save one drive for each Virtual Disk Group as a dynamic spare. And as the space in the “Performance” Virtual Disk Group is filled, add an expansion chassis of either “Performance” drives or “Archive” drives. SSD (Performance) drives are recommended in the main enclosure. Expansion is best served with standard or archive.

Boot from SAN

As a general rule, when booting from SAN with vSphere ESX servers, the Virtual Volume used for booting the server should be mapped to only one ESX Server system. The exception to this rule is when recovering from a system crash. If a new system is used for the server or the system’s FC HBAs have been changed (or the IP Address for iSCSI interfaces) then updating the mapping to the boot Virtual Volume is appropriate.

For more information regarding vSphere installation and boot-from-SAN configurations, refer to the vSphere Installation and Setup guides for vSphere.

Balancing controller ownership

The SAN administrator’s goal is to drive as much performance out of the array as possible. This is why it was recommended earlier, creating a balanced HPE MSA Storage Array environment. Both controllers in the HPE MSA are active at the same time. That is why 2 Storage Pools with the same configuration was created. To keep this configuration balanced, the VM Host’s goal is to keep storage needs balanced. A good rule to help keep the storage array balanced is to alternate VM storage requirements between the two Virtual Volumes created for the storage pool. This can be simplified through the administration of storage within vCenter. (See the [HPE MSA Considerations and Best Practices for vCenter](#) section).



Best practice: For production virtual environments, VM storage should be balanced between controller A and B. This can be administered most effectively through the creation of Virtual Volumes mapped to Storage Groups alternating between controllers and by using vCenter Datastore Clusters.

Changing LUN Response

Before creating volumes/LUNs for use by vSphere hosts, it is **essential** to change the setting for LUN Response on the HPE MSA Storage to ILLEGAL REQUEST. The following two VMware knowledgebase articles discuss this topic which relates to this subject.

[VMware KB article 1003433](#): SCSI events that can trigger ESX server to fail a LUN over to another path.

[VMware KB article 1027963](#): Understanding the storage path failover sequence in VMware® ESX®/ESXi 4.x and 5.x.

To change the LUN RESPONSE setting through the CLI:

1. Shut down the host.
2. Log into either the controller A or B CLI.

3. When logged into the CLI, enter the following command:

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

4. Restart the host.

Volume Mapping

Virtual machines access data using two methods: VMFS (.vmdk files in a VMFS file system) and raw device mapping (RDM). The only difference is RDM contains a mapping file inside VMFS that behaves like a proxy to the raw storage device. RDM is recommended in cases where a virtual machine needs to simulate interacting with a physical disk on the SAN. If using RDM is anticipated with the virtual machines, make sure the HBAs are supported for the ESX hosts. Not all HBAs are supported by VMware for this feature on the HPE MSA Storage. See the [SPOCK](#) compatibility matrix on the HPE website.

Caution

VMware vSphere versions prior to 5.5 had a virtual machine data (VMDK) size limitation of 2 TB. If dedicated volumes for VMDKs are planned, this will require the creation of multiple virtual volumes on the HPE MSA. This creates a less than optimal performance model and a more complex management for a cluster of ESX hosts that need to access these virtual volumes.

VMware recommends the following practices for volume mapping:

- Use explicit LUN mapping.
- Make sure that a shared LUN is mapped with the same LUN number to all vSphere ESX servers sharing the LUN.
- Make sure that the LUN is mapped through all the same controller ports for all mapped server WWNs, so that each server has the same number of paths to the LUN.
- Map the LUNs to the ports on the controller that own the disk group. Mapping to the non-preferred path may result in performance degradation.

The HPE MSA Storage version 3 SMU simplifies this process by providing virtual volumes and hosts configurations. The HPE MSA also supports ULP so only mapping between hosts and virtual volumes is needed—not each server WWN to each HPE MSA controller port and the virtual volume owner.

With version 3 of the SMU software, administrator can use aliases for initiators, hosts, and host groups to simplify the mapping and management of volumes. If the new features are not used, volumes must be individually mapped to the WWN or IDQ of each ESX server's interface and path to the array. Each server must include the exact same LUN number and mapping assignments to the array for every shared volume.

This process is very simple with the latest software which allows volumes to be mapped by cluster, servers in the cluster, or individual adapters of the servers in the cluster.



Best practice: When creating Mappings in the SMU, ensure there are no conflicting LUN numbers being exposed to vSphere Hosts and Clusters. And never identify a LUN number using 0 (zero).

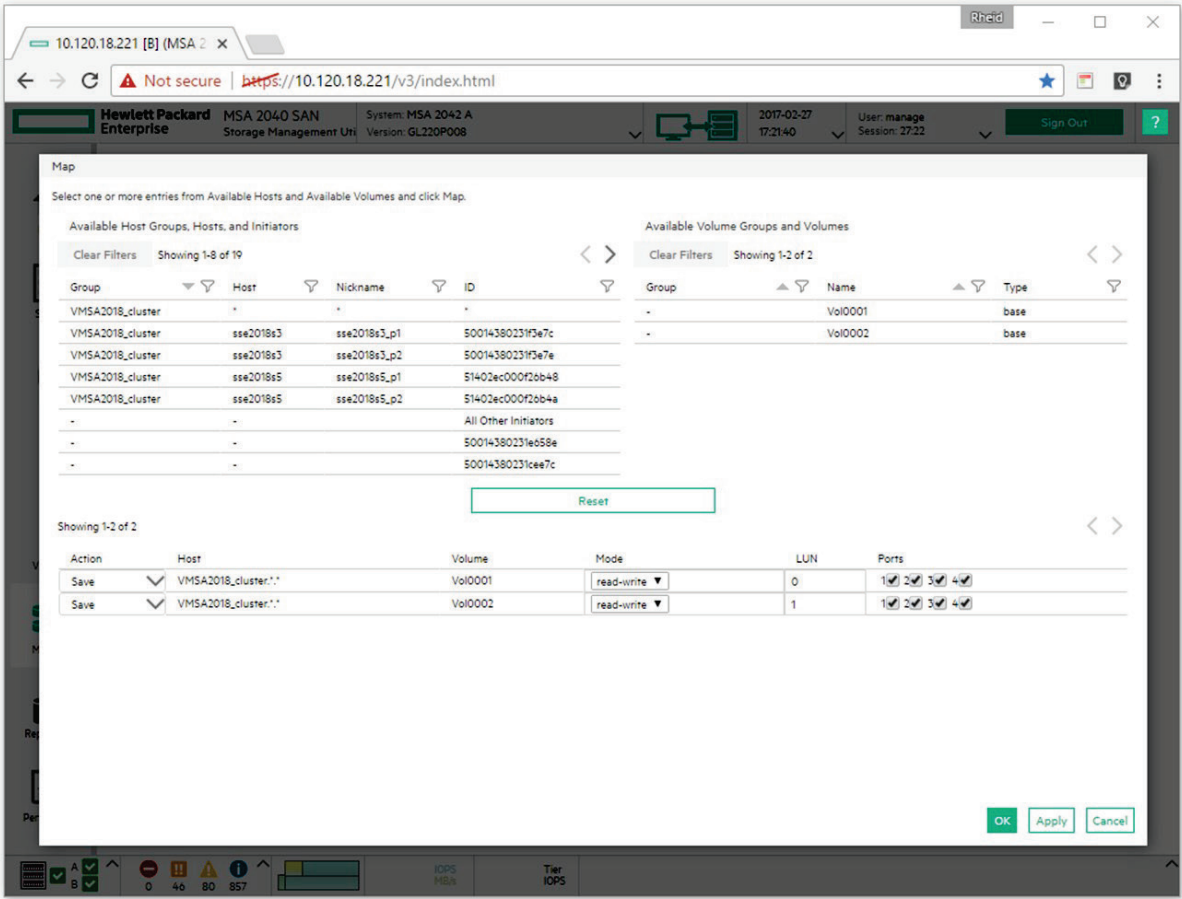


Figure 19. HPE MSA SMU Volume Mapping Screen

Common configuration error

The HPE MSA Storage enclosure is presented to vSphere as a Fibre Channel enclosure using LUN 0. This means all volumes mapped to the vSphere hosts or clusters must not use LUN 0. The SMU software doesn't automatically detect this and for each mapping created, it defaults to LUN 0. (See the Device view of the ESX host below to see how the HPE MSA controller FC ports show up as unknown type and capacity).

For example, in the previous screen shot two volumes were mapped to the VMMSA2018 cluster—Vol0001 assigned LUN 0 and Vol0002 assigned LUN 1. Because the HPE MSA enclosure is exposed as LUN 0, only Vol0002 could be seen in the vCenter management software.



Best Practice: HPE recommends using the new alias features and create Nicknames for each HBA for each vSphere server then adding a host name for the server. For ESX clusters, it is recommended to create a Host Group and managing the shared volumes through a Host Group mapping.

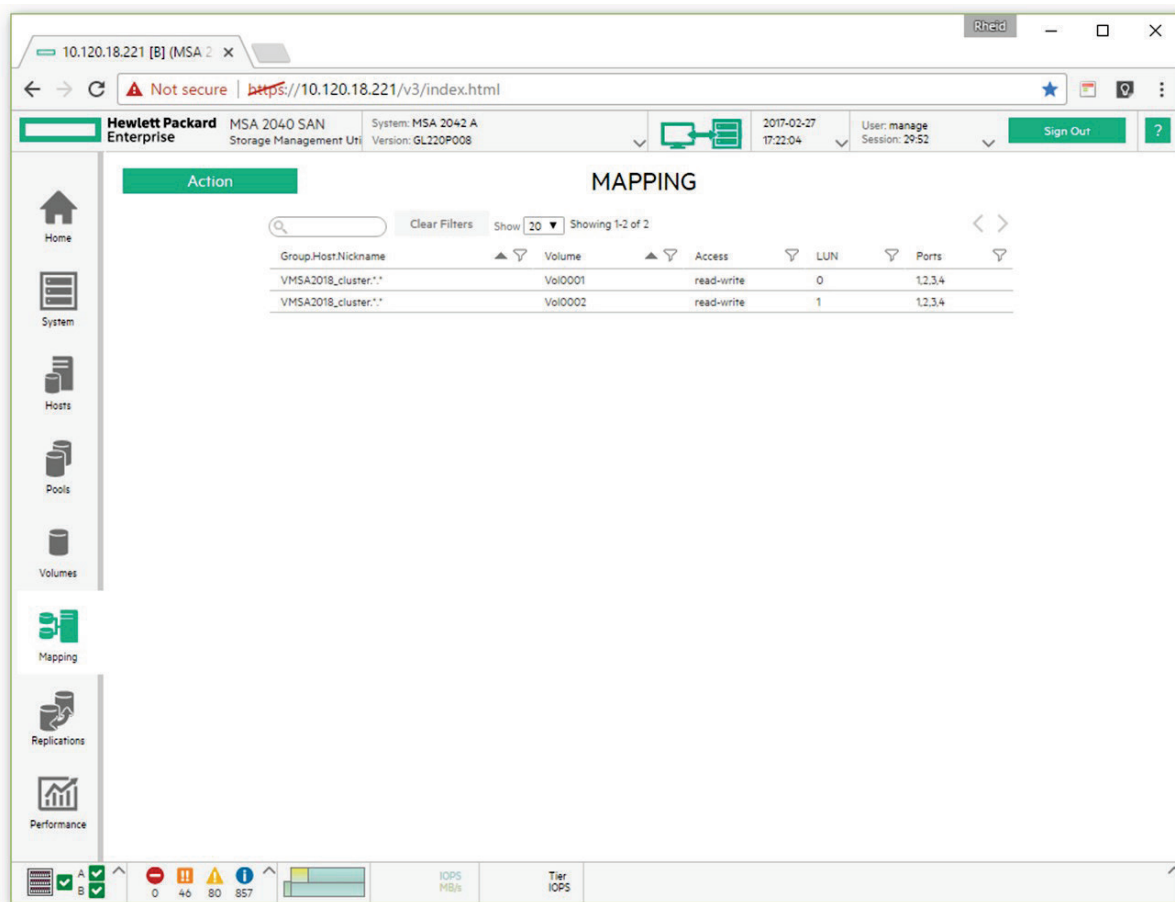


Figure 20. HPE MSA SMU Volume Mapping for a Cluster

Note

The simplification of mapping seen in the screen above actually represents 4 server connections to the two volumes through 4 ports on the array. When viewing the HPE MSA Storage configuration through the CLI interface on the array, the alias name associations will not be seen. Only the volume and WWN associations—all 16 connections will be seen.

Presenting Storage to Virtual Machines

Understanding the HPE MSA architecture, the vSphere cluster features, and the applications being virtualized are all essential when planning the creation of virtual volumes in the vSphere environment. For example, a single VM providing a large database accessed by multiple users can be adequately serviced by a single virtual volume. This single virtual volume when assigned to a storage pool with a disk group made up of RAID 6 SAS drives will provide adequate performance as well as fault-tolerant storage space for the database application.

The following sections highlight the best practices when configuring the HPE MSA and vSphere for the vSphere virtual environment.

Default Volume Mapping

Each volume created in vSphere has a default host-access setting called a default mapping. Default mappings allow all hosts specified in the mapping to connect to the controller host port(s) to access the volume. By default, these mapping tables are created such that all hosts connected to the specified ports have access to the volume. Specifying explicit host mappings during the creation of a volume map will restrict the visibility of the volume to the specified hosts.

The advantage of using a default mapping is that all connected hosts can discover the volume with no additional action by the administrator. The disadvantage is that all connected hosts can discover and access the volume without restrictions.



Best practice: Explicitly map virtual volumes to VMware vSphere hosts. Do not use the default mapping when creating any volume requiring restricted access.

ULP and vSphere

Hypervisors such as VMware vSphere use ALUA to communicate with backend storage arrays. ALUA provides multipathing (**two or more storage networking paths**) to the same LUN on a storage array and marks one path “**Active**” and the other “**Passive**.” The status of the paths may be changed either manually by the user or programmatically by the array.

VMware vSphere 5 is also ALUA-compliant. This was one of the major features added to the vSphere 5 architecture, allowing the hypervisor to:

- Detect that a Storage system is ALUA-capable—use ALUA to optimize I/O processing to the controllers
- Detect LUN failover between controllers

vSphere supports the following ALUA modes:

- Not supported
- Implicit
- Explicit
- Both implicit and explicit support

Additionally, vSphere 5 also supports all ALUA access types:

- Active-optimized—The path to the LUN is through the managing controller.
- Active-non-optimized—The path to the LUN is through the non-managing controller.
- Standby—The path to the LUN is not an active path and must be activated before I/O can be issued.
- Unavailable—The path to the LUN is unavailable through this controller.
- Transitioning—The LUN is transitioning from and to any one of the types defined above.

VMware vSphere 5 supports round robin load balancing, along with Most Recently Used (MRU) and Fixed I/O path policies. Round robin and MRU I/O path policies are ALUA-aware, meaning that both round robin and MRU load balancing will first attempt to schedule I/O requests to a LUN using a path through the managing controller. For more details, see the Multipath Considerations for vSphere section.

Multipath Considerations for vSphere

To maintain a constant connection between a vSphere host and storage, ESX software supports multipathing. To take advantage of this feature, the ESX host requires multiple FC, iSCSI, or SAS adapters and the HPE MSA virtual volumes need to be mapped to these adapters.

This can be accomplished easily on the HPE MSA Storage by creating a host definition as outlined in the previous section and associating the World Wide Names (WWNs) of the multiple interfaces (HBA ports) on the host server to this new host object. When mapping a Virtual Volume to the host object in the SMU, all the path mappings are automatically created to support multipath to the host. To do this in the CLI an entry for each path would need to be created or use the Host/Host Groups with wildcards.

As recommended in the previous section, HPE recommends configuring the HPE MSA Storage to use a Host Group for a vSphere cluster and use the cluster object when mapping Virtual Volumes. This will create all the mappings to all the adapters to support multipath in the cluster in one step.

VMware vSphere supports an active/active multipath environment to maintain a constant connection between the ESX host and the HPE MSA Storage Array. The latest version of vSphere offers 3 path policies: “Fixed,” “Most Recently Used,” and “Round Robin.”

HPE recommends using the “Round Robin” preferred selection path (PSP) policy for best performance and load balancing on the HPE MSA Storage.

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

With vSphere 6.x the default storage array type for the HPE MSA Storage Array is “VMW_SATP_ALUA.” By default the path selection policy is set to use the Most Recently Used path (VMW_PSP_MRU). This selects the first working path discovered during boot up. If the path becomes unavailable it moves to another path.



Best practice: For the HPE MSA Storage Array, HPE recommends for optimal performance to set the PSP policy to Round Robin for all vSphere hosts with multiple paths to the HPE MSA Storage. This will include all ESX hosts as long as the HPE MSA has more than one connection to the SAN and the SAN fabric is exposing those connections to all hosts.

Although path selection policies can be viewed and modified in many ways, the following example screen shots show the vCenter method for viewing and configuring the path selection policy. Since the multipath policy is specific to the ESX Host, each volume that supports multipath on each ESX Host must be changed to the Round Robin policy. See [Appendix E](#) for more specific multipathing details for older version of vSphere.

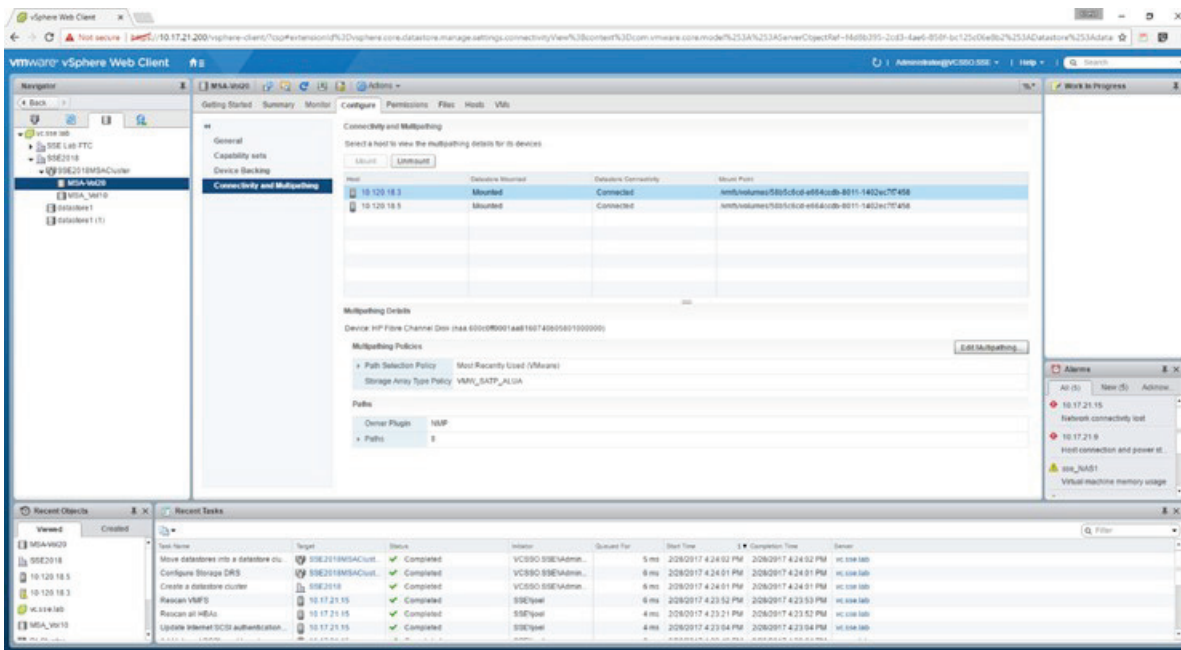


Figure 21. VMware's vCenter Volume Multipathing Screen

By selecting the Connectivity and Multipath tab for the storage volume in a storage cluster, you can edit the setting for the path selection policy to Round Robin as shown below. This will allow requests to utilize all available paths to the HPE MSA Storage Array.

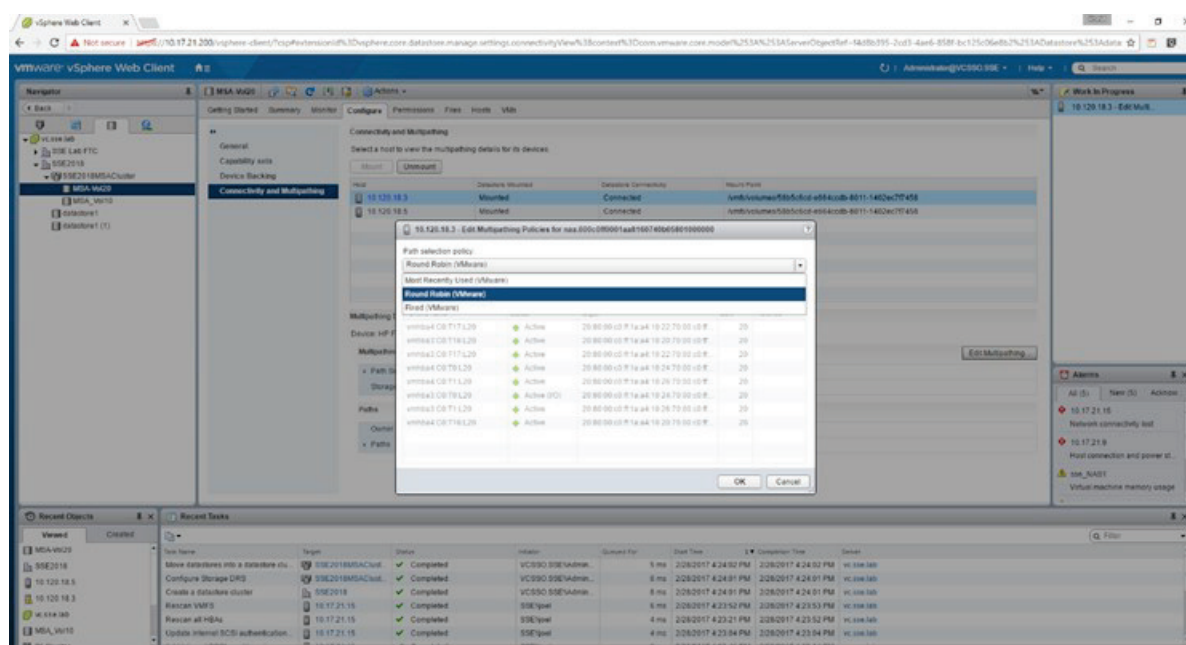


Figure 22. VMware vCenter Volume Multipathing Policy Screen

HPE MSA Considerations and Best Practices for vCenter

Datstore 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 virtual volumes allows for many of the vCenter automations to help balance the storage load. For example, when creating a VM, the initial storage creation will balance between the volumes of the datastore cluster.

Another benefit of the datastore cluster is during the datastore creation, 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 mappings on the MSA have been configured properly.



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

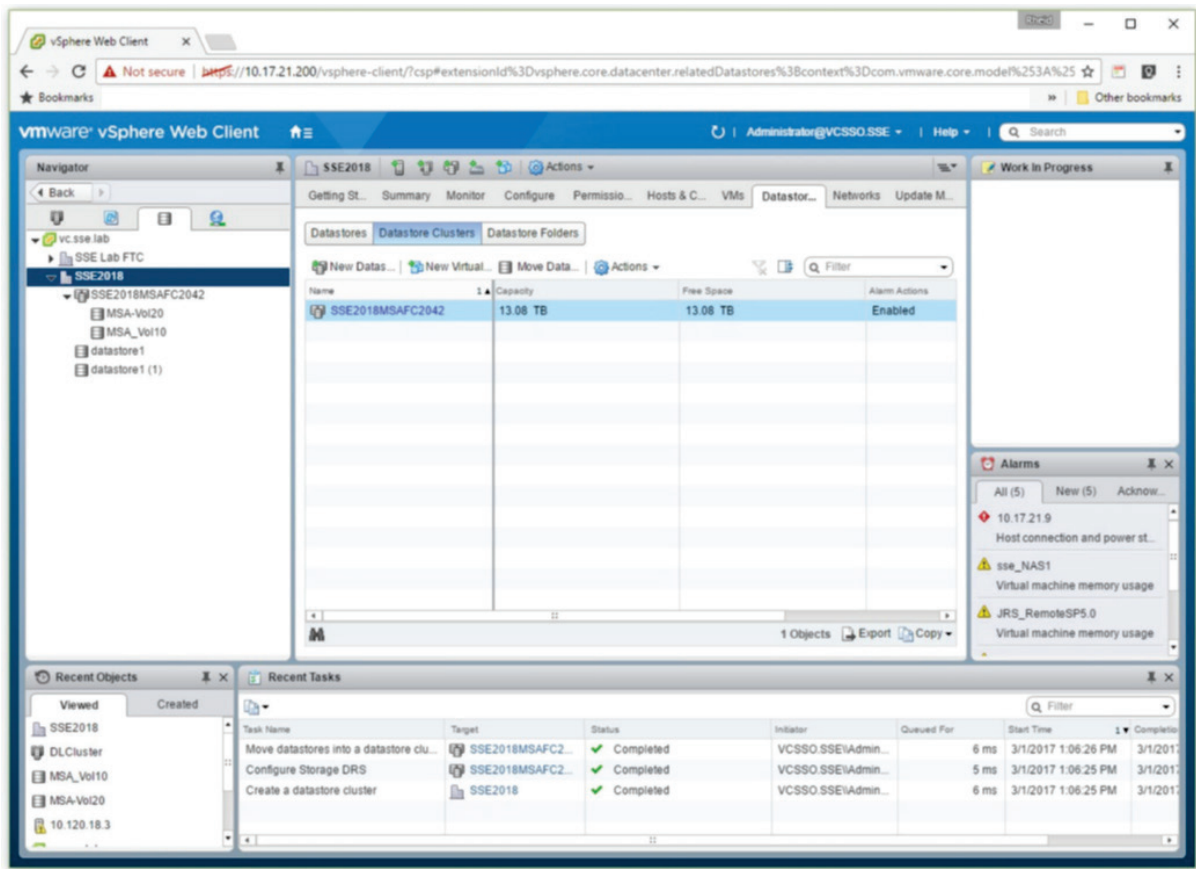


Figure 23. VMware vCenter Datastore Cluster Screen

Distributed Resource Scheduler

VMware vCenter Datastore Clusters provide a feature called Storage Distributed Resource Scheduler (SDRS) 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 MSA Storage Array controller and virtual volumes provide the benefits of DRS and I/O balancing dynamically, the VMware’s DRS and I/O load balancing are not needed with virtual volumes as long as all hosts in the cluster share the same MSA volume mappings.



Best practice: In the vCenter administration tools, do not enable DRS and I/O load balancing on data storage clusters that are solely based on MSA Virtual Volumes.

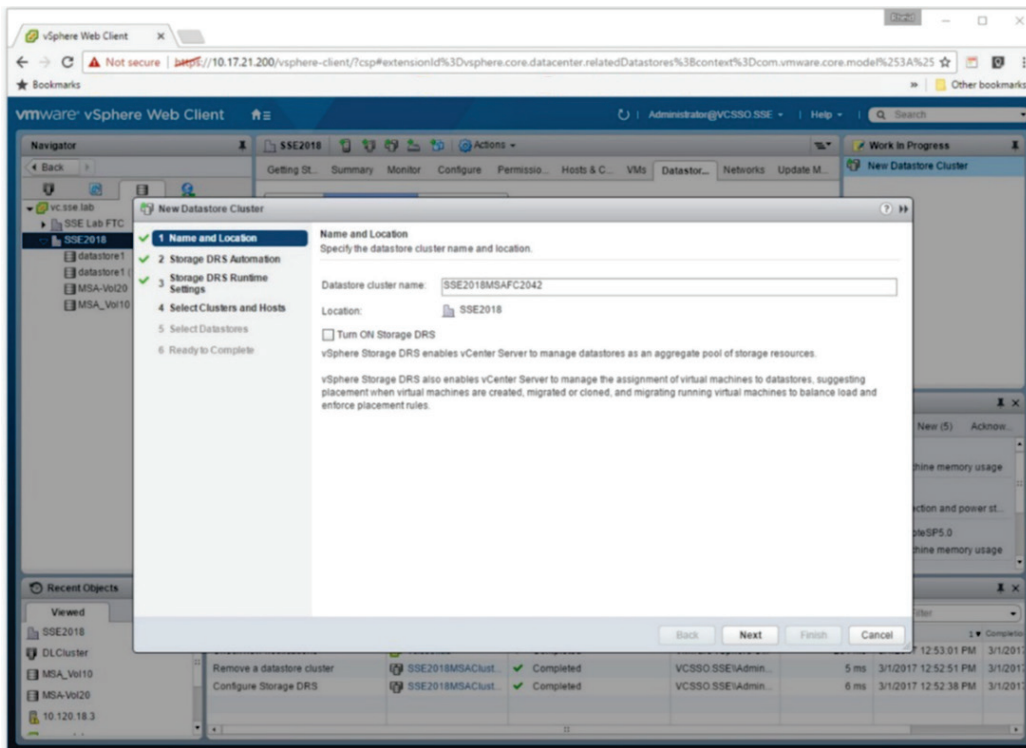


Figure 24. VMware vCenter Datastore Cluster Creation

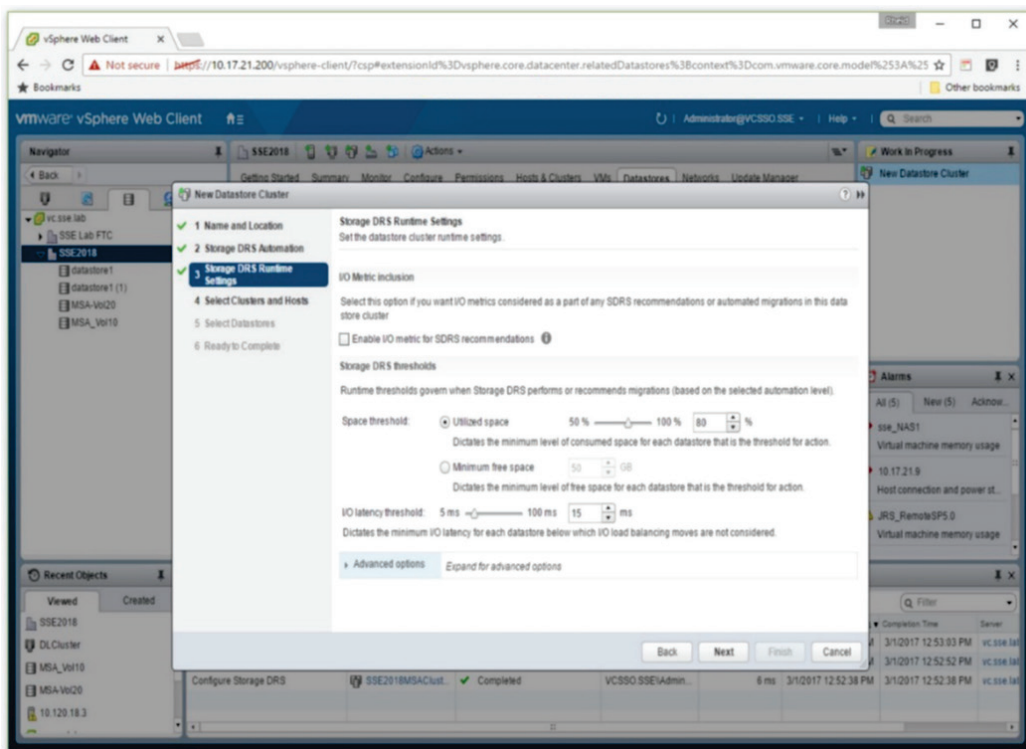


Figure 25. VMware vCenter Storage DRS Configuration Screen

Multipath and Datastore Clusters

The datastore cluster provides a simplified way of managing the multipath configuration for the volumes in the datastore cluster. With the latest version of vCenter, multipath setup and configuration of PSP can be done for any of the ESX hosts from the configuration view for any of the cluster volumes.



Best practice: For the HPE MSA Storage Array, HPE recommends for optimal performance to set the PSP policy to Round Robin for all vSphere hosts with multiple paths to the HPE MSA Storage.

Site Recovery Manager and the HPE MSA Storage

Along with the escalating costs of maintaining business-critical applications, for many small to midsize enterprises, the cost of replicating their application servers to recover from disasters can be prohibitive. Implementing a cost-effective data recovery (DR) solution using a VMware-certified SAN infrastructure that supports the failover between vSphere environments and virtual machine instances can be accomplished with the HPE MSA Storage Replication Adapter (SRA) and VMware's SRM.

A typical SRM configuration involves two geographically separated sites with TCP/IP connectivity between the protected site and the recovery site. The DR protected site is the site that is being replicated to the recovery site. Each site contains an HPE MSA Storage Array, vSphere hosts, a VMware® vCenter Server™, and an SRM server.

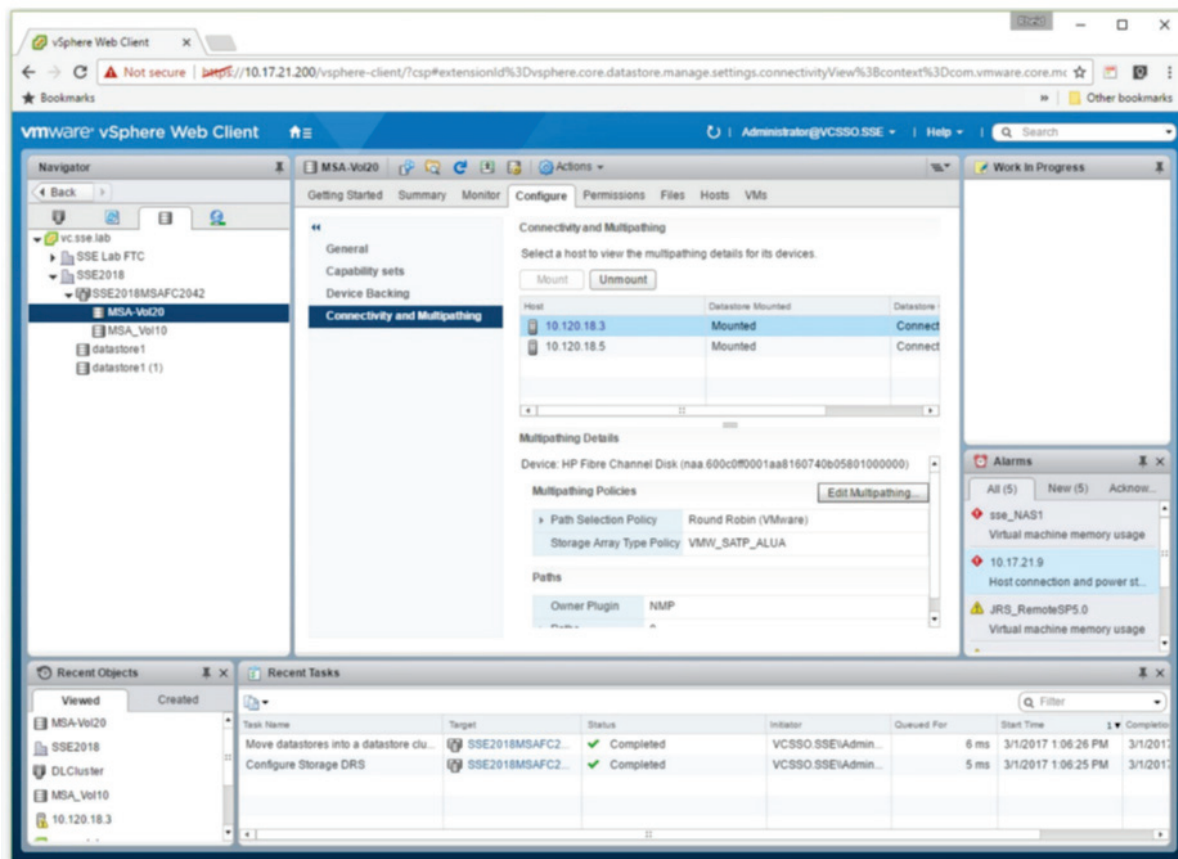


Figure 26. VMware vCenter Volume Multipathing Policies Screen

vCenter's single SRM interface displays and manages all protected recovery sites simplifying administration and the need to personally manage multiple sites separately.

The virtual machines used for DR failover can all reside on a small number of DR servers running VMware vSphere. Although the virtual machine instances may not offer the bandwidth to accommodate peak workloads, the solution avoids downtime. If necessary, use VMware vSphere® Storage vMotion® to migrate these virtual machines to more powerful servers as required. The VMware SRM does not replace the functionality of Remote Snap; rather it provides the mechanism for replicating virtual machine data between the protected site and the recovery site.

The HPE MSA SRA provides a platform for a seamless integration of HPE MSA Storage with SRM and aids in the automated DR of virtual machines. It enables SRM to coordinate the recovery process with the underlying replication mechanisms, so the virtual machines at the protected site are shut down properly and the replicated virtual machines can be powered up.

A recovery plan specifies the order in which the protected virtual machines are to be recovered. The recovery plan also specifies network parameters, such as IP addresses, and consists of user-specified scripts that perform custom recovery actions.

After a recovery, the virtual machines that are up and running are no longer protected. To reactivate protection, SRM supports a re-protect operation for virtual machines backed by HPE MSA storage systems. The re-protect operation reverses the roles of the two sites after the original protected site is back up. The site formerly the recovery site becomes the protected site, and the protected site becomes the recovery site.

SRM also enables recovery plan testing using native functionality without disruption to production environments. This is accomplished by using a temporary copy of the replicated data to conduct tests. After the re-protect operation is used to confirm that the new protected recovery site configuration is valid.

Building an effective disaster-tolerant solution can often be a very complex and time-consuming task. Most disaster-tolerant solutions implemented at customer sites are often untested and fail to protect the application when a failure occurs.

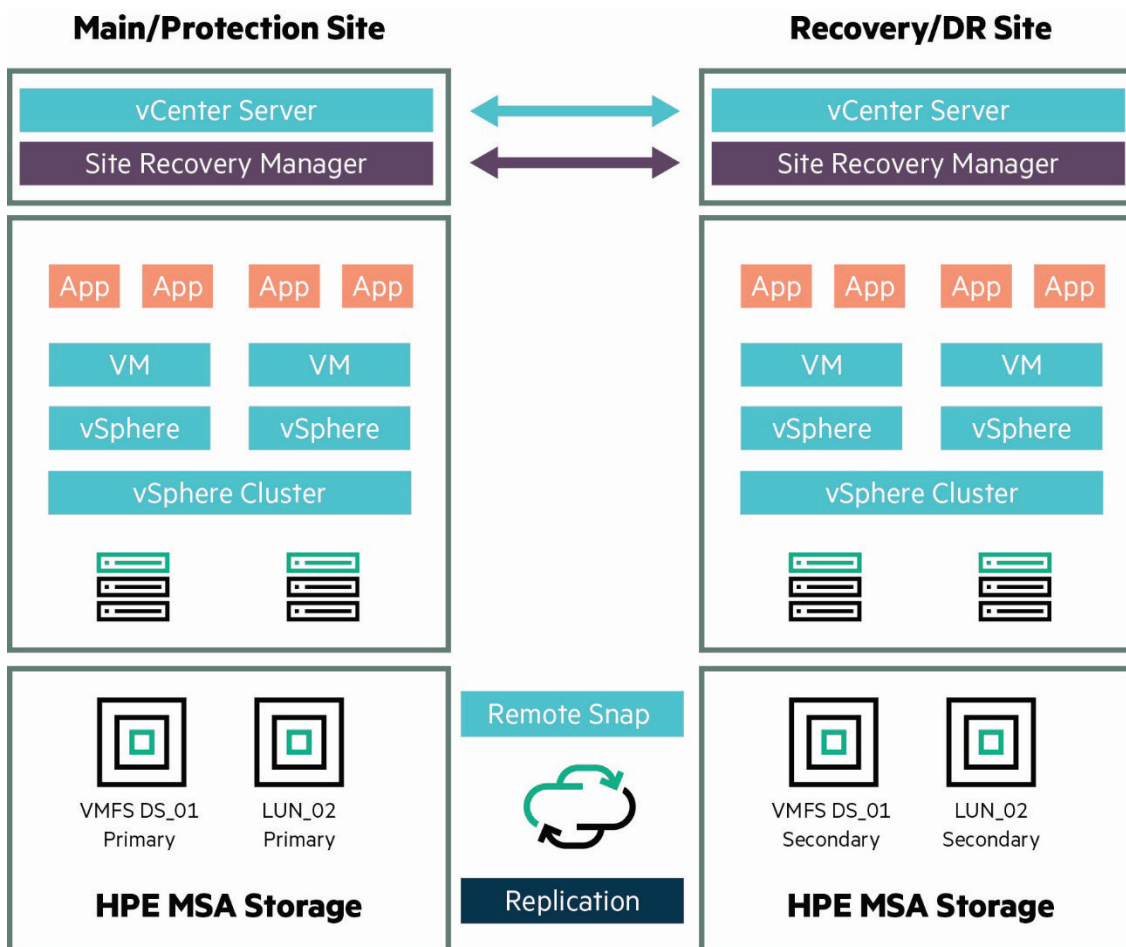


Figure 27. Multi-site Recover Model

Based on the data center solution or application, the recovery point objective (RPO) and recovery time objective (RTO) may differ from application to application and from customer to customer. Any disaster-tolerant solution must be able to accommodate both planned and unplanned downtime.

Taking full advantage of the encapsulation and isolation of virtual machines, VMware's Storage Recovery Manager (SRM) can use the HPE MSA SRA to automate a simplified disaster recovery environment. The HPE MSA SRA helps VMware's SRM meet customer's DR expectations by simplifying recovery time objectives, reducing the costs of traditional business continuance plans, and achieving low-risk, predictable results for a virtual recovery environment.



Best practices: When using the HPE MSA SRA with VMware SRM:

1. Prepare a plan for re-establishing replication schedules in the event of a site failover. After performing a reverse replication operation, setup replication schedules to ensure periodic replication of data from the new source volumes back to the original source site. Alternately, initiate replication manually, if appropriate.
2. Group virtual machines with related backup requirements or schedules on the same datastore volume, because replication occurs on a per-volume basis. For example, if some virtual machines do not need to be replicated to a remote site, or need to be replicated less frequently, do not store them on the same datastore volume as virtual machines which must be replicated frequently, to avoid replicating data unnecessarily.
3. Each array must have a unique alphanumeric "System Name" assigned that does not use any non-alphanumeric characters other than "." Or "-". No spaces are allowed in the system name.
4. Each array must have remote systems defined for each of the remote arrays to or from which data is being replicated. The SRA depends on these remote system names being defined for basic functionality.
5. The SRA only supports replication between identical hardware models. For example, replication between an iSCSI-only system and a Combo FC/iSCSI system is not supported.
6. Avoid mapping replication volumes to LUN 0. This avoids issues with dynamically mapping and unmapping LUNs, due to special management functionality assigned to LUN 0. Mapping volumes to LUN 0 can be done if those volumes are not expected to be mapped and unmapped automatically the way replication volumes are, such as local datastores that are not replicated.
7. Map replication volumes with the same LUN number on all hosts.
8. Do not use the same LUN number for different volumes that are mapped to different hosts.
9. Failover operations cause read-write host mappings for replication volumes to be converted to read only. Restoring replication converts all read-only mappings for the same volume to read/write. Be careful not to create read-only mappings of replication volumes, for instance, for data mining purposes. If a read-only mapping of a replication volume is required, consider creating a non-replicated hardware or software snapshot of the volume.

For more specifics about the SRA software see the [MSA Storage Replication Adapter Software User Guide](#).

Configuring SRM and the Storage Replication Adapter

Site Recover Manager is an extension of the vCenter Server. When configuring your SRM virtual DR environment that contain HPE MSA Storage arrays, SRM requires a specific Storage Replication Adapter for the HPE MSA. This adapter will need to be installed on both vCenter systems. Before configuring SRM, ensure the latest SRA software bundle has been downloaded and available by visiting the [VMware SRA Download site](#).

HPE MSA Configuration Requirements

Configuration of the Storage Replication Adapter for the HPE MSA Storage Array requires the setup of two HPE MSA Storage arrays—one at the protected location and one at the remote location.

An important aspect of the configuration of this environment is to ensuring the MSA connectivity to the hosts and vCenter resources are the same for both storage arrays. As a best practice, for production environments use Fibre Channel connectivity to both HPE MSA Storage arrays and the vSphere hosts.

The HPE MSA Storage arrays should also be configured on their own private IP Network for replication communications. This is not a necessity; however, isolating the replication traffic to their own private network will prevent potential network threats and network contention.

Configure both HPE MSA Storage arrays to use the same Network Time Protocol (NTP) servers.

Features of the SRA adapter require both HPE MSA Storage arrays to have licenses for Replication and Snapshots enabled.

Both MSA's need to also have the "Missing LUN Response" property set to "Illegal Request." See the [Changing LUN Response](#) section of this document for more details.

The HPE MSA Storage arrays need to be paired using the MSA's replication services.

For more specifics about the SRA software see the [HPE MSA Storage Replication Adapter Software User Guide](#).

Note

The Storage Replication Adapter can create extra host entries on the array to help keep track of remote IP or FC addresses. Do not delete these host entries. They will have names that start with "SRA."

vCenter Configuration Requirements

In order to activate the Site Recovery Manager, download the vCenter extension and install it to both vCenter Servers. The HPE MSA SRA adapter bundle will need to be installed on both vCenter Servers as well.

Configure both vCenter servers where SRM is going to be installed to use the same Network Time Protocol (NTP) servers that the HPE MSA Storage arrays are using.

After both SRM and the HPE MSA SRA are installed, access the "Getting Started" tab of the main SRM window in vCenter. For detailed SRM configuration instructions, refer to the Installation and Configuration, and Administration documentation available on the [Site Recovery Manager](#) website.

HPE OneView for VMware vCenter

HPE OneView for VMware vCenter includes a storage management integration plug-in for storage. HPE OneView for VMware vCenter provides the ability to create storage array devices, map them to array ports, and create datastores and virtual machines directly from vCenter.

HPE OneView for VMware vCenter can:

- Provision storage
- Obtain physical and logical end-to-end views of storage devices
- View reports on device capacity and usage
- Monitor storage information

HPE OneView for VMware vCenter saves time and resources, allowing administrators to manage both VMware vSphere and HPE Storage from a common interface.

HPE OneView for VMware vCenter integrates with the vCenter management server and the vSphere Client software from VMware. VMware vCenter is the single point of management for VMware virtual environments. These virtual environments can consist of many standalone or clustered vSphere configurations. The virtual environment administrator accesses the vCenter management capabilities using the vSphere Client software.

HPE OneView for VMware vCenter software can be installed on a management server or a VM and configured to connect and register with vCenter. After registering OneView for VMware vCenter, all vSphere clients connected to vCenter can use the HPE OneView Software to access the OneView for VMware vCenter software.

HPE OneView for VMware vCenter provides several ways to access storage information, based on the selected VMware object in the vCenter navigation tree. Users can select different views of the detailed storage information for the selected object or objects.

- **Cluster View:** Provides summary and detailed information about the HPE storage deployed in the selected cluster. Summary information describes the used and unused storage available to the cluster, and includes details describing the total amount of storage provisioned by the cluster as virtual disks (VMDKs or RDMs). The storage administrator receives, mapping information for each of the physical HPE Storage arrays that are presenting disks to the selected cluster.
- **ESX Server (host) level View:** Provides summary and detailed information about the HPE storage deployed to the selected host. Summary information describes the amount of storage available (provisioned) to the host and the amount of storage the host is currently using. Detailed information includes virtual and physical disks (including virtual to physical mapping), HBA's, data paths (including virtual to physical mapping, and local and remote copy replication status).
- **Virtual Machine View:** Provides summary information about the total amount of storage used by the VM and available to the VM. Provides a list of the arrays that provide storage to the selected VM and detailed mapping to the Virtual Disks (physical) Storage Disks, local and remote copy replication status, data paths and (physical) hosts, and datastores.
- **Datastore View:** Provides information about the datastores used by a selected VM, or the datastores located on the selected (physical) host. Summary information includes the provisioned and available capacity within the datastore. Detailed information includes mappings to (physical) storage disks, virtual disks, (physical) hosts, data paths, and replication status (local/remote copy).

HPE OneView for VMware vCenter is available online from this [HPE website](#) and can be installed and used at no charge. Acceptance of the standard HPE software licensing terms is part of the download and installation sequence. An HPE Passport sign-in will be required.

Use cases for vSphere and HPE MSA Storage

HPE MSA Storage products with VMware vSphere can accommodate a variety of dynamic business needs and business models. The following use cases, or real world scenarios, use HPE MSA Storage and VMware vSphere to provide for critical business needs while minimizing risk.

Mixed Physical and Virtual Computing Environment

Due to operating budget, time, or staffing constraints, many businesses are still virtualizing their core systems. Customer demand for services can also impact migration or upgrade projects. As a result, many businesses have a mixed computing environment, with traditional physical servers sharing workload and storage with virtualized systems. The HPE MSA storage solution supports the storage needs of both physical servers and vSphere hosts. As more systems become virtualized and consolidated, the HPE MSA can grow and scale to suit changing business needs.

Figure 28 is a diagram of a mixed or hybrid environment, with physical systems hosting critical applications and virtualized systems with virtualized applications, utilizing HPE MSA storage.




Figure 28. Mixed Physical and Virtual Computing Environment

Multiple vSphere clusters on HPE MSA Storage

Figure 29 illustrates a recommended best practice configuration for a vSphere 5.x clustered environment, including multiple volumes and LUNs presented to the vSphere hosts and formatted as VMFS datastores.




Figure 29. Multiple vSphere clusters on HPE MSA Storage

Distributed vSphere Environments

Remote Office/Branch Office (ROBO) deployments offer a decentralized approach for system deployments. Distributed VMware vSphere environments spanning multiple sites can have storage needs which can be cost-prohibitive. With a relatively small physical footprint, entry-level pricing, the HPE MSA Storage is a practical choice for a storage array solution for a remote office environment. Rapid Provisioning software eliminates array planning by delivering instant, fine-grained, application-tailored storage provisioning. Administrators don't spend their time managing storage because the simple and intelligent provisioning automatically adds more storage to support a growing virtual client environment.




Figure 30. Distributed vSphere Environments

Data Protection with HPE MSA Remote Snap

Cost-effective data protection can be a challenge for today's small to midsize businesses. The HPE MSA Storage Array and its Remote Snap feature, an affordable, dual storage-array with point-in-time data protection and data replication provides a disaster-tolerant solution. The array-based, Remote Snap replication feature, implemented with two HPE MSA storage system solutions, provides a business site with ultimate data protection and redundancy.

Remote Snap Software is an optional, licensed disaster-recovery feature that performs asynchronous (batch) replication of block-level data from a volume on a primary HPE MSA, to a secondary volume on the same system or on a separate, independent HPE MSA. These two HPE MSA Storage systems connect through switches and are on the same fabric or network.

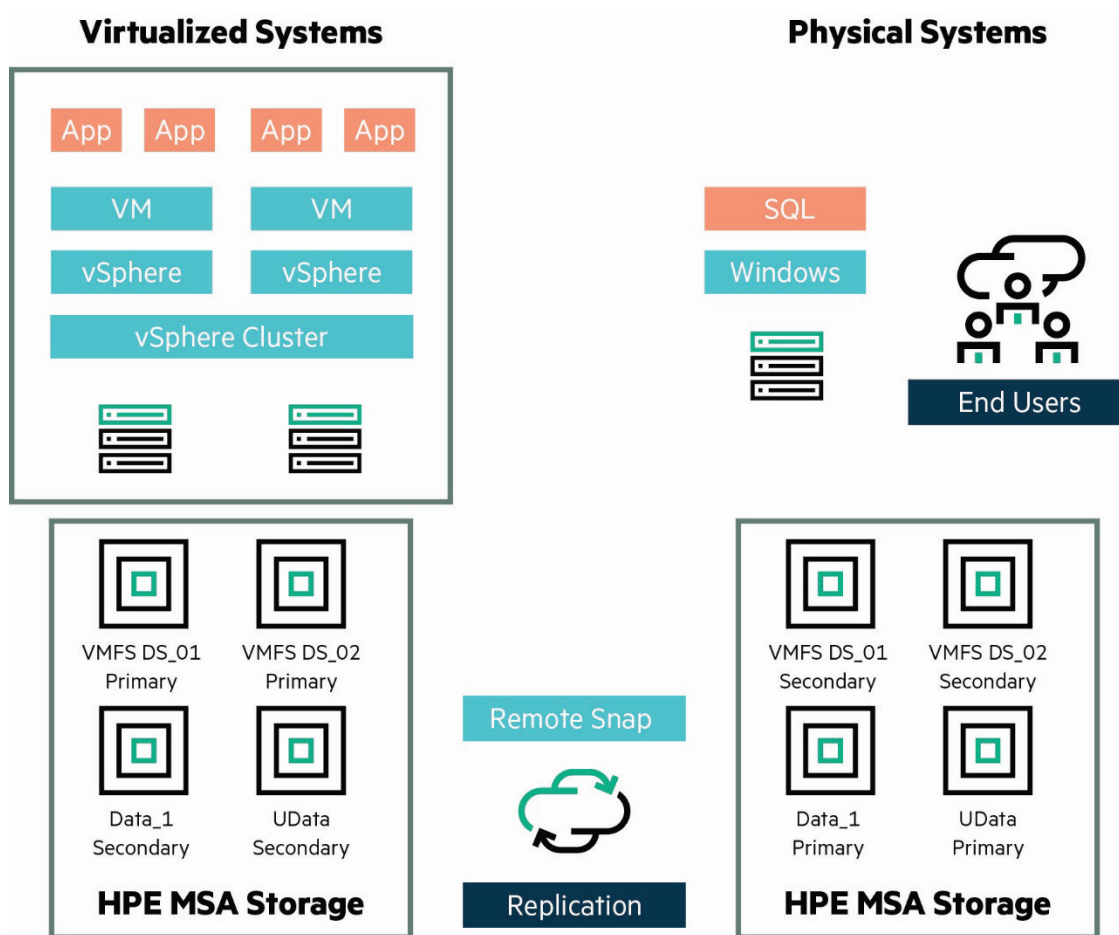


Figure 31. Data Protection with HPE MSA Remote Snap

Multi-Site Disaster Recovery

The VMware Site Recovery Manager solution provides virtual data center site disaster recovery. When implemented with HPE MSA Storage arrays for data redundancy, SRM becomes a valuable asset to a business with production demands that require minimal downtime.

VMware vCenter SRM uses the HPE MSA SRA to test DR scenarios and create different recovery plans.

HPE MSA Remote Snapshots must be created on the HPE MSA array before implementing VMware SRM. The HPE MSA SRA uses the Remote Snap feature provided from the HPE MSA. This combination provides the foundation for the SRM solution.

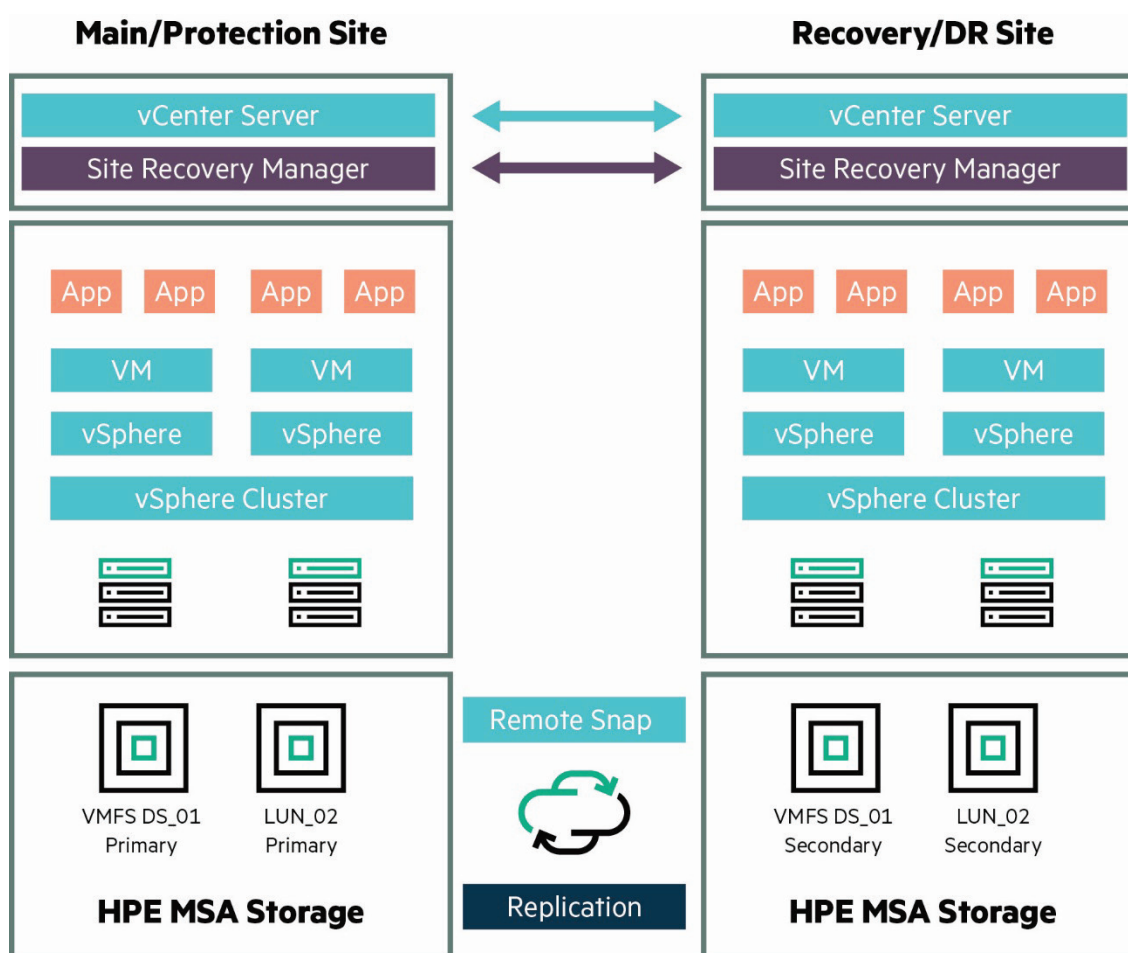


Figure 32. Multi-Site Disaster Recovery

The HPE MSA Storage Replication Adapter requires each Storage Array to have a licensed feature for Replication and Snapshot.

Summary and benefits

The HPE MSA Storage Array is an exceptional bargain for any businesses. It's rich feature set rivals enterprise storage array offerings making it very attractive as an entry-level storage solution; yet, it has the capacity to scale to meet any business needs of a growing virtualized server environment. The HPE MSA's Fibre Channel, iSCSI, and SAS connectivity options with Round Robin based Multipath policies makes the HPE MSA Storage Array ready to support test and development environments to full production vSphere solution. The array's Storage Pool, Disk Groups, and Virtual Volume ULP complement the VMware features of vCenter and the vSphere environment. All of these features combined with the HPE MSA's array-based snapshots and replication services create a very versatile storage array for VMware vSphere environments.

When following these best practices and configuration recommendations, the HPE MSA Storage will become an integral part of any IT Solution because it makes storage administration easy.

Appendix A

HPE MSA Firmware Update

When updating the firmware for the HPE MSA Storage Array family of products, the following recommendations should be considered:

- As with any other firmware upgrade, verify that a full backup has been performed before the upgrade.
- Before upgrading the firmware, make sure that the storage system configuration is stable and is not being reconfigured or changed in any way. If any configuration changes are in progress, monitor them using the SMU or CLI and wait until they are completed before proceeding with the upgrade.
- Do not power cycle or restart devices during a firmware update. If the update is interrupted or there is a power failure, the module could become inoperative. Should this happen, contact HPE customer support.
- After the device firmware update process is completed, confirm that the new firmware version is displayed correctly using one of the HPE MSA management interfaces—SMU GUI, MSA CLI, and so on.

HPE MSA Array Controller Firmware Update

- The array controller (or I/O module) firmware can be updated in an online mode only in redundant controller systems.
- When planning for a firmware upgrade, schedule an appropriate time to perform an online upgrade.
- For single domain systems, I/O must be halted.
- For dual domain systems, because the online firmware upgrade is performed while host I/Os are being processed, I/O load can impact the upgrade process. Select a period of low I/O activity to ensure the upgrade completes as quickly as possible and avoids disruptions to hosts and applications due to timeouts.
- When planning for a firmware upgrade, allow sufficient time for the update.
- In single-controller systems, it takes approximately 10 minutes for the firmware to load and for the automatic controller restart to complete.
- In dual-controller systems, the second controller usually takes an additional 20 minutes, but may take as long as one hour.
- When reverting to a previous version of the firmware, verify that the Management Controller (MC) Ethernet connection of each storage controller is available and accessible before starting the downgrade.
- When using a Smart Component firmware package, the Smart Component process automatically first disables Partner Firmware Update (PFU) and then performs a downgrade on each of the controllers separately (one after the other) through the Ethernet ports.
- When using a Binary firmware package, first disable the PFU option and then downgrade the firmware on each of the controller separately (one after the other).
- When performing firmware updates to HPE MSA70 drive enclosures, perform a power cycle on each enclosure.

Disk Drive Firmware Update

- Disk drive upgrades on the HPE P2000/MSA2000 storage systems are an off-line process. All host and array I/O must be stopped before the upgrade.
- If the drive is in a virtual disk, verify that it is not being initialized, expanded, reconstructed, verified, or scrubbed. If any of these tasks are in progress, before performing the update, wait for the task to complete or terminate it. Also verify that background scrub is disabled so that it doesn't start. The SMU or CLI interfaces can be used to check that the background scrub is disabled. If using a firmware smart component, it would fail and report if any of the above prerequisites are not being met.
- Disk drives of the same model in the storage system must have the same firmware revision. If using a firmware smart component, the installer would ensure that all the drives are updated.

Appendix B

Array/Volume Replication and Remote Snap

- A good rule of thumb is to make the secondary or replication target volume larger than the source volume being replicated (for example: by a measure of a third larger).
- While setting up the master volumes, note the size of the disk group and the primary/secondary volume. After they are part of a replication set, the sizes of the primary/secondary volume cannot be changed.
- Limit of one master volume per each disk group. Only make a master volume for replication and snapshot use.
- When replicating, put both the volume and snap-pool on the same disk group, but try to keep any other volumes off that disk group.
- Use a dual-controller array to try to avoid a failure of one controller. If one controller fails, replication continues through the second controller.
- Create no more than four volumes (master volume or snap pools) on a single disk group when used for snapshots or replication.
- To ensure that replication schedules are successful, schedule no more than three volumes to start replicating simultaneously.
- Users can set the replication image retention count to a preferred value. A best practice is to set the count to a maximum value of 32. Replication images can decrease the volume count per disk group. We recommend monitoring the number of replication images created.

Appendix C

Disk Background Scrubbing

The system continuously analyzes disks in Disk groups to detect, report, and store information about disk defects. You can enable or disable this utility.

Disk group-level errors reported include:

Hard errors, medium errors, and bad block replacements (BBRs). Disk-level errors reported include:

Metadata read errors, SMART events during scrub, bad blocks during scrub, and new disk defects during scrub. For RAID 3, 5, 6, and 50, the utility checks all parity blocks to find data-parity mismatches.

For RAID 1 and 10, the utility compares the primary and secondary disks to find data inconsistencies. For NRAID (Non-RAID, non-striped) and RAID 0, the utility checks for media errors.

A Disk group can be used while it is being scrubbed because the background Disk group scrub runs with a low priority, reducing activity if CPU usage increases above a certain percentage, or if I/O is occurring on the Disk group being scrubbed. A Disk group scrub may be in process on multiple Disk groups at once.

A new Disk group will first be scrubbed 20 minutes after creation and again after the interval specified by the Disk group Scrub Interval (hours) option. When a scrub is complete, the number of errors found is reported with event code 207 in the event log.

Utility priority levels

- High: Use when your highest priority is to get the system back to a fully fault-tolerant state. This causes heavy I/O with the host to be slower than normal and is the default.
- Medium: Use when you want to balance data streaming with data redundancy.
- Low: Use when streaming data without interruption, (for example a web media server), is more important than data redundancy. This enables operations like Reconstruct to run at a slower rate with minimal effect on host I/O.
- Background: Utilities run only when the processor has idle cycles.



Best practice: Leave the default setting of Background Scrub ON in the background priority for both Disk groups and available disks.

Scheduling drive spin down for all disk groups

For all Disk groups that are configured to use drive spin down (DSD), configure times to suspend and resume DSD so that Disk groups remain spun-up during hours of frequent activity. Also, configure DSD for available disks and global spares.

Best Practice: Set DSD for un-configured disks, spares, and configured Disk groups that do not perform a read/write operation at least once every 24 hours.

Important

Depending on which types of hard drives are in the HPE MSA array, using background disk scrub can have an impact on performance.

Appendix D

vSphere SAN troubleshooting

Volume Mapping

If the vSphere host does not see the LUN, initiate a rescan of the vSphere host's storage adapters. Then check for zone configurations, fabric switches, or fibre cables for any damaged or non-functional 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 WWN on ESX hosts

When creating hosts and host groups on the HPE MSA it is necessary to identify the WWNs associated with the initiators. Locating the WWNs in vCenter is found by selecting the host configuration for storage adapters as seen in the screen shot below.

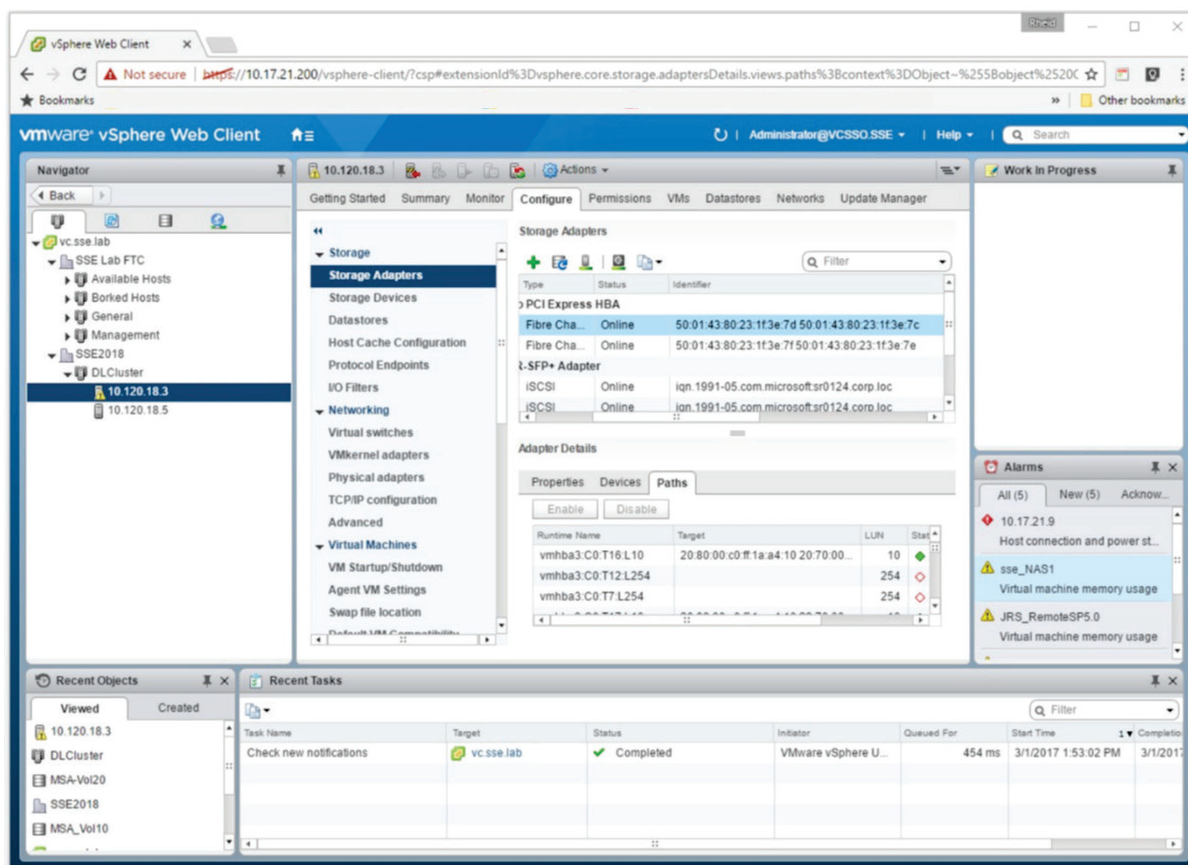


Figure 33. VMware vCenter ESX Host Fibre Channel Adapter Configuration Screen

If vCenter is not being used, finding WWNs in the vSphere ESX Web Client is not very easy. The following command can help find these WWN from the ESX console.

To identify the WWN for the HBAs in the ESX Servers use the following vSphere commands.

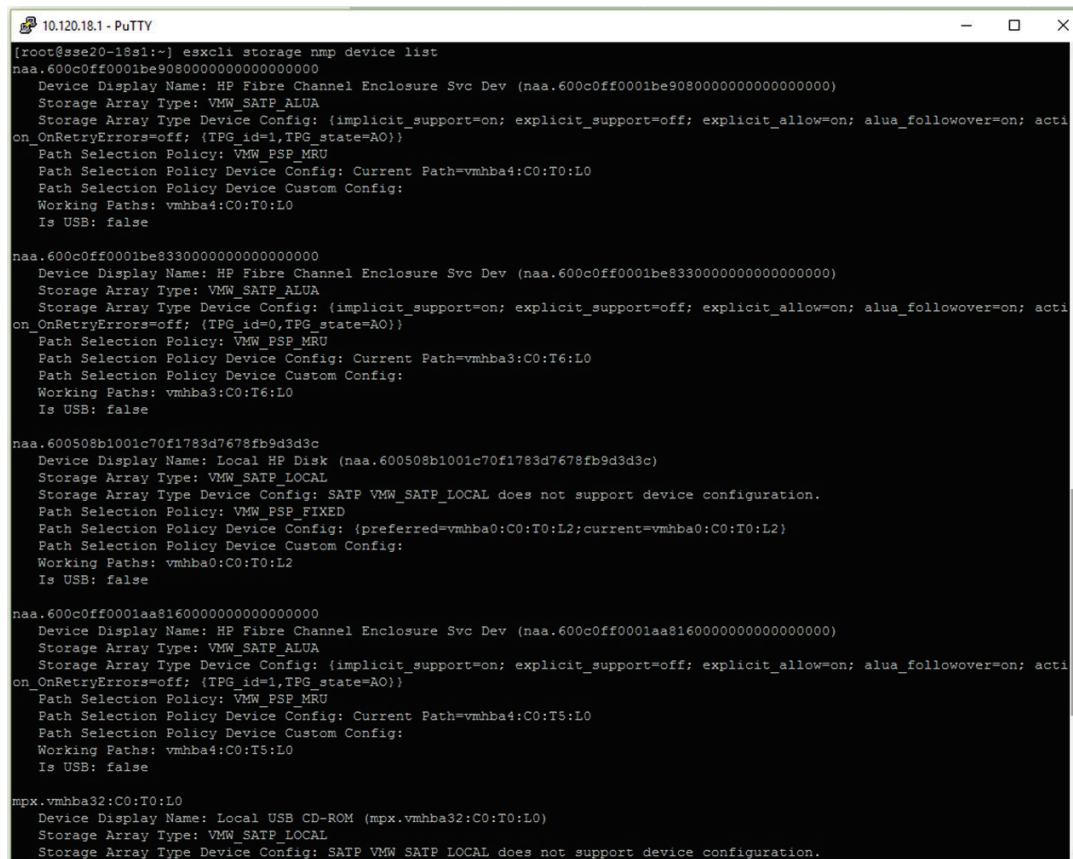
```
# esxcli storage core adapter list
```

```
[root@sse20-18s3:~] esxcli storage core adapter list
HBA Name  Driver      Link State  UID                               Capabilities
Description
-----
vmhba0    elxiscsi    online     iscsi.vmhba0                     (0000:04:00.2) Emulex Corporation HP FlexFabric 10Gb 2-port 556FLR-SFP+ Adapter
vmhba1    elxiscsi    online     iscsi.vmhba1                     (0000:04:00.3) Emulex Corporation HP FlexFabric 10Gb 2-port 556FLR-SFP+ Adapter
vmhba2    nhpsa      link-n/a   sas.500143803646d100             (0000:03:00.0) Hewlett-Packard Company Smart Array P440ar
vmhba3    qlnativefc link-up    fc.50014380231f3e7d:50014380231f3e7c Data Integrity, Second Level Lun ID
(0000:88:00.0) QLogic Corp 2600 Series 16Gb Fibre Channel to PCI Express HBA
vmhba4    qlnativefc link-up    fc.50014380231f3e7f:50014380231f3e7e Data Integrity, Second Level Lun ID
(0000:88:00.1) QLogic Corp 2600 Series 16Gb Fibre Channel to PCI Express HBA
[root@sse20-18s3:~]
```

Figure 34. HPE MSA Console Command for Identifying WWN

Or

```
# esxcli storage san fc list
```



```
10.120.18.1 - PuTTY
[root@sse20-18s1:~] esxcli storage nmp device list
naa.600c0ff0001be9080000000000000000
Device Display Name: HP Fibre Channel Enclosure Svc Dev (naa.600c0ff0001be9080000000000000000)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: (implicit_support=on; explicit_support=off; explicit_allow=on; alua_follower=on; acti
on_OnRetryErrors=off; (TPG_id=1,TPG_state=AO))
Path Selection Policy: VMW_PSP_MRU
Path Selection Policy Device Config: Current Path=vmhba4:C0:T0:L0
Path Selection Policy Device Custom Config:
Working Paths: vmhba4:C0:T0:L0
Is USB: false

naa.600c0ff0001be8330000000000000000
Device Display Name: HP Fibre Channel Enclosure Svc Dev (naa.600c0ff0001be8330000000000000000)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: (implicit_support=on; explicit_support=off; explicit_allow=on; alua_follower=on; acti
on_OnRetryErrors=off; (TPG_id=0,TPG_state=AO))
Path Selection Policy: VMW_PSP_MRU
Path Selection Policy Device Config: Current Path=vmhba3:C0:T6:L0
Path Selection Policy Device Custom Config:
Working Paths: vmhba3:C0:T6:L0
Is USB: false

naa.600508b1001c70f1783d7678fb9d3d3c
Device Display Name: Local HP Disk (naa.600508b1001c70f1783d7678fb9d3d3c)
Storage Array Type: VMW_SATP_LOCAL
Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
Path Selection Policy: VMW_PSP_FIXED
Path Selection Policy Device Config: (preferred=vmhba0:C0:T0:L2;current=vmhba0:C0:T0:L2)
Path Selection Policy Device Custom Config:
Working Paths: vmhba0:C0:T0:L2
Is USB: false

naa.600c0ff0001aa8160000000000000000
Device Display Name: HP Fibre Channel Enclosure Svc Dev (naa.600c0ff0001aa8160000000000000000)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: (implicit_support=on; explicit_support=off; explicit_allow=on; alua_follower=on; acti
on_OnRetryErrors=off; (TPG_id=1,TPG_state=AO))
Path Selection Policy: VMW_PSP_MRU
Path Selection Policy Device Config: Current Path=vmhba4:C0:T5:L0
Path Selection Policy Device Custom Config:
Working Paths: vmhba4:C0:T5:L0
Is USB: false

mpx.vmhba32:C0:T0:L0
Device Display Name: Local USB CD-ROM (mpx.vmhba32:C0:T0:L0)
Storage Array Type: VMW_SATP_LOCAL
Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
```

Figure 35. HPE MSA Console Command for Identifying Storage Devices

Boot from SAN

If booting from SAN, use the follow steps on the ESX 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
```

To check the storage device properties:

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

```
# esxcli storage core
```

Appendix E

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. 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.

The HPE MSA Storage supports the following VAAI primitives, which were first introduced in vSphere 4.1:

- **Hardware-Assisted Locking:** Also known as Atomic Test & Set (ATS). As illustrated in Figure 36, 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 VMDK, 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 having the host wait for the operation to complete, the array is able to signal that the operation has completed right away, handling the process on its own without involving the vSphere host.

The HPE MSA Storage Array natively supports these three primitives, which means an additional HPE MSA plug-in is no longer required with vSphere 5.0 and higher to receive the benefits of VAAI integration.

- VMware Native Multipathing (NMP) with HPE MSA firmware.
- HPE recommends VMW_PSP_RR for path failover policy.
- Refer to [VMware Compatibility Guide](#) for more information.

Remember to verify the version of HPE MSA firmware installed on the array; consult the [VMware Compatibility Guide](#) for detailed compatibility information regarding path failover, such as round-robin, and VAAI plug-in support for older versions of vSphere (hp_vaaip_p2000).

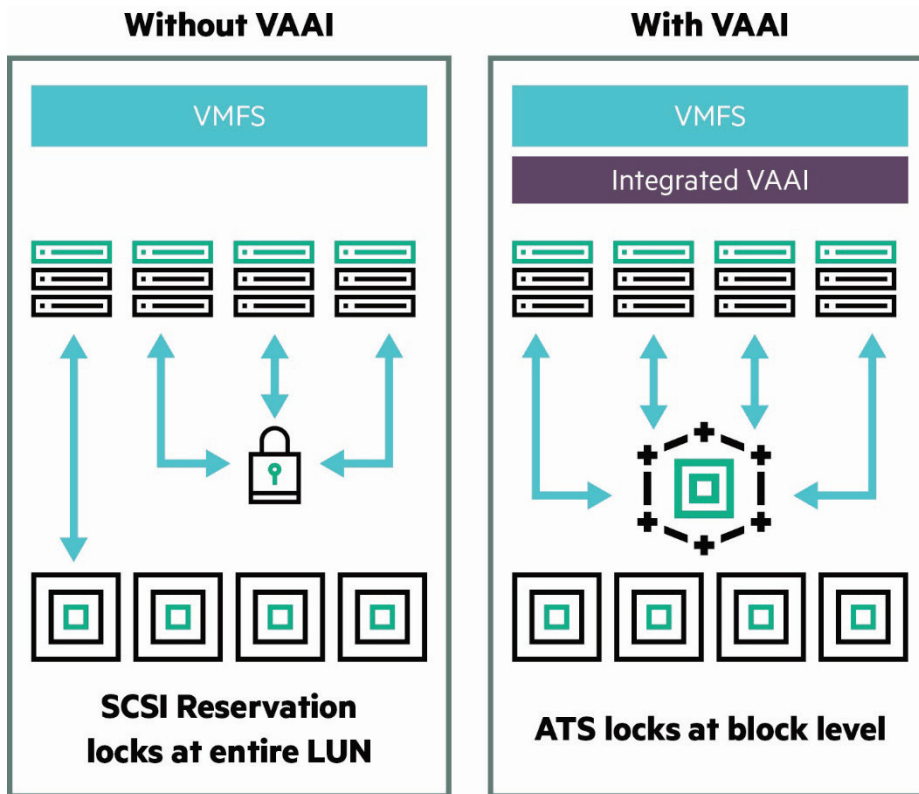


Figure 36. Comparison of Storage Array using VAAI and not using VAAI

VAAI benefits and use cases

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

The following use cases address the three VAAI features.

Use case 1: VM Migration

Using 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 quicker migration of VMs between clustered hosts.

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, users can increase the number of VMs for per server (also known as VM Density).

VMware Storage I/O Control

With vSphere 4.1, VMware introduced a new feature called Storage I/O control (SIOC), which enables you to perform attenuation of the I/O for each virtual disk you choose. The SIOC feature in vSphere 5.x is disabled by default.

Storage I/O Control provides I/O prioritization of virtual machines running on a cluster of vSphere hosts that access a shared storage pool. 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. For more information, refer to the technical white paper: vmware.com/files/pdf/techpaper/VMW-Whats-New-vSphere41-Storage.pdf

VMware Storage I/O Control and the HPE MSA Storage system combine to provide a more performance-optimized storage solution. Enabling Storage I/O control 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. Storage I/O control is not dependent on the array; it is more of a VMware vSphere infrastructure solution.

ALUA compliance in vSphere 5.x and the support for round robin load balancing are significant improvements for ESX 4.0 multipathing. These two simple features have eliminated all of the intricate configuration steps administrators carried out with ESX 3.x and older versions. These improvements also help to guarantee a much more balanced system configuration than administrators could achieve through manual preferred path configuration. With round robin I/O path policy, I/O can be queued to multiple controller ports on the G3, HPE MSA Storage, providing an instant performance boost.

VMware introduced a new multipathing framework with ESX 4.0. The components that comprise the framework are:

- Native Multi-Pathing (NMP)
- Storage Array Type Plugin (SATP)
- Path Selection Plugin (PSP)

Now ALUA-complaint, vSphere (ESXi) 5.x does not require the same intricate configuration process as with older vSphere generations. Historically, administrator tasks include:

- Configuring the HPE MSA 2000 volume and selecting the controller access policy
- Powering on or rebooting the ESX 5.x server or performing a rescan

The new HPE MSA 2050/2052 storage systems simplify the process. At boot up or after a rescan, vSphere 5.x detects the optimal access paths automatically. As long as MRU and/or round robin is the employed I/O path policy, vSphere 5.x will give these paths to the volume or LUN with the higher preference for I/O queuing.

Multipath

With server virtualization environments using vSphere 5.x, HPE recommends using the Round Robin (VMW_PSP_RR) failover policy with all HPE MSA Storage arrays. For optimal default system performance with HPE MSA Storage, configure the round robin load balancing selection to IOPS with a value of 1 for every LUN.

```
esxcli nmp round robin setconfig --type "iops" --iops 1 --device naa.xxxxxxxx
```

For environments with only HPE MSA Storage LUNs connected to vSphere 5.x hosts, the following simple script can be used to iterate through all LUNs and automatically set their I/O path access policy to round robin.

```
for i in `ls /vmfs/devices/disks/ | grep naa.600` ; do esxcli nmp roundrobin setconfig --type "iops" -  
-iops=1--device $i; done
```

For environments that have other array models in addition to HPE MSA arrays attached to vSphere 5.x hosts using the VMW_PSP_RR, change "grep naa.600" so that it matches the pattern to devices on the desired arrays only.

Remember to verify the version of HPE MSA firmware installed on the array; consult the [VMware Compatibility Guide](#) for detailed compatibility information regarding path failover, such as round-robin, and VAAI plug-in support for older versions of vSphere.

Appendix F

Changing the LUN RESPONSE with SMU v2

For backward compatibility purposes, the version 2 of the web based SMU is still available and can be used to change the LUN RESPONSE setting; however, it is not recommended. For administrators still using the v2 version of the SMU, the following steps identify how to set the HPE MSA Storage LUN RESPONSE setting.

1. Shut down any host accessing the HPE MSA.
2. Log into either controller A or B using the SMU using the v2 http path.
Ie. `https://10.120.18.221/v2/index.html`
3. In the **View** panel, right-click the system and select **Configuration > Advanced Settings > Cache**.
4. Change the LUN RESPONSE from **NOT READY** to **ILLEGAL REQUEST**.
5. Click **Apply**.
6. Restart hosts.

Appendix G

Terminology

Disk Group: A Disk Group is a collection of disks in a given redundancy mode (RAID 1, 5, 6, or 10 for Virtual Disk Groups and NRAID and RAID 0, 1, 3, 5, 6, 10, or 50 for Linear Disk Groups). A Disk Group is equivalent to a Disk group in Linear Storage and utilizes the same proven fault tolerant technology used by Linear Storage. Disk Group RAID level and size can be created based on performance and/or capacity requirements. With GL200 and greater firmware multiple Virtual Disk Groups can be allocated into a Storage Pool for use with the Virtual Storage features; while Linear Disk Groups are also in Storage Pools, there is a one-to-one correlation between Linear Disk Groups and their associated Storage Pools.

Initiator Port: A port is a connection from a device into the SAN. Each node in the SAN—host, storage device, and fabric component (router or switch)—has one or more ports that connect it to the SAN. Port identification can be by WWN, WWPN, or IQN.

iQN: iSCSI Qualified Name is a globally unique identifier for an initiator or a target node (not ports). It is UTF-8 encoding with human-readable format of up to 233 bytes. This address is not used for routing. The extended version is called the Extended Unique Identifier (EUI). In-depth information on SAN ports can be found at the website of the Storage Networking Industry Association, snia.org.

Linear Storage: Linear Storage is the traditional storage that has been used for the four HPE MSA generations. With Linear Storage, the user specifies which drives make up a RAID Group and all storage is fully allocated.

Logical Unit Number (LUN): The HPE MSA 1040/2040/2042/2050/2052 arrays support 512 volumes and up to 512 snapshots in a system. All of these volumes can be mapped to LUNs. Maximum LUN sizes are up to 128 TB and the LUNs sizes are dependent on the storage architecture: Linear vs. Virtualized. Thin Provisioning allows the user to create the LUNs independent of the physical storage.

Page: An individual block of data residing on a physical disk. For Virtual Storage, the page size is 4 MB. A page is the smallest unit of data that can be allocated, de-allocated, or moved between virtual disk groups in a tier or between tiers.

Read Cache: Read Cache is an extension of the controller cache. Read Cache allows a lower cost way to get performance improvements from SSD drives.

Sub-LUN Tiering: Sub-LUN Tiering is a technology that allows for the automatic movement of data between storage tiers based on access trends. In the HPE MSA 1040/2040/2042/2050/2052, Sub-LUN Tiering places data in a LUN that is accessed frequently in higher performing media while data that is infrequently accessed is placed in slower media.

Storage Pools: The GL200 firmware and greater introduces Storage Pools which are comprised of one or more Virtual Disk Groups or one Linear Disk Group. For Virtual Storage, LUNs are no longer restricted to a single disk group as with Linear Storage. A volume's data on a given LUN can now span all disk drives in a pool. When capacity is added to a system, users will benefit from the performance of all spindles in that pool.

When leveraging Virtual Storage Pools, the HPE MSA 1040/2040/2042/2050/2052 supports large, flexible volumes with sizes up to 128 TB and facilitates seamless capacity expansion. As volumes are expanded data automatically reflows to balance capacity utilization on all drives.

Thin Provisioning: Thin Provisioning allows storage allocation of physical storage resources only when they are consumed by an application. Thin Provisioning also allows over-provisioning of physical storage pool resources allowing ease of growth for volumes without predicting storage capacity upfront.

Thick Provisioning: All storage is fully allocated with Thick Provisioning. Linear Storage always uses Thick Provisioning.

Tiers: Disk tiers are comprised of aggregating 1 or more Disk Groups of similar physical disks. The HPE MSA 1040/2040/2042/2050/2052 support 3 distinct tiers:

1. A Performance tier with SSDs
2. A Standard SAS tier with Enterprise SAS HDDs
3. An Archive tier utilizing Midline SAS HDDs

Prior to GL200 firmware, the HPE MSA 1040/2040 operated through manual tiering, where LUN level tiers are manually created and managed by using dedicated Disk groups and volumes. LUN level tiering requires careful planning such that applications requiring the highest performance be placed on Disk groups utilizing high performance SSDs. Applications with lower performance requirements can be placed on Disk groups comprised of Enterprise SAS or Midline SAS HDDs. Beginning with GL200 and greater firmware, the HPE MSA 1040/2040/2042/2050/2052 supports Sub-LUN Tiering and automated data movement between tiers.

The HPE MSA 1040/2040/2042/2050/2052 automated tiering engine moves data between available tiers based on the access characteristics of that data. Frequently accessed data contained in pages will migrate to the highest available tier delivering maximum I/O's to the application. Similarly, "cold" or infrequently accessed data is moved to lower performance tiers. Data is migrated between tiers automatically such that I/O's are optimized in real time.

The Archive and Standard Tiers are provided at no charge on the HPE MSA 2040 platform beginning with GL200 and greater firmware. The HPE MSA 1040 requires a license when using the Archive and Standard Tiers. A Performance Tier utilizing a fault tolerant SSD Disk Group is a paid feature that requires a license for both the HPE MSA 1040 and HPE MSA 2040. Without the Performance Tier license installed, SSDs can still be used as Read Cache with the Sub-LUN Tiering feature. Sub-LUN Tiering from SAS MDL (Archive Tier) to Enterprise SAS (Standard Tier) drives is provided at no charge for the HPE MSA 2040.

Note

The HPE MSA 1040 requires a license to enable Sub-LUN Tiering and other Virtual Storage features such as Thin Provisioning.

Virtual Disk Group: The vDisk nomenclature is being replaced in version 3 of the SMU with Disk Group. For Linear Storage in the Storage Management Utility (SMU) Version 2 you will still see references to vDisks for Virtual Storage. In the Version 3 SMU, Disk Groups are now part of Storage Pools which allow for greater flexibility of the physical drives. Disk groups and vDisks are essentially the same and have the similar administration functions. On differences is Disk Groups for Linear disks have additional RAID types; NRAID, RAID 0, but are only available in the CLI.

Virtual Storage: Virtual Storage is an extension of Linear Storage where data is virtualized not only across a single disk group, as in the linear implementation, but also across multiple disk groups with different performance capabilities and use cases.

Virtual Volume: Virtual Volumes are equivalent to traditional LUNs but are not limited to the imposed barriers of physical disk grouping. Virtual Volumes are dynamic storage volumes visible to hosts. These abstract storage volumes allow flexibility of the Virtual Storage pool allowing for efficient use of the physical disks.

WWN: The World Wide Node Name is a globally unique identifier for a Fibre Channel HBA. Each FC HBA can have multiple ports, each with its own unique WWPN.

WWPN: This World Wide Port Name is a globally unique identifier for a port on an FC HBA. The FC switches discover the WWPN of a device or host and assign a port address to the device.

For more information

HPE resources

HPE Storage arrays: hpe.com/storage

HPE MSA home page: hpe.com/storage/msa

HPE MSA 2042 QuickSpecs: hpe.com/h20195/v2/GetDocument.aspx?docname=c05183822

HPE MSA 2052 QuickSpecs: hpe.com/h20195/V2/Getdocument.aspx?docname=a00008277enw

HPE Storage Networking: hpe.com/storage/san

HPE Servers: hpe.com/servers

For complete software requirements and compatibility lists, refer to the HPE SPOCK website: hpe.com/storage/spock

VMware resources

VMware vSphere Storage/SAN guide

VMware Compatibility Guide—Storage/SAN


VMware Storage Solutions from HPE

HPE Virtualization with VMware: hpe.com/storage/vmware

HPE OneView for VMware vCenter: h20392.www2.hpe.com/portal/swdepot/displayProductInfo.do%3FproductNumber%3DHPVPR

Learn more at

hpe.com/storage



Sign up for updates



© Copyright 2013, 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

VMware, VMware vSphere, VMware vSphere 5.1, VMware vSphere VMFS, VMware ESX, VMware vCenter, VMware vCenter Server, VMware vSphere Storage vMotion, and VMware ESXi are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Windows is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. All other third-party trademark(s) is/are property of their respective owner(s).

4AA4-7060ENW, August 2017, Rev. 2