

HPE MSA Considerations and Best Practices for vSphere Setup and Installation

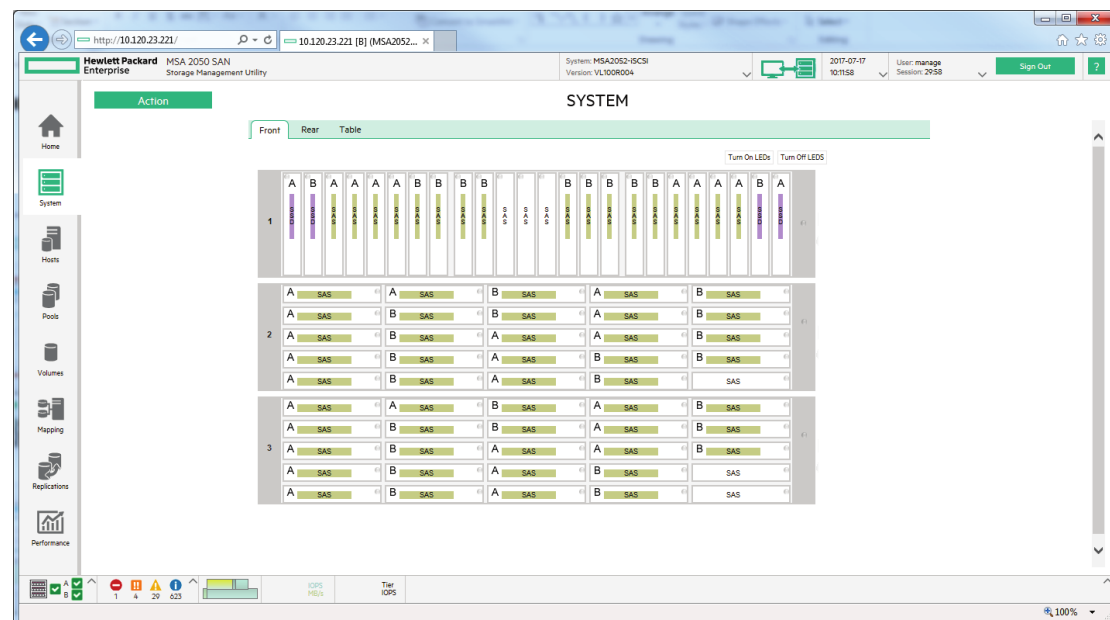
Storage plays a critical role in the success of VMware vSphere deployments. The following section highlights recommended practices for setup and configuration of the HPE MSA Storage Array best suited for virtualization environments running VMware vSphere.

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.

Storage configuration

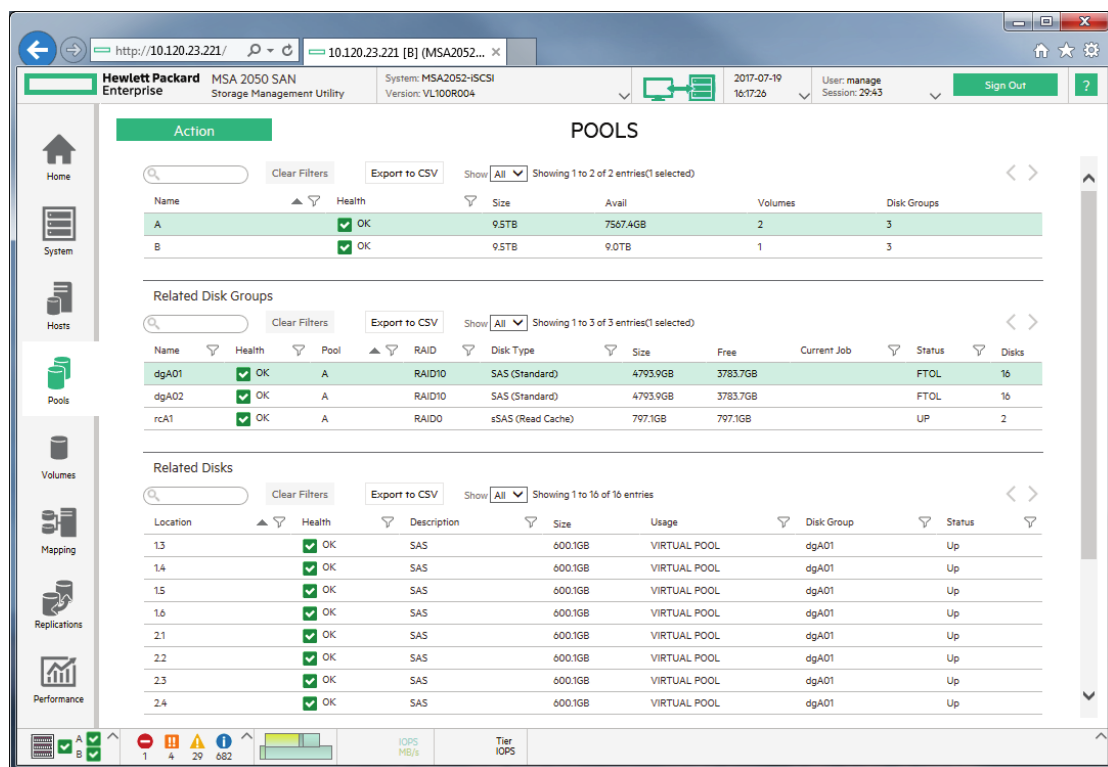
The storage demands of a typical virtualized environment using VMware vSphere are not I/O intensive once a virtual machine is up and running. There are exceptions; however, generally the user wants the storage to be performant upon startup of VMs, generally fast for interaction, and then smart enough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



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.



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 SAS 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.

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 article).

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.

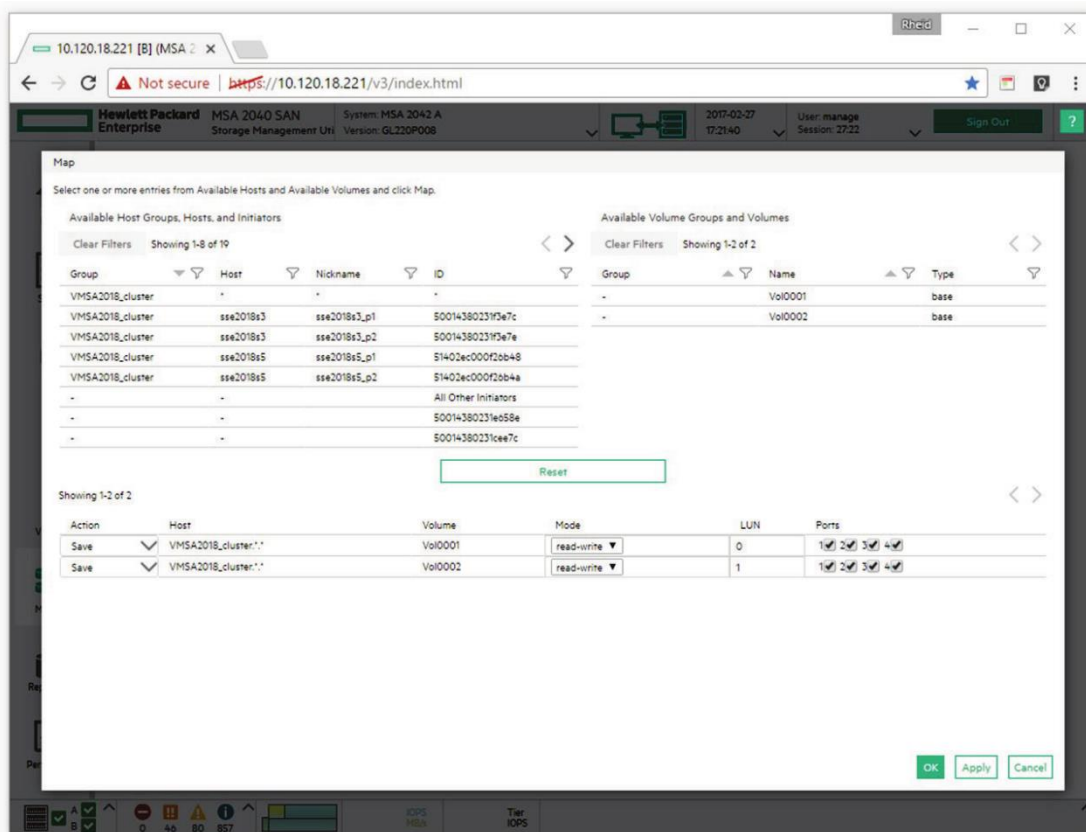
VMware recommends the following practices for volume mapping:

- 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.

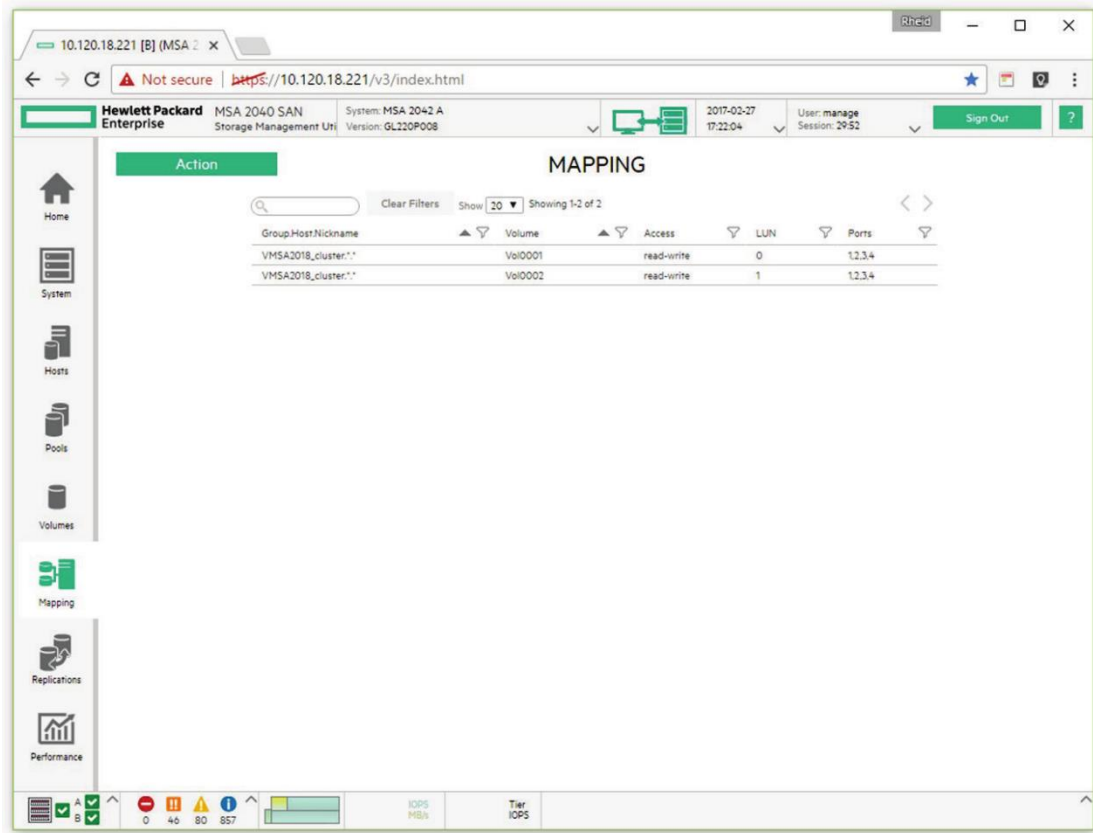
With version 3 of the SMU software, administrator can use aliases for initiators, hosts, and host groups to simplify the mapping and management of volumes. If the new features are not used, volumes must be individually mapped to the WWN or IDQ of each ESX server’s interface and path to the array. Each 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).



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.



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.

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.

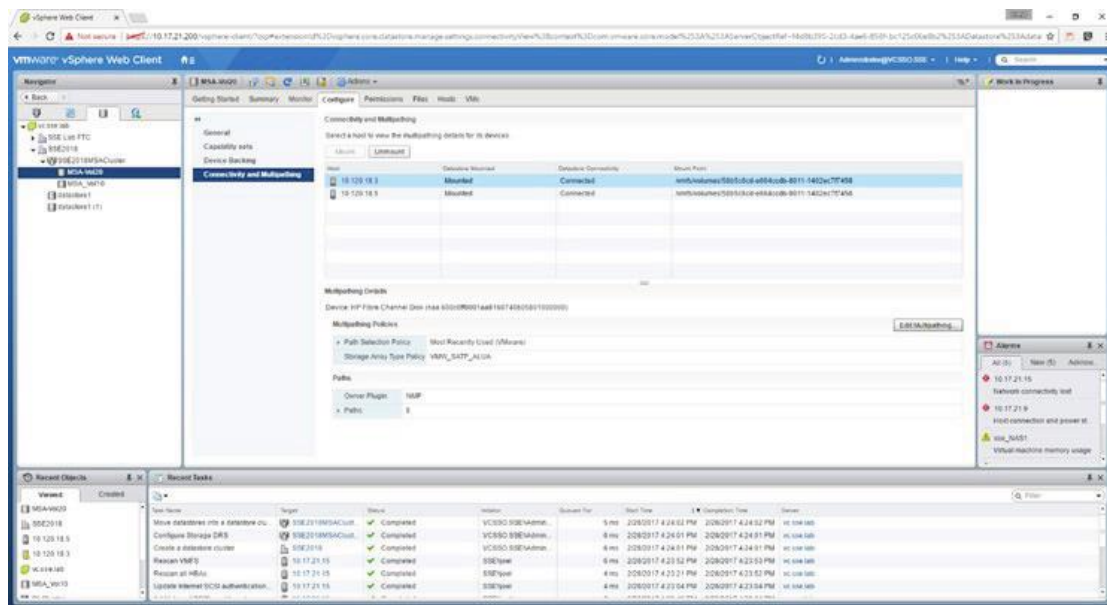
As recommended in the previous section, HPE recommends configuring the HPE MSA Storage to use a Host Group for a vSphere cluster and use the cluster object when mapping Virtual Volumes. This will create all the mappings to all the adapters to support multipath in the cluster in one step.

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.”

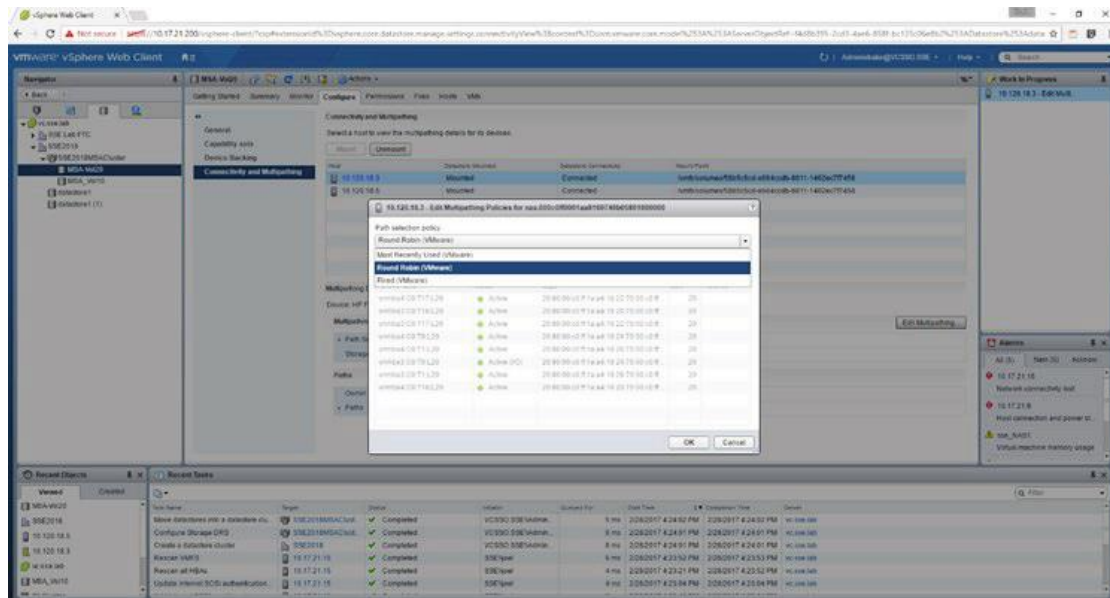
HPE recommends using the “Round Robin” preferred selection path (PSP) policy for best performance and load balancing on the HPE MSA Storage.

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. 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.



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.



Additional Information

[HPE MSA 1040 Storage Data Sheet](#)

[HPE MSA 2040 Storage Data Sheet](#)

[HPE MSA 2050 Storage Data Sheet](#)

[HPE MSA 2052 Storage Data Sheet](#)

Where to buy

To order HPE MSA storages, you can visit: [HPE MSA Storages Models List](#)

Contact us: +1-626-239-8066 (USA) / +852-3050-1066; sales@router-switch.com

About Us

Router-switch.com, founded in 2002, is one of the biggest Global Network Hardware Supplier. We are a leading provider of network products with 14,500+ customers in over 200 countries. We provide original new and used network equipments (Cisco, Huawei, HPE, Dell, Hikvision, Juniper, EMC, etc.), including Routers, Switches, Servers, Storage, Telepresence and Videoconferencing, Video surveillance, IP Phones, Firewalls, Wireless APs & Controllers, EHWIC/HWIC/VWIC Cards, SFPs, Memory & Flash, Hard Disk, Cables, and all kinds of network solutions related products.