



HP MSA 2040 Storage Configuration and Best Practices for VMware vSphere

Table of contents

Intended Audience.....	3
Executive summary	3
Product Overview	3
HP MSA 2040 Storage	3
Solution Overview	5
HP Storage Networking.....	6
HP MSA Storage Concepts and VMware vSphere.....	7
Virtual disk (vdisk)	7
Volumes	8
Default Mappings (Volumes).....	8
Unified LUN Presentation (ULP)	8
Unified LUN Presentation (ULP) Failover.....	10
RAID Level Considerations.....	11
VMware vSphere Storage Integration.....	12
VAAI Integration	12
VAAI Benefits and Use Cases.....	13
HP Insight Control Storage Module for VMware vCenter Server	13
Disaster Recovery with vCenter Site Recovery Manager and HP MSA 2040 Storage	14
Use Cases for vSphere and HP MSA 2040 Storage.....	15
Mixed Physical and Virtual Computing Environment	15
Multiple vSphere clusters on HP MSA Storage.....	16
ROBO deployments for distributed vSphere environments	17
Data Protection with HP MSA Remote Snap.....	18
Multi-Site Disaster Recovery	19

MSA Storage Installation for VMware vSphere	20
Additional VMware vSphere Storage Configuration Guides	20
Storage Network Connectivity	20
Cabling the MSA Storage Array	20
Best Practices: MSA Initial Configuration	21
Storage Administration	23
MSA Storage Best Practices.....	26
Maximize capacity and disk usage	26
Using multiple enclosures	26
Selecting the controller ownership for the vdisk	26
Storage configurations with several disks.....	26
RAID levels.....	26
Changing the default name of the volumes.....	26
Volume mapping	27
Naming hosts	27
Boot from SAN support/considerations.....	28
Multipath Considerations for vSphere.....	28
Changing LUN Response before LUN Creation	30
Additional VMware considerations.....	31
VMware Storage I/O Control	31
Setting HP MSA 2040 Storage MSA active-active arrays	32
Using vCenter Site Recovery Manager with HP MSA Remote Snap.....	33
Configuring VMware vCenter SRM with HP MSA 2040 SRA	34
Best Practices for MSA 2000 SRA and Remote Snap with VMware SRM.....	35
Summary and Benefits.....	35
Appendix A: Firmware Update Best Practices for 2040 MSA	36
P2000 MSA Array Controller Firmware Update Best Practices	36
Disk Drive Firmware Update Best Practices	36
Appendix B: Array/Volume Replication and Remote Snap Best Practices	37
Appendix C: HP MSA 2040 Disk Background Scrub.....	38
For more information	39
VMware.....	39

Intended Audience

Target audience: IT Administrators, VMware vSphere Administrators, and Solution Architects planning a server virtualization deployment with HP MSA 2040 storage. This and other documents pertaining to virtualization with HP and VMware are available at hp.com/go/vmware and at hp.com/go/storage.

If you are a VMware vSphere/ESXi administrator planning to set up hosts to work with Storage Area Networks (SANs), you should have a working knowledge of SAN concepts. You can find information about SANs in print and on the Internet. Because this industry changes constantly, check these resources frequently. If you are new to SAN technology, familiarize yourself with the basic terminology.

Executive summary

When supported with the correct underlying storage platform, server virtualization delivers increased infrastructure consolidation, administrative efficiency, business continuity, and cost savings. As a result, server virtualization is not only transforming the data center, but also the businesses that those data centers fuel. However, these transformative results depend on enterprise class storage to deliver the performance, availability, and flexibility to keep up with the dynamic and consolidated nature of virtualized server environments.

Storage is a critical component of the virtualization infrastructure. The storage platform must deliver high performance, use its capacity efficiently and scaling easily. Storage must also be modular and scalable to enable purchasing only what is needed, when it is needed. No single point of failure should exist, and failed components must be replaceable without interruption of service.

The latest generation of HP 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 HP MSA Storage with VMware vSphere. In addition, the paper covers the best practice settings and configurations needed to optimize HP MSA Storage for deployment with VMware vSphere.

When deployed together, HP 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.

Product Overview

Clustered, fault-tolerant, virtualization environments such as VMware vSphere rely heavily upon centrally-managed, scalable, storage area network (SAN) resources. The HP MSA 2040 Storage system provides a versatile entry-level SAN solution for vSphere host clusters. The following section provides an overview on the HP MSA 2040 Storage system.

HP MSA 2040 Storage

The HP MSA 2040 Storage array delivers cost-effective, consolidated storage solutions to meet departmental and small-to-mid-sized business (SMB) requirements. The HP MSA 2040 delivers 8Gb and 16 Gb Fiber Channel (FC) host connectivity and accommodates large form factor (LFF) and small form factor (SFF) SAS, SAS Midline (SAS MDL), and Solid State drives (SSD). The arrays ship standard with 64 Snapshots and Volume Copy enabled, which can be upgraded to 512 snaps.

The HP MSA 2000 array family helps customers with growing storage requirements across multiple servers to easily manage a versatile storage solution for their vSphere environment.

Figure 1. The HP MSA 2040 Storage array



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

Table 1. Common features of HP MSA models

Features	MSA family
Host protocols	8 Gb FC, 16 Gb FC, 1 GbE, 10 GbE, SAS
Max host ports	8 (4xFC, 4xFC+4x1 GbE, 8x1 GbE, 4x10 GbE, 8xSAS)
Cache support	4 GB standard (2 GB/controller)
Max LUNs	512
RAID supported	0, 1, 3, 5, 6, 10, 50
Rack/Physical	2U (for controllers and expansions)
Snapshots	64 Snapshots standard
Volume Copy	Bundled
vSphere Support	VAAI, VASA, SRM

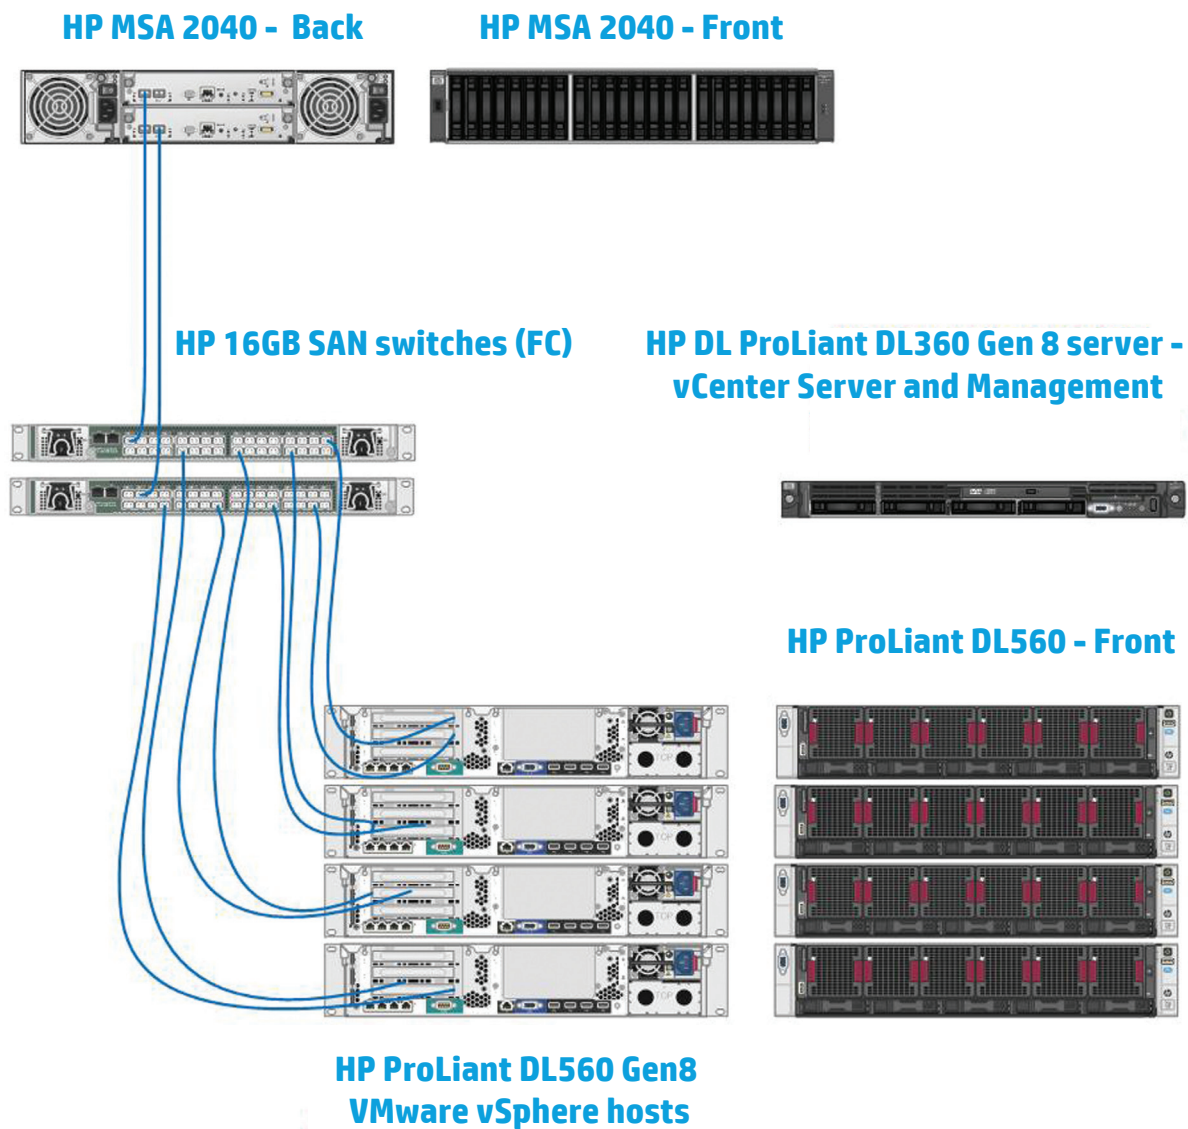
For more in-depth information on MSA storage system models, visit the [HP MSA Storage webpage](#).

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.

The following diagram (Figure 2) illustrates a recommended configuration for a fault-tolerant HP MSA 2040 Storage system supporting a VMware vSphere environment. In this configuration, two HP 16 GB FC switches are used in the storage fabric for redundancy. Two 16 GB HBA FC adapters in each HP DL560 server—deployed as VMware vSphere hosts—provide redundant paths to the HP 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 HP MSA Storage.

Figure 2. Example of a typical configuration with HP MSA 2040 Storage and HP ProLiant servers. Direct connection of host servers to FC host ports also possible, without the need for a FC switch



HP Storage Networking

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

A typical complement to an HP MSA Storage and VMware vSphere deployment is the HP SN6000B Fiber Channel Switch, which offers the features necessary to meet the demands of hyper-scale, private cloud storage environments by delivering market-leading 16 Gb/sec Fiber Channel technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the SN6000B Switch is configurable in 24, 36, or 48 ports and supports 4, 8, 10, or 16 Gb/sec Fiber 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.

HP 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.

HP 16Gb Host Bus Adapters (HBAs)

The HP SN1000E 16Gb single-port and dual-port Fiber Channel Host Bus Adapters are a step into the 16Gbit/sec environment with greater performance and advanced, embedded support for virtualized environments. The HP SN1000E 16Gb Fiber Channel HBA purchased today is backward compatible with 8Gb and 4Gb 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 HP SN1000E 16Gb FC HBA enables more rapid storage and retrieval of critical information. Designed for environments with greater virtual machine density and bandwidth requirements, the HP SN1000 16Gb 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 datacenter productivity by delivering twice the data throughput of 8Gb 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 Fiber Channel port, resulting in reduced cabling and a higher return on IT investment
- Includes PCIe 2.0 – Gen2 x8 lanes

For more information about HP Storage Networking solutions and products, visit hp.com/go/san.

HP MSA Storage Concepts and VMware vSphere

Growing storage needs for virtualized servers now require greater levels of storage performance and functionality at a lower cost of ownership. The new HP MSA 2040 Storage arrays are positioned to provide an excellent value for SMB customers needing increased performance to support initiatives such as consolidation and virtualization. The MSA 2040 delivers higher performance by leveraging new controller architecture with a new processor, four 8 Gb or 16 Gb FC ports, 4 GB cache per controller, easy Snap and Volume copy, and the latest drive technologies, such as SSDs.

Built for smaller-scale, affordable virtualization, HP MSA 2040 Storage delivers the following key advantages for a virtualization environment with VMware vSphere:

- High performance to meet the peak virtualization demands of cloning or migrating multiple VMs
- Cost effectiveness, allowing for affordable storage scaling
- Non-disruptive scalability to easily support data growth
- Integration with key VMware vSphere Storage APIs, including VAAI, VASA, and SRM
- Simplified management with easy-to-use storage administration
- Integration with the HP Insight Control Storage Module for VMware vCenter Server
- Support for large VMFS datastores
- Storage redundancy in a small footprint

HP MSA 2040 Storage delivers these advantages through a number of unique architectural features. Some of the virtualization benefits are detailed in the following sections, and for more in-depth information on product features and benefits, visit the [HP MSA Storage webpage](#).

Virtual disk (vdisk)

An HP MSA 2040 Storage array vdisk is the largest storage object within the array, combining the capacity of one or more physical disks. The maximum number of drives used in a RAID 1 vdisk is 2; RAID types 0, 3, 5, 6, and 10 uses 16; and RAID 50 uses 32 maximum drives.

When configuring a vdisk on HP MSA 2040 Storage for vSphere, an administrator must consider two factors:

- The application being virtualized
- The storage optimization objective

All disks in a vdisk must be the same type (SAS, MDL SAS, or SSD, small or large form-factor). Each controller can have a maximum of 16 vdisks.

A volume is a logical subdivision of a vdisk, mapped to controller host ports for access by vSphere hosts. The storage system presents only volumes, not vdisks, to the vSphere hosts, to be used as VMFS datastores or Raw Disk Mappings (RDM). A common misconception about server virtualization is that when an application is virtualized, its storage requirement can be reduced or changed; this is not the case.

With large storage configurations, consider creating fewer vdisks, each containing many drives, versus many vdisks containing few drives. For example, one 12-drive RAID 5 vdisk will have one parity drive and 11 data drives, compared to four 3-drive RAID 5 vdisk, each one having a parity drive and two data drives. Supporting large storage capacities requires advanced planning because it requires using large virtual disks with several volumes each or many virtual disks. Here is where performance and capacity come into effect with your planning. To maximize capacity, combine physical disks into a large vdisk, then subdivide that vdisk into several volumes with a capacity less than 2TB, because, as of VMware vSphere 5.1, VMFS has a 2TB limit on each "extent" of a datastore.

Volumes

The HP MSA 2040 can have a maximum of 128 volumes per vdisk. Only volumes map to vSphere host systems. A mapped volume provides the storage for your VMFS file system partitions.

Single-volume vdisks work well in VMware vSphere environments that need one large, fault-tolerant storage space. An example would be a large database accessed by users on a single host.

Multiple-volume vdisks work well when you have very large disks and you want to make the most efficient use of disk space for fault tolerance (parity and spares). However, I/O to multiple volumes containing vSphere VMFS datastores on the same vdisk can have an adverse effect on overall system performance at the vSphere cluster level.

Default Mappings (Volumes)

Each volume has default host-access settings when created; these settings are called default mappings. Default mapping allows all hosts connected to the controller host ports to access a newly created volume. Specifying host mappings during creation allows only those mapped hosts to access the volume.

The advantage of 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. Do not use default mapping for specialized volumes requiring restricted access. It is a recommended best practice to use explicit mapping of volumes for VMware vSphere hosts. When creating volumes, see the section on [RAID Level Considerations](#).

Unified LUN Presentation (ULP)

The HP MSA storage supports Unified LUN Presentation (ULP), which enables the array to present all LUNs to all host ports, eliminating the need for the interconnect path between the storage controller units. ULP presents the same World-Wide Node Name (WWNN) for both controllers. There is no duplicate LUN allowed between controllers, and either controller can use any unused logical unit number. As noted in the Asymmetric Logical Unit Access (ALUA) specifications, the preferred path, which is identified by the reported target port groups (RTPG), indicates the owning controller and gives it the optimal path to the array. The owning controller always performs the I/O to disk.

ULP appears to the host as an active-active storage system where the host can choose any available path to access a LUN regardless of vdisk ownership.

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

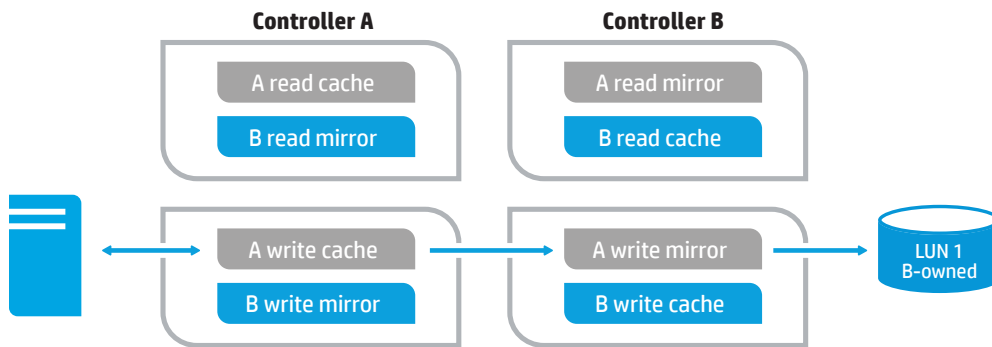
The ability to identify and alter the LUN controller ownership is defined by the ALUA extensions in the SPC3 standard. The HP MSA 2040 array supports implicit ALUA modes. This ability means that the array can assign and change the managing controller for the LUN but LUN ownership cannot be assigned to one particular MSA controller.

Overview of ULP Operation

- ULP presents all LUNs to all host ports
 - Removes the need for controller interconnect path
 - Presents the same World Wide Node Name (WWNN) for both controllers

- Shared LUN number between controllers
 - No duplicate LUNs allowed between controllers
 - Either controller can use any unused logical unit number
- ULP recognizes which paths are “preferred”
 - The preferred path indicates which is the owning controller per ALUA specifications
 - “Report Target Port Groups” identifies preferred path
 - Performance is slightly better on preferred path
- Write I/O Processing with ULP
 - Write command to controller A for LUN 1 owned by Controller B
 - The data is written to Controller A cache and broadcast to Controller A mirror
 - Controller A acknowledges I/O completion back to host
 - Data written back to LUN 1 by Controller B from Controller A mirror

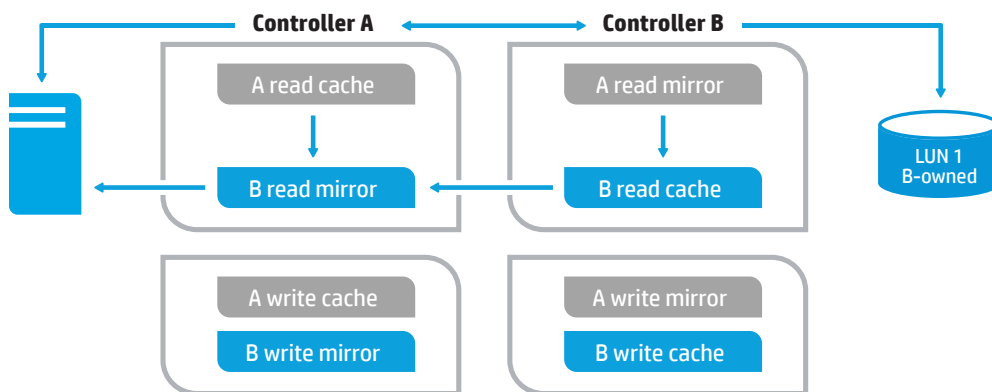
Figure 3. Mechanics of write I/O processing with ULP



Read I/O Processing with ULP

- Read command to controller A for LUN 1 owned by Controller B:
 - Controller A asks Controller B if data is in Controller B cache
 - If found, Controller B tells Controller A where in Controller B read mirror cache it resides
 - Controller A sends data to host from Controller B read mirror, I/O complete
 - If not found, request is sent from Controller B to disk to retrieve data
 - Disk data is placed in Controller B cache and broadcast to Controller B mirror
 - Read data sent to host by Controller A from Controller B mirror, I/O complete

Figure 4. Mechanics of read I/O processing with ULP



Unified LUN Presentation (ULP) Failover

If a controller unit fails on the HP MSA 2040 Storage array with dual controllers, the ownership transfers from the failing controller to the secondary, or backup, controller in the array. Single controller configurations do not work in this scenario. The vdisk ownership transfers to the secondary controller unit; the same single World Wide Node (WWN) is still presented, and all LUNs are now presented through the surviving controller unit. The multipathing software continues providing I/O and the surviving controller reports all paths as preferred.

The ability to identify and alter the LUN controller ownership is defined by the ALUA extensions in the SPC3 standard. The HP MSA 2040 Storage MSA array supports implicit ALUA modes. This ability means that the array can assign and change the managing controller for the LUN, but LUN ownership cannot be assigned to one particular 2040 controller.

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 all four 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.

RAID Level Considerations

This section contains considerations for RAID levels on the HP MSA 2040.

Table 2. Overview of supported RAID implementations for HP MSA 2040

RAID level	Cost	Performance	Protection level
RAID 0 Striping	N/A	Highest	No data protection
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	High	Protects against any individual drive failure; 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

Types of RAID sets

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

- RAID 0 (block level striping)

RAID 0 is also known as a stripe set or a 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.

- RAID 1+0 (mirroring + striping)

RAID 1+0 (also known as RAID 10) 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.

- RAID 5+0 (RAID 5 distributed parity+striping).

RAID 5 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 block 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 identical to RAID 5 except for the addition of a second parity block. It does not have a performance penalty for read operations; it does have a performance penalty for write operations, due to the overhead incurred by parity calculations.

VMware vSphere Storage Integration

The HP MSA 2040 Storage integrates with key VMware vSphere features and vSphere Storage APIs to provide additional benefits around performance, protection, and ease-of-use.

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.

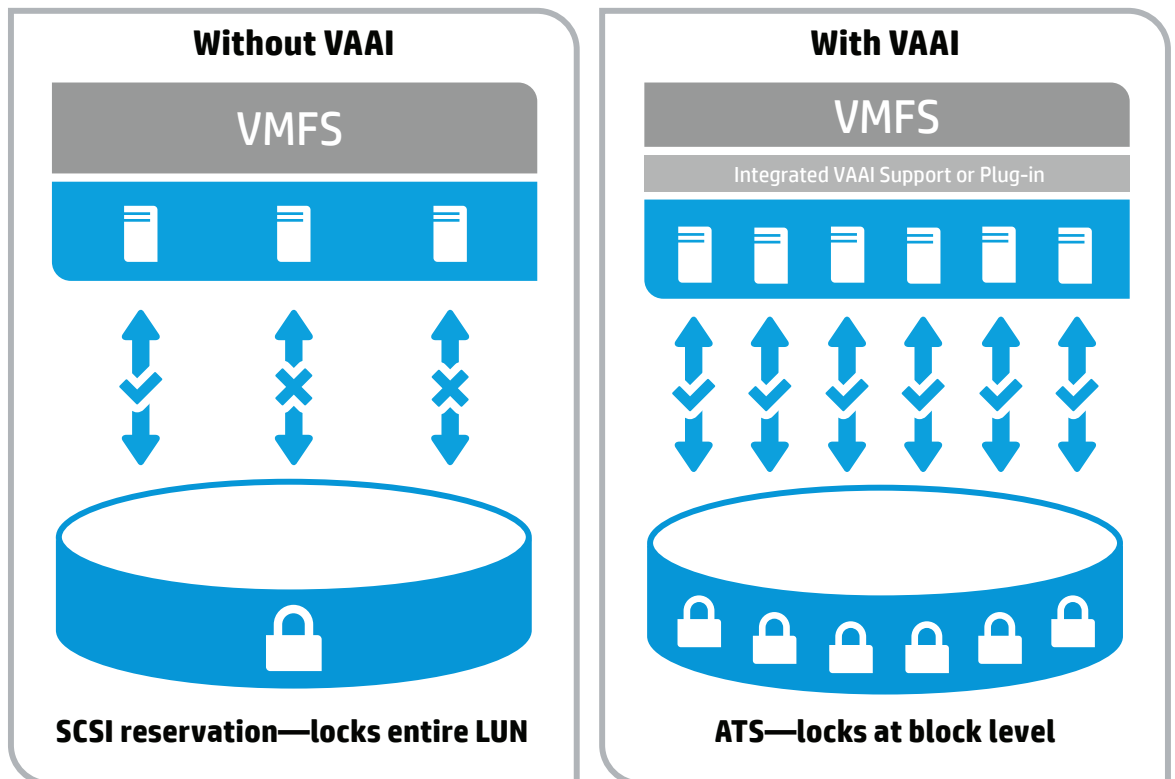
Primitives are specific functions used with VAAI that serve as integration points to storage arrays. When supported by the array, primitives in VAAI allow the hypervisor to communicate directly with storage arrays to offload storage functionality traditionally handled by the hypervisor. Storage arrays 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 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 3, 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 virtual machine data (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 HP MSA 2040 array natively supports these three primitives, which means an additional MSA plug-in is no longer required with vSphere 5.0 and higher to receive the benefits of VAAI integration.

Figure 5. Comparison of vSphere host contention with SCSI reservations/LUN locking and vSphere with VAAI, without LUN locking issue



VAAI Benefits and Use Cases

VAAI helps reduce the storage bandwidth consumed by a vSphere host and improves datacenter 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 HP MSA 2040 Storage system.

VAAI Use cases

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 IO performance impact to those hosts. Because the chance for SCSI reservation contention is greatly reduced, users can increase the number of VMs per server (also known as VM Density).

HP Insight Control Storage Module for VMware vCenter Server

HP Insight Control Storage Module for VMware vCenter Server is a storage management integration plug-in. [HP Insight Control for vCenter](#) provides the ability to create storage array devices, map them to array ports, and create datastores and virtual machines directly from vCenter Server.

HP Insight Control Storage Module 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

HP Insight Control for vCenter Server saves time and resources, allowing administrators to manage both VMware vSphere and HP Storage from a common interface.

HP Insight Control for vCenter integrates with the vCenter management server (also known as vCenter 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.

HP Insight Control for vCenter software can be installed on a management server or a VM and configured to connect and register with a vCenter Server. After registering Insight Control for vCenter with a vCenter server, all vSphere clients connected to the vCenter server can use the HP Insight Software tab to access the Insight Control for vCenter software.

The HP Insight Control Storage Module for 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 HP 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 HP storage arrays that are presenting disks to the selected cluster.

- **ESX Server (host) level View:** Provides summary and detailed information about the HP 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).

The HP Insight Control Storage Module for vCenter is available online from [this HP website](#) and can be installed and used at no charge. You will need to accept the standard HP software licensing terms as part of the download and installation sequence. An HP Passport sign-in will be required.

Disaster Recovery with vCenter Site Recovery Manager and HP MSA 2040 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 of each business-critical application helps to control costs.

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, you can use VMware 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 HP MSA P2040 Storage Replication Adapter (SRA) provides a platform for a seamless integration of HP MSA P2040 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 that the virtual machines at the protected site are shut down properly (when the protected site virtual machines are available) and the replicated virtual machines can be powered up. SRA helps in executing the DR plans.

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 reprotect operation for virtual machines protected on an HP MSA 2040 storage system. The reprotect operation reverses the roles of the two sites after the original protected site is backed up. The site that was formerly the recovery site becomes the protected site, and the site that was formerly the protected site becomes the recovery site.

SRM also enables recovery plan testing using native functionality, without disruption to production environments, by using a temporary copy of the replicated data. Conduct tests after a reprotect operation to confirm that the new protected recovery site configuration is valid. For more information about the integration of vSphere vCenter Site Recovery Manager with HP MSA 2040 Storage array, refer to [this whitepaper](#).

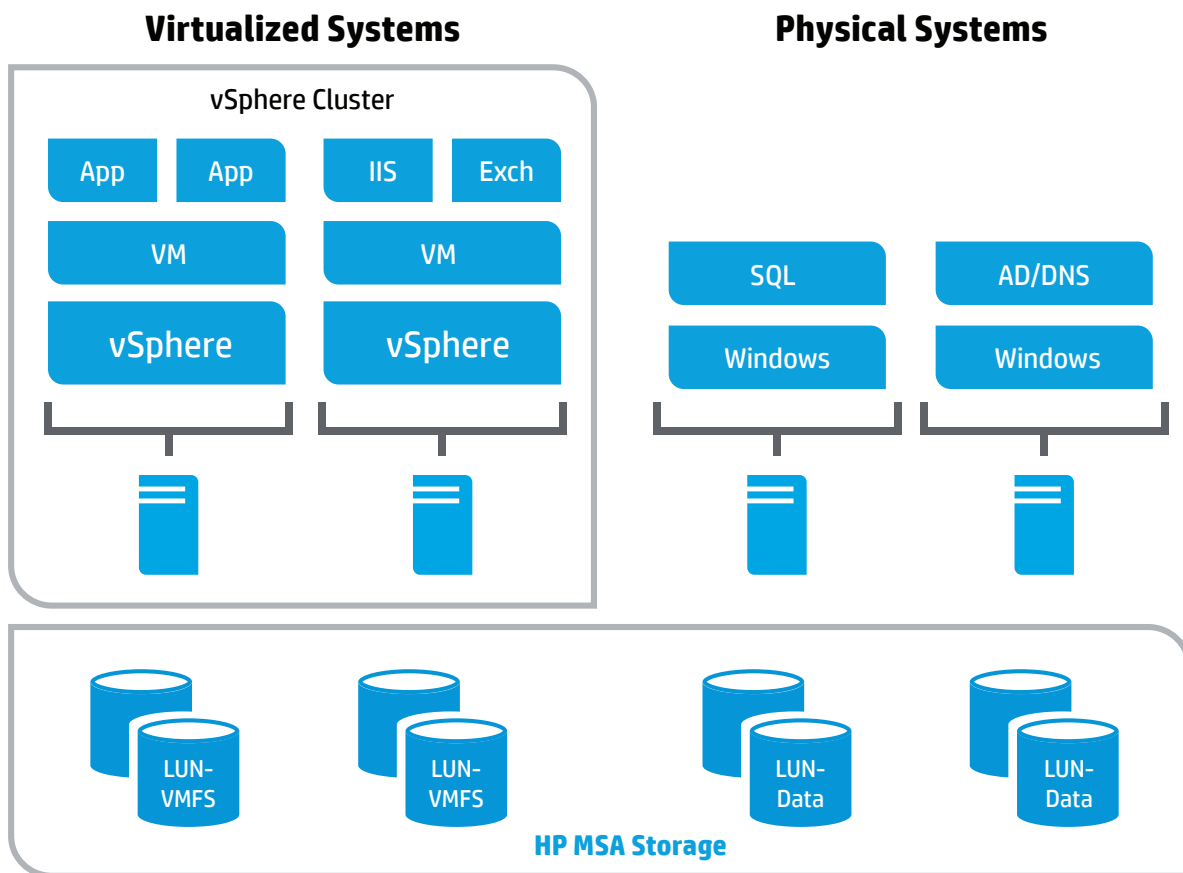
Use Cases for vSphere and HP MSA 2040 Storage

HP 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 HP MSA 2040 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 HP MSA 2040 storage solution supports the storage needs of both physical servers and vSphere hosts. As more systems become virtualized and consolidated, the HP MSA can grow and scale to suit changing business needs.

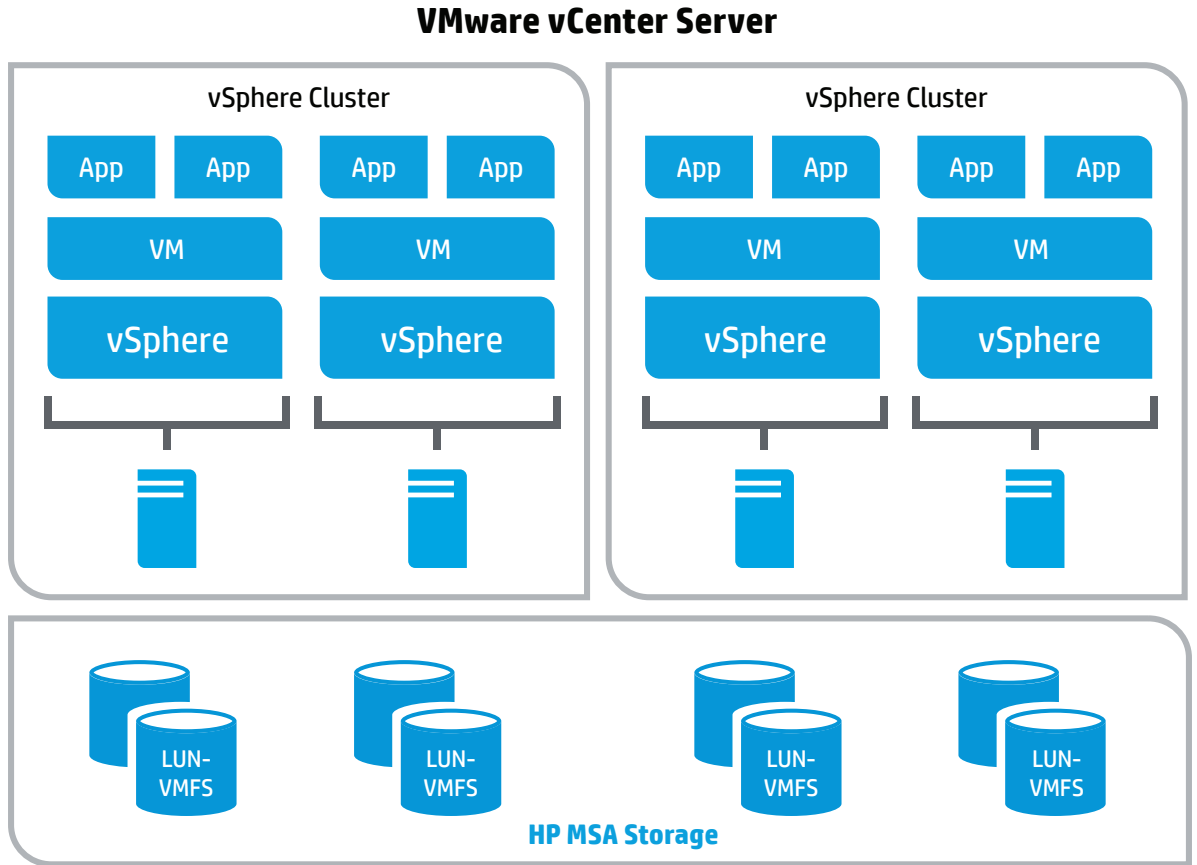
Figure 6. Diagram of a mixed or hybrid environment, with physical systems hosting critical applications and virtualized systems with virtualized applications, utilizing HP MSA 2040 storage



Multiple vSphere clusters on HP MSA Storage

Figure 5 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 7. Recommended vSphere cluster deployment, with physical servers running vSphere and hosting virtualized servers and applications, using HP MSA 2040 storage array



ROBO deployments for distributed vSphere environments

ROBO (remote office/branch office) offers a decentralized approach for systems deployment. Distributed VMware vSphere environments spanning multiple sites can have storage needs which are cost-prohibitive. With a relatively small physical footprint and affordable, entry-level pricing, the HP MSA 2040 is a practical choice for storage in a remote office location. Rapid Provisioning software eliminates array planning by delivering instant, fine-grained, application-tailored storage provisioning. Administrators manage the time-consuming task of storage provisioning intelligently and automatically, easily adding more storage to support a growing virtual client environment.

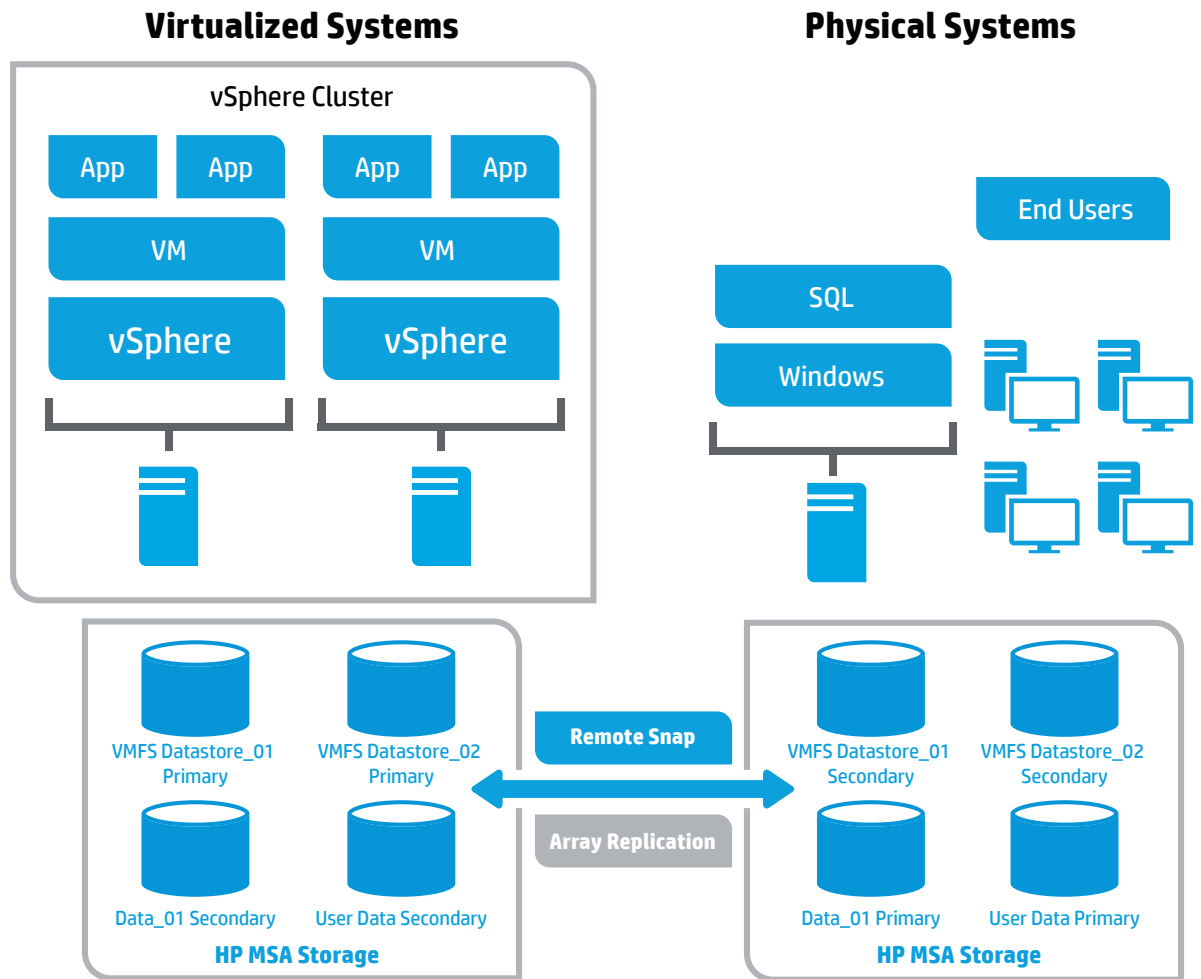
Figure 8. Configuration diagram of a ROBO environment with VMware vCenter Server, SRM and HP MSA 2040 Storage



Data Protection with HP MSA Remote Snap

Cost-effective data protection can be a challenge for today's small to midsize businesses. The HP MSA 2040 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 HP MSA 2040 storage system solutions, provides a business site with ultimate data protection and redundancy.

Figure 9. Configuration diagram with VMware vSphere environment, HP MSA Storage, and Remote Snap replication between two storage arrays



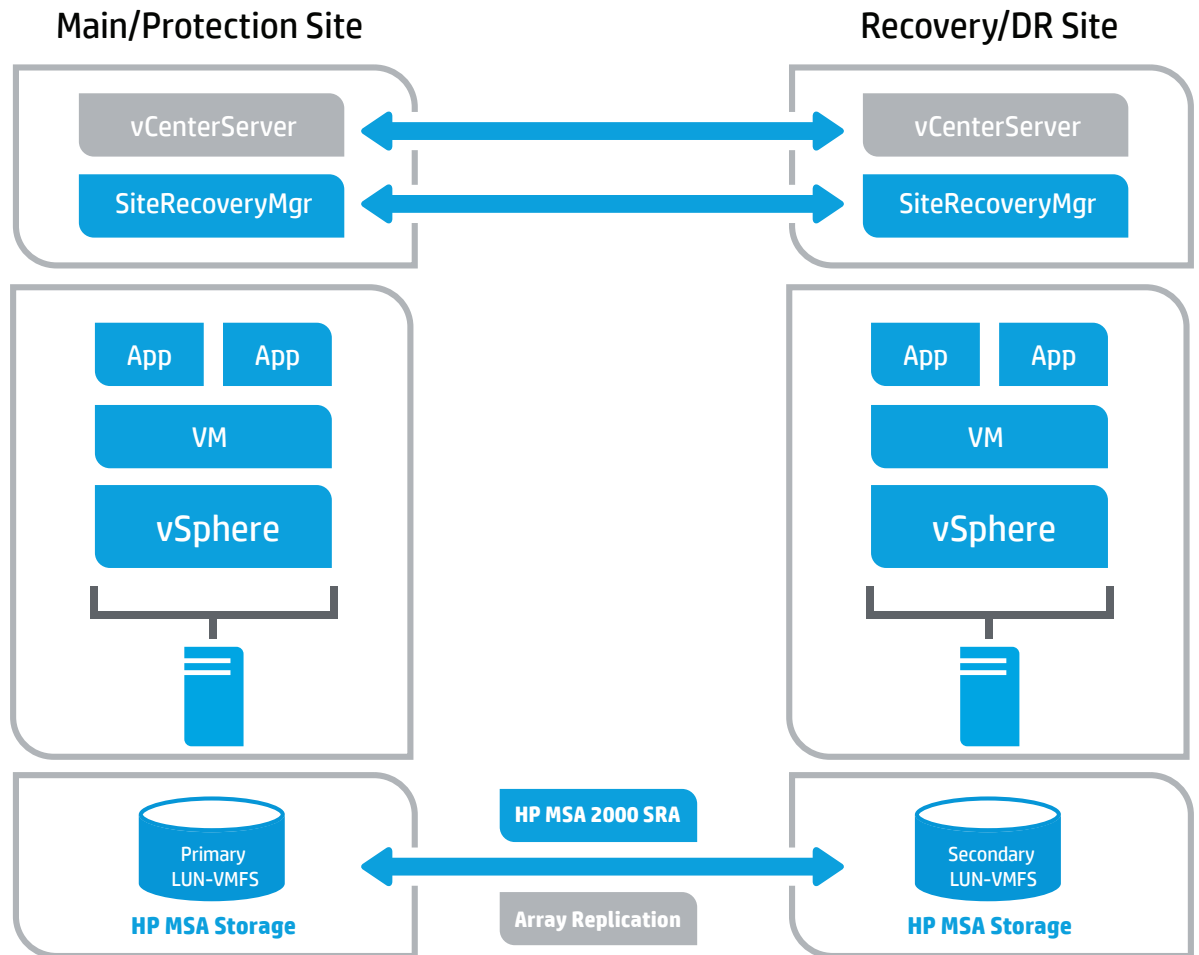
Remote Snap Software is an optional, licensed disaster-recovery feature that performs asynchronous (batch) replication of block-level data from a volume on a primary MSA 2040, to a secondary volume on the same system or on a separate, independent MSA 2040. These two MSA 2040 systems connect through switches and are on the same fabric or network.

Multi-Site Disaster Recovery

The VMware Site Recovery Manager solution provides virtual datacenter site disaster recovery. When implemented with HP MSA 2040 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 HP MSA P2040 SRA to test DR scenarios and create different recovery plans.

Figure 10. Configuration diagram of a multi-site environment with VMware SRM and HP MSA 2040 MSA



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

The HP MSA 2040 SRA is a licensed feature for use with the HP MSA 2040 Storage array.

MSA Storage Installation for VMware vSphere

Because storage plays a critical role in the success of a VMware vSphere deployment, it is important to properly install and configure the HP MSA storage. The following section highlights some key installation and configuration notes for the HP MSA 2040 storage array. This paper focuses on the best practices for deploying and configuring HP MSA Storage that are specific to VMware vSphere; complete installation steps for the MSA hardware are not covered in detail here.

Other resources, such as HP MSA Configuration Guides, provide detailed installation information. Refer to the following resources, which are available on the [HP MSA 2040 webpage](#):

- The HP MSA 2040 User Guide
- The HP MSA 2040 Cable Configuration Guide
- The HP MSA 2040 SMU Reference Guide

Additional VMware vSphere Storage Configuration Guides

Refer to the [VMware Storage Guide](#) for guidance on using vSphere with HP MSA storage, iSCSI, or Fiber-Channel storage, and related best practices not covered in this paper. For information specific to storage protocols and VMware vSphere, the following resources are available:

- [VMware iSCSI SAN Configuration Guide](#)
- [VMware Fiber Channel SAN Configuration Guide](#)

Storage Network Connectivity

The MSA 2040 model supports 16 Gb (as well as 8 Gb) FC host connections and four FC host ports per controller. By using the latest 16 Gb Fiber Channel technology, users benefit from:


- Fast and reliable access to data
- Host ports doubled in speed from the previous 8 Gb FC
- Faster controller architecture has double the cache verify in QS
- Overall performance benefits from the new 6 Gb SAS access to the drives
- Scalability has increased substantially over the previous model, but still allows you to start small and grow as needed.

Cabling the MSA Storage Array

HP MSA 2040 Storage supports both fault-tolerant and straight-through SAS cabling. 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.

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

Figure 11. HP MSA 2040 backplane with interfaces/ports



1. Power supplies
2. 16 Gb or 8 Gb Fiber Channel ports
3. CLI port (mini-USB)
4. Reserved for future use
5. Management Ethernet port
6. Expansion port

A configuration with dual FC switches for host connectivity to the MSA array is a best practice, offering excellent redundancy. For the HP MSA 2040 Storage cabling requirements and options, consult the “HP MSA 2040 User Guide” at hp.com/go/msa.

FC Switches and SFP transceivers

The FC switches that HP offers use differing models of Small Form-Factor Pluggable (SFP) transceivers. The correct SFPs must be loaded into the correct 16 Gb switches when connecting to the HP MSA 2040 Storage FC ports. Likewise, the correct SFPs must be used in the supported 16 Gb FC switches when connecting to the MSA 2040 array. If the wrong SFPs are used in an unsupported FC switch, the storage on the HP MSA 2040 Storage FC MSA will not be visible. Consult the MSA 2040 QuickSpecs for more information on supported SFP transceivers and FC switches at hp.com/go/san.

For more guidance regarding HP MSA 2040 Storage cabling, consult the “HP MSA 2040 Storage MSA System Cable Configuration Guide” on the [HP MSA2040 webpage](http://hp.com/go/msa).

Best Practices: MSA Initial Configuration

During the installation and initial configuration, follow these best practices to ensure a successful implementation. Review these best practices for configuring your HP MSA storage solution specifically for deployment into a VMware vSphere environment.

Fastest throughput optimization

The following guidelines list the general best practices to follow when configuring your storage system for fastest throughput for your VMware vSphere environment:

- Configure host ports for 16 Gb/sec on the HP MSA 2040 Storage FC MSA.
- Balance Virtual disks between the two controllers.
- Balance Disk drives between the two controllers.
- Set cache settings to match (optimizing performance for your application).
- To get the maximum sequential performance from a vdisk, create only one volume per vdisk. Otherwise, you will introduce randomness into the workload when multiple volumes on the vdisk are being exercised.
- Map volumes to dedicated host ports; this ensures optimal performance for virtualization workloads.
- Distribute the load across as many drives as possible.

Note

Create vdisks with the appropriate RAID level for the specific data protection requirements and virtualization demands. Use RAID 5 to lower storage-capacity costs and to provide adequate redundancy for vSphere deployments. Use RAID 10 when storage performance is a high priority. Refer to the MSA Storage Best Practices for more guidance.

Highest fault tolerance optimization for 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 Input/Output (MPIO) software; VMware Native Multipathing (NMP) with HP MSA firmware. HP recommends VMW_PSP_RR for path failover policy. Refer to [VMware Compatibility Guide](#) for more information.

Remember to verify the version of 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).

Global Spares (Drives)

You can designate a maximum of 16 global spare drives for the system. If a disk in any redundant Vdisk (RAID 1, 3, 5, 6, 10, and 50) fails, a global spare is automatically used to reconstruct the Vdisk.

Give careful consideration to the minimum number of global spares, based on the number of planned vdisks and the high-availability requirements of the vSphere hosts. Using global spares with vSphere VMFS datastores ensures a storage-based layer of high availability to complement vSphere's HA feature. When creating a vdisk using RAID-5 or RAID-6, a minimum of 2 global spares is recommended. The disk capacity of the MSA 2040 and the planned number of vdisks will determine the number of global spares needed. Refer to section on MSA Storage Best Practices for recommendations on vdisk creation.

At least one Vdisk must exist before you can add a global spare. A spare must have sufficient capacity to replace the largest disk in an existing Vdisk. If a drive in the virtual disk fails, the controller automatically uses the Vdiskglobal spare for reconstruction of the critical virtual disk to which it belongs. A spare drive must be the same type (SSD, SAS or SATA) as other drives in the virtual disk. If two drives fail in a RAID 6 virtual disk, two properly sized spare drives must be available before reconstruction can begin. For RAID 50 virtual disks, if more than one sub-disk becomes critical, reconstruction and use of Vdisk spares occurs in the order sub-Vdisks are numbered.

You can designate a global spare to replace a failed drive in any virtual disk, or a Vdisk spare to replace a failed drive in only a specific virtual disk. Alternatively, you can enable dynamic spares in the HP Storage Management Utility (SMU). Dynamic sparing enables the system to use any drive that is not part of a virtual disk to replace a failed drive in any virtual disk.

To change the system's global spares

1. In the Configuration View panel, right-click the system and select Provisioning > Manage Global Spares. The main panel shows information about available disks in the system. Existing spares are labeled GLOBAL SP. In the Disk Sets table, the number of white slots in the Disks field shows how many spares you can add. In the Graphical or Tabular view, only existing global spares and suitable available disks are selectable.
2. Select spares to remove, disks to add as spares, or both.
3. Click "Modify Spares". The panel will update to indicate which disks are now global spares.

Storage Administration

After installation, the MSA array is configurable and managed from the Storage Management Utility (SMU), a web browser interface. Using SMU, administrators can create vdisks and volumes, map volumes to hosts, replicate and delete volumes and vdisks, and monitor the health of the system components. SMU gives storage administrators a simple way to perform day-to-day administration tasks. In addition, a command line utility does everything the GUI does and offers performance monitoring counters.

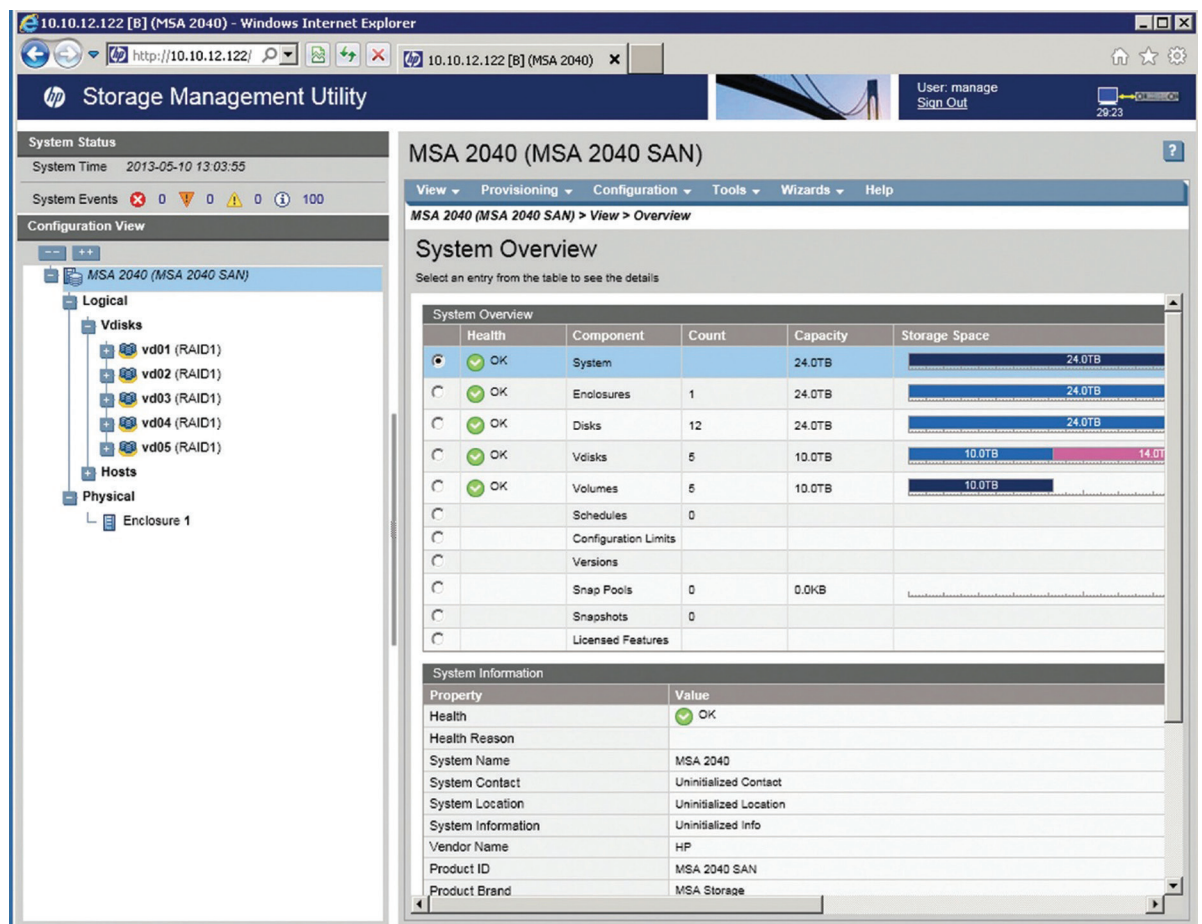
Administering HP MSA 2040 with HP Storage Management Utility

Open your web browser, and enter the IP address of the controller module's network port in the address field, then press Enter. To sign in to SMU, use the default user name **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 HP MSA 2040 Storage System 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 12. HP MSA 2040 Storage Management Utility (SMU) interface



Configuring a new HP MSA 2040 Storage system

Follow these steps to prepare to deploy a new system for VMware vSphere.

To configure and provision a storage system for the first time:

1. Configure your web browser for SMU and sign in.
2. Set the system date and time.
3. Use the Configuration Wizard to configure other system settings.
4. Use the Provisioning Wizard to create virtual disk (vdisk) containing storage volumes, and optionally to map the volumes to hosts.
5. Use the Replication Setup Wizard to configure replication for a primary volume to a remote system volume.
6. If you mapped volumes to hosts, verify the mappings by mounting/presenting the volumes from each host and performing simple read/write tests to the volumes.
7. Verify that controller modules and expansion modules have the latest firmware.

Complete the initial installation and setup of the MSA array before provisioning your vSphere hosts. Provisioning vdisks, volumes, and LUNs for vSphere environments requires forethought and planning.

VMware Hosts: Storage Presentation

Configuring vSphere hosts is not discussed extensively in this paper. For more details, see the *VMware Fiber Channel SAN Configuration Guide*. Configuration is also documented on several VMware blogs and other Storage related websites.

During configuration of your vSphere host, place two FC ports on separate fabrics.

When scanned by the vSphere hosts, the HP MSA 2040 Storage array will always show LUN 0 in the controllers. When planning your LUN's with the HP MSA 2040 Storage, remember LUN 0 is already used by the HP MSA 2040 array for its **Storage** controllers.

If you have no fiber HBA drivers installed on the host, make sure you have installed the HBA drivers following the manufacturer's directions for the HBA model installed in your system.

In our example in figure 13, we used HP HBAs to test our proof-of-concept. Appendix A describes how to install fiber HBA drivers after vSphere 5.x has been installed.

As vSphere hosts are deployed or installed with the correct fiber HBA drivers, the Storage adapters and their corresponding WWN names will be displayed. In our example, we used a Brocade 425/825 8G FC HBA (2 port HBA), which is the HP 82B PCI-e 8Gb FC dual port, host-based controller, as shown in figure 13.

Figure 13. A vSphere 5.x host with Emulex HBAs and their corresponding WWN

Storage Adapters
Add... Remove

Device	Type	WWN
Smart Array P410i		
vmhba4	Block SCSI	
LPe12000 8Gb Fibre Channel Host Adapter		
vmhba0	Fibre Channel	20:00:00:00:c9:90:5d:20 10:00:00:00:c9:90:5d:20
vmhba1	Fibre Channel	20:00:00:00:c9:90:5d:21 10:00:00:00:c9:90:5d:21
vmhba2	Fibre Channel	20:00:00:00:c9:90:5f:02 10:00:00:00:c9:90:5f:02
vmhba3	Fibre Channel	20:00:00:00:c9:90:5f:03 10:00:00:00:c9:90:5f:03

Details

vmhba0

Model: LPe12000 8Gb Fibre Channel Host Adapter

WWN: 20:00:00:00:c9:90:5d:20 10:00:00:00:c9:90:5d:20

Targets: 1 Devices: 6 Paths: 7

View:
Devices
Paths

Runtime Name	Target	LUN	Status
vmhba0:C0:T0:L22	2a:21:00:02:ac:00:14:e7 22:21:00:02:ac:00:14:e7	22	◆ Active (I/O)
vmhba0:C0:T0:L0	2a:21:00:02:ac:00:14:e7 22:21:00:02:ac:00:14:e7	0	◆ Active (I/O)
vmhba0:C0:T0:L4	2a:21:00:02:ac:00:14:e7 22:21:00:02:ac:00:14:e7	4	◆ Active (I/O)

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 fiber cables for any damaged or non-functional components. With the correct drivers installed, you can provision the LUNs for the vSphere hosts in your cluster or datacenter. To verify storage connectivity, use the SMU option that shows discovery of the host port by the vSphere host.

Virtual machines access data using two methods: VMFS (.vmdk files in a VMFS file system) and Raw Device Mapping (RDM), which is a mapping file inside the VMFS that behaves like a proxy to the raw storage device. RDM is recommended in cases when a virtual machine needs to interact with a physical disk on the SAN.

Mapping volumes to an vSphere host

For all volumes in all vdisks or a selected vdisk, you can change access to those volumes by a specific host. When multiple volumes are selected, LUN values are sequentially assigned, starting with a LUN value that you specify. For example, if the starting LUN value is 1 for 30 selected volumes, the first volume’s mapping is assigned LUN 1 and so forth, and the last volume’s mapping is assigned LUN 30. For LUN assignment to succeed, verify that no value in the sequence is already in use. When you specify access through specific ports, the ports and host must be the same type (for example, FC).

Caution

Volume mapping changes take effect immediately. Make changes that limit access to volumes when the volumes are not in use. Before changing a volume’s LUN, be sure to unmount/unpresent/unmap the volume.

MSA Storage Best Practices

The following section covers general best practices for HP MSA Storage in a VMware vSphere environment.

Maximize capacity and disk usage

To maximize capacity and disk usage, create vdisks larger than 2TB to increase the usable capacity of storage configurations by reducing the total number of parity disks required when using parity-protected RAID levels. This method differs from using a volume larger than 2 TB, which requires specific support by the host operating system, I/O adapter, and application and does not maximize performance.

Using multiple enclosures

When using multiple enclosures, stripe across shelf enclosures to enable data integrity in the event of an enclosure failure.

Selecting the controller ownership for the vdisk

When creating vdisks, be sure to add them evenly across the two controllers. Because both controllers are active, you must have at least one vdisk assigned to each controller. If you use the default value of Auto on, this will be done for you. In addition, each controller should own a similar number of vdisks.

When optimizing for performance, the SAN administrator's goal is to drive as much performance out of the array as possible. This has implications for usable storage capacity, due to the configuration decisions.

Storage configurations with several disks

When an array has many physical disks available, create a few vdisks, each containing many disks, instead of many vdisks each containing a few disks. I/O to multiple volumes containing vSphere VMFS datastores on the same vdisk can have an adverse effect on overall system performance at the vSphere cluster level. Supporting large storage capacities requires advanced planning because it requires using large virtual disks with several volumes each or many virtual disks. Here is where performance and capacity come into effect with your planning. To maximize capacity, combine physical disks into a large vdisk, then subdivide that vdisk into several volumes with a capacity less than 2TB, because, as of VMware vSphere 5.1, VMFS has a 2TB limit on each "extent" of a datastore.

RAID levels

Use RAID 5 to lower storage-capacity costs and to provide adequate redundancy for vSphere deployments. Use RAID 10 when storage performance is a high priority.

Refer to the "HP StorageWorks P2000 G3 MSA best practices" technical white paper for more information: h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA3-2141ENW&cc=us&lc=en

Changing the default name of the volumes

Change the default name of the volume to identify the purpose of the volume. For example, a volume used for storing datastores belonging to ESX cluster1 can be named 'esx_cluster1_datastores'.

To rename a volume in the SMU, follow these steps:

- Log in to the SMU and locate the vdisk which contains the volume.
- Right-click on the volume. Select the Configuration menu option and select "Modify Volume Name".
- Rename the volume to something the SAN Administrator can understand and manage.

Volume mapping

For best performance, follow the VMware-specific practices for volume mapping:

- Use explicit mapping.
- Make sure that a shared LUN is mapped with the same LUN number to all vSphere 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 volumes to the ports on the controller that own the vdisk. Mapping to the non-preferred path may result in performance degradation

Note

Configuring vSphere hosts is not discussed extensively in this paper. For more details, see the [“VMware Fiber Channel SAN configuration Guide”](#).

When configuring your vSphere host, you should place two fiber ports on separate fabrics.

Caution

When scanned by the vSphere hosts, the HP MSA 2040 storage array will always show LUN 0 in the controllers. When planning your LUN's with the HP MSA 2040 storage; remember LUN 0 is already used by the array for its storage controllers.

Naming hosts

To provide better administration of your SAN environment, always rename hosts to something meaningful. Datacenter administrators may have a naming convention in use within the datacenter to keep track of servers and their components.

To rename a host in the SMU, follow these steps:

- Locate the WWN from the vSphere client.
- Log in to the SMU and locate the WWN which correspond to the ESXi host's fiber HBA.
- Rename the hosts to something the SAN Administrator can understand and manage.

In many cases, because the HP MSA 2040 Storage is targeted to small businesses, you 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 fiber switch level.

Boot from SAN support/considerations

As a rule, when you're using boot from SAN, each boot LUN should be seen only by the ESX Server system that is booting from that LUN. However, if you are trying to recover from a crash by pointing a second ESX Server system to the same LUN, an exception is made. In this case, the SAN LUN in question is not really a boot from SAN LUN. No ESX Server system is booting from it because it is corrupted. The SAN LUN is a regular, non-boot LUN that is made visible to an ESX Server system.

For more information regarding vSphere installation and boot-from-SAN configurations, refer to the vSphere Installation and Setup guides for [vSphere 5.0](#) and [vSphere 5.1](#).

Follow these steps to identify boot LUNS for ESXi hosts:

To obtain the 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 device list -d naaID
```

Multipath Considerations for vSphere

To maintain a constant connection between an ESXi host and its storage, ESXi supports multipathing. To take advantage of this support, virtual volumes should be exported to multiple paths to the host server. To do this, create a host definition on the HP MSA Storage system that includes the World Wide Names (WWNs) of multiple HBA ports on the host server and then export the VLUNs to that host definition. For a vSphere cluster, the VLUNs must be exported to all of the host definitions for the cluster nodes, or a host set may be created containing all of the servers and the VLUNs can be exported to the host set.

VMware vSphere includes active/active multipath support to maintain a constant connection between the ESXi host and the HP MSA Storage array. Three path policies are available, "Fixed", "Most Recently Used," and "Round Robin". For HP MSA storage, Round Robin is the recommended policy for best performance and load balancing; however, it may not be enabled by default. The path policies can be viewed and modified from the vSphere Web Client on a per datastore basis as follows:

1. In the vSphere Web Client, select the datastore.
2. Select the Manage tab, then the Settings tab, and then click Connectivity and Multipathing.
3. Select one of the ESXi hosts, and then click the Edit Multipathing button.
4. In the pop-up window, select Round Robin from the Path Selection drop-down menu.
5. Click OK to save the new setting.
6. Repeat steps 3 through 5 for each ESXi host.

In the context of this document, a port is the connection from a device into the SAN. Each node in the SAN—each host, storage device, and fabric component (router or switch)—has one or more ports that connect it to the SAN. Ports can be identified in a number of ways:

- **WWN:** The World Wide Node Name is a globally unique identifier for a Fiber 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. To view the WWPN using the VI Client, click the host's Configuration tab and choose storage Adapters. You can then select the storage adapter for which you want to see the WWPN.
- **iSCSI Qualified Name (iqn)** – 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 web site of the Storage Networking Industry Association, snia.org.

A given host might be able to access a volume on a storage array through more than one path. Having more than one path from a host to a volume is called multipathing. By default, VMware ESX systems use only one path from the host to a given volume at any time. 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.

VMware vSphere multipathing framework

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

The new HP MSA 2040 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.

Selecting Best Path Failover (PSP) policy

With server virtualization environments using vSphere 5.x or later, VMW_PSP_RR for path failover policy is recommended.

As a best practice, change the default PSP option for VMW_SATP_ALUA to VMW_PSP_RR for HP MSA 2040 Storage SAN environments. For optimal default system performance with HP MSA 2040 Storage, configure the round robin load balancing selection to IOPS with a value of 1 for every LUN using the command:

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

For environments with only HP MSA 2040 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 HP 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 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).

Changing LUN Response before LUN Creation

Before creating volumes/LUNs for use by vSphere hosts, it is essential to change a setting for LUN Response on the MSA 2040. The following are two VMware knowledgebase articles on 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. Change the LUN RESPONSE setting through the HP System Management Unit (SMU) or CLI using either controller A or B. To change the LUN RESPONSE setting through the SMU:

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

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.

Additional VMware considerations

Configuring HP MSA Storage involves considerations of some VMware vSphere storage-related features, such as Storage I/O Control (SIOC) and multipathing policies (round robin, MRU).

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 HP MSA 2040 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.

Starting with ESX 4.0, the round robin load balancing policy is supported; both MRU and round robin path policy are ALUA-aware and have the following characteristics:

MRU

- Gives preference to an optimal path to the LUN
- Uses a non-optimized path if all optimal paths are unavailable
- Fails back to the optimal path when an optimal path becomes available
- Uses a single controller port for LUN access per ESX hosts

Round Robin

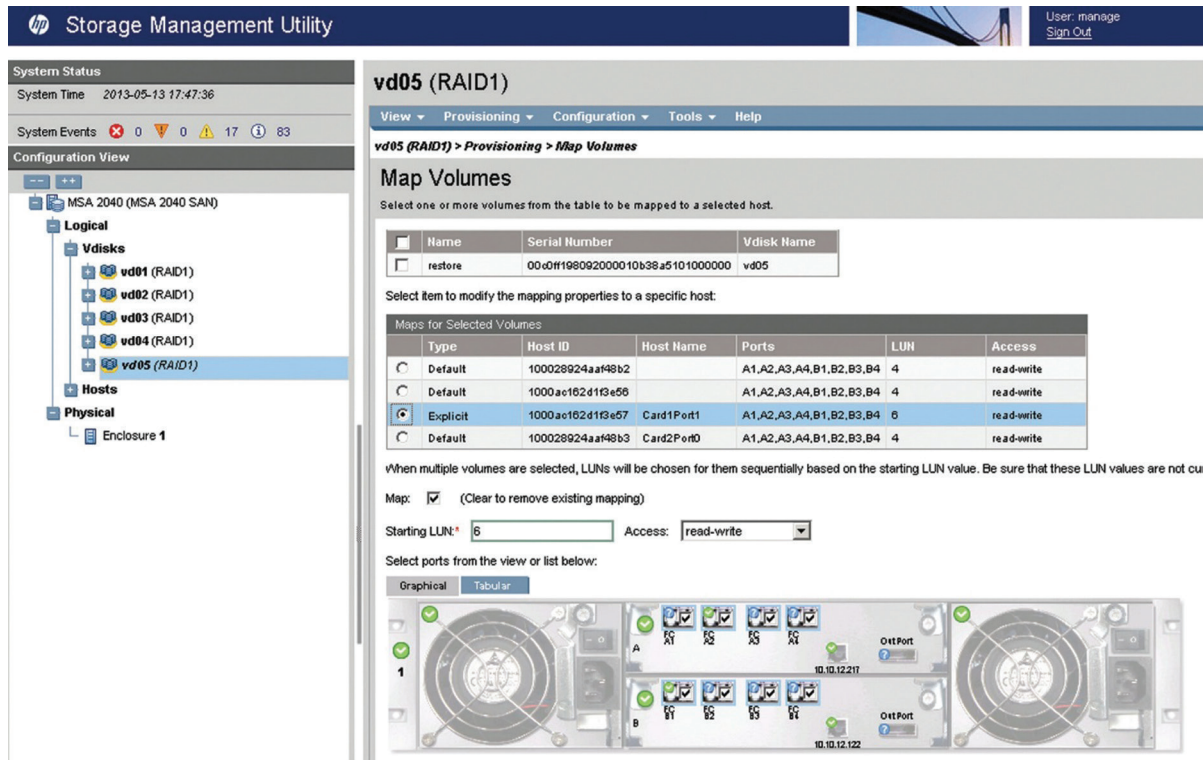
- Queues I/O to the LUNs on all ports of the owning controller in a round robin fashion, providing instant bandwidth improvement.
- Continues queuing I/O in a round robin fashion to the optimal controller ports until none are available; will then failover to the non-optimized path.
- When an optimal path returns, it will failback to the original storage path.

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, HP MSA 2040 Storage, providing an instant performance boost.

Setting HP MSA 2040 Storage MSA active-active arrays

Round Robin I/O path policy is the recommended setting for the HP MSA 2040 Storage MSA array. MRU is the default setting; it is also suitable if Round Robin is not the desired setting in the specific environment.

Figure 14. Mapping an explicit volume in HP MSA SMU



To create an explicit volume mapping

1. In the Maps for Volume table, select a host.
2. Select Map.
3. Set the LUN and select the ports and access type.
4. Click Apply. A message confirms or denies the change.
5. Click OK. The mapping becomes explicit with the new settings.

After you have created your explicit mappings and have installed vSphere/ESX 5.x on your host server, open the vSphere client and connect to the vCenter server and managed hosts. Select a vSphere host, open the configuration tab, and rescan all/scan for new storage devices. When the new LUN has been detected by the vSphere host, you can create the new VMFS datastores on the LUN or storage device.

To explicitly map multiple volumes

1. In the Configuration View panel, right-click vdisks or a vdisk and then select Provisioning > Map Volumes. In the main panel, a table shows all the volumes for the selected vdisk.
2. In the table, select the volumes to change. To select up to 100 volumes or clear all selections, toggle the checkbox in the heading row.
3. In the Maps for Selected Volumes table, select the host to change access for.
4. Select Map.
5. Either map volumes to the host or mask the volumes from the host as follows.
 - Map the volumes to the host by setting a starting LUN, selecting ports, and setting access to read-only or read-write.
 - Mask the volumes from the host by setting a starting LUN, selecting ports, and setting access to no-access.
6. Click Apply. A message specifies whether the change succeeded or failed.
7. Click OK.

Using vCenter Site Recovery Manager with HP MSA Remote Snap

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 may fail to protect customers when an application failure occurs. Based on the data center solution or application, the recovery point objective (RPO) and recovery time objective (RTO) may differ 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, SRM uses the HP MSA 2040 SRA software to enable simplified disaster recovery automation. SRM with HP MSA helps meet recovery time objectives, reduces the costs that are traditionally associated with business continuance plans, and achieves low-risk and predictable results for the recovery of a virtual environment.

SRM uses the HP 2040 SRA to test a DR plan by executing a failover in an isolated environment. This environment is quickly and automatically cleaned up following the testing to avoid disrupting production operations. Adequate disk array and server sizing mitigates the impact of testing failovers that include application recovery.

Use SRM to initiate a DR plan manually, which can be used for data migration. SRM displays each step of the progress and records each successful completion. In the event of a failure, it provides event and error logging.

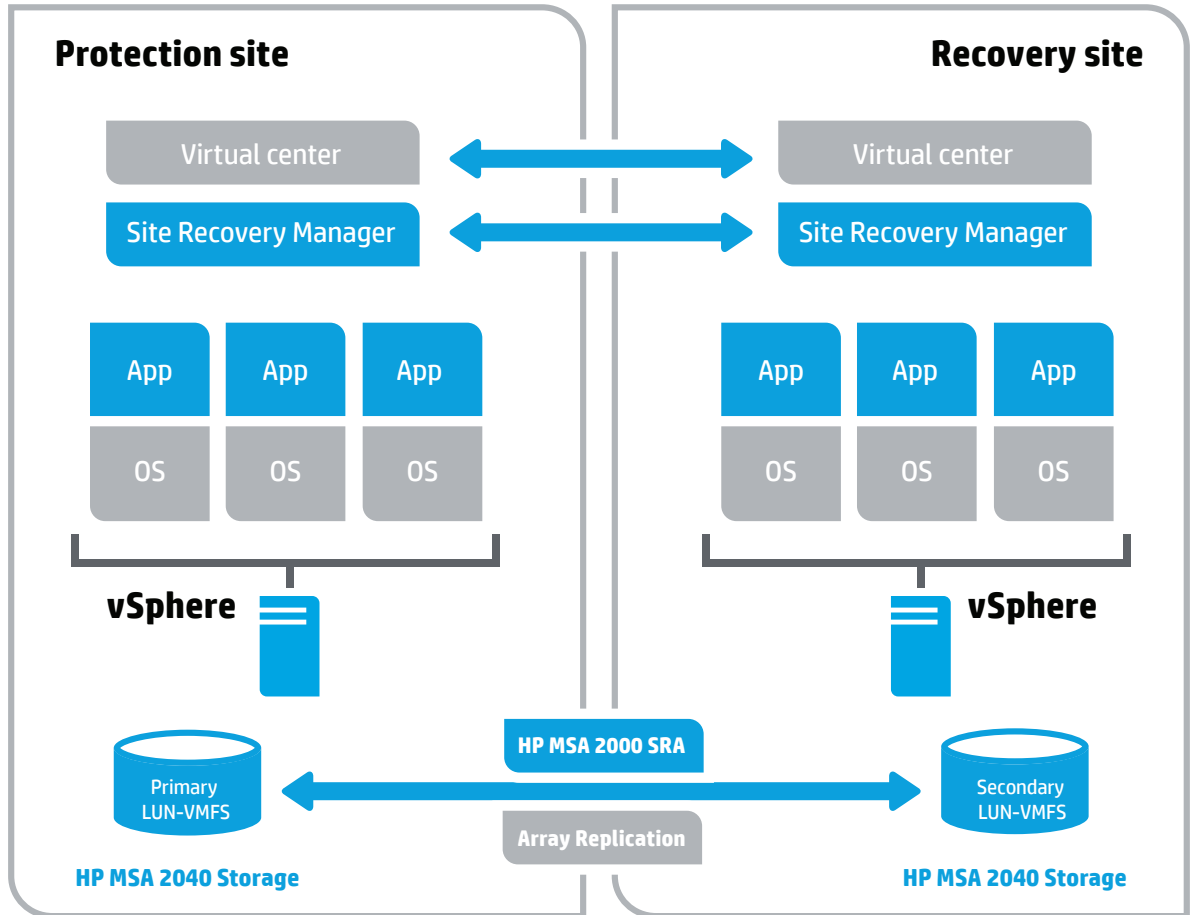
A single SRM interface displays and manages both the protected and recovery site installations of SRM. A user logged into either site can still see both sites. Administrators can perform all management tasks, eliminating the need to manage the two sites through separate interfaces.

Configuring VMware vCenter SRM with HP MSA 2040 SRA

After you have both SRM and SRA installed, the Getting Started tab of the main SRM window guides you through the steps necessary to configure your environment. For detailed SRM configuration instructions, refer to the Site Recovery Manager administration guide that is available on VMware’s website.

A typical SRM configuration involves two geographically separated sites with TCP/IP connectivity, the protection site, and the recovery site. The protection site is the site that is being replicated to the recovery site for DR. Each site contains an HP MSA 2040 Storage Array, vSphere hosts, a vCenter Server, and an SRM server.

Figure 15. VMware SRM with HP MSA 2040 SRA configuration with a primary “protection” site and a secondary “recovery” site



Best Practices for MSA 2000 SRA and Remote Snap with VMware SRM

Specific recommendations for using the HP MSA 2040 SRA and Remote Snap software in conjunction with a VMware SRM DR solution include these best practices.

1. Prepare a plan for re-establishing replication schedules in the event of a site failover. After performing a reverse replication operation, you must set up replication schedules to ensure periodic replication of data from the new source volumes back to the original source site. Alternately, you can 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. You can map volumes to LUN 0 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.

The SRA may create host entries on the array to keep track of remote IP or FC addresses. Do not delete host entries whose name starts with SRA.

Summary and Benefits

The HP P2000 and MSA 2040 storage arrays are well suited for small to medium businesses looking for a cost-effective, entry-level storage solution with the capacity to scale to meet business needs and virtualization growth. With support for various connectivity options, array-based snapshots, and VMware VAAI storage APIs out-of-the-box, HP MSA provides a versatile storage array for VMware vSphere environments. When these best practices and configuration recommendations are followed, the HP MSA 2040 makes it easy to deploy, scale, and maintain storage to ensure that your crucial business data remains available.

Appendix A: Firmware Update Best Practices for 2040 MSA

Appendix A details common array firmware update best practices for all generations of the HP MSA storage family, including general best practices and considerations when planning a firmware update on the P2000 array.

- As with any other firmware upgrade, verify that you have a full backup 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 HP customer support.
- After the device firmware update process is completed, confirm that the new firmware version is displayed correctly using one of the MSA management interfaces—SMU GUI, MSA CLI, and so on.

P2000 MSA Array Controller Firmware Update Best Practices

- 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 MSA70 drive enclosures, perform a power cycle on each enclosure.

Disk Drive Firmware Update Best Practices

- Disk drive upgrades on the HP 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. You can determine this using SMU or CLI interfaces. 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 Best Practices

- 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 vdisk 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 vdisk. Only make a master volume for replication and snapshot use.
- When replicating, put both the volume and snap-pool on the same vdisk, but try to keep any other volumes off that vdisk.
- 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 vdisk 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 vdisk. We recommend monitoring the number of replication images created.

Appendix C: HP MSA 2040 Disk Background Scrub

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

Vdisk-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-stripped) and RAID 0, the utility checks for media errors.

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

A new Vdisk will first be scrubbed 20 minutes after creation and again after the interval specified by the Vdisk 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 Vdisks and available disks

Scheduling drive spin down for all vdisks

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

Best Practice: Set DSD for unconfigured disks, spares, and configured Vdisks 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 MSA array, using background disk scrub can have an impact on performance.

For more information

HP Storage Arrays: [HP Storage](#)

MSA 2040 home page: hp.com/go/msa2040

For a list of QuickSpecs, go to: h18000.www1.hp.com/products/quickspecs/13551_div/13551_div.html

For all HP MSA 2040 manuals and user guides, go to: hp.com/go/msa2040

Storage Networking: hp.com/go/san

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

VMware

[VMware vSphere Storage/SAN guide](#)

[VMware Compatibility Guide – Storage/SAN](#)

[HP Servers](#)

HP Virtualization with VMware:

hp.com/go/vmware

VMware Storage Solutions from HP:

hp.com/go/storage/vmware

[HP Unified Sizer for Server Virtualization](#)

[HP Insight Control Storage Module for VMware vCenter Technical Documentation](#)

[HP Insight Control Storage Module for VMware vCenter Download](#)

To help us improve our documents, please provide feedback at hp.com/solutions/feedback.

Sign up for updates
hp.com/go/getupdated



Share with colleagues



Rate this document

