



Table of contents

About this document	3
Intended audience	3
Prerequisites	3
Related documentation	3
Introduction	3
General best practices	4
Become familiar with the array by reading the manuals	4
Stay current on firmware	4
Use tested and supported configurations	4
Understand what a host is from the array perspective	4
Rename hosts to a user friendly name	4
Vdisk initialization	5
Best practice for monitoring array health	5
Configure email, SNMP, and Syslog notifications	5
Setting the notification level for email, SNMP, and Syslog	6
Sign up for proactive notifications for the HP MSA 2040 array	6
Best practices when choosing drives for HP MSA 2040 Storage	7
Drive types	7
Best practices to improve availability	7
Volume mapping	7
Redundant paths	7
Multipath software	8
Dual power supplies	8
Dual controllers	8
Reverse cabling of expansion enclosures	8
Create vdisks across expansion enclosures	9
Drive sparing	9
Best practices to enhance performance	10
Cache settings	10
Other methods to enhance array performance	12

- Best practices for SSDs 14
 - Use SSDs for randomly accessed data..... 14
 - SSD and performance 14
 - SSD wear gauge 15
- Best practices for virtual disk expansion 15
 - Vdisk expansion capability for supported RAID levels 15
 - Vdisk expansion recommendations..... 16
 - Re-create the vdisk with additional capacity and restore data 17
- Best practices for firmware updates..... 17
 - General MSA 2040 device firmware update best practices 17
 - MSA 2040 array controller or I/O module firmware update best practices 17
 - MSA 2040 disk drive firmware update best practices 18
- Miscellaneous best practices 18
 - Boot from storage considerations 18
 - 8 Gb/16 Gb switches and small form-factor pluggable transceivers 18
 - MSA 2040 iSCSI considerations..... 18
 - IP address scheme for the controller pair 19
- Summary 20

About this document

This white paper highlights the best practices for optimizing the HP MSA 2040, and should be used in conjunction with other HP Modular Smart Array manuals. Modular smart array (MSA) technical user documentations can be found at hp.com/go/msa2040.

Intended audience

This white paper is intended for entry-level and mid-range HP MSA 2040 administrators with previous storage area network (SAN) knowledge. It offers MSA practices that can contribute to an MSA best customer experience.

This paper is also designed to assist MSA users for best practices in the deployment of the HP MSA 2040 array.

Prerequisites

Prerequisites for using this product include knowledge of:

- Networking
- Storage system configuration
- SAN management
- Connectivity methods such as direct attach storage (DAS), Fibre Channel, and serial attached SCSI (SAS)
- Internet SCSI (iSCSI) and Ethernet protocols

Related documentation

In addition to this guide, please refer to other documents for this product:

- HP MSA System Racking Instructions
- HP MSA 2040 Installation Guide
- HP MSA 2040 System Cable Configuration Guide
- HP MSA 2040 User Guide
- HP MSA 2040 SMU Reference Guide
- HP MSA 2040 CLI Reference Guide
- HP MSA 2040 Troubleshooting Guide

You can find these documents on the HP MSA 2040 home page: hp.com/go/msa2040

Introduction

The HP MSA 2040, a high-performance storage array designed for entry-level HP customers desiring 8 and/or 16 Gb Fibre Channel, 6 Gb SAS and/or 12 Gb* SAS, and 1GbE and/or 10GbE iSCSI connectivity with 4 host ports per controller. This next generation MSA 2040 Storage array provides an excellent value for customers needing performance balanced with price to support initiatives such as consolidation and virtualization. The MSA 2040 delivers this performance by offering:

- New controller architecture with a new processor
- 4 GB cache per controller
- Support for solid state drives (SSDs)
- 4 host ports per controller
- 4 Gb/8 Gb/16 Gb FC connectivity
- 6 Gb/12 Gb* SAS
- 1GbE/10GbE iSCSI

*12 Gb SAS in a future release

The HP MSA 2040 Storage System brings the performance benefits of SSDs to MSA array family customers. This array has been designed to maximize performance by using high-performance drives across all applications sharing the array.

The HP MSA 2040 Storage arrays are positioned to provide an excellent value for customers needing increased performance to support initiatives such as consolidation and virtualization.

The HP MSA 2040 Storage ships standard with a license for 64 Snapshots and Volume Copy for increased data protection. There is also an optional license for 512 Snapshots. The HP MSA 2040 can also replicate data between arrays (P2000 G3 and/or MSA 2040 SAN model only using FC or iSCSI protocols) with the optional Remote Snap feature.

General best practices

Become familiar with the array by reading the manuals

The first recommended best practice is to read the HP MSA 2040 User Guide and the HP MSA 2040 SMU Reference Guide or the HP MSA 2040 Command Line Interface (CLI) Reference Guide, depending on the interface you will use to configure the array. Always operate the array in accordance with the user manual. In particular, never exceed the environmental operation requirements.

Other HP MSA 2040 guides of importance to review are:

- The HP MSA 2040 Troubleshooting Guide located at: hp.com/support/msa2040/Troubleshooting
- The MSA Remote Snap Guide can be found at: h20195.www2.hp.com/v2/GetPDF.aspx/4AA1-0977ENW.pdf

Stay current on firmware

Use the latest controller, disk, and expansion enclosure firmware to benefit from the continual improvements in the performance, reliability, and functionality of the HP MSA 2040. For additional information, see the release notes and release advisories for the respective MSA products.

This information can be located at: hp.com/go/msa2040

Use tested and supported configurations

Deploy the MSA array only in supported configurations. Do not risk the availability of your critical applications to unsupported configurations. HP does not recommend nor provide HP support for unsupported MSA configurations.

HP's primary portal used to obtain detailed information about supported HP Storage product configurations is single point of connectivity knowledge (SPOCK). An HP Passport account is required to enter the SPOCK website.

SPOCK can be located at: hp.com/storage/spock

Understand what a host is from the array perspective

From the perspective of the MSA 2040 array, each individual unique initiator is considered a host. A host is analogous to an external port on a host bus adapter (HBA). A host port does not equate to a physical server, but rather a unique host connection on that server. For example, a dual port FC HBA has two ports and therefore there are two unique initiators.

Rename hosts to a user friendly name

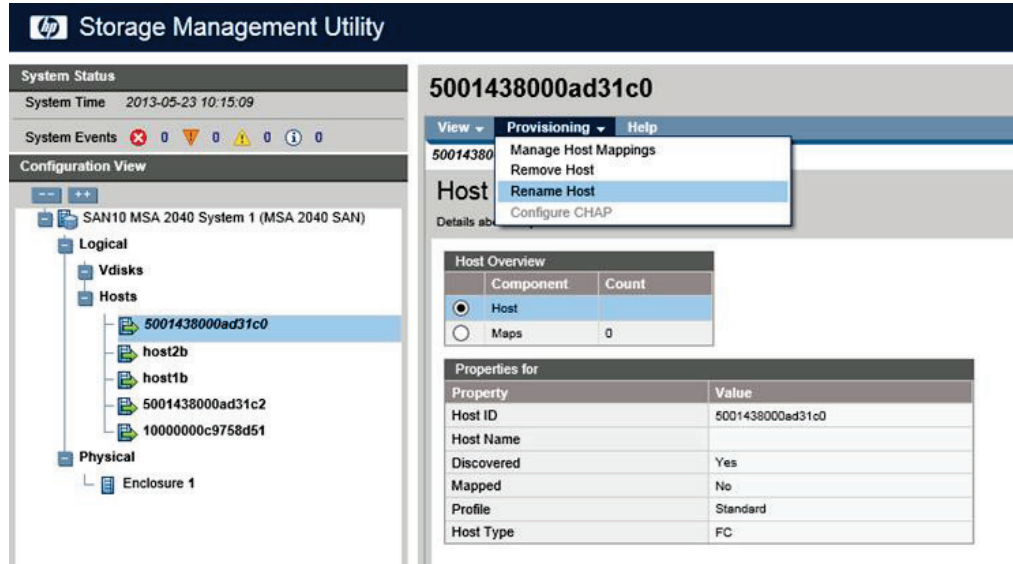
Applying friendly names to the hosts enables easy identification of which hosts are associated with servers and operating systems. It also allows for applying any special host profiles that may need to be set, such as HP-UX. The best practice for acquiring and renaming Worldwide Name (WWN) is to connect one cable at a time and then rename the WWN to an identifiable name.

The procedure below outlines the steps needed to rename hosts.

1. Log into the SMU and click "+" next to "Hosts" from the left frame. This will expand the list to show all connected hosts.
2. Highlight the host in the list that you want to rename by clicking the WWN name.
3. On the right window frame, click Provisioning -> Rename Host.
4. Type in the host nickname and choose the Profile and then click Modify Name.

Repeat for additional host connections.

Figure 1. Renaming hosts




Vdisk initialization

During the creation of a virtual disk (vdisk), the user has the option to create a vdisk in online mode (default) or offline mode.

If the “online initialization” option is enabled, you can use the vdisk while it is initializing. Online initialization takes more time because parity initialization is used during the process to initialize the vdisk. Online initialization is supported for all HP MSA 2040 RAID levels except for RAID 0 and NRAID. Online initialization does not impact fault tolerance.

If the “online initialization” option is unchecked, which equates to “offline initialization,” you must wait for initialization to complete before using the vdisk, but the initialization takes less time to complete.

Figure 2. Choosing online or offline initialization



Best practice for monitoring array health

Setting up the array to send notifications is important for troubleshooting and log retention.

Configure email, SNMP, and Syslog notifications

The storage management utility (SMU) is the recommended method for setting up email, SNMP, and Syslog notifications. Setting up these services is easily accomplished by using a Web browser; to connect; type in the IP address of the management port of the HP MSA 2040.

Email notifications can be sent to up to as many as three different email addresses. In addition to the normal email notification, enabling managed logs notifications, with the “Include Logs” option enabled is recommended. When the “Include Logs” feature is enabled, the system automatically attaches the system log files to the managed logs email notifications sent. The managed logs email notification is sent to an email address which will retain the logs for future diagnostic investigation.

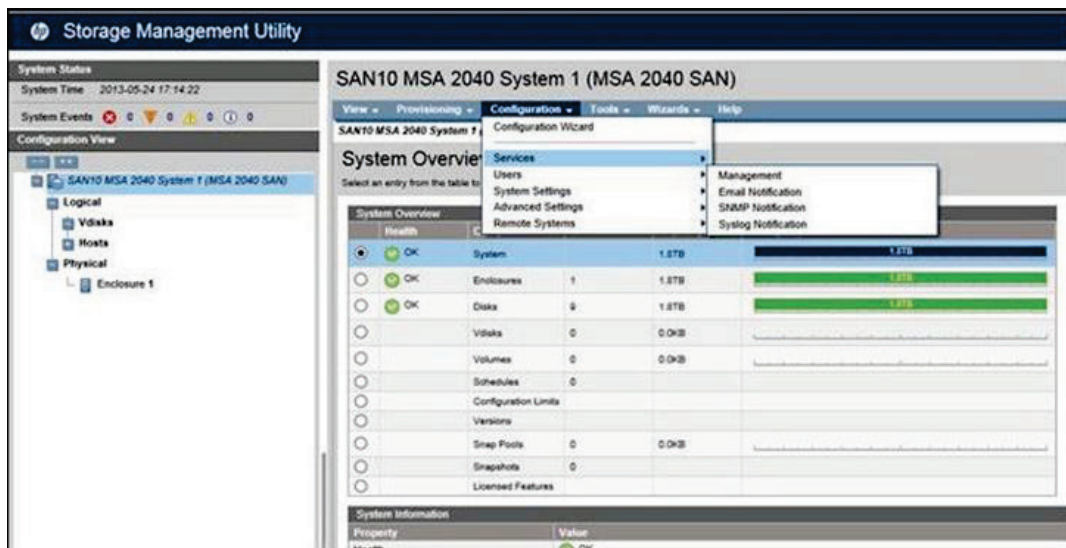
The MSA 2040 system has a limited amount of space to retain logs. When this log space is exhausted, the oldest entries in the log are overwritten. For most systems this space is adequate to allow for diagnosing issues seen on the system. The managed logs feature notifies the administrator that the logs are nearing a full state and that older information will soon start to get overwritten. The administrator can then choose to manually save off the logs. If “Include Logs” is also checked, the segment of logs which is nearing a full state will be attached to the email notification. Managed logs attachments can be multiple MB in size. In a typical system a managed log event should only occur every few weeks. If managed logs events are occurring more frequently an investigation of the system health should be done.

Enabling the managed logs feature allows log files to be transferred from the storage system to a log-collection system to avoid losing diagnostic data. The “Include Logs” option is disabled by default.

HP recommends enabling SNMP traps. SNMP traps can be sent to up to three host trap addresses (i.e., HP SIM Server or other SNMP server). SNMP traps can be useful in troubleshooting issues with the MSA 2040 array.

To configure these settings in the SMU, click **Configuration -> Services**. Enter the correct information for email, SNMP, and Syslog.

Figure 3. Management services



Setting the notification level for email, SNMP, and Syslog

Setting the notification level to Warning, Error, or Critical on the email, SNMP, and Syslog configurations will ensure that events of that level or above are sent to the destinations (i.e., SNMP server, SMTP server) set for that notification. HP recommends setting the notification level to Warning or above.

HP MSA 2040 notification levels:

- Warning will send notifications for all Warning, Error, or Critical events.
- Error will only send Error and Critical events.
- Critical will only send Critical events.

Sign up for proactive notifications for the HP MSA 2040 array

Sign up for proactive notifications to receive MSA product advisories. Applying the suggested resolutions can enhance the availability of the product.

Sign up for the notifications at: hp.com/go/myadvisory

Best practices when choosing drives for HP MSA 2040 Storage

The characteristics of applications and workloads are important when selecting drive types for the HP MSA 2040 array.

Drive types

The HP MSA 2040 array supports SSDs, SAS enterprise drives, and SAS Midline (MDL) drives. The HP MSA 2040 array does not support Serial ATA (SATA) drives. Choosing the correct drive type is important; drive types should be selected based on the workload and performance requirements of the volumes that will be serviced by the array. For sequential workloads, SAS enterprise drives or SAS MDL drives provide a good price-for-performance tradeoff. SAS enterprise drives offer higher performance than SAS MDL and should also be considered for random workloads when performance is not a premium. For high performance random workloads, SSDs would be appropriate.

SAS MDL drives are not recommended for constant high workload applications. SAS MDL drives are intended for archival purposes.

Best practices to improve availability

There are many methods to improve availability when using the HP MSA 2040 array. High availability is always advisable to protect your assets in the event of a device failure. Outlined below are some options that will help you in the event of a failure.

Volume mapping

Using volume mapping correctly can provide high availability from the hosts to the array. For high availability during a controller failover, a volume must be mapped to at least one port accessible by the host on both controllers. Mapping a volume to ports on both controllers ensures that at least one of the paths is available in the event of a controller failover, thus providing a preferred/optimal path to the volume.

In the event of a controller failover, the surviving controller will report that it is now the preferred path for all vdisks. When the failed controller is back online, the vdisks and preferred paths switch back to the original owning controller.

Best practice is to map volumes to two ports on each controller to take advantage of load balancing and redundancy to each controller.

It is not recommended to enable more than 8 paths to a single host, i.e., 2 HBA ports on a physical server connected to 2 ports on the A controller and 2 ports on the B controller. Enabling more paths from a host to a volume puts additional stress on the operating system's multipath software which can lead to delayed path recovery in very large configurations.

Note:

In the SMU when a new volume is created, the volume mapping defaults to the "all other hosts read-write access" mapping (known as default mapping). Please refer to the HP MSA 2040 SMU Reference Guide for directions related to explicit and default mapping practices.

Note:

Volumes should not be mapped to multiple servers at the same time unless the operating systems on the servers are cluster aware. However, since a server may contain multiple unique initiators, mapping a volume to multiple unique initiators (that are contained in the same server) is supported and recommended.

Redundant paths

To increase the availability of the array to the hosts, multiple, redundant paths should be used along with multipath software. Redundant paths can also help in increasing performance from the array to the hosts (discussed later in this paper). Redundant paths can be accomplished in multiple ways. In the case of a SAN attach configuration, best practice would be to have multiple, redundant switches (SANs) with the hosts having at least one connection into each switch (SAN), and the array having one or more connections from each controller into each switch. In the case of a direct attach configuration, best practice is to have at least two connections to the array for each server. In the case of a direct attach configuration with dual controllers, best practice would be to have at least one connection to each controller.

Multipath software

To fully utilize redundant paths, multipath software should be installed on the hosts. Multipath software allows the host operating system to use all available paths to volumes presented to the host; redundant paths allow hosts to survive SAN component failures. Multipath software can increase performance from the hosts to the array. Table 1 lists supported multipath software by operating systems.

Note:

More paths are not always better. Enabling more than 8 paths to a single volume is not recommended.

See the HP MSA 2040 MPIO Guide:

bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c02021677/c02021677.pdf

Table 1. Multipath and operating systems

Operating system	Multipath name	Vendor ID	Product ID
Windows® 2008/2012	Microsoft® multipath I/O (MPIO)	HP	MSA 2040 SAN
Linux	Device mapper/multipath	HP	MSA 2040 SAN
VMware	Native multipath (NMP)	HP	MSA 2040 SAN

Dual power supplies

The HP MSA 2040 chassis and supported expansion enclosures ship with dual power supplies. At a minimum, connect both power supplies in all enclosures. For the highest level of availability, connect the power supplies to separate power sources.


Dual controllers

The HP MSA 2040 can be purchased as a single or dual controller system. Utilizing a dual controller system is best practice for increased reliability for two reasons. First, dual controller systems will allow hosts to access volumes during a controller failure or during firmware upgrades (given correct volume mapping discussed above). Second, if the expansion enclosures are cabled correctly, a dual controller system can withstand an expansion IO Module (IOM) failure, and in certain situations a total expansion enclosure failure.

Reverse cabling of expansion enclosures

The HP MSA 2040 firmware supports both fault tolerant (reverse cabling) and straight-through SAS cabling of expansion enclosures. Fault tolerant cabling allows any expansion enclosure to fail or be removed without losing access to other expansion enclosures in the chain. For the highest level of fault tolerance, use fault tolerant (reverse) cabling when connecting expansion enclosures.

Figure 4. Reverse cabling example



See the [HP MSA 2040 System Cable Configuration Guide](#) for more details on cabling the HP MSA 2040.

Create vdisks across expansion enclosures

HP recommendation is to stripe virtual disks across shelf enclosures to enable data integrity in the event of an enclosure failure. A virtual disk created with RAID 1, 10, 3, 5, 50, or 6 can sustain one or more expansion enclosure failures without loss of data depending on RAID type. Vdisk configuration should take into account MSA drive sparing methods such as dedicated, global, and dynamic sparing.

Drive sparing

Drive sparing, sometimes referred to as hot spares, is recommended to help protect data in the event of a disk failure in a fault tolerant vdisk (RAID 1, 3, 5, 6, 10, or 50) configuration. In the event of a disk failure, the array automatically attempts to reconstruct the data from the failed drive to a compatible spare. A compatible spare is defined as a drive that has sufficient capacity to replace the failed disk and is the same media type (i.e., SAS SSD, Enterprise SAS, or Midline SAS). The HP MSA 2040 supports dedicated, global, and dynamic sparing. The HP MSA 2040 will reconstruct a critical or degraded vdisk.

Important

An offline or quarantined vdisk is not protected by sparing.

Supported spare types:

- **Dedicated spare**—reserved for use by a specific vdisk to replace a failed disk. This method is the most secure way to provide spares for vdisks. The array supports up to 4 dedicated spares per vdisk.
- **Global spare**—reserved for use by any fault-tolerant vdisk to replace a failed disk. The array supports up to 16 global spares per system. At least one vdisk must exist before you can add a global spare.
- **Dynamic spare**—all available drives are available for sparing. If the MSA has available drives and a vdisk becomes degraded any available drive can be used for vdisk reconstruction.

Sparing process

When a disk fails in a redundant vdisk, the system first looks for a dedicated spare for the vdisk. If a dedicated spare is not available or the disk is incompatible, the system looks for any compatible global spare. If the system does not find a compatible global spare and the dynamic spares option is enabled, the system uses any available compatible disk for the spare. If no compatible disk is available, reconstruction cannot start.

During reconstruction of data, the effected vdisk will be in either a degraded or critical status until the parity or mirror data is completely written to the spare, at which time the vdisk returns to fault tolerant status. For RAID 50 vdisks, if more than one sub-vdisk becomes critical, reconstruction and use of spares occurs in the order sub-vdisks are numbered. In the case of dedicated spares and global spares, after the failed drive is replaced, the replacement drive will need to added back as a dedicated or global spare.

Best practice for sparing is to configure at least one dedicated spare for every fault tolerant vdisk in the system.

Drive replacement

In the event of a drive failure, replace the failed drive with a compatible drive as soon as possible. As noted above, if dedicated or global sparing is in use, mark the new drive as a spare (either dedicated or global), so it can be used in the future for any other drive failures.

Implement Remote Snap replication

The HP MSA 2040 Array System Remote Snap feature is a form of asynchronous replication that replicates block-level data from a volume on a local system to a volume on the same system or on a second independent system. The second system may be at the same location as the first, or it may be located at a remote site.

Best practice is to implement Remote Snap replication for disaster recovery.

Note:

Remote Snap requires a purchasable license in order to implement.

To obtain a Remote Snap license, go to:

h18004.www1.hp.com/products/storage/software/p2000rs/index.html

See the MSA 2040 Remote Snap Guide:

h20195.www2.hp.com/v2/GetPDF.aspx/4AA1-0977ENW.pdf

Use VMware Site Recovery Manager with Remote Snap replication

VMware vCenter Site Recovery Manager (SRM) is an extension to VMware vCenter that delivers business-continuity and disaster-recovery solution that helps you plan, test, and execute the recovery of vCenter virtual machines. SRM can discover and manage replicated datastores, and automate migration of inventory from one vCenter to another. Site Recovery Manager integrates with the underlying replication product through a storage replication adapter (SRA).

For best practices with SRM and MSA Remote Snap replication, see the “Integrate VMware vCenter SRM with HP MSA Storage” technical white paper: h20195.www2.hp.com/v2/GetPDF.aspx/4AA4-3128ENW.pdf

Best practices to enhance performance

This section outlines configuration options for enhancing performance for your array.

Cache settings

One method to tune the array is by choosing the correct cache settings for your volumes. Controller cache options can be set for individual volumes to improve a volume’s I/O performance.

Caution:

Only disable write-back caching if you fully understand how the host operating system, application, and adapter move data. If used incorrectly, you might hinder system performance.

Using write-back or write-through caching

By default, volume write-back cache is enabled. Because controller cache is backed by super-capacitor technology, if the system loses power, data is not lost. For most applications, write-back caching enabled is the best practice. With the transportable cache feature, write-back caching can be used in either a single or dual controller system.

You can change a volume's write-back cache setting. Write-back is a cache-writing strategy in which the controller receives the data to be written to disks, stores it in the memory buffer, and immediately sends the host operating system a signal that the write operation is complete, without waiting until the data is actually written to the disk. Write-back cache mirrors all of the data from one controller module cache to the other. Write-back cache improves the performance of write operations and the throughput of the controller. This is especially true in the case of random I/O, where write-back caching allows the array to coalesce the I/O to the vdisks.

When write-back cache is disabled, write-through becomes the cache-writing strategy. Using write-through cache, the controller writes the data to the disks before signaling the host operating system that the process is complete. Write-through cache has lower write operation and throughput performance than write-back, but all data is written to non-volatile storage before confirmation to the host. However, write-through cache does not mirror the write data to the other controller cache because the data is written to the disk before posting command completion and cache mirroring is not required. You can set conditions that cause the controller to change from write-back caching to write-through caching. Please refer to the HP MSA 2040 User Guide for ways to set the auto write through conditions correctly. In most situations, the default settings are acceptable.

In both caching strategies, active-active failover of the controllers is enabled.

Caution:

Only change read-ahead cache settings if you fully understand how the host operating system, application, and adapter move data so that you can adjust the settings accordingly.

Optimizing read-ahead caching

You can optimize a volume for sequential reads or streaming data by changing its read ahead, cache settings. Read ahead is triggered by sequential accesses to consecutive LBA ranges. Read ahead can be forward (that is, increasing LBAs) or reverse (that is, decreasing LBAs). Increasing the read-ahead cache size can greatly improve performance for multiple sequential read streams. However, increasing read-ahead size will likely decrease random read performance.

- **Adaptive**—this option works well for most applications: it enables adaptive read-ahead, which allows the controller to dynamically calculate the optimum read-ahead size for the current workload. This is the default.
- **Stripe**—this option sets the read-ahead size to one stripe. The controllers treat non-RAID and RAID 1 vdisks internally as if they have a stripe size of 512 KB, even though they are not striped.
- **Specific size options**—these options let you select an amount of data for all accesses.
- **Disabled**—this option turns off read-ahead cache. This is useful if the host is triggering read ahead for what are random accesses. This can happen if the host breaks up the random I/O into two smaller reads, triggering read ahead.

Optimizing cache modes

You can also change the optimization mode for each volume.

- **Standard**—this mode works well for typical applications where accesses are a combination of sequential and random; this method is the default. For example, use this mode for transaction-based and database update applications that write small files in random order.
- **No-mirror**—in this mode each controller stops mirroring its cache metadata to the partner controller. This improves write I/O response time but at the risk of losing data during a failover. Unified LUN presentation (ULP) behavior is not affected, with the exception that during failover any write data in cache will be lost. In most conditions No-mirror is not recommended, and should only be used after careful consideration.

Parameter settings for performance optimization

You can configure your storage system to optimize performance for your specific application by setting the parameters as shown in table 2. This section provides a basic starting point for fine-tuning your system, which should be done during performance baseline modeling.

Table 2. Optimizing performance for your application

Application	RAID level	Read-ahead cache size	Cache write optimization
Default	5 or 6	Adaptive	Standard
High-Performance Computing (HPC)	5 or 6	Adaptive	Standard
Mail spooling	1	Adaptive	Standard
NFS_Mirror	1	Adaptive	Standard
Oracle_DSS	5 or 6	Adaptive	Standard
Oracle_OLTP	5 or 6	Adaptive	Standard
Oracle_OLTP_HA	10	Adaptive	Standard
Random 1	1	Stripe	Standard
Random 5	5 or 6	Stripe	Standard
Sequential	5 or 6	Adaptive	Standard
Sybase_DSS	5 or 6	Adaptive	Standard
Sybase_OLTP	5 or 6	Adaptive	Standard
Sybase_OLTP_HA	10	Adaptive	Standard
Video streaming	1 or 5 or 6	Adaptive	Standard
Exchange database	5 for data; 10 for logs	Adaptive	Standard
SAP	10	Adaptive	Standard
SQL	5 for data; 10 for logs	Adaptive	Standard

Other methods to enhance array performance

There are other methods to enhance performance of the HP MSA 2040. In addition to the cache settings, the performance of the HP MSA 2040 array can be maximized by using the following techniques.

Place higher performance SSD and SAS drives in the array enclosure

The HP MSA 2040 controller is designed to have a single SAS link per drive in the array enclosure and only four SAS links to expansion enclosures. Placing higher performance drives (i.e., SSD and Enterprise SAS drives) in the array enclosure allows the controller to utilize the performance of those drives more effectively than if they were placed in expansion enclosures. This process will help generate better overall performance.

Fastest throughput optimization

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

- Host ports should be configured to match the highest speed your infrastructure supports.
- Virtual disks should be balanced between the two controllers.
- Disk drives should be balanced between the two controllers.
- Cache settings should be set to match table 2 (“Optimizing performance for your application”) for the application.
- In order to get the maximum sequential performance from a vdisk, you should only create one volume per vdisk. Otherwise you will introduce randomness into the workload when multiple volumes on the vdisk are being exercised concurrently.
- Distribute the load across as many drives as possible.
- Distribute the load across multiple array controller host ports.

Creating virtual disks

When creating virtual disks, best practice is to add them evenly across both controllers. With at least one virtual disk assigned to each controller, both controllers are active. This active-active controller configuration allows maximum use of a dual-controller configuration’s resources.

Choosing the appropriate RAID levels

Choosing the correct RAID level when creating virtual disks can be important for performance. However, there are some trade-offs with cost when using the higher fault tolerant RAID levels.

See table 3 below for the strengths and weaknesses of the supported HP MSA 2040 RAID types.

Table 3. HP MSA 2040 RAID levels

RAID level	Minimum disks	Allowable disks	Description	Strengths	Weaknesses
NRAID	1	1	Non-RAID, non-striped mapping to a single disk	Ability to use a single disk to store additional data	Not protected, lower performance (not striped)
0	2	16	Data striping without redundancy	Highest performance	No data protection: if one disk fails all data is lost
1	2	2	Disk mirroring	Very high performance and data protection; minimal penalty on write performance; protects against single disk failure	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required
3	3	16	Block-level data striping with dedicated parity disk	Excellent performance for large, sequential data requests (fast read); protects against single disk failure	Not well-suited for transaction-oriented network applications; write performance is lower on short writes (less than 1 stripe)
5	3	16	Block-level data striping with distributed parity	Best cost/performance for transaction-oriented networks; very high performance and data protection; supports multiple simultaneous reads and writes; can also be optimized for large, sequential requests; protects against single	Write performance is slower than RAID 0 or RAID 1
6	4	16	Block-level data striping with double distributed parity	Best suited for large sequential workloads; non-sequential read and sequential read/write performance is comparable to RAID 5; protects against dual disk failure	Higher redundancy cost than RAID 5 because the parity overhead is twice that of RAID 5; not well-suited for transaction-oriented network applications; non-sequential write performance is slower than RAID 5
10 (1+0)	4	16	Stripes data across multiple RAID 1 sub-vdisks	Highest performance and data protection (protects against multiple disk failures)	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required; requires minimum of four disks
50 (5+0)	6	32	Stripes data across multiple RAID 5 sub-vdisks	Better random read and write performance and data protection than RAID 5; supports more disks than RAID 5; protects against multiple disk failures	Lower storage capacity than RAID 5

Volume mapping

For increased performance, map volumes to the ports on the controller that owns the vdisk. Mapping to the non-preferred path results in a slight performance degradation.

Optimum performance with MPIO can be achieved with volumes mapped to multiple paths on both controllers. When the appropriate MPIO drivers are installed on the host, only the preferred (optimized) paths will be used. The non-optimized paths will be reserved for failover.

Best practices for SSDs

SSDs are supported in the MSA 2040 and with their performance capabilities, SSDs are a great alternative to traditional spinning hard disk drives (HDD). SSDs cost more \$/GB than spinning hard drives; however, SSDs cost much less in \$/IOPS. Keep this in mind when choosing the numbers of SSDs per MSA 2040 array.

Use SSDs for randomly accessed data

The use of SSDs can greatly enhance the performance of the array. Since there are no moving parts in the drives, data that is random in nature can be accessed much faster.

Data such as database indexes and TempDB files would best be placed on a volume made from an SSD based vdisk since this type of data is accessed randomly.

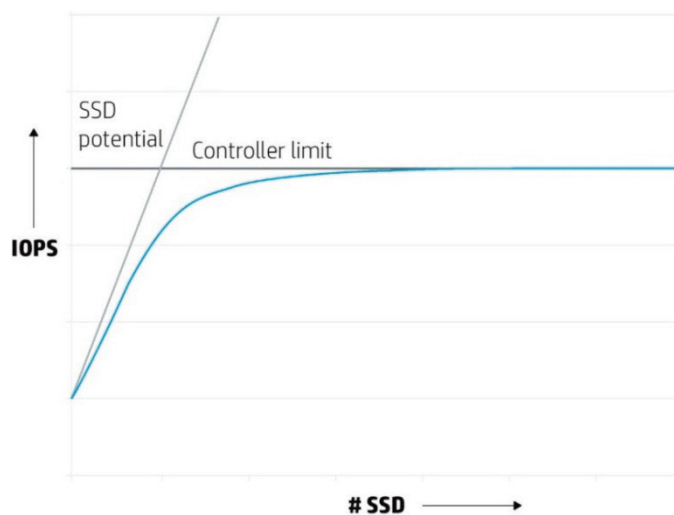
Another good example of a workload that would benefit from the use of SSDs is desktop virtualization, for example, virtual desktop infrastructure (VDI) where boot storms require high performance with low latency.

SSD and performance

There are some performance characteristics which can be met with linear scaling of SSDs. There are also bandwidth limits in the MSA 2040 controllers. There is a point where these two curves intersect. At the intersecting point, additional SSDs will not increase performance. See figure 5.

The MSA 2040 reaches this bandwidth at a low number of SSDs. For the best performance using SSDs on the MSA 2040, use at least a minimum of 4 SSDs with 1 mirrored pair of drives (RAID 1) per controller. RAID 5 and RAID 6 are also good choices for SSDs, but require more drives using the best practice of having one vdisk owned by each controller. This would require 6 SSDs for RAID 5 and 8 SSDs for RAID 6. All SSD volumes should be contained in fault tolerant vdisks for data integrity.

Figure 5. SSD performance potential vs. MSA 2040 controller limit



Note:

The HP MSA 2040 array is not limited by the number of SSDs in the system

SSD wear gauge

SSDs have a limited number of times they can be written and erased due to the memory cells on the drives. The SSDs in the HP MSA 2040 come with a wear gauge as well as appropriate events that are generated to help detect the failure. Once the wear gauge reaches 0%, the integrity of the data is not guaranteed. Best practice is to replace the SSD when the events and gauge indicate <5% life remaining to prevent data integrity issues.

Best practices for virtual disk expansion

With the ever changing storage needs seen in the world today, there comes a time when storage space gets exhausted quickly. The HP MSA 2040 gives you the option to grow the size of a LUN to keep up with your dynamic storage needs.

A vdisk expansion allows you to grow the size of a vdisk in order to expand an existing volume or create volumes from the newly available space on the vdisk. Depending on several factors, vdisk expansion can take a significant amount of time to complete. For faster alternatives, see the “vdisk expansion recommendations” section on page 16.

The factors that should be considered with respect to vdisk expansion include but are not limited to:

- Physical disk size
- Number of disks to expand (1-4)
- I/O activity during vdisk expansion

Note:

During vdisk expansion, other disk utilities are disabled. These utilities include vdisk Scrub and Rebuild.

Vdisk expansion capability for supported RAID levels

The chart below gives information on the expansion capability for the HP MSA 2040 supported RAID levels.

Expansion capability for each RAID level

RAID level	Expansion capability	Maximum disks
NRAID	Cannot expand	1
0, 3, 5, 6	Can add 1–4 disks at a time	16
1	Cannot expand	2
10	Can add 2 or 4 disks at a time	16
50	Can expand the vdisk one RAID 5 sub-vdisk at a time. The added RAID 5 sub-vdisk must contain the same number of disks as each original sub-vdisk	32

Important:

If during the process of a vdisk expansion one of the disk members of the vdisk fails, the reconstruction of the vdisk will not commence until the expansion is complete. During this time, data is at risk with the vdisk in a DEGRADED or CRITICAL state.

If an expanding vdisk becomes DEGRADED (e.g., RAID 6 with a single drive failure) the storage administrator should determine the level of risk of continuing to allow the expansion to complete versus the time required to backup, re-create the vdisk (see page 16 “vdisk expansion recommendations”) and restore the data to the volumes on the vdisk.

If an expanding vdisk becomes CRITICAL (e.g., RAID 5 with a single drive failure) the storage administrator should immediately employ a backup and recovery process. Continuing to allow the expansion places data at risk of another drive failure and total loss of all data on the vdisk.

Vdisk expansion can be very time consuming. There is no way to reliably determine when the expansion will be complete and when other disk utilities will be available.

Follow the procedure below.

1. Backup the current data from the existing vdisk.
2. Using the WBI or CLI, start the vdisk expansion.
3. Monitor the vdisk expansion percentage complete.

Note:

Once a vdisk expansion initiates it will continue until completion or until the vdisk is deleted.

Vdisk expansion recommendations

Before expanding a vdisk, review the information below to understand the best alternative method for allocating additional storage to hosts.

Allocate “quiet” period(s) to help optimize vdisk expansion

Vdisk expansion can take a few hours with no data access for smaller capacity hard drives and may take several days to complete with larger capacity hard drives. This procedure can be paused by reconnecting the host side cables and restarting hosts. Priority is given to host I/O or data access over the expansion process during normal array operation. While the system is responding to host I/O or data access requests, it may seem as if the expansion process has stopped. When expanding during “quiet” periods, expansion time is minimized and will allow quicker restoration of other disk utilities.

This method of expansion utilizes the expand capability of the system and requires manual intervention from the administrator. The procedure below outlines the steps to expand a vdisk during a “quiet” period.

In this context, a “quiet” period indicates a length of time when there is no host I/O or data access to the system. Before starting the vdisk expansion:

1. Stop I/O to existing volumes on the vdisk that will be expanded.
2. Backup the current data from the existing volumes on the vdisk.
3. Shutdown all hosts connected to the HP MSA 2040 system.
4. Label and disconnect host side cables from the HP MSA 2040 system.

Start and monitor vdisk expansion

1. Using the WBI or CLI, start the vdisk expansion.
2. Monitor the vdisk expansion percentage complete.

When expansion is complete or data access needs to be restored:

1. Re-connect host side cables to the HP MSA 2040 system.
2. Restart hosts connected to the HP MSA 2040 system.

If additional “quiet” periods are required to complete the vdisk expansion:

1. Shutdown all hosts connected to the HP MSA 2040 system.
2. Label and disconnect host side cables from the HP MSA 2040 system.
3. Monitor the vdisk expansion percentage complete.

Re-create the vdisk with additional capacity and restore data

This method is the easiest and fastest method for adding additional capacity to a vdisk. The online vdisk initialization allows a user to access the vdisk almost immediately and will complete quicker than the expansion process on a vdisk that is also servicing data requests. The procedure below outlines the steps for recreating a vdisk with additional capacity and restoring data to that vdisk.

Procedure:

1. Stop I/O to existing volumes on the vdisk that will be expanded.
2. Backup the current data from the existing volumes on the vdisk.
3. Delete the current vdisk.
4. Using the WBI or CLI, create a new vdisk with the available hard drives using online initialization.
5. Create new larger volumes as required.
6. Restore data to the new volumes.

Best practices for firmware updates

The sections below detail common firmware update best practices for the MSA 2040.

General MSA 2040 device firmware update best practices

- As with any other firmware upgrade, it is a recommended best practice to ensure that you have a full backup prior to 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 configurations 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 the new firmware version is displayed correctly via one of the MSA management interfaces—e.g., SMU or CLI.

MSA 2040 array controller or I/O module 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 controller systems, I/O must be halted.
 - For dual controller 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 avoid 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, ensure 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 will automatically first disable partner firmware update (PFU) and then perform 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).

MSA 2040 disk drive firmware update best practices

- Disk drive upgrades on the HP MSA 2040 storage systems is an off line process. All host and array I/O must be stopped prior to 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 is 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 pre-requisites 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 all the drives are updated.

Miscellaneous best practices

Boot from storage considerations

When booting from SAN, construct a separate virtual disk and volume that will be used only for the boot from SAN. Do not keep data and boot from SAN volumes on the same vdisk. This can help with performance. If there is a lot of I/O going to the data volume on a vdisk that shares a boot from SAN volume, there can be a performance drop in the I/O to the operating system drives.

8 Gb/16 Gb switches and small form-factor pluggable transceivers

The HP MSA 2040 uses specific small form-factor pluggable (SFP) transceivers which are not qualified in the HP 8 Gb and 16 Gb switches. Likewise, the HP Fibre Channel switches use SFPs which will not operate in the HP MSA 2040.

The HP MSA 2040 controllers do not include SFPs. Qualified SFPs for the HP MSA 2040 are available for separate purchase in 4 packs. Both 8G and 16G SFPs are available to meet the customer need and budget constraints. All SFPs in an HP MSA 2040 should conform to the installation guidelines given in the product QuickSpecs. SFP speeds and protocols can be mixed, but only in the specified configurations.

In the unlikely event of an HP MSA 2040 controller or SFP failure, a field replacement unit (FRU) is available. SFPs will need to be moved from the failed controller to the replacement controller.


Please see the HP Transceiver Replacement Instructions document for details found at hp.com/support/msa2040/manuals.

MSA 2040 iSCSI considerations

When using the MSA 2040 SAN controller in an iSCSI configuration, it is a best practice to use at least three network ports per server, two for the storage (Private) LAN and one or more for the Public LAN(s). This will ensure that the storage network is isolated from the other networks.

The Private LAN is the network that goes from the server to the MSA 2040 SAN controller. This Private LAN is the storage network and the Public LAN is used for management of the MSA 2040. The storage network should be isolated from the Public LAN to improve performance.

Figure 6. MSA 2040 SAN iSCSI Network



IP address scheme for the controller pair

The MSA 2040 SAN controller in iSCSI configurations should have ports on each controller in the same subnets to enable preferred path failover. The suggested means of doing this is to vertically combine ports into subnets. See examples below.

Example with a netmask of 255.255.255.0:

- Controller A port 1: 10.10.10.100
- Controller A port 2: 10.11.10.120
- Controller A port 3: 10.10.10.110
- Controller A port 4: 10.11.10.130
- Controller B port 1: 10.10.10.140
- Controller B port 2: 10.11.10.150
- Controller B port 3: 10.10.10.160
- Controller B port 4: 10.11.10.170

Jumbo frames

A normal Ethernet frame can contain 1500 bytes whereas a jumbo frame can contain a maximum of 9000 bytes for larger data transfers. If you are using jumbo frames, make sure to enable jumbo frames on all network components in the data path.

Summary

HP MSA 2040 administrators should determine the appropriate levels of fault tolerance and performance that best suits their needs. Understanding the workloads and environment for the MSA SAN is also important. Following the configuration options listed in this paper can help optimize the HP MSA 2040 array accordingly.

Learn more at

hp.com/go/MSA

Sign up for updates

hp.com/go/getupdated



Share with colleagues



Rate this document

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

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle and/or its affiliates.

