Hewlett Packard Enterprise

HPE MSA 1050/2050/2052

Best practices

Technical white paper

Contents

About this document

This white paper highlights the best practices for optimizing the HPE 1050/2050/2052, and should be used in conjunction with other HPE Modular Smart Array (MSA) manuals. The HPE MSA 1050/2050/2052 is the fifth generation MSA SAN. HPE MSA technical user documentations can be found at: hpe.com/storage/msa.

Intended audience

This white paper is intended for HPE MSA 1050/2050/2052 administrators with previous storage area network (SAN) knowledge. It offers HPE MSA practices that can contribute to an HPE MSA best customer experience.

This paper is also designed to convey best practices in the deployment of the HPE MSA 1050/2050/2052 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)
- iSCSI and Ethernet protocols

Related documentation

In addition to this guide, please see other documents or materials for this product:

- HPE MSA System Racking Instructions
- HPE MSA 1050 User Guide
- HPE MSA 2050 User Guide
- HPE MSA 1050/2050 SMU Reference Guide
- HPE MSA 1050/2050 CLI Reference Guide
- HPE MSA 1050 Quick Start Instructions
- HPE MSA 2050/2052 Quick Start Instructions

You can find the HPE MSA 1050/2050/2052 documents at: hpe.com/storage/msa.

Introduction

The HPE MSA 1050 is designed for entry-level market needs featuring 8 Gb Fibre Channel, 6/12 Gb SAS, 1GbE, and 10GbE iSCSI protocols.

An outline of the MSA 1050 features:

- 6 GB cache per controller (Data [Read/Write] cache = 4 GB and Metadata and System OS memory = 2 GB)
- Support for SFF and LFF solid state drives (SSDs)
- 6 Gb/12 Gb SAS connectivity
- Support for MSA Fan-out SAS cables
- 2 host ports per controller
- 4 Gb/8 Gb FC connectivity
- 1GbE/10GbE iSCSI connectivity

- Support for up to 4 disk enclosures including the array enclosure
- Support for up to 96 small form factor (SFF) drives and 48 large form factor (LFF) drives
- Support for Thin Provisioning
- Support for Automated Tiering
- Support for Read Cache
- Support for Performance Tier[1]
- Wide Striping: Wide Striping allows more hard drives behind a single volume to improve performance (e.g., >16 drives for a volume).
- Support for replication snapshot history and queueing for Remote Snap; Remote Snap requires a license
- Remote Snap for both FC and iSCSI; Remote Snap requires a license

The HPE MSA 2050 is a high-performance Storage System designed for HPE customers desiring 8 Gb and/or 16 Gb Fibre Channel, 6 Gb and/or 12 Gb SAS, and 1GbE and/or 10GbE iSCSI connectivity with four host ports per controller. The HPE MSA 2050 Storage System provides an excellent value for customers needing performance balanced with price to support initiatives such as consolidation and virtualization.

The HPE MSA 2050 delivers this performance by offering

- 2X I/O performance from fourth generation
- 8 GB cache per controller (Data [Read/Write] cache = 4 GB and metadata and system OS memory = 4 GB)
- Support for SFF and LFF solid-state drives (SSDs)
- Four host ports per controller
- 4 Gb/8 Gb/16 Gb FC connectivity
- 1GbE/10GbE iSCSI connectivity
- 6 Gb/12 Gb SAS connectivity
- Support for both FC and iSCSI in a single controller
- Support for up to eight Disk Enclosures including the Array Enclosure
- Support for up to 192 small form factor (SFF) drives and 96 large form factor (LFF) drives
- Support for Thin Provisioning
- Support for Automated Tiering
- Support for Read Cache
- Support for Performance Tier[2]
- Wide Striping: Wide Striping allows more hard drives behind a single volume to improve performance (for example, >16 drives for a volume)
- Support for replication snapshot history and queueing for Remote Snap; Remote Snap requires a license
- Remote Snap for both Fibre Channel and iSCSI; Remote Snap requires a license
- Volume Copy across pools
- Support for Full Drive Encryption (FDE) using Self-Encrypting Drives (SEDs)[3]
- Support for replication snapshot history and queueing for Remote Snap; Remote Snap requires a license
- Remote Snap for both FC and iSCSI

[1, 2] A license is required for the Performance Tier for mixed SSD and HDD systems. For a system with only SSDs, the Performance Tier license is not required.
[3] SED drives are only supported in HPE MSA 2050.

HPE MSA 2052 SAN Storage offers an entry-level storage platform with built-in hybrid flash for application acceleration and high-performance. It is ideal for performance-hungry applications and includes 1.6 TB of solid-state drive (SSD) capacity.

An outline of HPE MSA 2052 features:

- All the features of the MSA 2050
- The industry's fastest entry array, now with 1.6 TB of SSD capacity, standard
- Advanced Data Services (ADS) Software Suite

The HPE MSA ADS Suite is included as a standard feature on the MSA 2052 at no extra charge. The standard MSA ADS Software Suite includes the following functionality:

- Performance Tiering and Archive Tiering Software
- 512 Snapshots and Volume Copy Software
- Remote Snap Software

Note
The HPE Advanced Data Services (ADS) Software Suite license key comes standard with the HPE MSA 2052 and must be redeemed and installed to enable the services.

The HPE MSA 1050/2050 Storage Systems ship standard with a license for 64 snapshots and Volume Copy for increased data protection. There is also an optional license for 512 snapshots. The HPE MSA 1050/2050/2052 can also replicate data between arrays (HPE MSA 1040/2040/2042 using iSCSI virtual volumes with GL firmware and HPE MSA 1050/2050/2052 SAN model using FC or iSCSI with the optional Remote Snap feature).

Terminology

Virtual Storage: Data is virtualized not only across a single Disk Group but also across multiple Disk Groups with different performance capabilities and use cases.

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

Disk Group: A Disk Group is a collection of disks in a given redundancy.

Storage Pools: Storage Pools are comprised of one or more Virtual Disk Groups. A volume's data on a given LUN can span all disk drives in a pool. When capacity is added to a system, users will benefit from the performance of all spindles in that pool.

When leveraging Storage Pools, HPE MSA 1050/2050/2052 supports large, flexible volumes with sizes up to 128 TiB and facilitates seamless capacity expansion. As volumes are expanded, data automatically reflows to balance capacity utilization on all drives.

Logical Unit Number (LUN): HPE MSA 1050/2050/2052 arrays support 512 volumes and up to 512 snapshots in a system. All of these volumes can be mapped to LUNs. Maximum LUN sizes are up to 128 TiB.

Thin Provisioning: Thin Provisioning allows storage allocation of physical storage resources only when they are consumed by an application. Thin Provisioning also allows overprovisioning of physical Storage Pool resources allowing ease of growth for volumes without predicting storage capacity up front.

Tiers: Disk tiers are comprised of aggregating one or more Disk Groups of similar physical disks. HPE MSA 1050/2050/2052 support three distinct tiers:

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

LUN-level tiering requires careful planning such that applications requiring the highest performance be placed on Disk Groups utilizing high-performance SSDs. Applications with lower performance requirements can be placed on Disk Groups comprised of Enterprise SAS or Midline SAS HDDs.

HPE MSA 1050/2050/2052 Automated Tiering engine moves data between available tiers based on the access characteristics of that data. Frequently accessed data contained in pages will migrate to the highest available tier delivering maximum I/Os to the application. Similarly, cold or infrequently accessed data is moved to lower performance tiers. Data is migrated between tiers automatically such that I/Os are optimized in real time.

The Archive and Standard Tiers are provided at no charge on HPE MSA 2050 platform. Performance Tier utilizing a fault-tolerant SSD Disk Group is a paid feature that requires a license for HPE MSA 1050/2050. Without the Performance Tier license installed, SSDs can still be used as Read Cache with the Sub-LUN Tiering feature. Sub-LUN Tiering from SAS MDL (Archive Tier) to Enterprise SAS (Standard Tier) drives is provided at no charge for HPE MSA 1050/2050/2052.

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

Automated Tiering: Automated Tiering is a technology that allows for the automatic movement of data between storage tiers based on access trends. In HPE MSA 1050/2050/2052, Automated Tiering places data in a LUN that is accessed frequently in higher performing media while data that is infrequently accessed is placed in slower media.

Array Enclosure: This is the array head or chassis of HPE MSA that includes HPE MSA controllers.

Disk Enclosure: This is the expansion shelf that is connected to the Array Enclosure.

Storage System: This is the whole HPE MSA system that includes the Array Enclosure and Disk Enclosure(s).

General best practices

Become familiar with the array by reading the manuals

The first recommended best practice is to read the corresponding guides for HPE MSA 1050/2050/2052. These documents include the User Guide, the Storage Management Utility (SMU) Reference Guide, or the Command Line Interface (CLI) Reference Guide. The appropriate guide will depend on the interface that you will use to configure the storage array. Always operate the array in accordance with the user manual. In particular, never exceed the environmental operation requirements.

Other HPE MSA 1050/2050/2052 materials of importance to review are:

- The MSA Remote Snap Software technical white paper located at hpe.com/h20195/v2/Getdocument.aspx?docname=4AA1-0977ENW.

Stay current on firmware

Use the latest controller, disk, and Disk Enclosure firmware to benefit from the continual improvements in the performance, reliability, and functionality of the HPE MSA 1050/2050/2052. For additional information, see the release notes and release advisories for the respective HPE MSA products.

This information can be located at hpe.com/storage/msa

Use tested and supported configurations

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

HPE primary portal used to obtain detailed information about supported HPE storage product configurations is HPE Single Point of Connectivity Knowledge (SPOCK). An HPE Passport account is required to enter the SPOCK website.

SPOCK can be located at hpe.com/storage/spock.

Understand what a host is from the array perspective

An initiator is analogous to an external port on a host bus adapter (HBA). An initiator port does not equate to a physical server but rather a unique connection on that server. For example, a dual port FC HBA has two ports and therefore there are two unique initiators, and the array will show two separate initiators for that HBA.

A host is a collection of one or more initiators. HPE MSA firmware can support 512 hosts with multiple initiators per host. HPE MSA can manage 1024 initiators.

The array supports the grouping of initiators under a single host and grouping hosts into a host group. Grouping of initiators and hosts allows simplification of the mapping operations.

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. A recommended method 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 following procedure outlines the steps needed to rename hosts using the SMU.

1. Log in to the SMU and click Hosts from the left frame.
2. Locate and highlight the WWN (ID) you want to name.
3. From the Action button, click Modify Initiator.
4. Type in the initiator nickname and click OK.
5. Repeat for additional initiator connections.

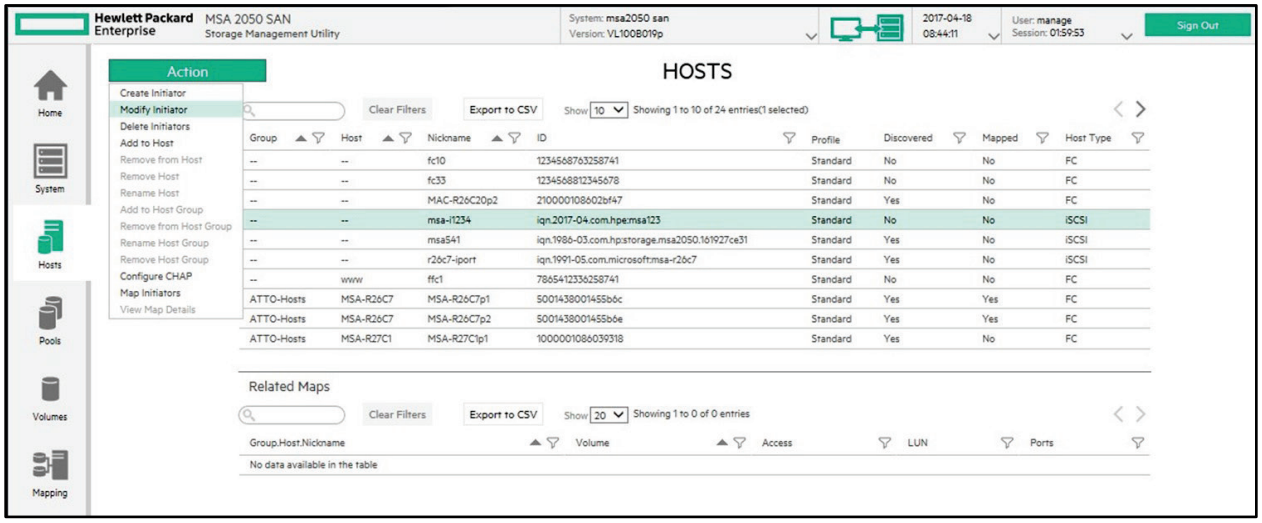

Figure 1. Renaming hosts

The recommended practice would be to use initiator nicknaming as outlined in Figure 1—host aggregating of initiators and the grouping of hosts using the SMU.

Best practice for monitoring array health

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

Configure email and SNMP notifications

The Storage Management Utility (SMU) is the recommended method for setting up email and SNMP 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 HPE MSA 1050/2050/2052.

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 with the **Include logs as an email attachment** option enabled is recommended. When the Include logs as an email attachment feature is enabled, the system automatically attaches the system log files to the managed logs' email notifications that is sent. The managed logs email notification is sent to an email address, which will retain the logs for future diagnostic investigation.

HPE MSA 1050/2050/2052 Storage 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 as an email attachment** is also checked, the segment of logs that is nearing a full state will be attached to the email notification. Managed logs attachments can be multiple MB in size.

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 as an email attachment** option is disabled by default.

Hewlett Packard Enterprise recommends enabling SNMP traps. Version 1 SNMP traps can be sent to up to three host trap addresses (that is, an SNMP server). To send version 3 SNMP traps, create an SNMPv3 user with the trap target account type. Use SNMPv3 traps rather than SNMPv1 traps for greater security. SNMP traps can be useful in troubleshooting issues with HPE MSA 1050/2050/2052 array.

To configure email and version 1 SNMP settings in the SMU, click **Home -> Action -> Set Up Notifications**. Enter the correct information for Email, SNMP, and Managed Logs. See Figure 3.

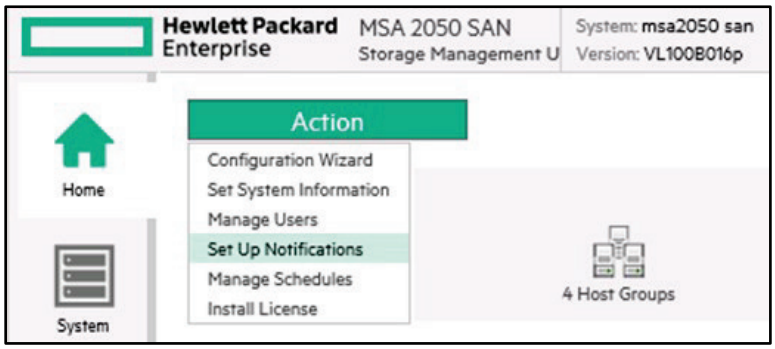

Figure 2. Setting up management services

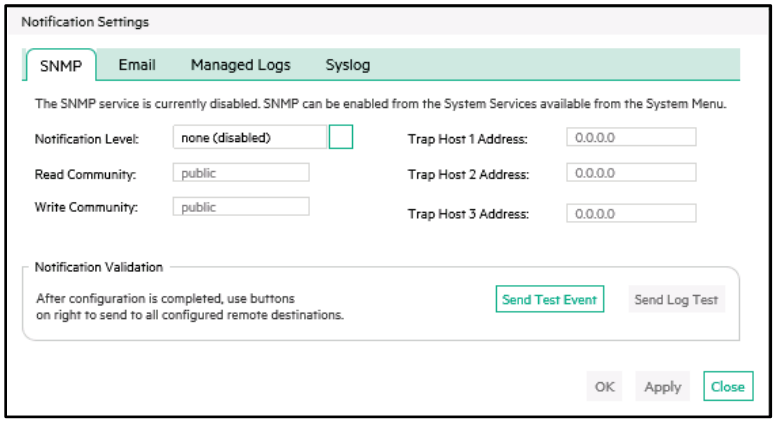

Figure 3. SNMP, Email, and Managed Logs Notification Settings

To configure SNMPv3 users and trap targets, click **Home -> Action -> Manage Users**. See Figure 4.

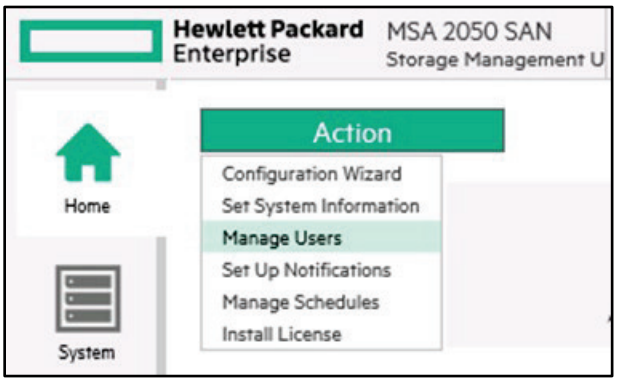

Figure 4. Manage Users

Enter the correct information for SNMPv3 trap targets. See Figure 5.

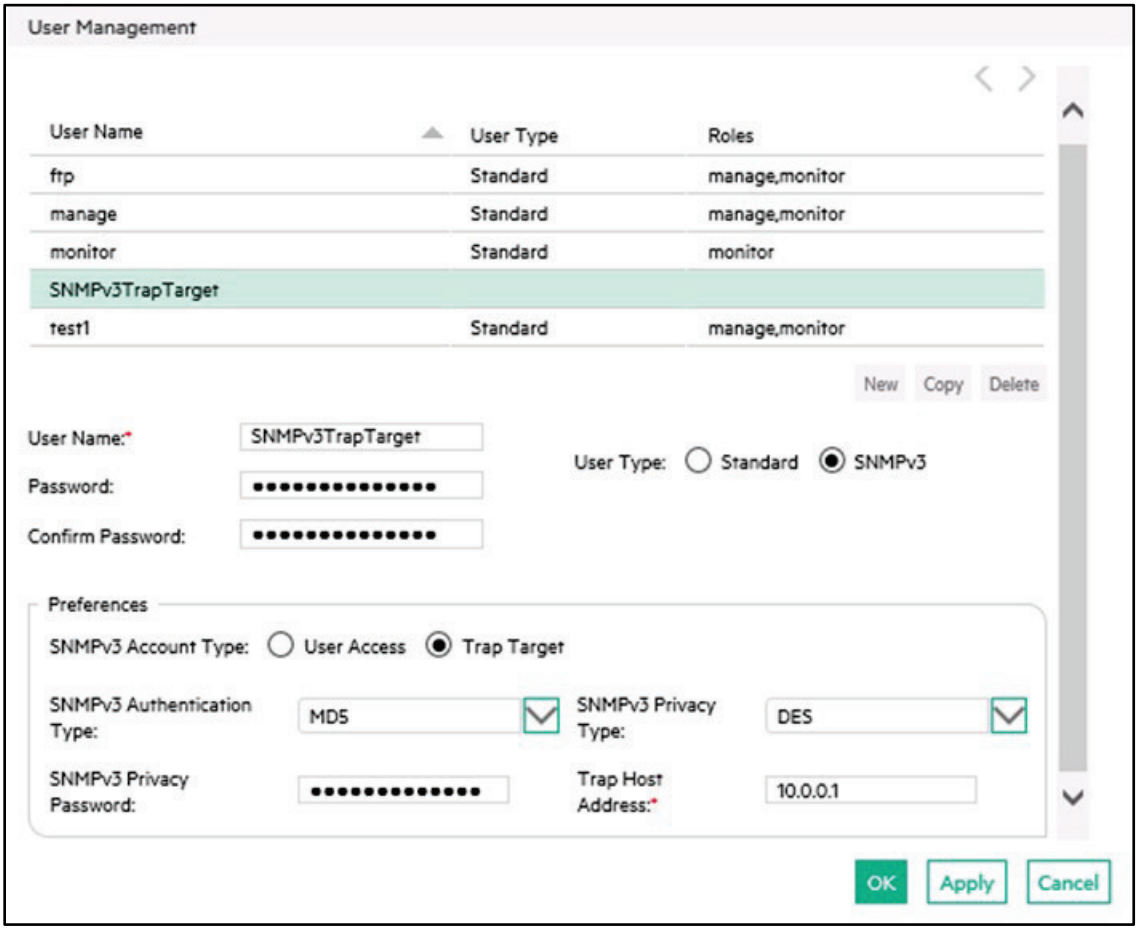

Figure 5. User Management

Setting the notification level for email and SNMP

Setting the notification level to Critical, Error, Warning on the Email and Warning on SNMP configurations will ensure that events of that level or above are sent to the destinations (that is, SNMP server, SMTP server) set for that notification. Hewlett Packard Enterprise recommends setting the notification level to Warning.

HPE MSA 1050/2050/2052 notification levels:

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

Sign up for proactive notifications for HPE MSA 1050/2050/2052 array

Sign up for proactive notifications to receive HPE MSA product advisories. Applying the suggested resolutions can enhance the availability of the product. Sign up for the notifications at h41360.www4.hpe.com/alerts-signup.php.

Best practices for provisioning storage on HPE MSA 1050/2050/2052

The following section will assist in the best methods for optimizing Virtual Storage features such as Thin Provisioning, Wide Striping, and Automated Tiering for HPE MSA 1050/2050/2052.

Thin Provisioning

Thin Provisioning is a storage allocation scheme that automatically allocates storage as your applications need it.

Thin Provisioning dramatically increases storage utilization by removing the need to match purchased capacity to allocated capacity. Traditionally, application administrators purchased storage based on the capacity required now and for future growth. This resulted in over purchasing capacity and unused space.

With Thin Provisioning, applications can be provided with all of the capacity to which they are expected to grow but can begin operating on a smaller amount of physical storage. As the applications fill their storage, new storage can be purchased as needed and added to the array's Storage Pools. This results in a more efficient utilization of storage and a reduction in power and cooling requirements.

Overcommit is enabled by default. The overcommit setting lets the user oversubscribe the physical storage (that is, provision volumes in excess of physical capacity). If a user disables overcommit, they can only provision virtual volumes up to the available physical capacity.

Overcommit is performed on a per pool basis and using the **Change Pool Settings** option. To change the **Pool Settings** to overcommit disabled:

1. Open the SMU and select **Pools**.
2. Click **Action** > **Change Pool Settings.**
3. Uncheck the **Enable overcommitment of pool?** by clicking the box. See Figures 6 and 7.

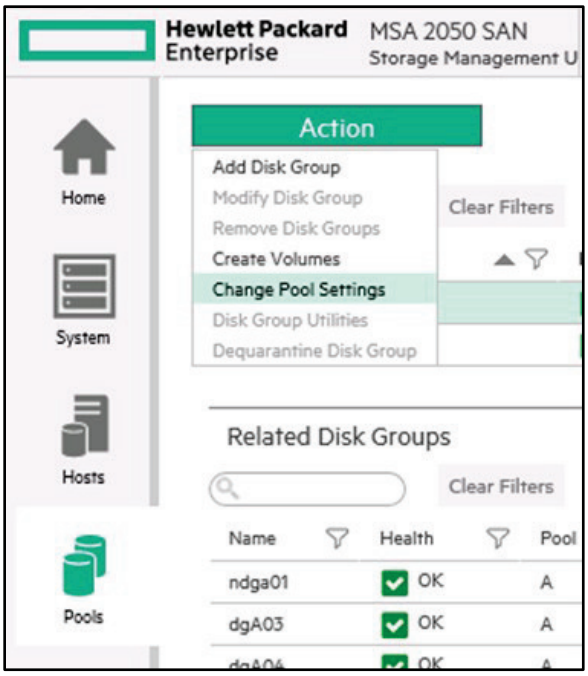

Figure 6. Changing pool settings

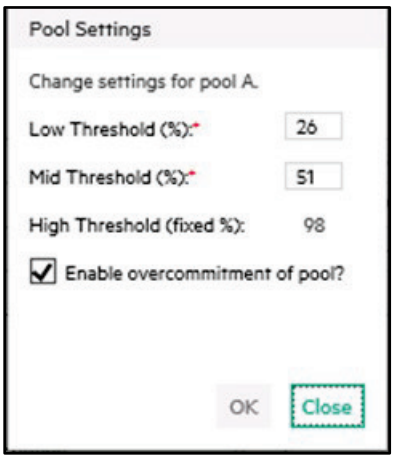

Figure 7. Setting overcommit for the pool

Thresholds and notifications

If you use Thin Provisioning, monitor space consumption and set notification thresholds appropriately for the rate of storage consumption. The following thresholds and notifications can help determine when more storage needs to be added.

Users with a manage role can view and change settings that affect the thresholds and corresponding notifications for each Storage Pool.

- **Low Threshold**—When this percentage of virtual pool capacity has been used, informational event 462 will be generated to notify the administrator. This value must be less than the Mid Threshold value. The default is 25%.

- **Mid Threshold**—When this percentage of virtual pool capacity has been used, event 462 will be generated to notify the administrator to add capacity to the pool. This value must be between the Low Threshold and High Threshold values. The default is 50%. If the pool is not overcommitted, the event will have Informational severity. If the pool is overcommitted, the event will have Warning severity.
- **High Threshold**—When this percentage of virtual pool capacity has been used, event 462 will be generated to alert the administrator to add capacity to the pool. This value is automatically calculated based on the available capacity of the pool minus 200 GB of reserved space. If the pool is not overcommitted, the event will have Informational severity. If the pool is overcommitted, the event will have Warning severity and the system will use write-through cache mode until virtual pool usage drops back below this threshold.

See Figures 6 and 7 on how to set the thresholds.

T10 UNMAP for Thin Reclaim

UNMAP is the ability to reclaim thinly provisioned storage after the storage is no longer needed. There are procedures to reclaim UNMAP space when using Thin Provisioning and VMware® ESX®.

The user should run the UNMAP command with ESX 5.0 Update 1 or higher to avoid performance issues. In ESX 5.0, UNMAP is automatically executed when deleting or moving a virtual machine.

In ESX 5.0 Update 1 and greater, the UNMAP command was decoupled from auto reclaim; therefore, use the VMware vSphere® CLI command to run UNMAP command.

See VMware® knowledge base for further details on the UNMAP command and reclaiming space.

Pool balancing

Creating and balancing Storage Pools properly can help with performance of HPE MSA array. Hewlett Packard Enterprise recommends keeping pools balanced from a capacity utilization and performance perspective. Pool balancing will leverage both controllers and balance the workload across the two pools.

Assuming symmetrical composition of Storage Pools, create and provision storage volumes by the workload that will be used. For example, an archive volume would be best placed in a pool with the most available Archive Tier space. For a high performance volume, create the Disk Group on the pool that is getting the least amount of I/O on the Standard and Performance Tiers.

Determining the pool space can easily be viewed in the SMU. Simply navigate to **Pools** and click the name of the pool.

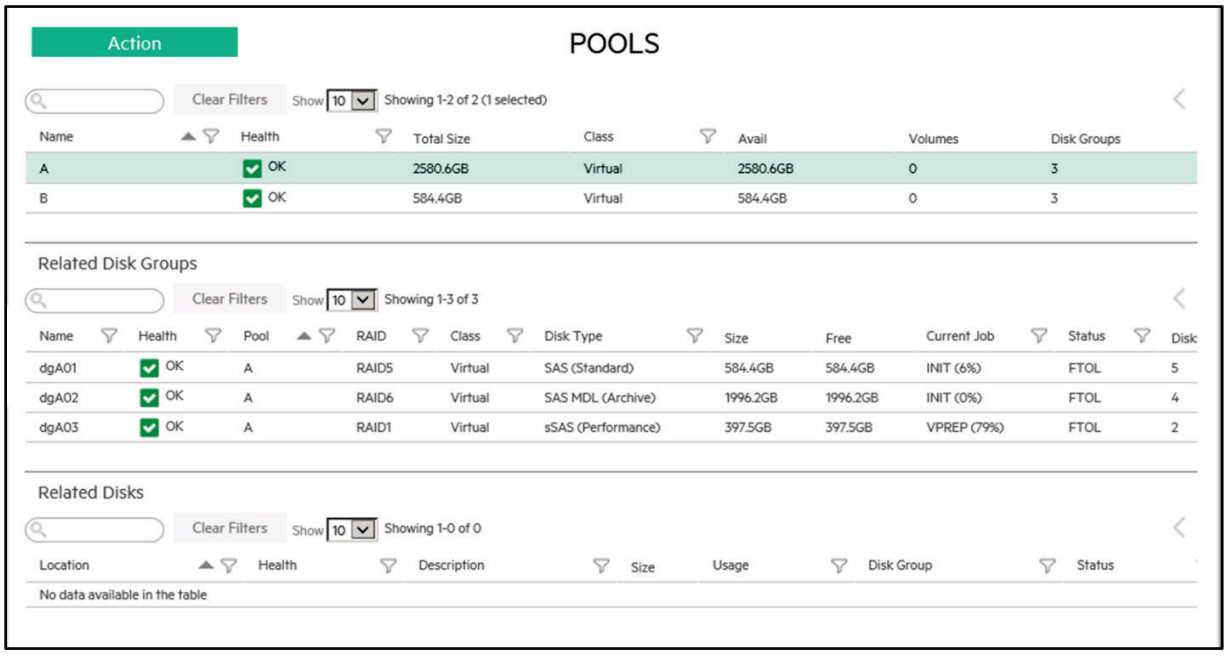

Figure 8. HPE MSA Pool A screen

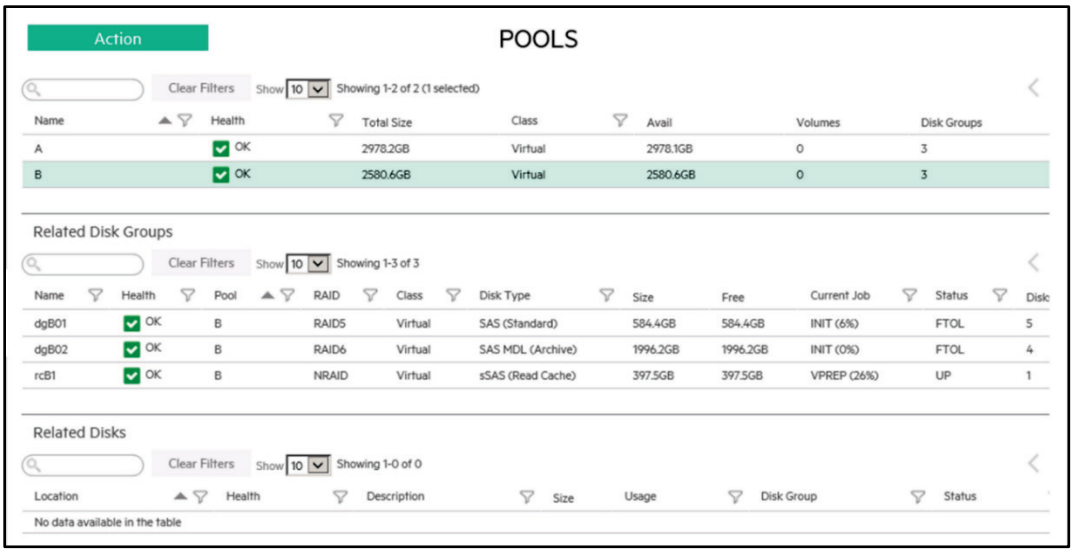

Figure 9. HPE MSA Pool B screen

Viewing the performance of the pools or Virtual Disk Groups can also assist in determining where to place the Archive Tier space.

From the SMU, navigate to Performance, and then click **Virtual Pools** from the **Show**: drop-down box. Next, click the pool and for real-time data, click **Show Data**. For historical data, click the **Historical Data** box and **Set time range**.

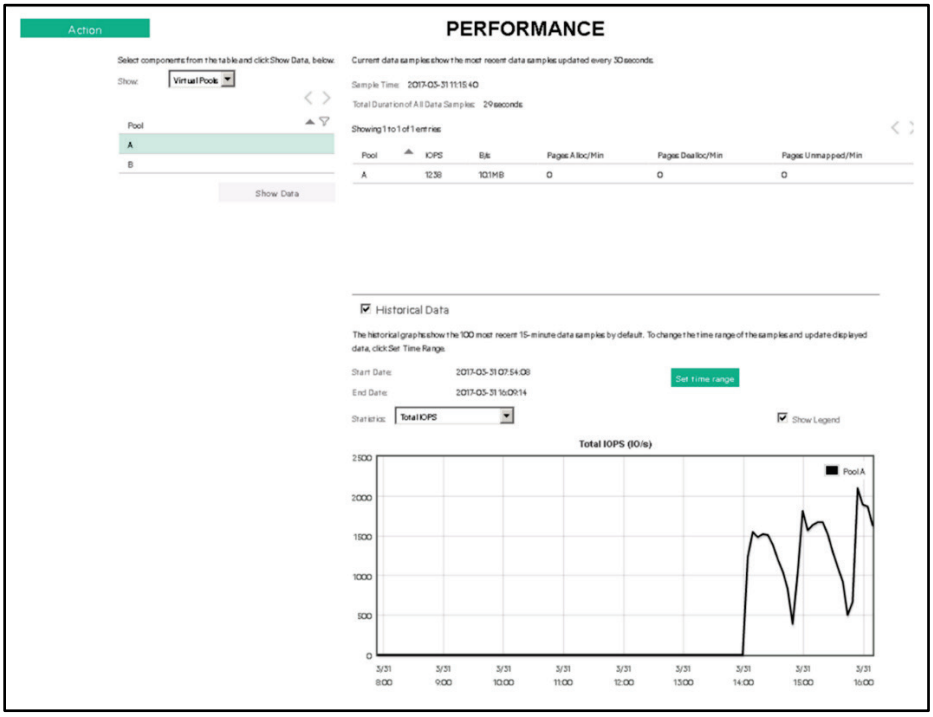

Figure 10. HPE MSA Virtual Pools Performance screen

The HPE MSA 1050/2050/2052 also has the ability to copy volumes owned by one pool to the other pool. Use this volume copy feature if the storage system becomes unbalanced.

Wide Striping

HPE MSA 1050/2050/2052 support the Wide Striping concept for Virtual Storage.

With Virtual Disk Groups in a Storage Pool, the HPE MSA algorithm evenly distributes the allocated Pages of a volume across all Disk Groups in the Storage Pool.

Wide Striping also allows for rapid expansion of volumes in the case of consuming all pool capacity. By adding one or more Virtual Disk Groups into the Storage Pool, the volumes within the pool will immediately benefit from the additional capacity. The leveling process will begin automatically and will redistribute data evenly across all Disk Groups in the Storage Pool. In essence, you will increase the storage space and storage performance by adding Disk Groups.

Note
The rebalancing happens automatically; no user interaction is required.

Expanding virtual volumes

There might come a time when the Virtual Disk Group in a pool will start to fill up. To easily add more space, HPE MSA implements Wide Striping, as mentioned above, to increase the size of the virtual volumes. The recommended method to increase the volume size is to add a new Virtual Disk Group with the same amount of drives and RAID type as the existing Virtual Disk Group has.

For example, a Virtual Disk Group in Pool A is filling up. This Disk Group is a five, 900 GB drive, 10K rpm, RAID 5 Disk Group. The recommended procedure would be to create a new Virtual Disk Group on Pool A that also has five, 900 GB 10K rpm disk drives in a RAID 5 configuration.

Using HPE MSA 2052 embedded SSDs

HPE MSA 2052 further drives HPE MSA 2050 family into the world of flash acceleration with a set of HPE MSA models, which include flash/SSD drives in SFF and LFF configurations. In addition to including two SSDs in the base configuration as a standard feature for flash acceleration, the HPE MSA 2052 also includes a rich set of software features as standard that includes 512 snapshots, Remote Snap Replication, and Performance Tiering capabilities. Each HPE MSA 2052 model comes standard with 1.6 TB (2 x 800 GB) mixed-use SSDs, which can be used as Read Cache or as a start to building a fully tiered configuration. With the inclusion of the Advanced Data Services software license, customers can choose how to best utilize these SSDs to provide the optimal flash acceleration for their environment. The MSA 2052 utilizes a real-time tiering algorithm, which works without user intervention to place the hottest pages of an array on the fastest medium at the right time. The tiering engine moves hot pages up to SSD for flash acceleration and moves cooler, less used, pages back down to spinning disks as workloads change in real time. All of which are all done hands free and without any management overhead from the IT manager.

The following section will assist in the best methods for using HPE MSA 2052.

Using HPE MSA 2052 SSDs for Read Cache

With the inclusion of two SSDs standard on HPE MSA 2052, each Storage Pool can have one Read Cache SSD. The method to construct this setup is as follows:

1. Create a Virtual Disk Group on Pool A using Enterprise SAS or Midline SAS drives.

Important
Make sure to choose Pool A from the drop-down box.

2. Next, create the Read Cache for Pool A using SSD number 1.

Important
Make sure to choose Pool A from the drop-down box.

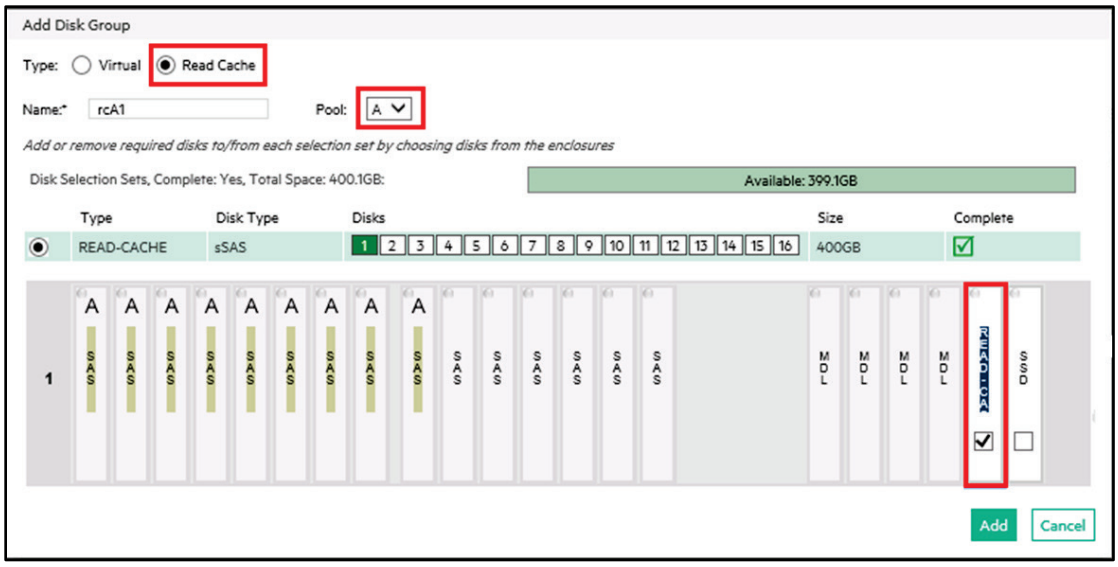

Figure 11. Creating Read Cache for Pool A

3. Create a Virtual Disk Group on Pool B.

> **Important**
> Make sure to choose Pool B from the drop-down box.

4. Next, create the Read Cache for Pool B using SSD number 2.

> **Important**
> Make sure to choose Pool B from the drop-down box.

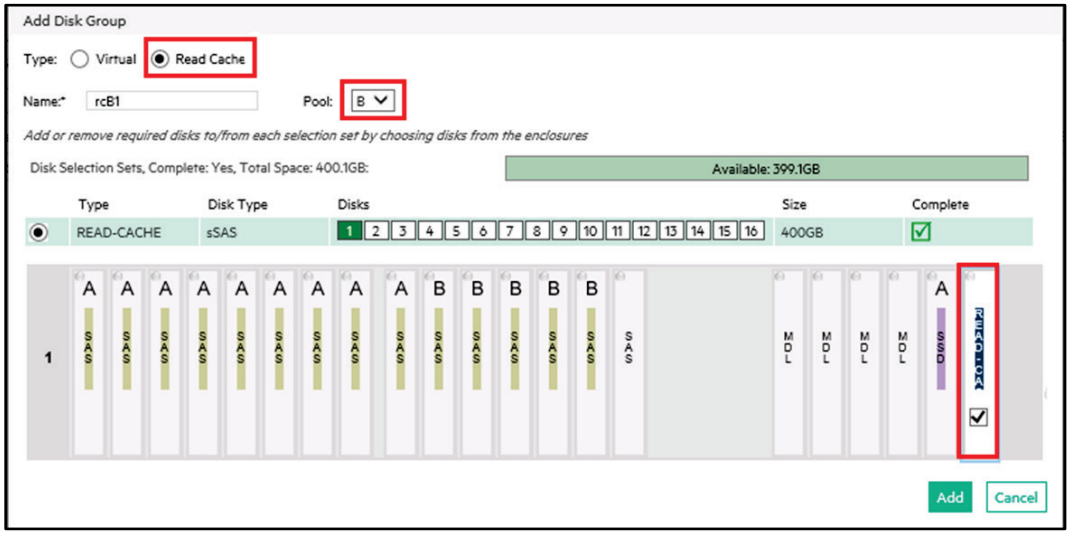

Figure 12. Creating Read Cache for Pool B

Now, you have utilized each SSD of the HPE MSA 2052 as Read Cache.

Using HPE MSA 2052 SSDs for Performance Tiering

With the inclusion of two SSDs standard on HPE MSA 2052, only one Storage Pool can have a Performance Tier Disk Group without purchasing additional SSD drives. This is because you need a minimum of two SSDs to construct a RAID 1 Virtual Disk Group to use for auto tiering to the Performance Tier. If additional SSD drives are purchased, then RAID 1, 5, or 6 disk groups may be created.

The method to construct this setup is as follows:

1. Create a Virtual Disk Group for Pool A using the SSDs.

> **Important**
> Select the correct RAID level and make sure to choose Pool A from the drop-down box.

Now, you have utilized SSDs installed in HPE MSA 2052 as a Virtual Disk Group to be used as a Performance Tier.

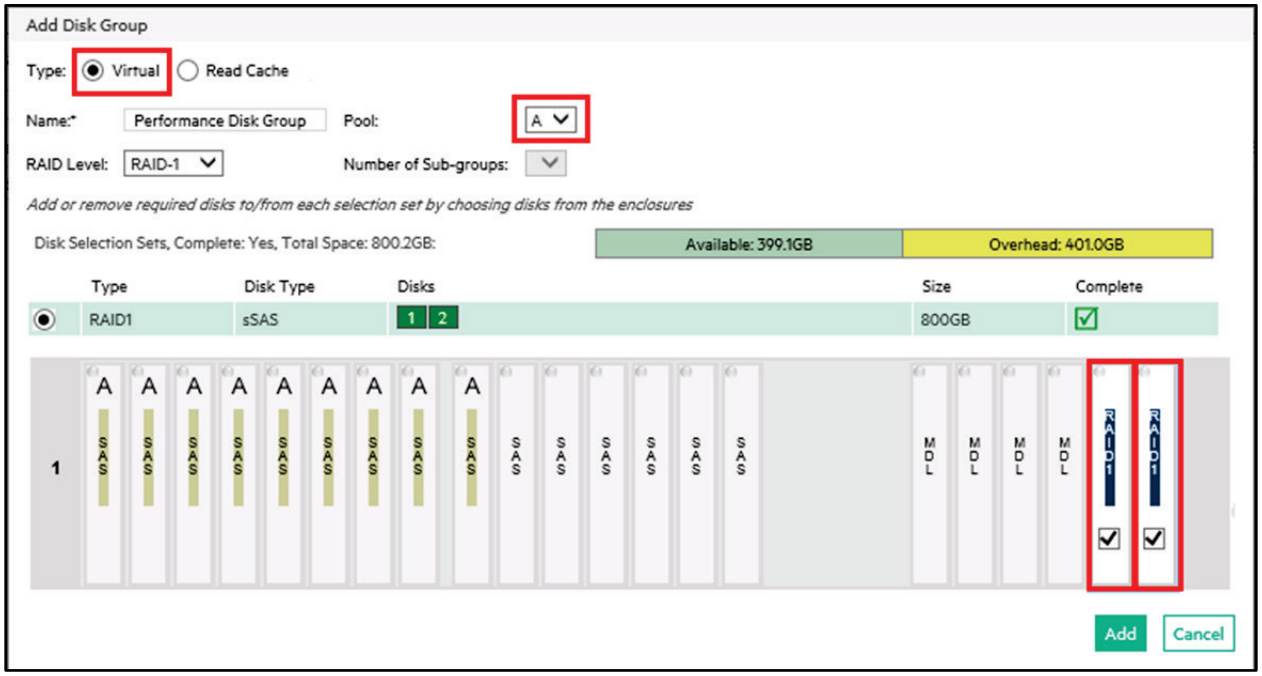

Figure 13. Creating a Performance Tier for Pool A

Hewlett Packard Enterprise recommends balancing Disk Groups across both pools as seen in the Creating Disk Groups section. For Performance Tiering, this would require the installation of a minimum of two more 800 GB SSDs in addition to the two 800 GB SSDs that come standard with HPE MSA 2052. With additional SSDs, different RAID levels can be chosen and SSD Disk Groups can be created for each pool.

Configuring the two SSDs into a single pool is supported, but will result in an unbalanced system because only one pool will be utilized with SSDs. Customers will not get the full performance benefits of HPE MSA 2052 in an unbalanced system; only the pool with the SSDs would get the full performance benefit. For example, if you configure two SSDs in Pool A but no SSDs in Pool B, then any applications with data in Pool B will not have the performance benefit of the SSDs in Pool A.

Automated Tiering

Automated Tiering is a Virtual Storage feature that automatically moves data residing in one class of disks to a more appropriate class of disks based on data access patterns:

- Frequently accessed, hot data can move to disks with higher performance, typically lower capacity, and typically higher costs.
- Infrequently accessed, cool data can move to disks with higher capacity, lower performance, and typically lower costs per GB.

Each Virtual Disk Group, depending on the type of disks it uses, is automatically assigned to one of the following tiers:

- **Performance**—This highest tier uses SAS SSDs, which provide the best performance, but also the highest cost per GB.
- **Standard**—This middle tier uses enterprise-class 10K rpm/15K rpm spinning SAS disks, which provide good performance with mid-level cost per GB and capacity.
- **Archive**—This lowest tier uses Midline 7.2K rpm spinning SAS disks, which provide the lowest performance with the lowest cost per GB and highest capacity.

How Automated Tiering works

Automated Tiering uses virtualized storage and is accomplished by paging. HPE MSA virtual volumes allocate data into small, 4 MB chunks (Pages) from within the Virtual Storage Pool. These Pages are ranked based on a sophisticated algorithm. The Page rank is used to efficiently select appropriate Pages to move between tiers. The result is that Pages can be migrated between tiers automatically such that I/Os are optimized in real time.

In contrast to data movement at the LUN level, Automated Tiering at the Sub-LUN level provides highly efficient data movement. Only a minimum amount of CPU and memory resources are needed to support the data movement between tiers; therefore, movement can happen in real time rather than in offline batch movements.

Automated Tiering concepts

HPE MSA tiering algorithm runs every five seconds and Pages are ranked, scanned, and migrated during this period.

1. Pages are ranked by access patterns.
2. A scan looks for highly ranked Pages.
3. These highly ranked Pages are then migrated up tier levels.
   a. Pages are only migrated down a tier if space is needed for a highly ranked Page.
   b. Only 80 MB of data is migrated every five seconds to avoid degrading the system throughput.
4. HPE MSA tiering algorithm is tuned to avoid thrashing, or moving Pages back and forth between tiers in brief amounts of time.
5. Infrequently accessed or cold data is only moved to lower tiers as capacity is required for more frequently accessed data. This keeps as much of your data in a higher performing tier as possible.
6. In order to optimize workloads, sequential writes initially are written to the fastest spinning drive tier with free capacity. A random write workload will initially be written to the fastest tier with free capacity. In a three-tiered HPE MSA, the random write workload would be written to the Performance Tier.

Note

The information in Steps 5 and 6 explain the No Affinity volume setting. Consult the Volume Tier Affinity section to review different virtual volume settings and how they affect quality of service (QoS).

Automated Tiering components

To get the benefits of Automated Tiering, a license is needed on the arrays. Once the proper license(s) are installed, the Automated Tiering function will be enabled.

Consult the following table to determine which license is needed.

Table 1. HPE MSA 1050/2050/2052 Automated Tiering components

	HPE MSA 1050	HPE MSA 2050	HPE MSA 2052
Automated Tiering component			
Archive Tiering (Enterprise SAS to SAS Midline)	Standard	Standard	Standard
Performance Tiering (SSD to Enterprise SAS)	Optional*	Optional*	Standard**

* Requires Advanced Data Services Software Suite

** All HPE MSA 2052 models come standard with the HPE Advanced Data Services (ADS) Software Suite. Software titles included in the ADS Software Suite include:

- Performance Tiering and Archive Tiering Software
- 512 Snapshots and Volume Copy Software
- Remote Snap Software

Note

The HPE Advanced Data Services (ADS) Software Suite license key comes standard with HPE MSA 2052 and must be redeemed and installed to enable the services.

Advantages of using Automated Tiering

- Because a virtual pool can have multiple Virtual Disk Groups, each belonging to a different tier, a virtual pool can provide multiple tiers of storage, which will lower total cost of ownership (TCO) by moving less accessed data to the lower cost per GB media.
- The I/O load is automatically balanced between components in a tier, which improves performance and allows for easy expansion.
- Virtual Disk Groups can be added or removed without disrupting I/O. Data in Virtual Disk Groups that are being removed are automatically migrated to other Disk Groups in a Storage Pool as long as the other disk Groups in the Storage Pool have enough storage space. If there is not enough space, the system will not delete the Disk Groups until enough data is removed.

Disk Group considerations

Allocated Pages are evenly distributed between Disk Groups in a tier; therefore, create all Disk Groups in a tier with the same RAID type and number of drives to ensure uniform performance in the tier.

Consider an example where the first Disk Group in the Standard Tier consists of five 15K Enterprise SAS drives in a RAID 5 configuration.

To ensure consistent performance in the tier, any additional Disk Groups for the Standard Tier should also be a RAID 5 configuration. Adding a new Disk Group configured with four 10K Enterprise SAS drives in a RAID 6 configuration will produce inconsistent performance within the tier due to the different characteristics of the Disk Groups.

For optimal write performance, parity-based Disk Groups (RAID 5 and RAID 6) should be created with “The Power of 2” method. This method means that the number of data (non-parity) drives contained in a Disk Group should be a power of 2. See Table 2.

Table 2. The Power of 2 method

RAID type	Total drives per Disk Group	Data drives	Parity drives
RAID 5	3	2	1
RAID 5	5	4	1
RAID 5	9	8	1
RAID 6	4	2	2
RAID 6	6	4	2
RAID 6	10	8	2

Due to the limitation of Disk Groups in a pool, which is 16, RAID type should be considered when creating new Disk Groups. For example, instead of creating multiple RAID 1 Disk Groups, consider using a larger RAID 10 Disk Group.

Drive type and capacity considerations when using tiering

All hard disk drives in a tier should be the same type. For example, do not mix 10K rpm and 15K rpm drives in the same Standard Tier. The SSD Performance Tier or Read Cache should be sized correctly to fit your active data size to gain the best performance boost. See the section on SSD Read Cache for more information on sizing Read Cache data.

HPE MSA 1050/2050/2052 supports all SSD drives shipped with previous HPE MSA 1040/2040/2042 products to support upgrades. SSDs can be utilized in the Read Cache or Performance Tier. All SSDs will be classified into the same tier and will interoperate together. Different SSD saleable parts will have differences in both performance and endurance and mixing of the different parts should be done with care.

The following table lists Disk Group configurations that can be used in HPE MSA 1050/2050/2052 array systems and their support information.

Table 3. Supported Disk Group types

Different drive types in the same Disk Group	Supported in HPE MSA 1050/2050/2052
LFF 15K and SFF 15K	Yes
SSD Mainstream Endurance[4] and SSD Mixed Use	Yes
SSD and 15K	No
SSD and 10K	No
10K and 15K	Yes, but not recommended

Disk Group RAID type considerations

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

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

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

Global spares with tiers

Using global spares is recommended for all tiers based on spinning media. When using these global spares, make sure to use the same drive types as the Disk Group. The drive size must be equal or larger than the smallest drive in the tier. For more information on sparing, go to the Drive sparing section.

Creating multiple tiers on one Storage Pool

By default when creating multiple Virtual Disk Groups, HPE MSA SMU attempts to alternate the Disk Group ownership between Pools A and B. Make sure to choose one pool only if you need multiple tiers on one pool.

Important

In this example, make sure to choose Pool A from the drop-down box when creating the Virtual Disk Groups.

[4] SSD Mainstream Endurance drives are only supported in HPE MSA 1050/2050/2052 after an upgrade from an HPE MSA 1040/2040/2042.

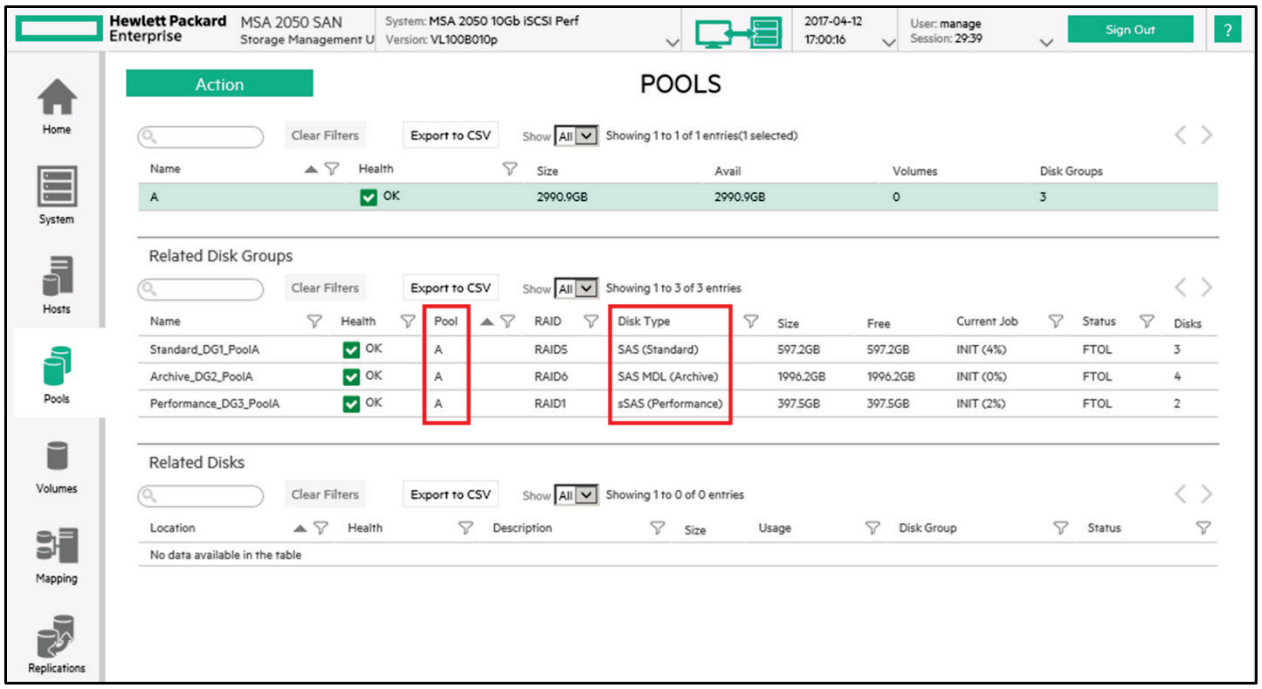

Figure 14. HPE MSA Virtual Pools screen

Volume Tier Affinity

Volume Tier Affinity is a settable attribute that allows a storage administrator to define quality of service (QoS) preferences for virtual volumes in a tiered environment. There are three Volume Tier Affinity options—Archive, Performance, and No Affinity. A setting of Archive will prefer the lowest tier of service, Performance will prefer the higher tiers of service, and **No Affinity** will use the Standard Tiering strategy. Tier Affinity is not the same as Tier Pinning and does not restrict data to a given tier and capacity. Data on a volume with **Archive** affinity can still to be promoted to a Performance Tier if that data becomes in demand to the host application.

Note

The Performance affinity does not require an SSD tier and will use the highest Performance Tier available.

Mechanics of Volume Tier Affinity

Volume Tier Affinity acts as a guide to the system on where to place data from a given volume in the available tiers.

The standard strategy is to prefer the highest spinning disk (non-SSD) tiers for new sequential writes and the highest tier available (including SSD) for new random writes. As data is later accessed by the host application, data will be moved to the most appropriate tier based on demand with hot data being promoted up toward the highest Performance Tier and cold data being demoted downward to the lower spinning disk based tiers. This standard strategy will be followed for data on volumes set to **No Affinity**.

For data on volumes set to the **Performance** affinity, the standard strategy will be followed for all new writes; however, subsequent access to that data will have a lower threshold for promotion upward making it more likely for that data to be available on the higher Performance Tiers. Preferential treatment will be provided to **hot** data that has **Performance** affinity at the SSD tier making it more likely for **Archive** or **No Affinity** data to be demoted out of the SSD tier to make room. This is useful for volumes where you know the data will be in demand and want to ensure that it has priority treatment for promotion to and retention in the highest Performance Tier.

For volumes that are set to the **Archive** affinity, all new writes will be initially placed in the Archive Tier as long as space is available. If no space is available, new writes will be placed on the next higher tier available. Subsequent access to that data will allow for its promotion to the

Performance Tiers as it becomes **hot**; however, the data will have a lower threshold for demotion and will be moved out of the highest performance SSD tier if there is a need to promote hot data up from a lower tier.

Volume Tier Affinity impact in existing environments

New virtual volumes will have a default setting of **No Affinity** and will continue to use the Standard Tiering strategy.

If the affinity of an existing volume is changed to the **Performance** affinity, there is no immediate change made. The current data on the volume will be promoted with the **Performance** strategy outlined in the previous section based on host application needs.

If the affinity of an existing volume is changed to **Archive**, a background operation will be performed to begin moving the data for the affected volume down to the Archive Tier. This process is done at a low priority to have as minimal effect on host I/O performance as possible. As data in a volume with an **Archive** affinity becomes **hot**, it will be promoted up to the higher tiers automatically. New data written to the volume will be targeted to the Archive Tiers based on the strategy outlined in the previous section.

Configuring Volume Tier Affinity

For new virtual volumes, at volume creation time, set the affinity from the **Preference** drop-down. Note that the default is **No Affinity**.

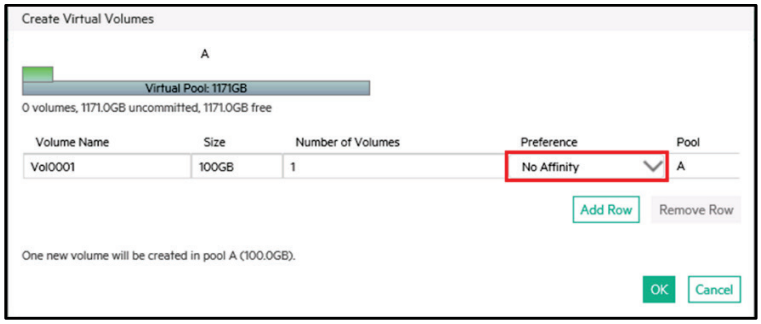

Figure 15. Setting Virtual Volume Tier Affinity

For existing virtual volumes, from HPE SMU select the **Volumes** tab, select the volume you want to set an affinity on and from the **Action** menu select **Modify Volume** and set the affinity from the **Preference** drop-down.

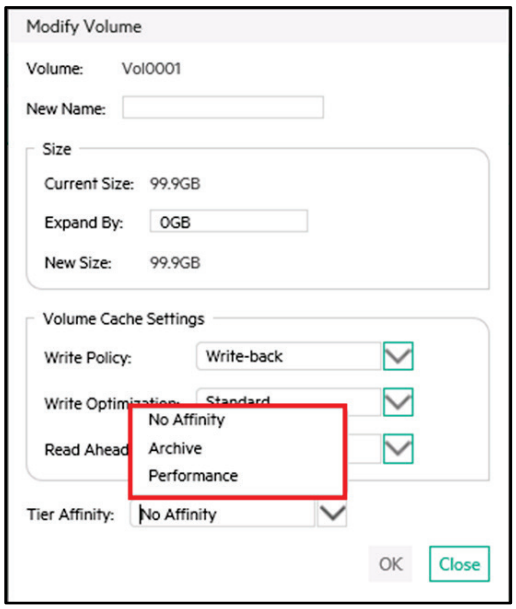

Figure 16. Modifying Virtual Volume Tier Affinity

Hewlett Packard Enterprise recommends the default **No Affinity** option for most configurations. This setting attempts to balance the frequency of data access, disk cost, and disk availability by moving this volume's data to the appropriate tier.

If the virtual volume uses mostly random or bursty low latency workloads such as online transaction processing (OLTP), virtual desktop infrastructure (VDI), or virtualization environments, Hewlett Packard Enterprise recommends setting the preference to **Performance**. This setting keeps as much of this volume's data in the Performance Tier for as long a period as possible.

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

Best practices when choosing drives for HPE MSA 1050/2050/2052 storage

The characteristics of applications and workloads are important when selecting drive types for HPE MSA 1050/2050/2052 array.

Drive types

HPE MSA 1050 and MSA 2052 arrays support SSDs, SAS Enterprise drives, and SAS Midline (MDL) drives. HPE MSA 2050 array supports SSDs, SAS Enterprise drives, SAS Midline (MDL) drives, and Self-Encrypting Drives (SEDs). See the Full Disk Encryption section for more information on SED 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 Storage System. For sequential workloads, SAS Enterprise drives or SAS MDL drives provide a good price-for-performance trade-off over SSDs. If more capacity is needed in your sequential environment, SAS MDL drives are recommended. SAS Enterprise drives offer higher performance than SAS MDL and should also be considered for random workloads when performance is 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 HPE MSA 1050/2050/2052 array. High availability is always advisable to protect your assets in the event of a device failure. Outlined 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 or 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 Disk Groups. When the failed controller is back online, the Disk Groups 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.

Mapping a port will make a mapping to each controller. Thus, mapping port 1 will map host ports A1 and B1. Mapping to port 2 will map host ports A2 and B2.

With this in mind, make sure that physical connections are set up correctly on HPE MSA so that a server has a connection to both controllers on the same port number. For example, on a direct attach HPE MSA 1050/2050/2052 array with multiple servers, make sure that ports A1 and B1 are connected to server 1, ports A2 and B2 are connected to server 2, and so on.

HPE MSA 1050 SAS with Fan-out cables

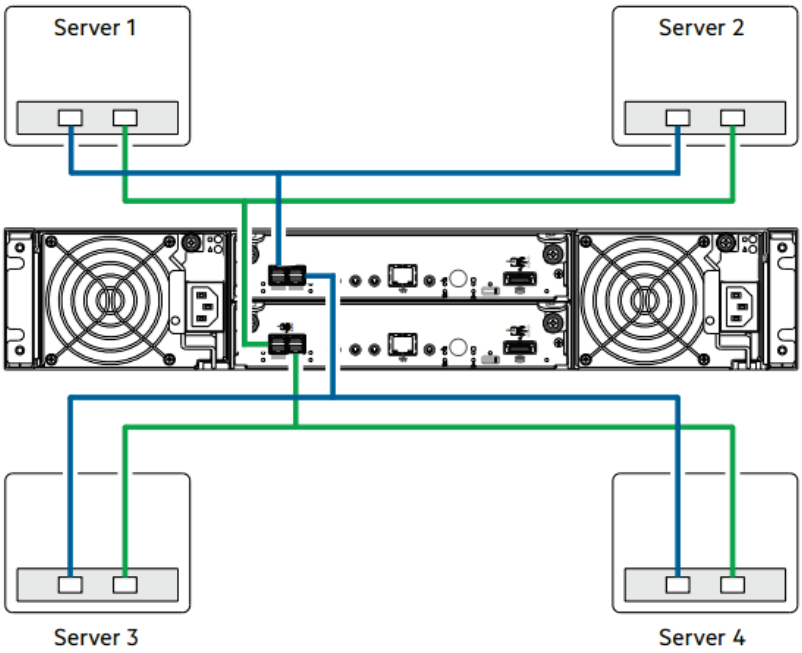

HPE MSA 2050 SAN

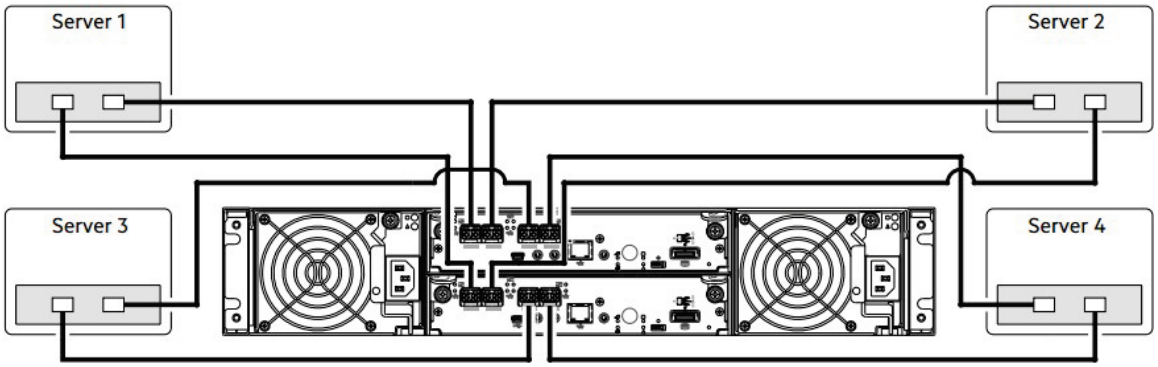

HPE 2050 SAS

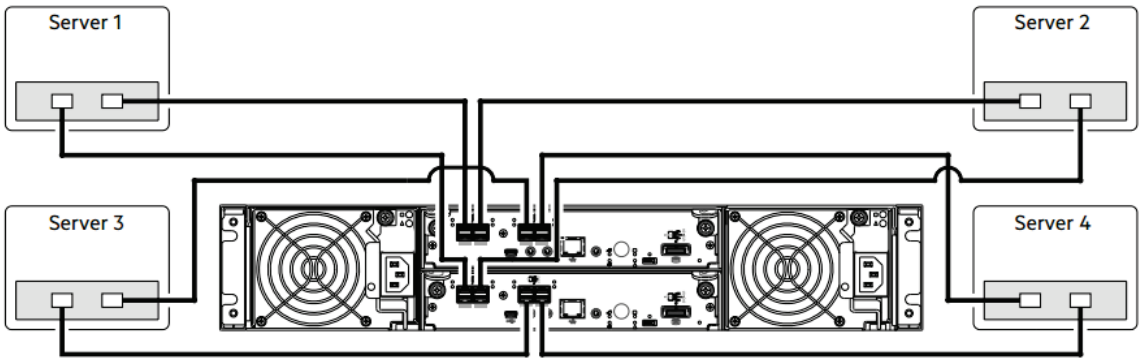

Figure 17. Direct attach cabling

It is not recommended to enable more than eight paths to a single host, that is, two HBA ports on a physical server connected to two ports on the A controller and two 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
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. Recommended practice is to put multiple initiators for the same host into a host and map the host to the LUNs, rather than individual maps to initiators.

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 4 lists supported multipath software by operating systems.

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

Table 4. Multipath and operating systems

Operating system	Multipath name	Vendor ID	Product ID
Windows Server® 2012/2016	Microsoft® Multipath I/O (MPIO)	HPE	MSA 1050 SAN MSA 1050 SAS MSA 2050 SAN MSA 2050 SAS
Linux®	Device mapper/multipath	HPE	MSA 1050 SAN MSA 1050 SAS MSA 2050 SAN MSA 2050 SAS
VMware	Native multipath (NMP)	HPE	MSA 1050 SAN MSA 1050 SAS MSA 2050 SAN MSA 2050 SAS

Installing MPIO on Windows Server 2012/2016

Microsoft has deprecated server manager cmd for Windows Server 2012/2016, so you will use the ocsetup command instead.

1. Open a command prompt window and run the following command:

```
C:\>ocsetup MultiPathIO /norestart
C:\>mpclaim -n -i -d "HPE     MSA 2050 SAN"
```

Figure 18. HPE MSA 1050/2050/2052 mpclaim command

Note
There are five spaces between HPE and MSA in the mpclaim command.

The mpclaim–n option avoids rebooting. Reboot is required before MPIO is operational.

The MPIO software is installed. When running the mpclaim command, type in the correct product ID for your HPE MSA product. See Table 4.

Once the MPIO DSM is installed, no further configuration is required; however, after initial installation, you should use Windows Server Device Manager to ensure that MPIO DSM has installed correctly as described in Managing MPIO LUNs.

Long failover times when using MPIO with large numbers of LUNs
Microsoft Windows Servers running MPIO use a default Windows® Registry **PDORemovePeriod** setting of 20 seconds. When MPIO is used with a large number of LUNs, this setting can be too brief, causing long failover times that can adversely affect applications.

The Microsoft Technical Bulletin Configuring MPIO Timers describes the following **PDORemovePeriod** setting:

"This setting controls the amount of time (in seconds) that the multipath pseudo-LUN will continue to remain in system memory, even after losing all paths to the device. When this timer value is exceeded, pending I/O operations will be failed, and the failure is exposed to the application rather than attempting to continue to recover active paths. This timer is specified in seconds. The default is 20 seconds. The max allowed is MAXULONG."

Workaround: If you are using MPIO with a large number of LUNs, edit your registry settings so that

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PDORemovePeriod is set to a higher value.

- If you are using a Fibre Channel connection to a Windows Server running MPIO, use a value of 90 seconds.
- If you are using an iSCSI connection to a Windows Server running MPIO, use a value of 300 seconds.

Managing MPIO LUNs
The Windows Server Device Manager enables you to display or change devices, paths, and load balance policies, and enables you to diagnose and troubleshoot the DSM. After initial installation of MPIO DSM, use Device Manager to verify that it has installed correctly.

If MPIO DSM was installed correctly, each HPE MSA 1050/2050/2052 storage volume visible to the host will be listed as a multipath disk drive as shown in the following example.

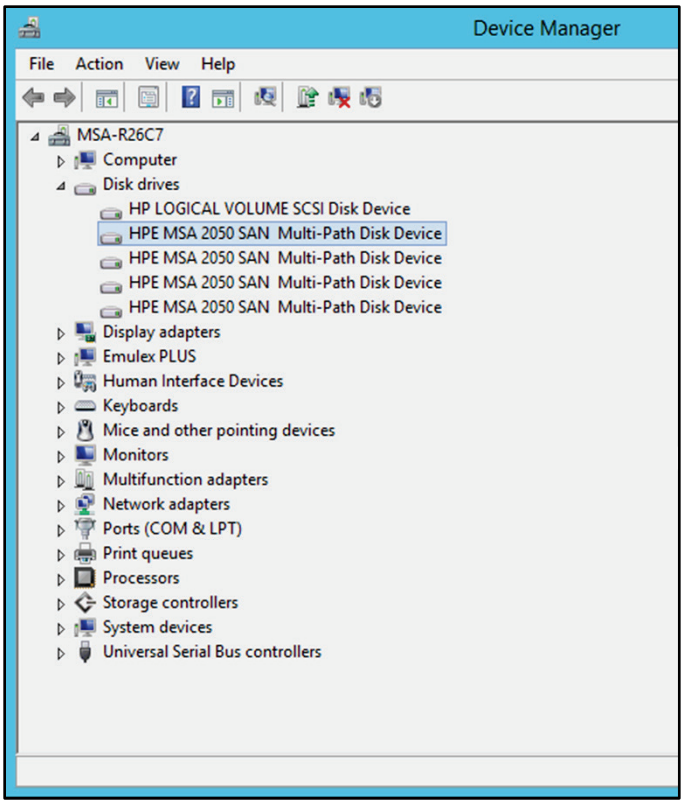

Figure 19. HPE MSA Multi-Path Disk Device

To verify that there are multiple, redundant paths to a volume, right-click the **Multi-Path Disk** Device and select **Properties**.

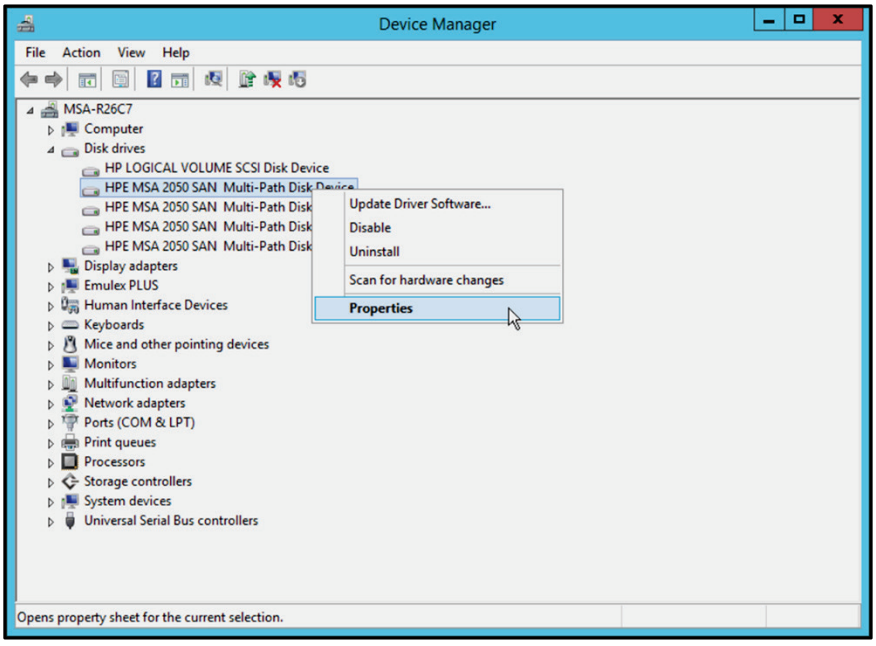

Figure 20. Selecting Properties of HPE MSA Multi-Path Disk Device

Click the MPIO tab, which enables you to view or change the load balance policy and view the number of paths and their status.

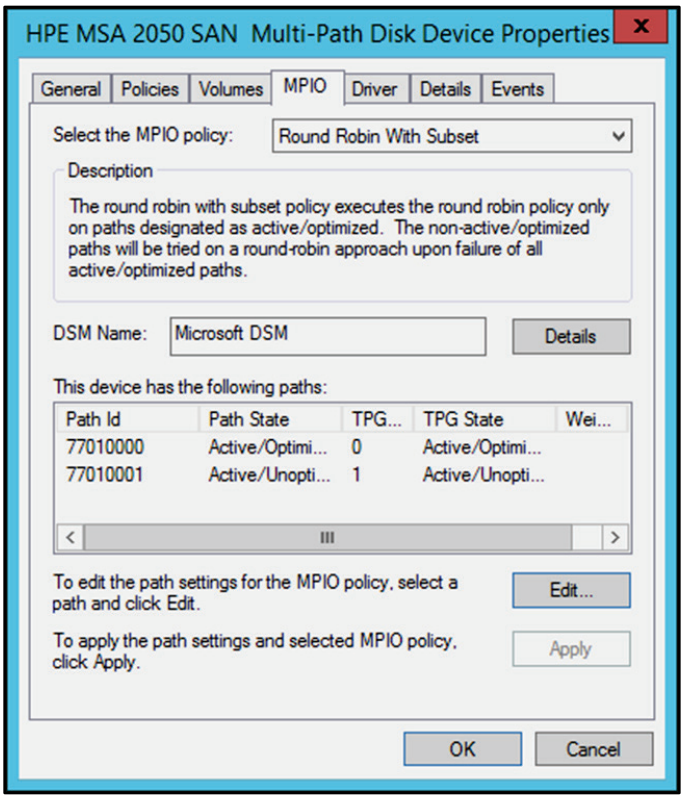

Figure 21. HPE MSA Multi-Path Disk Device Properties—MPIO tab

The Details tab shows additional parameters.

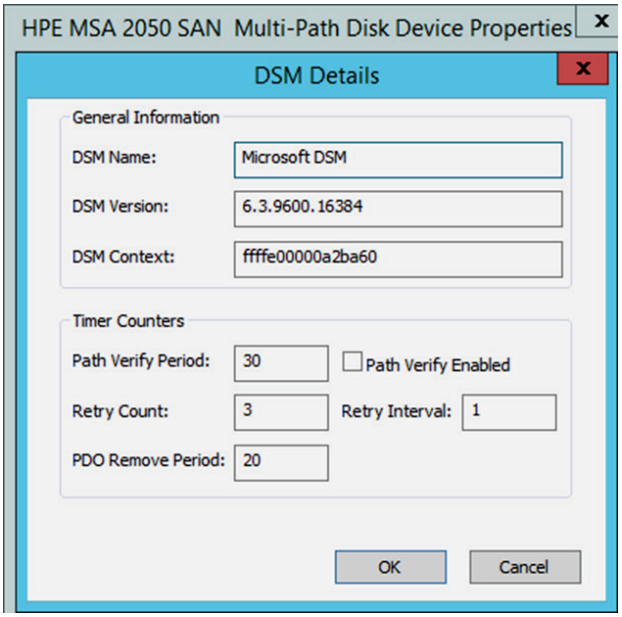

Figure 22. HPE MSA Multi-Path Disk Device Properties—DSM Details

Snapshots

HPE MSA snapshot services enable increased data protection by creating recovery points for your data by taking a picture of your data at a specific point in time. Snapshots are then maintained even as data continues to change. In the event of a failure, you can recover to any previous snapshot. Snapshots are a great complement to tape or disk backup strategy.

HPE MSA snapshot functionality is controller-based, so host resources are not used. HPE MSA 1050/2050/2052 snapshot services utilize redirect-on-write capabilities.

Hewlett Packard Enterprise recommends utilizing the snapshot functionality for data protection. Review the following snapshot space management guidelines when using HPE MSA snapshots.

Snapshots on virtual volumes

How do you use snapshots on virtual volumes?

Occasionally—If you already maintain your snapshot space, then you probably do not need to change anything. The system automatically sets the limit at 10% of the pool and only notifies you if a threshold is crossed.

Regularly—If you want the system to remove old snapshots, you might want to consider changing the limit policy to delete. Only unmapped snapshots that are leaves of a snapshot tree are considered for deletion. The oldest and lowest priority snapshots are deleted first.

What is the rate of change of your data?

Based on the rate of change in your data and the desired snapshot retention, adjust the snapshot space limit and thresholds accordingly.

Note

Retention policies apply to individual volumes and are inherited. If you set the retention level of a base volume, this will not affect existing snapshots of that volume but the snapshots created after setting the retention level.

Dual power supplies

HPE MSA 1050/2050/2052 Array Enclosure and supported Disk 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.

Reverse cabling of Disk Enclosures

HPE MSA 1050/2050/2052 firmware supports both fault-tolerant (reverse cabling) and straight-through SAS cabling of Disk Enclosures. Fault-tolerant cabling allows any Disk Enclosure to fail or be removed without losing access to other Disk Enclosures in the chain. For the highest level of fault tolerance, use fault-tolerant (reverse) cabling when connecting Disk Enclosures.

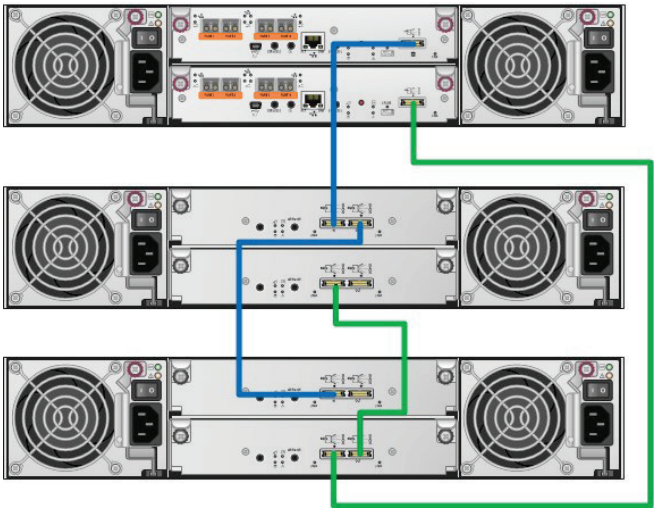

Figure 23. Reverse cabling example using the HPE MSA 2050 system

See the MSA Cable Configuration Guide for more details on cabling the HPE MSA 1050/2050/2052.

HPE MSA 1050/2050 Cable Configuration Guides can be found on HPE MSA support pages.
For the HPE MSA 1050, visit: hpe.com/support/msa1050. For the HPE MSA 2050, visit hpe.com/support/msa2050/.

Create Disk Groups across Disk Enclosures

HPE recommendation is to stripe Disk Groups across shelf enclosures to enable data integrity in the event of an enclosure failure. A Disk Group created with RAID 1, 10, 5, or 6 can sustain one or more Disk Enclosure failures without loss of data depending on RAID type. Disk Group configuration should take into account HPE MSA drive sparing methods such as global and dynamic sparing.

Drive sparing

Drive sparing is recommended to help protect data in the event of a disk failure in a fault-tolerant Disk Group (RAID 1, 5, 6, or 10) 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 (that is, SAS SSD, Enterprise SAS, Midline SAS, or SED drives). HPE MSA 1050/2050/2052 supports global and dynamic sparing. HPE MSA 1050/2050/2052 will reconstruct a critical or degraded Disk Group.

Important
An offline or quarantined Disk Group is not protected by sparing.

Supported spare types:

- **Global spare**—Reserved for use by any fault-tolerant Disk Group to replace a failed disk. The array supports up to 16 global spares per system. At least one Disk Group must exist before you can add a global spare.
- **Dynamic spare**—All available drives are available for sparing. If HPE MSA has available drives and a Disk Group becomes degraded, any available drive can be used for Disk Group reconstruction.

Sparing process

When a disk fails in a redundant Disk Group, the system first looks for a 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 Disk Group will be in either a degraded or critical status until the parity data is completely written to the spare, at which time the Disk Group returns to fault tolerant status. In the case of global spares, after the failed drive is replaced, the replacement drive needs to be added back as a global spare.

Best practice for sparing is to configure at least one spare for every fault-tolerant Disk Group in the system. Dynamic sparing is enabled by default.

Note
Warnings alerts are sent out when the last global spare is used in a system.

Drive replacement

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

Implement Remote Snap replication

HPE MSA 1050/2050/2052 Storage 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 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.

The MSA 2050/2052 supports up to four peer connections; the MSA 1050 supports only one.

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

Use the secured web access (HTTPS) when utilizing Remote Snap replication on HPE MSA.

Note
Remote Snap requires a purchasable license in order to implement. On HPE MSA 2052, the ADS Software Suite includes HPE MSA Remote Snap Software, but will require a license key from Hewlett Packard Enterprise and must be installed on HPE MSA 2052 to enable Remote Snap.

To obtain a Remote Snap license, go to myenterpriselicense.hpe.com/.

See MSA Remote Snap software technical white paper at hpe.com/h20195/v2/Getdocument.aspx?docname=4AA1-0977ENW.

Best practices to enhance performance

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

Cache settings

One method to tune the Storage System 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, enabling write-back cache is the best practice. With the transportable cache feature, write-back caching can be used in either a single or dual controller system. See HPE MSA 1050 User Guide or HPE MSA 2050 User Guide for more information on the transportable cache feature.

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 unless cache optimization is set to no-mirror. 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 Disk Groups.

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 nonvolatile 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 see HPE MSA 1050/2050 SMU Reference 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.

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 Logical Block Addressing (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 Disk Groups 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.

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 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 5. This section provides a basic starting point for fine-tuning your system, which should be done during performance baseline modeling.

Table 5. 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

Host port utilization for maximum performance

Maximum array performance may require the use of multiple host ports per controller, as well as utilizing multiple ports on the hosts accessing the array. For example, a single 16 Gb FC port can only produce roughly 1.6 GB/s throughput. Likewise a 1 Gb iSCSI port can only produce roughly 110 MB/s, which not only limits sequential throughput, but will limit IOPS per port as well.

A 1 Gb host port running at 110 MB/s unidirectionally with an 8 KB I/O size will only achieve roughly 14,200 IOPS on that single port, and a maximum of only 113,600 IOPS with all host ports being used.

Other methods to enhance array performance

There are other methods to enhance performance of HPE MSA 1050/2050/2052. In addition to the cache settings, the performance of HPE MSA 1050/2050/2052 array can be maximized by using the following techniques.

Place higher performance SSD and SAS drives in the Array Enclosure

The HPE MSA 1050/2050/2052 controller is designed to have a single SAS link per drive in the Array Enclosure and only four SAS links to Disk Enclosures. Placing higher performance drives (that is, 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 Disk 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.
- Disk Groups should be balanced between the two controllers.
- Disk drives should be balanced between the two controllers.
- Cache settings should be set to match Table 5 for the application.
- In order to get the maximum sequential performance from a Disk Group, you should only create one volume per Disk Group. Otherwise, you will introduce randomness into the workload when multiple volumes on the Disk Group are being exercised concurrently.
- Distribute the load across as many drives as possible.
- Distribute the load across multiple array controller host ports.

Creating Disk Groups

When creating Disk Groups, best practice is to add them evenly across both pools. With at least one Disk Group 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 Disk Groups can be important for performance. However, there are some trade-offs with cost when using the higher fault-tolerant RAID levels.

See Table 6 for the strengths and weaknesses of the supported HPE MSA 1050/2050/2052 RAID types.

Table 6. HPE MSA 1050/2050/2052 RAID levels

RAID level	Minimum disks	Allowable disks	Description	Strengths	Weaknesses
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.
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 disk failure	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-Disk Groups	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.

Note
RAID types NRAID and RAID 0 can only be created when setting Read Cache. NRAID and RAID 0 are used with Read Cache as the data in the Read Cache SSDs is duplicated on either the Standard or Archive Tier.

Volume mapping

For increased performance, access the volumes from the ports on the controller that owns the Disk Group, which would be the preferred path. Accessing the volume on 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

The performance capabilities of SSDs are a great alternative to traditional spinning hard disk drives (HDDs) in highly random workloads. SSDs cost more in terms of dollars per GB throughput than spinning hard drives; however, SSDs cost much less in terms of dollars per IOPS. Keep this in mind when choosing the numbers of SSDs per HPE MSA 1050/2050/2052 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.

If you have the Performance Tier license, you can use SSDs for virtual disk groups. When combined with virtual disk groups that consist of other classes of disks, improved read and write performance is possible through automated tiered storage. Alternatively, you can use one or two SSDs in read-cache disk groups to increase read performance for pools without a Performance tier. The application workload of a system determines the percentage of SSDs of the total disk capacity that is needed for best performance.

Data such as database indexes and TempDB files would benefit from SSDs 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 HPE MSA 1050/2050/2052 controllers. There is a point where these two curves intersect. At the intersecting point, additional SSDs will not increase performance. See Figure 24.

HPE MSA 1050/2050/2052 reaches this bandwidth at a low number of SSDs. For the best performance using SSDs on HPE MSA 1050/2050/2052, use a minimum of four SSDs with one 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 Disk Group owned by each controller. This would require six SSDs for RAID 5 and eight SSDs for RAID 6. All SSD volumes should be contained in fault-tolerant Disk Groups for data integrity.

Base the number of SSDs to use on the amount of space that is needed for your highly random, high performance data set. For example, if the amount of data that is needed to reside in the SSD volumes exceeds a RAID 1 configuration, use a RAID 5 configuration.

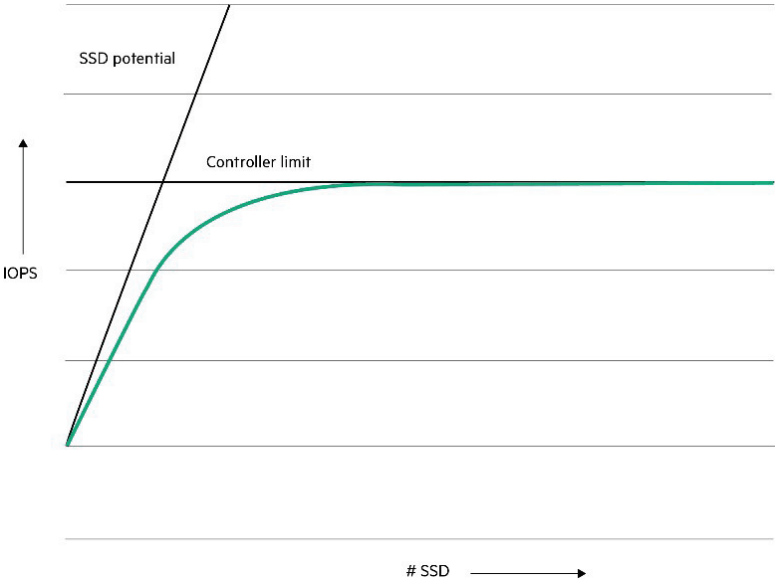

Figure 24. SSD performance potential vs. HPE MSA 1050/2050/2052 controller limit

Note

There is no limit to the number of SSDs that can be used in HPE MSA 1050/2050/2052 array system.

SSD Read Cache

SSD Read Cache is a feature that extends HPE MSA 1050/2050/2052 controller cache.

Read Cache is most effective for workloads that are high in random reads. The user should size the Read Cache capacity based on the size of the hot data being randomly read. A maximum of two SSD drives per pool can be added for Read Cache.

Hewlett Packard Enterprise recommends beginning with one SSD assigned per Storage Pool for Read Cache. Monitor the performance of the Read Cache and add more SSDs as needed.

There is a 4 TB maximum limit per pool for Read Cache.

Note

You can have SSDs in a fault-tolerant Disk Group as a Performance Tier or as a non-fault-tolerant (up to two disks) Disk Group as Read Cache. However, neither pool can have both a Performance Tier and a Read Cache. For example, Pool A can have a Performance Tier and Pool B can have a Read Cache.

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 HPE MSA 1050/2050/2052 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.

Full Disk Encryption

Full Disk Encryption (FDE) is a data security feature used to protect data on disks that are removed from a storage array. The FDE feature uses special Self-Encrypting Drives (SEDs) to secure user data. FDE functionality is available on HPE MSA 2050.

The SED is a drive with a circuit built into the drive's controller chipset, which encrypts/decrypts all data to and from the media automatically. The encryption is part of a hash code, which is stored internally on the drive's physical medium. In the event of a failure of the drive or the theft of a drive, a proper key sequence needs to be entered to gain access to the data stored within the drive.

Full Disk Encryption on HPE MSA 2050

HPE MSA 2050 Storage System uses a passphrase to generate a lock key to enable securing the entire Storage System. All drives in a Full Disk Encryption (FDE) secured system are required to be SED (FDE capable). By default, a system and SED drive are not secured and all data on the disk may be read or written by any controller. The encryption on the SED drive conforms to FIPS 140-2.

To secure an HPE MSA 2050, you must set a passphrase to generate a lock key and then FDE secures the system. Simply setting the passphrase does not secure the system. After an HPE MSA 2050 system has been secured, all subsequently installed disks will automatically be secured using the system lock key. Non-FDE capable drives will be unusable in a secured HPE MSA 2050 system.

Note
The system passphrase should be saved in a secure location. Loss of the passphrase could result in loss of all data on HPE MSA 2050 Storage System.

All HPE MSA 2050 Storage Systems will generate the same lock key with the same passphrase. It is recommended that you use a different passphrase on each FDE secured system.

If you are moving the entire Storage System, it is recommended to clear the FDE keys prior to system shutdown. This will lock all data on the disks in case of loss during shipment. Only clear the keys after a backup is available and the passphrase is known. Once the system is in the new location, enter the passphrase and the SED drives will be unlocked with all data available.

SED drives which fail in an FDE secured system can be removed and replaced. Data on the drive is encrypted and cannot be read without the correct passphrase.

Best practices for firmware updates

The following sections detail common firmware update best practices for HPE MSA 1050/2050/2052.

General HPE MSA 1050/2050/2052 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 configuration changes are in progress, monitor them using the SMU or CLI and wait until they are completed before proceeding with the upgrade.
- Do not power cycle or restart devices during a firmware update. If the update is interrupted or there is a power failure, the module could become inoperative. Should this happen, contact HPE customer support.
- After the device firmware update process is completed, confirm the new firmware version is displayed correctly via one of the HPE MSA management interfaces—for example, SMU or CLI.

HPE MSA 1050/2050/2052 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. You can check the firmware upgrade readiness on your system with the CLI command: `check firmware-upgrade-health`.
- When planning for a firmware upgrade, schedule an appropriate time to perform an online upgrade.
  - 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 time-outs.
- When planning for a firmware upgrade, allow sufficient time for the update.
  - Plan for firmware upgrade to take at least 30 minutes; large configurations or heavy I/O workload during upgrade can extend this time.
  - When upgrading or downgrading 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 use best practices to upgrade the controllers in the system. Using a Smart Component is recommended.

HPE MSA 1050/2050/2052 disk drive firmware update best practices

- Disk drive upgrades on the HPE MSA 1050/2050/2052 Storage Systems is an offline process. All host and array I/O must be stopped prior to the upgrade.
- If the drive is in a Disk Group, 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 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 all the drives are updated.

Miscellaneous best practices

The following sections detail miscellaneous best practices for HPE MSA 1050/2050/2052.

Boot from storage considerations

When booting from storage on HPE MSA 1050/2050/2052, set the Volume Tier Affinity to Performance even if an SSD Performance Tier is not available. Choosing this option will keep operating system boot latencies to a minimum and will avoid delays on the operating system.

Even though the Volume Tier Affinity is set to Performance, the data on the boot drive can move to other tiers if not accessed.

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

HPE MSA 2050/2052 Storage System uses specific small form-factor pluggable (SFP) transceivers that will not operate in HPE 8 Gb and 16 Gb switches. Likewise, HPE Fibre Channel switches use SFPs, which will not operate in the HPE MSA 2050/2052.

HPE MSA 2050/2052 controllers do not include SFPs. Qualified SFPs for HPE MSA 2050/2052 are available for separate purchase in four packs. Both 8 Gb and 16 Gb SFPs are available to meet the customer need and budget constraints. All SFPs in an HPE MSA 2050/2052 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 HPE MSA 2050/2052 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 HPE Transceiver Replacement Instructions document at: h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-c02010667&lang=en-us&cc=us.

HPE MSA 1050/2050/2052 iSCSI considerations

When using HPE MSA 1050/2050/2052 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 HPE MSA 1050/2050/2052 SAN controller. This Private LAN is the storage network and the Public LAN is used for management of HPE MSA 1050/2050/2052. The storage network should be isolated from the Public LAN to improve performance.

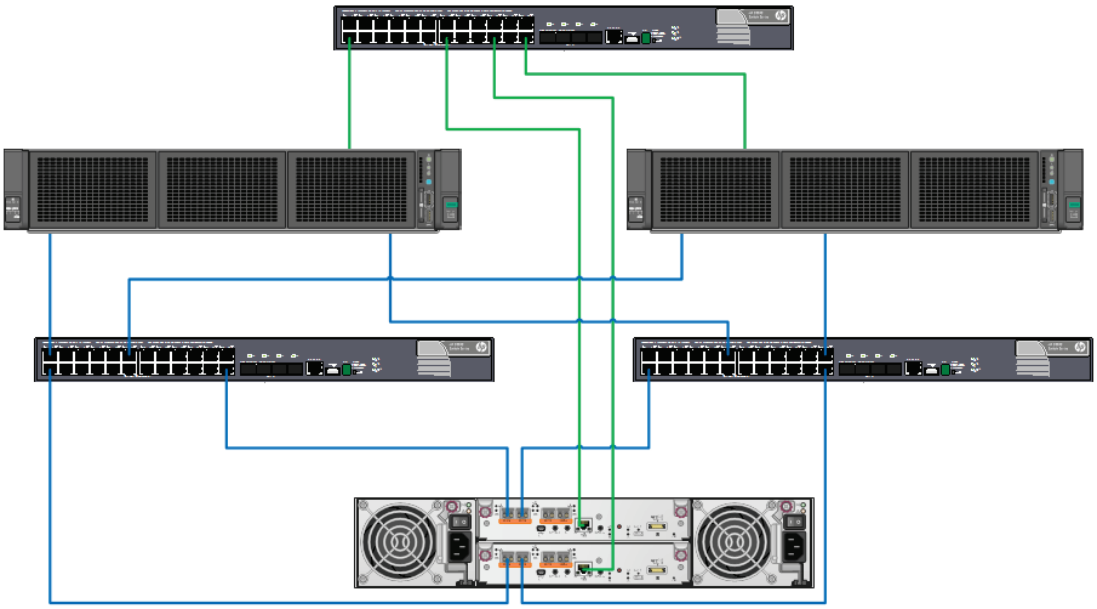

Figure 25. HPE MSA 1050/2050/2052 SAN iSCSI network

IP address scheme for the controller pair

The MSA 2050 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: MSA 2050 SAN:

Controller A port 1: 10.10.10.100

Controller A port 2: 10.11.10.110

Controller A port 3: 10.10.10.120

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. HPE MSA reserves some of this frame size; the current maximum frame size is 1400 for a normal frame and 8900 for a jumbo frame. This frame

maximum can change without notification. If you are using jumbo frames, make sure to enable jumbo frames on all network components in the datapath.

Hot adding Disk Enclosures

This section details the necessary steps required for upgrading both fault-tolerant and straight-through configurations while HPE MSA 1050/2050/2052 Storage System is online and operational. This is accomplished without removing the storage from use, provides a transparent upgrade path with no disruption of service, and assures data integrity throughout the process.

Adding additional storage to existing HPE MSA 1050/2050/2052 Storage Systems will not present a major impact to current data center operations and should be performed during a maintenance window where I/O to the affected Storage System(s) will be minimized.

Differences in Disk Enclosures

There are two different types of Disk Enclosures that you may choose for HPE MSA 1050/2050/2052 Storage System depending on whether you are using small or large form factor drives. The two Disk Enclosure examples are shown in Figure 26 and Figure 27. Note that the backs of the small and large form factor disk enclosures look the same.

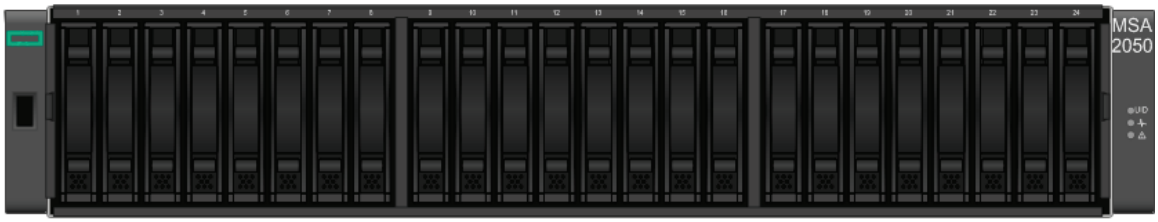

Figure 26. HPE MSA 2050 SFF Disk Enclosure (without the bezel)

Figure 27. HPE MSA 2050 LFF Disk Enclosure (without the bezel)

The common feature to note in these examples is that both Disk Enclosures have two SAS ports per I/O module. The left port on each module is the **In** port and the right port is the **Out** port.

Determining the current cabling configuration

If the Storage System to be upgraded is an HPE MSA 1050/2050/2052 single-shelf system, addition of the second shelf requires no special considerations. This section is intended to address Disk Enclosure upgrades where:

- You are moving from an HPE MSA 1050/2050/2052 with one or more Disk Enclosures to add additional Disk Enclosures.
- You are moving from a straight-through cabling configuration to a fault-tolerant cabling design. To determine the current cabling configuration, see **Figure 28** and **Figure 29**.

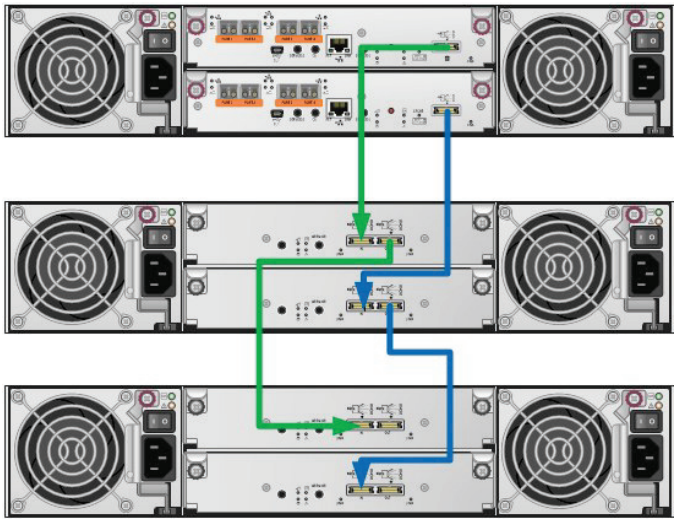

Figure 28. Straight-through cabling configuration

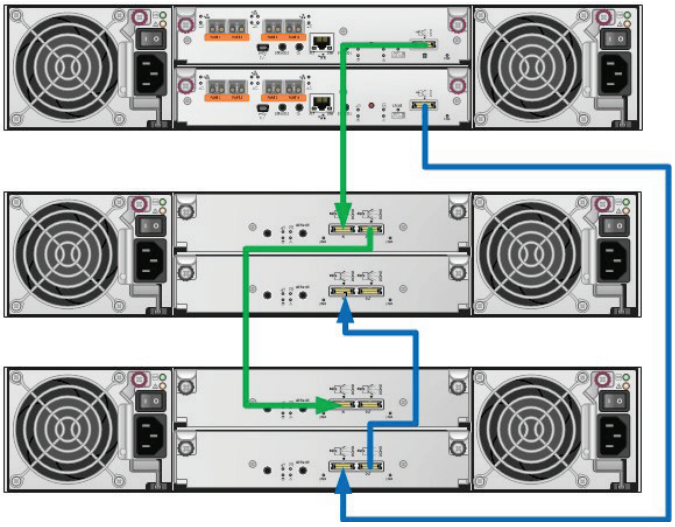

Figure 29. Fault-tolerant cabling configuration

Implementing the upgrade

To ensure data integrity, it is important to determine the type of configuration that is currently implemented and what type of end configuration will be needed moving forward. The following steps will refer you to the correct section to ensure that during the upgrade, there is no risk to data integrity.

If you are currently using:

- An HPE MSA 1050/2050/2052 with a straight-through cabling configuration (**Figure 28**) and wish to add one or more additional Disk Enclosures of storage while continuing to use the same cabling configuration, see the section entitled Extending the straight-through cabling configuration.
- An HPE MSA 1050/2050/2052 with a straight-through cabling configuration (**Figure 28**) and requires the system to be reconfigured into a fault-tolerant cabling configuration or is adding an additional Disk Enclosure to your current configuration and want to implement a fault-tolerant configuration, see the section entitled Upgrading from straight-through cabling to a fault-tolerant cabling configuration.

- You have an existing HPE MSA 1050/2050/2052 Storage System currently configured for fault-tolerant design (**Figure 29**) and need additional Disk Enclosures for the existing configuration; see the section entitled Adding additional storage to an existing fault-tolerant environment.

Extending the straight-through cabling configuration

Use the following steps to complete the upgrade:

1. Install and mount the new Disk Enclosure and then install all drives.
2. Power on the new Disk Enclosure.
3. Ensure all currently connected Disk Enclosures are powered on and functioning properly. Verify this by using the Storage Management Utility (SMU).
4. Connect the new SAS cable (**Figure 30 green** cable) from the A port **Out** of the current HPE MSA Storage System to the A port **In** of the new storage enclosure.
5. Allow the system a few moments to initialize and recognize the addition of the new Disk Enclosure.
6. Restart or reload the SMU for the HPE MSA 1050/2050/2052 and verify that the system has recognized the new storage.
7. Connect the new SAS cable (**Figure 30 blue** cable) from the B port **Out** of the current HPE MSA Storage System to the B port **In** of the new Disk Enclosure.
8. Allow the system a few moments to initialize and recognize the addition of the new ports.
9. Ensure all firmware is at the latest levels for all Disk Enclosures.
10. Verify proper system operation by validating all drives in the new Disk Enclosure. Validate this by using the Storage Management Utility (SMU).

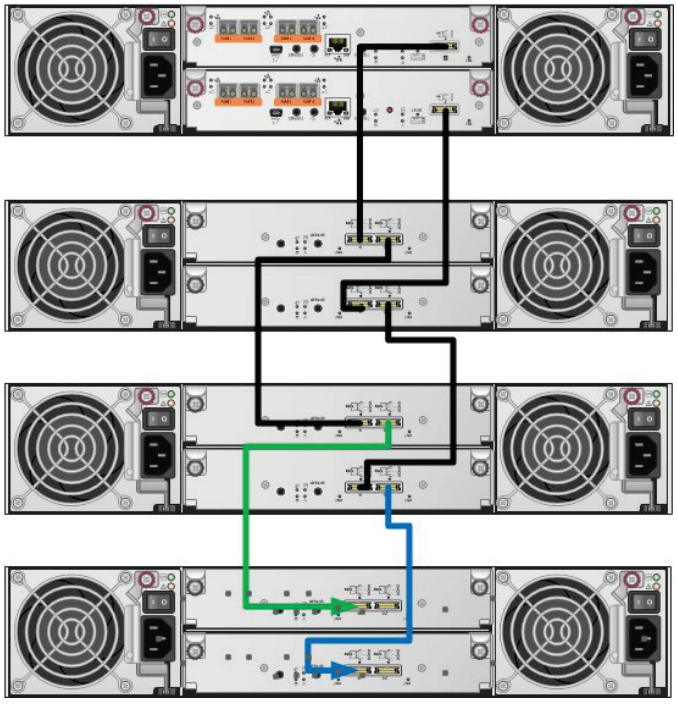

Figure 30. Extended straight-through cabling configuration

Upgrading from straight-through cabling to a fault-tolerant cabling configuration

Use the following steps to complete the upgrade:

1. Install and mount any new Disk Enclosures, if needed, and then install all drives.
2. Power on the new Disk Enclosures.
3. Ensure all Disk Enclosures are powered on and functioning properly. Verify this by using the SMU.
4. Using the SMU or the Command Line Interface (CLI), perform a shutdown of the B storage controller to ensure that any pending write or read requests are flushed from the controller B cache.
5. Remove existing cabling from the B channel chain as indicated in **Figure 31**.

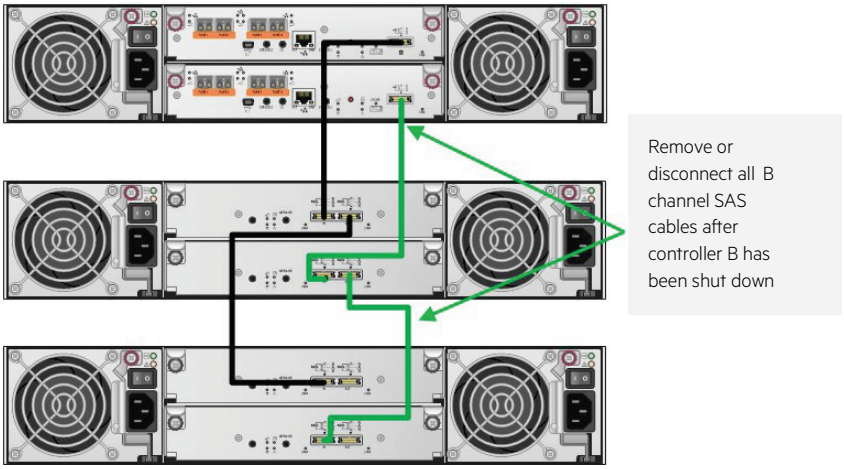

Figure 31. Straight-through cabling configuration

6. Connect an SAS interconnect cable from controller B on HPE MSA 1050/2050/2052 Array Enclosure to the last **In** port B channel on the Disk Enclosures as indicated in **Figure 32**.

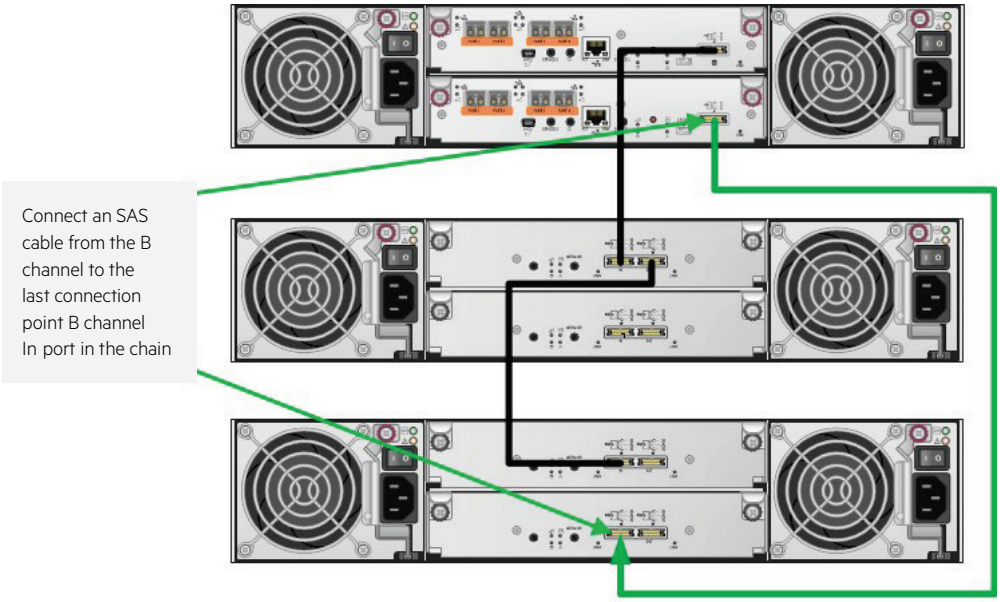

Figure 32. Initial B channel cabling configuration

7. Connect or reconnect an SAS interconnect cable from the last **Out** port B channel to the previous Disk Enclosure's B channel In port, as indicated in **Figure 33**.
8. Using the SMU or the CLI, restart the B storage controller.
9. Allow the Storage System a few moments to initialize and recognize the addition of the new ports.
10. Ensure all firmware is at the latest levels for all enclosures.
11. Verify proper system operation by validating that all drives in the new Disk Enclosure are seen.

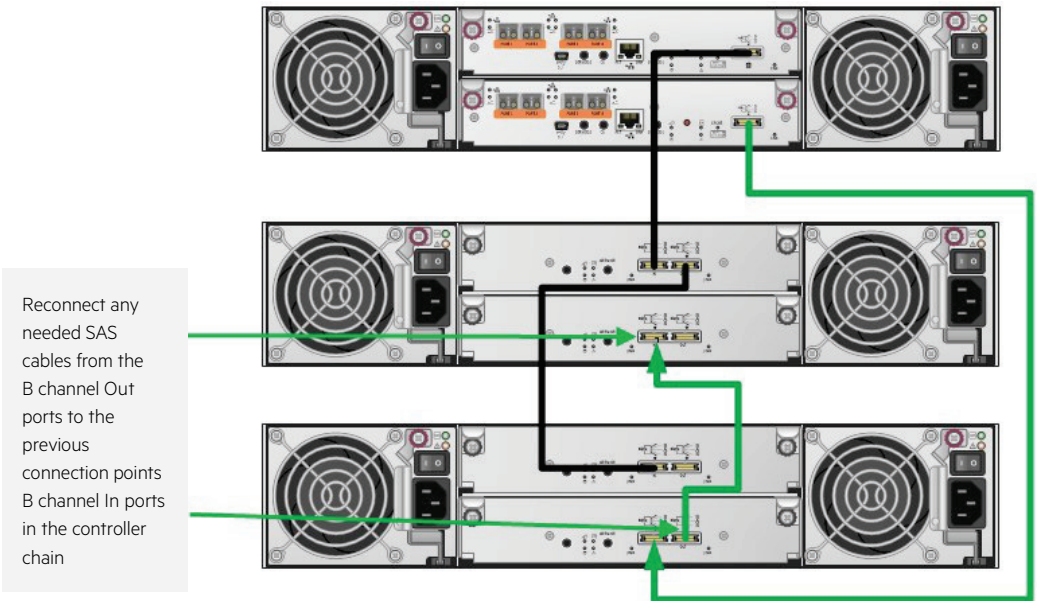

Figure 33. Completed B channel cabling configuration

Adding additional storage to an existing fault-tolerant environment

Use the following steps to complete the upgrade:

1. Install and mount the new Disk Enclosure and then install all drives.
2. Power on the new Disk Enclosures.
3. Ensure all Disk Enclosures are powered on and functioning properly. Verify this by using the SMU.
4. Connect the new SAS cable (**Figure 34 green** cable) from the A port **Out** of the current HPE MSA Storage System to the A port **In** of the new Disk Enclosure.
5. Allow the system a few moments to initialize and recognize the addition of the new Disk Enclosure.
6. Restart or reload the SMU for HPE MSA 1050/2050/2052 and verify that the system has recognized the new storage.
7. Using the SMU or the CLI, perform a shutdown of the B storage controller to ensure that any pending write or read requests are flushed from the controller B cache.

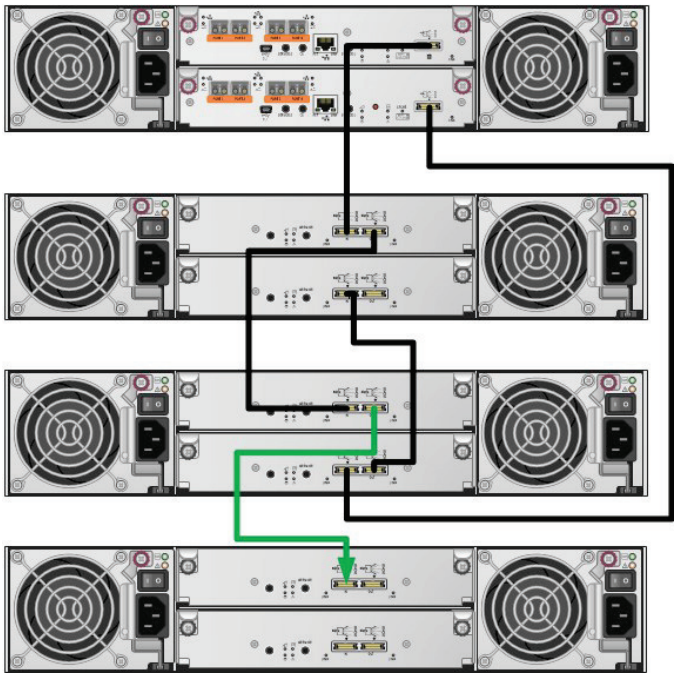

Figure 34. Adding new SAS cable to extend a channel to the additional shelf

8. Disconnect the existing SAS cable (**Figure 35** dashed **blue** cable) from controller B of the current HPE MSA 1050/2050/2052 Storage System to the B port In of the existing Disk Enclosure and reconnect the SAS cable to the last Disk Enclosure in the chain (**Figure 36 blue** cable).

Ensure B controller is shut down and then disconnect SAS cable at the last connection point in the chain

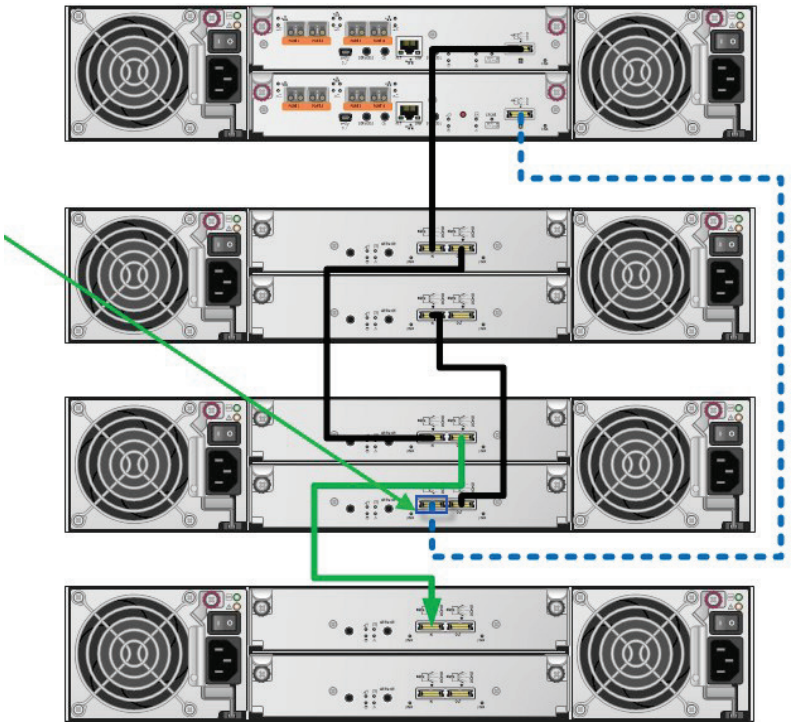

Figure 35. Breaking the B channel chain to extend B channel to additional shelf

9. Connect the new SAS cable (**Figure 36 blue** cable) from the B port **Out** of the new HPE MSA 1050/2050/2052 Disk Enclosure to the B port **In** of the previous Disk Enclosure.
10. Using the SMU or the CLI, restart the B controller.
11. Allow the system a few moments to initialize and recognize the addition of the new ports.
12. Ensure all firmware is at the latest levels for all Disk Enclosures.
13. Verify proper system operation by validating that all drives in the new Disk Enclosure are seen.

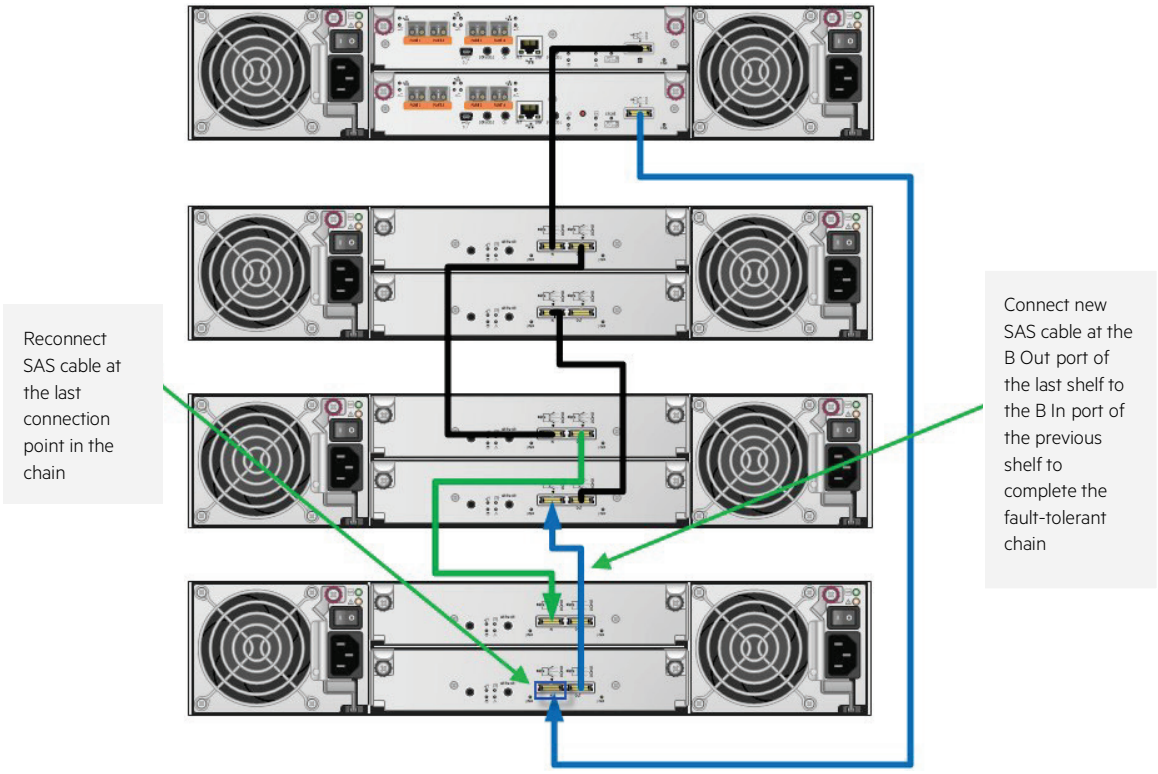

Figure 36. Completing the fault-tolerant chain on the B channel

Summary

HPE MSA 1050/2050/2052 administrators should determine the appropriate levels of fault tolerance and performance that best suits their needs. Understanding the workloads and environment for HPE MSA SAN is also important. Following the configuration options listed in this paper can help optimize HPE MSA 1050/2050/2052 array accordingly.

Learn more at

hpe.com/storage/msa

a00015961ENW, September 2017, Rev. 1