Technical white paper

HPE MSA 2070 and HPE MSA 2072 Storage best practices

Contents

Executive summary 3
Intended audience 3
Connectivity best practices 3
Naming hosts 3
iSCSI 3
MPIO 5
Maintaining supported configurations 7
Best practices for maintaining system health 7
Users 7
Firmware 8
System monitoring 10
Background scrubbing 11
Data protection 12
Periodic health checks 13
Storage best practices 13
Disk drives 13
Choosing disk group types 14
Sparing 15
Single vs. dual pools 20
Thin Provisioning 22
Full disk encryption 22
Capacity expansion 23
Volume mapping 24
Summary 26
Resources 26

Executive summary

This paper provides guidance on configuring HPE MSA Storage arrays to meet recommended best practices from Hewlett Packard Enterprise. These recommendations help improve application availability and performance, as well as improve system security. This paper is not a user guide but complements other official documentation that explains HPE MSA Storage technology and how to configure array settings. These best practices focus on providing clear recommendations rather than detailed information on the technologies they reference. Technology details in the best practices documents for previous generation arrays have migrated between the HPE MSA Gen7 virtual storage technical reference guide and technology-specific documents found in the HPE Support Center.

Intended audience

This paper is for those tasked with the installation, configuration, and ongoing maintenance of HPE MSA Storage systems. Additionally, this paper assists technical sales staff in designing optimal solutions.

Connectivity best practices

Connectivity is the fundamental connection between devices. This section covers topics regarding the management network, data network, and path management software.

Naming hosts

Best practices for naming hosts include:

– Best practice: Group initiators (IDs) as hosts and define friendly names for them

- System default: None
- Detail: The default HPE MSA Storage Management Utility (SMU) (web-based user interface) behavior is to not allow the mapping of a volume to a host without first creating a host of one or more initiators. Initiator names such as the World Wide Port Name (WWPN), which is applicable to Fibre Channel and SAS, and the iSCSI Qualified Name (IQN), which is applicable to iSCSI, are composed of long alphanumeric strings that are difficult to remember or recognize. HPE recommends providing port-based naming, which follows meaningful device inventory naming within an organization.
- Example: Host name: dl380_gen11_1

  ID #1 Nickname: 10009cdc7172690a → dl380_gen11_1_port_0

  ID #2 Nickname: 10009cdc71726909 → dl380_gen11_1_port_1

Note

It is only possible to map volumes to individual initiators through the CLI.

iSCSI

Best practices for configuring an iSCSI connection include:

– Best practice: Use three network ports per host

- System default: None

  Detail: To achieve isolation between management and application traffic, HPE recommends using separate networks for management and iSCSI traffic. Additionally, at least two physical connections should provide connectivity to the data networks. Actual network topography may vary depending on the use of VLANs and network switch virtualization.

Note

A host does not require connectivity to either management port to access data.

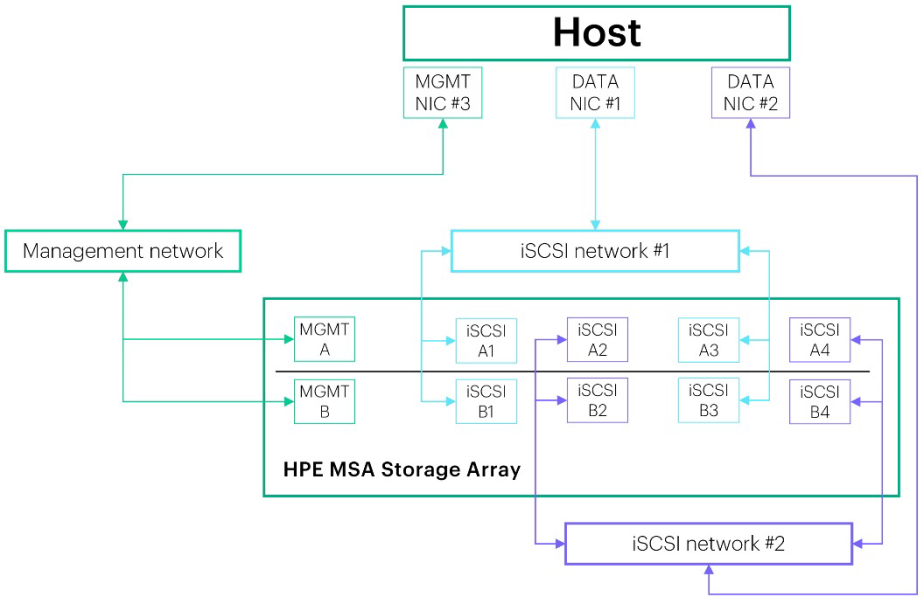


Figure 1. A sample iSCSI network topography

– Best practice: Use at least two isolated data networks when more than four array host ports are in use.

- System default: None
- Detail: To both improve performance and minimize application disruption, HPE recommends not configuring more than eight paths to a volume. The more paths configured, the longer the failover time can be when active paths are lost. Using two networks limits an initiator to a total of eight paths to a volume (four active/optimized, four active/unoptimized).
- Example:

  Controller A, Port 1 (A1): 10.10.10.10/24

  Controller B, Port 1 (B1): 10.10.10.11/24

  Controller A, Port 2 (A2): 10.10.20.10/24

  Controller B, Port 2 (B2): 10.10.20.11/24

  Controller A, Port 2 (A3): 10.10.10.12/24

  Controller B, Port 2 (B3): 10.10.10.13/24

  Controller A, Port 2 (A4): 10.10.20.12/24

  Controller B, Port 2 (B4): 10.10.20.13/24

  Host data NIC #1: 10.10.10.100/24

  Host data NIC #2: 10.10.20.100/24

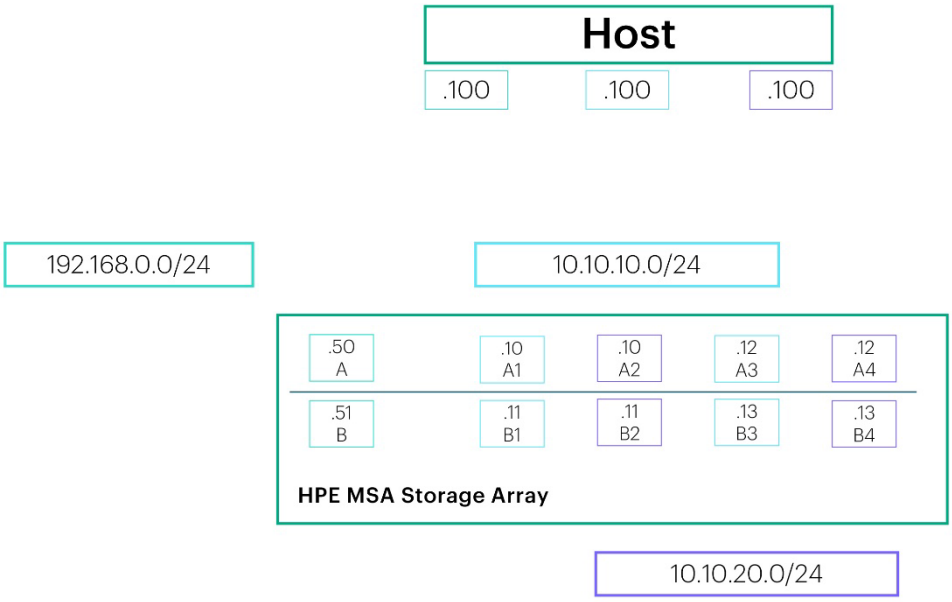


Figure 2. Example of iSCSI network subnet assignments

– Best practice: Set all devices on the same network to use jumbo frames when configured for an HPE MSA Storage array

- System default: Disabled (MTU = 1400)
- Detail: Jumbo frames increase the payload limit per Ethernet frame and can improve end-to-end performance. However, all devices within the datapath must also use jumbo frames to avoid packet fragmentation or loss. HPE MSA Storage arrays advertise a jumbo frame payload of 8900 bytes. Sending devices usually agree with the MTU advertised by the receiver. If a device is unable to adjust its MTU automatically, do so manually.
- Example: Issuing this command in the HPE MSA CLI enables jumbo frames:

```
set iscsi-parameters jumbo-frames enabled
```

MPIO

– Best practices for configuring an MPIO connection include:

- Best practice: Install and configure multipath software on connected hosts. Consult with the Single Point of Connectivity Knowledge (SPOCK) for HPE Storage products for current MPIO configuration recommendations.
- System default: Not installed
- Detail: Multipath software provides load balancing and tolerance to link failure between a host and a storage array. Without multipath software, a volume mapped to a host appears as multiple physical disks, each with a single path. With multipath software, a volume can appear as a single physical disk that has multiple paths. A path is a connection between an initiator (host bus adapter [HBA] port or iSCSI software initiator) and a target (HPE MSA Storage array host port). When there are multiple active paths to the owning controller/pool for a given volume, multipath software can improve performance by distributing traffic evenly among those paths.

  When configuring MPIO, the product string must be formatted correctly where product name is 2070:

  ```
  'MSA product name protocol'.
  ```

Note
HPE MSA 2072 Storage product name is the same as the HPE MSA 2070 Storage.

– Example: Microsoft Windows Server

Issue the following commands in Windows Server PowerShell to enable MPIO with an HPE MSA 2070 FC Storage array:

```
Install-WindowsFeature –name MultiPath-io
mpclaim –n –I –d “HPE MSA     2070 FC”
```

Note
There are five spaces between **HPE** and **MSA** in the sample command

– Example: Red Hat® Enterprise Linux®

Add the following to /etc/multipath.conf devices section:

```
device {
vendor “HPE"
product “MSA     2070 iSCSI"
path_grouping_policy “group_by_prio"
prio “alua"
path_selector “round-robin 0"
failback “immediate"
no_path_retry 18
}
```

Note
HPE MSA Storage arrays do not support link aggregation.

– Best practice: Modify MPIO timers on Microsoft Windows Server hosts when connecting to large numbers of logical unit numbers (LUNs).

- System default: 20 seconds
- Detail: Microsoft Windows Server has a default period of 20 seconds in which it retains a multipath pseudo-LUN in memory, even after losing all paths to the device. When this time has passed, pending I/O operations fail and the failure is exposed to applications, rather than continuing to recover active paths. When a Windows host has many volumes (LUNs) mapped, 20 seconds might be too brief a time to wait. This can cause long failover times and adversely affect applications.

HPE recommends modifying the PDORemovePeriod value within the system registry depending on the protocol used:

- Fibre Channel: 90 seconds
- iSCSI: 300 seconds

– Example: Issue the following command at the Microsoft Windows Server command prompt:

```
reg add HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters /t REG_DWORD /v
PDORemovePeriod /d 300 /f
```

Note
A reboot may be required on Microsoft Windows Server systems the first time a device is added to the multipath device list. The reboot may be circumvented by disabling and reenabling the disk device within device manager but is not advised if the disk device is the same as the boot volume.

Maintaining supported configurations

Best practices for maintaining supported configurations include:

– Best practice: Cross-reference support for a configuration through the Single Point of Connectivity Knowledge (SPOCK) for HPE Storage products before deploying HPE storage array or infrastructure upgrades

- System default: N/A
- Detail: HPE only provides support for configurations listed within SPOCK. SPOCK is a database of tested interoperability between all active components in the datapath, such as the operating system, HBAs, SAN switching infrastructure, and arrays, including all firmware and related software. As the storage vendor, HPE, not the operating system vendor, is responsible for issuing support statements. Where possible, HPE works with operating system vendors to provide matching support statements. In the event of differing statements, SPOCK is the final authority.
- Example: A SPOCK end-to-end support statement:
- Windows Server 2025 ? SN1610Q FC HBA ? Brocade SN3600B 32Gb FC switch ? MSA 2070 FC

Best practices for maintaining system health

The availability of application data is paramount to any organization. This section provides recommendations to help maintain the highest levels of availability.

Users

Best practices for maintaining system health for users include:

– Best practice: Disable unsecure protocols.

- System default: Disabled
- Detail: To minimize the possibility of both unauthorized access and vulnerability to attack, HPE recommends that unsecure protocols remain disabled. These include:
  - Telnet
  - FTP
  - SNMP (unless creating an SNMPv3 user)
  - HTTP
  - Debug

– Example: Issuing this command in the HPE MSA CLI or by deselecting the service from within the SMU, disables telnet

```
set protocols telnet disabled
```

– Best practice: Assign LDAP user groups to the standard role.

- System default: N/A
- Detail: For added security, users affiliated with the standard role can perform some array management tasks but cannot manage users or clear logs.
- Example: N/A

Firmware

Array firmware includes software that enables features and functionality and includes internal software components that enhance array stability and performance.

Maintaining current firmware is critical not just for the latest features, but to maintain application availability. Best practices include:

– Best practice: Configure an **Update Server** to automatically receive the latest information on available system firmware.
  - System default: Configured
  - Detail: The Update Server feature periodically downloads a manifest of current firmware and compares it to installed firmware. When more recent firmware is published, the system logs an alert and sends a notification. Consider hosting the manifest on a local server if internet connectivity is unavailable. The Update Server can only be functionally disabled by setting the resource URL to " ".
  - Example: Set the Update Server resource URL to hpe.com/support/MSAmanifest

Note

A valid network configuration, including a gateway and DNS server, must be configured for the Update Server feature to work. It may also be necessary to configure a proxy server.

– Best practice: Have a tested backup strategy before applying firmware.
  - System default: N/A
  - Detail: Applying a firmware update to any component introduces a risk of failure and subsequent data unavailability or loss, although such failures are highly unlikely. HPE recommends a tested backup solution is in place before the application of a firmware update. Strategies can vary greatly depending on the use case ranging from array-based snapshot replication (Remote Snap) to agent-based backup of applications to tape.
  - Example: HPE MSA Remote Snap Replication

– Best practice: Maintain current versions of all component firmware throughout the datapath.
  - System default: N/A
  - Detail: SPOCK support matrices are continuously updated to reflect changing conditions as new firmware is released. What is supported when first installed might not be supported months later. Firmware releases sometimes bring new features, but often also close security holes and fix bugs. Check SPOCK and the HPE MSA Storage Firmware support page periodically for new support streams and apply new firmware accordingly. HBAs, switches, storage arrays, and their connected disk enclosures and drives are all examples of firmware that is regularly maintained. Depending on the configuration, other components may also require updating.
  - Example: As of March 2025

    Operating system: Microsoft Windows Server 2025

    HBA: SN1610Q Fibre Channel HBA firmware version 9.15.00

    Switch: Brocade SN6600B 32 Gb Fibre Channel switch FOS version 9.2.1a

    Array: HPE MSA 2070 Fibre Channel firmware version IN300R004

    Disk: HPE MSA 1.92 TB SAS RI SFF SSD (KPM71RUG1T92) firmware version 0104

– Best practice: Update the HPE MSA Storage array and connected components through HPE MSA Storage smart component software.
  - System default: N/A
  - Detail: HPE MSA smart component software enforces best practices when updating firmware and provides monitoring of progress that is unavailable elsewhere. HPE MSA smart component software increases the probability of a successful firmware update compared to direct updates through the SMU or SFTP.
  - Example: N/A
– Best practice: Consult the HPE MSA 2070 / HPE MSA 2072 Storage Management Guide (SMG) for further best practices.
  - System default: N/A
  - Detail: The HPE MSA Gen7 SMG contains a list of additional best practices to follow before and during a firmware update. To avoid duplication or misalignment, refer to the HPE MSA 2070 / HPE MSA 2072 Storage Management Guide.
  - Example: N/A

Controller firmware

Best practices for configuring controller firmware include:

– Best practice: Keep the Partner Firmware Update (PFU) setting enabled so that both controllers run the same firmware version.
  - System default: Enabled
  - Detail: If a controller is replaced or updated outside of HPE MSA smart component software, there is a possibility of controller firmware mismatch. The PFU setting enables automated actions that both controllers are running the same firmware version, which is essential to maintaining system stability and functionality.
  - Example: Set within the SMU or by issuing this command enables the PFU setting.

    ```
    set advanced-settings partner-firmware-upgrade enabled
    ```
– Best practice: Schedule controller firmware updates.
  - System default: N/A
  - Detail: Updating controller firmware causes each controller to reboot, although at different times. While a controller is rebooting, it is unable to respond to I/O, and the remaining controller temporarily serves I/O for both pools. In dual-pool configurations, there could be a measurable drop in performance during this time. To minimize the impact on applications, schedule controller firmware updates to occur during quiet periods. Although performance should not degrade in single pool configurations, HPE nevertheless recommends scheduling firmware updates to minimize application impact if unexpected problems arise.

    The average time to update firmware is five minutes per controller and no more than 30.
  - Example: 19:00, Saturday

Disk drive firmware

Note
Gen7 HPE MSA Storage arrays support online drive firmware upgrades via the HPE Smart Component only and using bundles dated from February 2025 onwards.

– Best practice: Wait for all background tasks to complete before updating disk drive firmware.

- System default: N/A
- Detail: Some background tasks read and write data to disks. If disk drives are targets for firmware updates and are participating in these tasks, there is a possibility of unwanted interruption. Before proceeding with drive firmware updates, check via the activity monitor that the none of the following disk group tasks are active:
  - Initialization
  - Expansion (MSA-DP+ disk groups only)
  - Reconstruction
  - Verification
- Example: N/A

Note
The background scrubbing task runs almost continuously and does not need to be stopped.

System monitoring

The monitoring of array health is essential because, without notifications, an administrator would be unable to take the appropriate corrective action promptly. This section contains recommendations that help administrators receive appropriate and necessary array information at the right time.

– Best practice: Configure email notifications of system events.

- System default: Not configured
- Detail: Email notifications are essential so that administrators are made aware of pending issues with the array. When configured, the HPE MSA Storage array sends emails to up to three addresses with information on alerts as they happen. If more than three recipients are required, configure a distribution list of relevant email addresses in the email server software.
- Example: Issuing this sample command in the HPE MSA CLI configures email notification parameters.

```
set email-parameters server smtp.org.net domain org.net email-list admin@org.net
sender msa2070_1 sender-password Password123 port 25 security-protocol TLS alert-
notification-level all
```

– Best practice: Enable the managed logs feature and include logs

  - System default: Not configured
  - Detail: HPE MSA Storage arrays have a finite amount of storage for log data. The managed logs feature helps reduce the likelihood of losing log data due to wrapping by notifying defined email addresses that the log is nearly full. After a notification is received, an administrator should access the HPE MSA Storage array and download the logs (pull).

    However, HPE recommends also enabling the include logs option, which automatically attaches the logs to the notification email (push). Doing so lessens the possibility of losing historical log data.
  - Example: Issuing these sample commands in the HPE MSA CLI enables the managed logs feature.

    ```
    set email-parameters include-logs enabled email-list
    admin@org.net,,,logcollector@org.net
    ```

    ```
    set advanced-settings managed-logs enable
    ```

– Best practice: Configure SNMPv3.

  - System default: Not configured
  - Detail: SNMP is used to monitor events from managed devices centrally and minimizes downtime caused by unacknowledged system events such as component failure. HPE recommends creating an SNMPv3 user for added security.
  - Example: N/A

– Best practice: Sign up for proactive product advisory notifications.

  - System default: N/A
  - Detail: To increase availability, performance, and features, subscribe to HPE Support alerts through the HPE Email Preference Center.
  - Example: N/A

Background scrubbing

Best practices for background scrubbing include:

– Best practice: Do not disable the background disk group scrubbing feature.

  - System default: Enabled
  - Detail: The disk group scrubbing feature is important to maintaining array health and increasing application availability. In addition to both finding and attempting to fix disk errors within a disk group, disk group scrubbing can also reclaim zeroed pages and return capacity to the pool. Background disk group scrub runs continually and is designed to avoid contention with I/O; it is not recommended to create a background disk group scrub schedule.
  - Example: Issuing this command in the HPE MSA CLI enables the background disk group scrubbing feature:

    ```
    set advanced-settings background-scrub enabled
    ```

Data protection

HPE MSA Storage arrays provide various technologies to aid in the protection of application data and the retention of earlier copies. In addition to low-level mechanisms used to distribute data on disk, snapshots and replication provide further peace of mind. Best practices include:

– Best practice: Schedule volume snapshots and configure retention policies.

- System default: Not configured
- Detail: HPE MSA Storage arrays provide redirect-on-write (ROW) volume snapshots that enable the immediate recovery of data from a given point in time. HPE recommends that all volumes have at least one schedule in place for automatic taking of snapshots.

  Configure retention policies to make sure that snapshots for a given interval do not exceed a defined number. Multiple snapshots schedules allow for finer control over snapshot retention.
- Example: Volume_a001

  Schedule #1: Once per day, retention count = 7

  Schedule #2: Once per week, retention count = 4

  Schedule #3: Once per month, retention count = 12

– Best practice: Implement Remote Snap replication.

- System default: Not configured
- Detail: HPE MSA Storage arrays enable replication of volumes to remote arrays over switched Fibre Channel and iSCSI networks. In the event of a disaster at one site, Remote Snap provides an array-based solution to recover access to data as it was when last successfully replicated.
- Example: N/A

– Best practice: Maintain a tested backup solution.

- System default: N/A
- Detail: Snapshots are convenient and provide near-instantaneous recovery to a given point in time, but they are not an ideal solution to provide long-term retention of data and might not offer application consistency without the coordination of host-based software. Backup solutions take many forms, including the use of agents within virtual machines that back up data over the network to central locations and to various media. HPE does not recommend a specific backup solution but does recommend that a tested backup solution is configured together with array-based snapshots and replication.
- Example: N/A

– Best practice: Consider defining a fixed percentage of pool capacity for use by snapshots.

- System default: 10%
- Detail: To prevent snapshots consuming more capacity than desired or to allow them to consume more, consider defining the percentage of a pool's capacity reserved for snapshot data.

  Example: Issue the following sample command in the HPE MSA CLI to define 15% of pool capacity to be used by snapshots.

  ```
  set snapshot-space pool A limit 15% middle-threshold 85% limit-policy delete
  ```

– Best practice: Enable both the controller-failure and partner-notify settings through the CLI

- System default: Enabled
- Detail: When a single controller is unavailable, these settings instruct the remaining controller to change the cache mode to write-through, which facilitates the committal of data to disk at the cost of reduced write performance. After both controllers are operational, the cache mode is reverted to its previous setting, which is usually write-back mode.
- Example: Issue the following commands in the HPE MSA CLI to enable both settings:

```
set advanced-settings controller-failure enable
set advanced-settings partner-notify enable
```

Periodic health checks

HPE Best practices when performing periodic health checks include:

– Best practice: Perform regular checks of array health through the HPE MSA Health Check tool.

- System default: N/A
- Detail: The HPE MSA Health Check tool provides a free and easy way to check array health and conformity to best practices that improve array availability. HPE MSA Health Check is especially helpful in environments where an array does not have access to the internet and where a local manifest is not maintained because it provides a list of available firmware updates. HPE recommends putting a regular schedule in place to upload logs securely to HPE MSA Health Check and to review the subsequent report.
- Example: N/A

Storage best practices

Storage in this section refers to the underpinning technologies used to distribute and protect data across multiple drives and the data services that provide the resulting capacity, features, and performance. This section targets best practices regarding the configuration of disk drives, disk groups, and pools.

Disk drives

Best practices when configuring disk drives include:

– Best practice: Choose the correct drive types for the workload.

- System default: N/A
- Detail: Different drives provide differing ratios of performance, capacity, and cost. Consider workloads before building a solution and verify that the solution does not use inappropriate drive types.
  - SSD: Suited to workloads that require high random performance (IOPS) and low latency
  - Enterprise SAS: Optimized for around-the-clock low-end random performance (1K to 2K IOPS) and high-throughput and low-latency sequential I/O
  - Midline SAS: Optimized for archival data; not recommended for constant high-workload applications
- Example: N/A

– Best practice: Install SSDs in the array enclosure first.

- System default: N/A
- Detail: For optimal performance, install SSDs in the main array enclosure first, and then the closest enclosures as required.
- Example: N/A

– Best practice: Replace SSDs when their remaining life reaches 5%.

- System default: N/A
- Detail: SSD failure or wear-out is extremely rare. However, if an SSD reaches 0% life left, data loss occurs. To avoid this, the HPE MSA Storage array logs a warning-level event at 5%, 1%, and 0%. HPE recommends replacing SSDs no later than when 5% life of remains.
- Example: N/A

Choosing disk group types

A disk group is a collection of disks logically teamed together to provide capacity and performance to a pool. Best practices when selecting a disk type group include:

– Best practice: Choose an appropriate RAID protection level.

- System default: MSA-DP+
- Detail: RAID provides both resiliencies to drive failure and performance benefits. With the introduction of MSA-DP+, HPE recommends that all mechanical HDDs take advantage of its benefits and SSDs when in all-flash pool configurations.
- Example: N/A

Note

HPE does not recommend RAID 5 with mechanical HDDs.

Table 1. HPE recommended RAID levels

Disk type	Tier	RAID protection level	Pool type	Detail
SSD	Performance	RAID 1/10	– Hybrid or all-flash	– Performance-optimized solutions that require the best random write performance – Low-capacity solutions where RAID types requiring more drives (RAID 5 / MSA-DP+) would not be cost-effective
		RAID 5	– Hybrid	– Typical configurations that require very high random read performance and medium-high random write performance
		MSA-DP+	– All-flash	– Recommended for all-flash configurations
HDD	Standard or Archive	MSA-DP+	– Any	– Recommended for all configurations – Highly resilient to failure – Very fast rebuild times
		RAID 6		– Recommended when the capacity goal is too low for MSA-DP+ – Medium resiliency to failure – Medium performance – Long rebuild times

– Best practice: Follow the power of 2 rule.

- System default: N/A
- Detail: The power of 2 rule achieves an optimal data layout with parity-based disk groups (RAID 5 and 6). Not following this rule while using HDDs can reduce sequential write performance. See the HPE MSA Gen7 virtual storage technical reference guide for detailed information.
- Example: N/A

Table 2. Optimal drive configurations

RAID level	Number of driver in a disk group	Number of data chunks in a stripe	Number of parity chunks in a stripe
RAID 5	3	2	1
RAID 5	5	4	
RAID 5	9	8	
RAID 6	4	2	2
RAID 6	6	4	
RAID 6	10	8	

Note
MSA-DP+ disk groups have the power of 2 rule embedded into their design regardless of drive count.

Note
It is not strictly necessary to configure SSDs to follow the power of 2 rule.

Sparing

Sparing enables an array to automatically assign idle disks or capacity to rebuild a degraded disk group or stripe zone, thus reducing the chance of data loss. Best practices and array behavior vary depending on drive technology, RAID type, and drive form factor. Spare disks are consumable by any non-MSA-DP+ disk group in either pool where the drive type is a match. Unlike MSA-DP+ disk groups where disks store both data and distributed spare capacity, spare disks for non-MSA-DP+ disk groups are idle until required to rebuild a degraded disk group.

– Best practice: Assign global spares when using non-MSA-DP+ disk groups.

- System default: Dynamic sparing
- Detail: By default, the HPE MSA Storage array uses dynamic sparing, which consumes unassigned drives (AVAIL) as needed. However, to reserve disks as spares, HPE recommends allocating drives as global spares as shown in Table 3.

Note
When an SSD drive fails in a disk group, if no matching spare SSD is available for rebuilding the disk group, then to improve data reliability, the data on the SSD disk group drains to an HDD tier. Because dedicated SSD spares are costly and could reduce availability should the disk group degrade due to wear-out, HPE recommends, when using hybrid pools, allocating global spares for HDDs only. Additionally, global spares are neither required nor usable by MSA-DP+ disk groups, because sparing is an integral feature of this RAID group type.

Table 3. Recommended sparing strategy for non-MSA-DP+ disk groups

Drive type	Tier	RAID	Initial global spares	Recurring global spares
SSD	SSD read cache	n/a	None	None
	Performance	RAID 1 / 10 RAID 5		
	All-flash pool	RAID 6	1 for up to 24 configured drives	1 per 24 configured drives
Enterprise SAS	Standard			
Midline SAS	Archive		2 for up to 24 configured drives	

Example: Table 4 provides an example of recommended global spares when using non-MSA-DP+ disk groups.

Table 4. Examples of non-MSA-DP+ global spare assignment for a single dual-pool HPE MSA Storage system

Drive	Tier	Drive capacity	RAID	Number of drives in disk groups	Global spares	Total drives
SSD	Performance	1.92 TB	RAID 5	3	None	3
	All-flash pool		RAID 6	10	1	11
				40	2	42
Enterprise SAS	Standard	2.4 TB		20		22
Midline SAS	Archive	24 TB		30	3	33

– Best practice: Define an adequate target spare capacity when using MSA-DP+ disk groups.

- System default: Equal to the sum of the two largest drives in the disk group
- Detail: Instead of consuming idle spare drives as needed, MSA-DP+ disk groups include integrated spare capacity. By default, spare capacity equals the summed capacity of the two largest drives within the disk group

  There are two scenarios to consider regarding spare capacity for an MSA-DP+ disk group:

  - During the initial creation of the disk group—It is not possible to define the target spare capacity through the SMU. Therefore, if the quantity of disks is greater than one enclosure or is an SSD disk group, add it to the pool through the CLI. Refer to the HPE MSA 2070/ HPE MSA 2072 CLI reference guide for more information on the spare-capacity switch of the add disk-group command. Follow Table 5 for guidance on target spare space capacity.
  - The expansion of a disk group—Before expansion, modify the target spare capacity in multiples of the largest drive capacity as per Table 5, then add drives and expand the disk group.

Table 5. Recommended spare capacity for MSA-DP+ disk groups

Drive type	Tier	Initial spare capacity	Recurring spare capacity
SSD	Performance	1x drive capacity per 24 drives	1x drive capacity per 24 drives
Enterprise SAS	Standard	2x drive capacity per 24 drives	
Midline SAS	Archive		

- Example: Table 6 provides examples of recommended spare capacity when using MSA-DP+ disk groups.

Table 6. Example of MSA-DP+ spare capacity for a single dual-pool HPE MSA Storage system

Drive type	Tier	Drive capacity	Number of drives in a disk group	Sparing	Target spare capacity
SSD	Performance	1.92 TB	24	1x drive capacity	1.92 TB
Enterprise SAS	Standard	2.4 TB	48	2x drive capacity	3.84 TB
Midline SAS	Archive	12 TB	24		4.8 TB
			36	3x drive capacity	36 TB

Warning
It is only possible to define the target spare capacity through the CLI. If the target spare capacity is to be set higher or lower than the default, it must be defined during creation of the disk group or before adding new drives. Additionally, a combination of drive capacity and quantities sufficient to reach the target spare capacity must be included or added to the disk group. Not doing so results in a target spare capacity greater than what is available, thus placing the disk group into a degraded state.

Tiering

Tiering is the process of granularly distributing a volume across different drive types to provide the best balance of cost and performance. Tiering best practices include:

– Best practice: Maintain proportions across tiers.

- System default: N/A
- Detail: The HPE MSA Storage tiering engine works in near real time to deliver the best balance of performance to cost. If a deep combined knowledge of array tiering behavior and workload knowledge is present, it might be possible to successfully work outside of the recommended tier ratios. However, following HPE recommendations allows the tiering engine to deliver its value reliably and effectively and with minimal administrator intervention.

Note
Although it is supported, combining SSDs and Midline SAS drives within the same pool without an intermediary tier of Enterprise SAS drives is not recommended. Doing so can reduce the effectiveness of the tiering engine and reduce both performance and value. For example, even low-capacity configurations using MSA-DP+ would require more SSD read cache than can be configured. If SSDs were used in a performance tier instead, the solution would be neither as cost-effective nor as performant as when using three tiers. However, if known in advance that the working set is anticipated to remain sufficiently small to fit into the available SSD region, then it might be a practical solution. HPE recommends reevaluating performance and ratios over time.

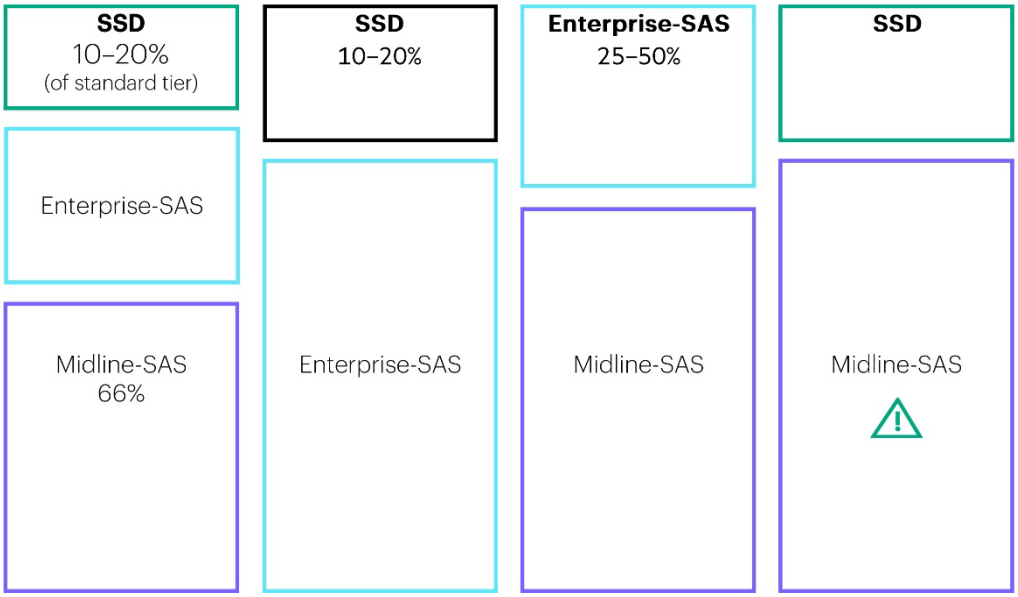


Figure 3. Recommended tier ratios

There are two possibilities regarding the recommended quantity of SSDs when combined in a pool with HDDs. In both cases, HPE recommends proportioning SSDs against the standard tier.

Table 7. Recommended SSD proportion rules

SSD use case	Flash ratio to maintain	Proportion
Performance tier	10%–20%	Percentage of the fastest HDD tier
SSD read cache		

- Example: Table 8 provides examples of disk tier proportions in a three-tier hybrid pool.

Table 8. Examples of disk tier proportions in a three-tier hybrid pool

Tier	RAID	Number of drives	Drive capacity	Number of disk groups	Capacity	Capacity ratio	Flash ratio (rounded)
Performance	RAID 5	5	1.92 TB	1	7.65 TB	5%	16.6% (7.65 / 45.99) x 100
Standard	MSA-DP+	24	2.4 TB		45.99 TB	31%	
Archive		12	12 TB		98.88 TB	64%	N/A

– Best practice: Configure SSDs as capacity within the performance tier for generic workloads.

- System default: N/A
- Detail: Mixing SSDs and Enterprise SAS drives as capacity tiers automatically enable tiering and delivers the best balance of cost and performance for common workloads. Refer to the HPE MSA Gen7 virtual storage technical reference guide for information on how the HPE MSA Storage tiering engine works and why it is an effective choice.
- Example:

  Performance tier: 1x RAID 10 disk group

  Standard tier: 1x MSA-DP+ disk group

– Best practice: Choose read cache only when a workload is known to have a very low percentage of random writes.
  - System default: N/A
  - Detail: SSD read cache accelerates random reads, and does not accelerate random writes or sequential I/O. If random writes are frequent, use performance tiering.
  - Example: Customer requirement
    - Random writes: <2K IOPS
    - Random reads: >2K IOPS
– Best practice: Choose single-tier HDD configurations for throughput-heavy workloads or data archiving.
  - System default: N/A
  - Detail: HDDs provide an ideal solution for workloads that are not dominantly random. Introducing performance tiering or SSD read cache for such workloads is unlikely to yield tangible benefits but does increase the cost.
  - Example: Customer requirement
    - Random I/O: <2K IOPS
    - Throughput: <8.7 GB/s reads
      <5.5 GB/s writes
– Best practice: Choose single-tier SSD configurations for extremely high throughput-heavy workloads or for large capacities of low-latency storage and datasets that are not candidates for data-reduction techniques.
  - System default: N/A
  - Detail: SSDs provide low latencies for random I/O and over 14 GB/s of sequential throughput. Because an HPE MSA Storage array does not offer deduplication or compression, it might be an economical solution for large datasets that are not eligible for compaction. For datasets that are eligible, HPE recommends considering other solutions within the portfolio such as HPE Alletra Storage arrays.
  - Example: Customer requirement
    - Random I/O: >5K IOPS
    - Throughput: >8.7 GB/s
    - Random latencies: <10 ms
    - Capacity: Many tens of terabytes
– Best practice: Periodically review and maintain sufficient SSDs capacity needed to absorb 80% of daily random I/O.
  - System default: N/A
  - Detail: Recommended tier proportions in this guide are for new installations, and a configuration applied to an HPE MSA Storage array when brought into service might not continue to be effective over the full duration of an array's useful life. For optimal performance, SSDs must have enough capacity to store 80% or more of daily random I/O, which typically accounts for a small fraction of a pool's total size.

    Located in the capacity area of the SMU, the I/O workload tool gives administrators a graph of historic daily I/O and a visual representation of the relationship to SSD capacity. HPE recommends using the I/O workload tool at defined intervals throughout the life of the array to drive changes in the configuration of a pool as needed.
  - Example: N/A

– Best practice: Do not change the volume tier affinity setting.

- System default: No Affinity
- Detail: HPE recommends that volumes use the default tier affinity setting, which is No Affinity. Modifying the affinity setting of a volume could result in the unnecessary degradation of performance for other volumes within the pool and might not yield the expected results.

  There are, however, workloads and scenarios where changing a volume's affinity is warranted. For example, the Archive Affinity setting is useful for infrequently accessed data because it frees capacity in the upper tiers for performance-sensitive applications. Refer to the HPE MSA Gen7 virtual storage technical reference guide for further guidance on how and when to use the tier affinity setting.
- Example:
  - Backups: Archive
  - HDD images: Archive
  - Boot volumes: Performance
  - SQL: No Affinity
  - Video streaming: No Affinity
  - General VM storage: No Affinity

Single vs. dual pools

HPE MSA Storage arrays ship with two array controllers, each supporting a single pool of storage. A pool contains disk groups that provide capacity and performance independently of the other pool. This section provides best practices regarding the use of a single (asymmetrical) or dual pool (symmetrical) configurations.

Important

The array defaults for controller-failure and partner-notify settings are enabled. If a single controller becomes unavailable, these settings cause the remaining controller cache policy to switch to write-through, which can cause a degradation in write performance in return for the assurance that written data is committed to disk. Although the full performance of the remaining controller cannot be realized during single-controller operation, it is still important to consider controller headroom so the overall impact to application performance can be minimized.

– Best practice: Use a single pool unless either the performance or capacity goal explicitly requires a second pool.

- System default: N/A
- Detail: Single-controller HPE MSA Storage performance and addressable capacity scales beyond the requirements of typical workloads. Additionally, single-controller configurations reduce the impact to performance in the unlikely event of controller unavailability during a peak in I/O demand.

  The official HPE tool for sizing is HPE Ninja Online for MSA, which automatically suggests single and dual pool configurations that match the intended performance and capacity goals.

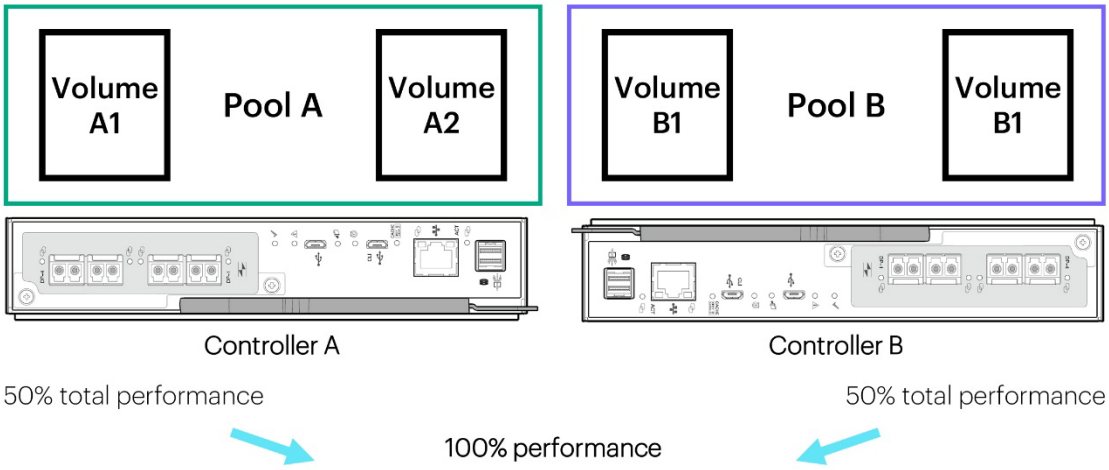


Figure 4. Representation of pool and array performance and capacity potential

- Example

  Dual pool: 780K IOPS 4K random read performance

  14.7 GB/s sequential read performance

  7.3 PB raw capacity

  Single pool: 390K IOPS 4K random read performance

  7.35 GB/s sequential read performance

  4 PiB usable capacity

– Best practice: Use the Volume Copy feature to rebalance an underperforming array.

- System default: N/A
- Detail: HPE MSA Gen7 Storage arrays can address up to 4 PiB of storage per pool and provide more than adequate performance such that two pools are unlikely necessary. However, for extremely demanding workloads, the array can provide, in most cases, double the capacity and performance by using its second pool.

  If a single-pool array configuration no longer meets requirements, use the Volume Copy feature to assist in the migration of a volume to the other pool. The Volume Copy feature might also help in rebalancing a dual pool configuration where one pool is underutilized.

Important
Volume Copy requires the volume to be unmapped from a host and application traffic for that volume halted. After the copy is complete, map the newly copied volume to the host, resume any applications, and consider deleting the source volume.

- Example: Issuing the following command in the HPE MSA CLI copies the volume SourceVol from Pool A to Pool B with a new name "DestVol":

```
copy volume SourceVol destination-pool B name DestVol
```

Thin Provisioning

Thin Provisioning is the practice of defining volume capacities that, when totaled, exceed the physical capacity of a pool. The principal goal of Thin Provisioning is to reduce the initial costs of owning an array by both reducing the number of drives that must be initially purchased and the power and cooling costs required to operate them.

– Best practice: Monitor pool usage and set appropriate thresholds and notifications.

- System default
  - Low threshold: 50%
  - Middle threshold: 75%
  - High threshold: Calculated (available pool capacity minus 400 GiB)
- Detail: When a pool is overcommitted and has no capacity remaining, incoming writes to previously unwritten areas of any volume are rejected. Additionally, when a pool reaches its high threshold, performance is reduced.

  Configure notifications and set thresholds that allow sufficient time to procure new physical capacity or remove unwanted data from the pool.
- Example: N/A

– Best practice: If using VMware ESXi™ 6.5 or later, periodically issue the UNMAP command.

- System default: N/A
- Detail: ESXi 6.5 introduced the automatic release of allocated block storage after the removal of files from a datastore. However, VMware® issues the UNMAP command at a 1 MB granularity, but the HPE MSA Storage operates with a 4 MB page size. As a result, pages do not become free automatically but can be released when the UNMAP command is invoked manually.
- Example: ESXi CLI

```
esxcli storage vmfs unmap -l MyDatastore
```

– Best practice: Leave the overcommit setting enabled for pools with a high number of snapshots.

- System default: Enabled
- Detail: Overcommit allows snapshots to consume a minimal amount of allocated capacity within the pool. When overcommit is disabled, snapshots consume available space equal to the size of the original volume. On average, disabling overcommit has an undesired impact on the available capacity of a pool and should be left enabled.
- Example: N/A

Full disk encryption

HPE MSA Storage arrays support data-at-rest encryption with self-encrypting drives (SEDs).

Note

The HPE MSA 2072 Storage array is encryption-capable, but the included 1.92 TB SSDs are not self-encrypting. Because full disk encryption requires every drive in the system to be an SED, it is neither economical nor recommended to use the HPE MSA 2072 Storage array for this purpose.

Important

It is not possible to recover a lost passphrase, and it is not possible to access data on a locked system without it. Therefore, HPE strongly recommends keeping a copy of all passphrases in a secure location.

– Best practice: Use a different passphrase for each FDE-secured system.
  - System default: N/A
  - Detail: All HPE MSA Storage arrays that support encryption generate the same lock key and the same passphrase. To increase security, HPE recommends locking each array by using a unique passphrase, which also generates a unique lock key.
  - Example: N/A
– Best practice: Clear FDE keys before shutdown when moving an entire HPE MSA storage system.
  - System default: Not cleared
  - Detail: To protect data after relocation, HPE recommends clearing the FDE keys. When the array is powered on again, the disks shall be in a secure, locked state, and the original passphrase must be entered to re-access data.
  - Example: Issue the following commands in the HPE MSA CLI to clear the FDE keys.

    Before powering down: `clear fde-keys current-passphrase myPassphrase`

    After power is reapplied: `set fde-lock-key passphrase myPassphrase`

Capacity expansion

Best practices regarding capacity expansion include:

– Best practice: Expand tiers comprising RAID 1/10, 5, or 6 with disk groups of equal proportions and rotational speeds.
  - System default: N/A
  - Detail: For performance to be consistent across a tier and pool, all disk groups within a tier should exhibit the same performance characteristic and provide equal capacity.
  - Example:

    Before expansion: 2x RAID 6 disk groups, each with ten 2.4 TB Enterprise SAS drives (Standard tier)

    After expansion: 3x RAID 6 disk groups, each with ten 2.4 TB Enterprise SAS drives (Standard tier)
– Best practice: Expand tiers comprising MSA-DP+ disk groups with drive capacities no greater than a factor of two than the smallest.
  - System default: N/A
  - Detail: Using drives of capacities greater than a factor of two can lead to the inefficient use of the capacity provided by the larger drive and reduce their value.
  - Example:

    2x 1.2 TB HDD + 1 x 2.4 TB HDD = OK

    12x 600 GB HDD + 1 x 2.4 TB HDD = Not recommended

Note

A running scrub task can be aborted to allow an MSA-DP+ disk group to be expanded.

– Best practice: Make sure a disk group is fault-tolerant before attempting to expand it.
  - System default: N/A
  - Detail: Before attempting a disk group expansion, a disk group must be healthy within the WBI or FTOL (Fault Tolerant On-Line) within the CLI. In rare cases, expanding a degraded disk group could conflict with the rebuild process.
  - Example: N/A

Volume mapping

Best practices to follow when mapping volumes include:

– Best practice: Do not use the default mapping feature unless there is a documented application-specific requirement.
  - System default: Disabled, uses explicit mapping.
  - Detail: Default mapping allows unrestricted access to a volume for all attached initiators. While default mapping may be convenient, its use can lead to locking issues and uncommitted writes. Explicit mapping reduces the likelihood of mistakes and lost data.
  - Example: N/A

– Best practice: Map volumes through sufficient host ports to meet performance requirements.
  - System default: N/A
  - Detail: If volumes are not attached to hosts through enough host ports, array performance might be limited. Consider how many active paths are required to meet the performance potential of the configured system.
  - Example: Table 9 provides examples of data rates per host port and protocol.

Table 9. Data rates per host port and protocol

Link speed	Throughput (MB/s)	IOPS
1 Gb iSCSI	110	14200
10 Gb iSCSI	1100	142000
25 Gb iSCSI	2750	355000
12 Gb SAS[1]	5280	681600
8 Gb FC	880	113600
16 Gb FC	1760	227200
32 Gb FC	3520	454400

[1] Each SAS port contains four 12 Gb SAS lanes.

– Best practice: Do not configure more than eight paths from a host to the array.
  - System default: N/A
  - Detail: The time for MPIO to recover from multiple path failures might increase unacceptably if there are too many paths to a volume. Because performance cannot benefit from more than eight paths, use Fibre Channel zoning or subnetting to limit available paths.
  - Example: N/A

– Best practice: Provision redundant paths from a host to the array.
  - System default: N/A
  - Detail: For application data to remain available during a controller or path outage, confirm that hosts have no less than two paths to a volume; one to Controller A and one to Controller B.

    Connect fabric-attached hosts to at least two independent Fibre Channel or Ethernet networks or SAS switches. Connect directly connected hosts with at least one path to each controller. Refer to HPE SPOCK for compatibility information.
  - Example: Figure 5 illustrates a host with redundant paths to a volume.

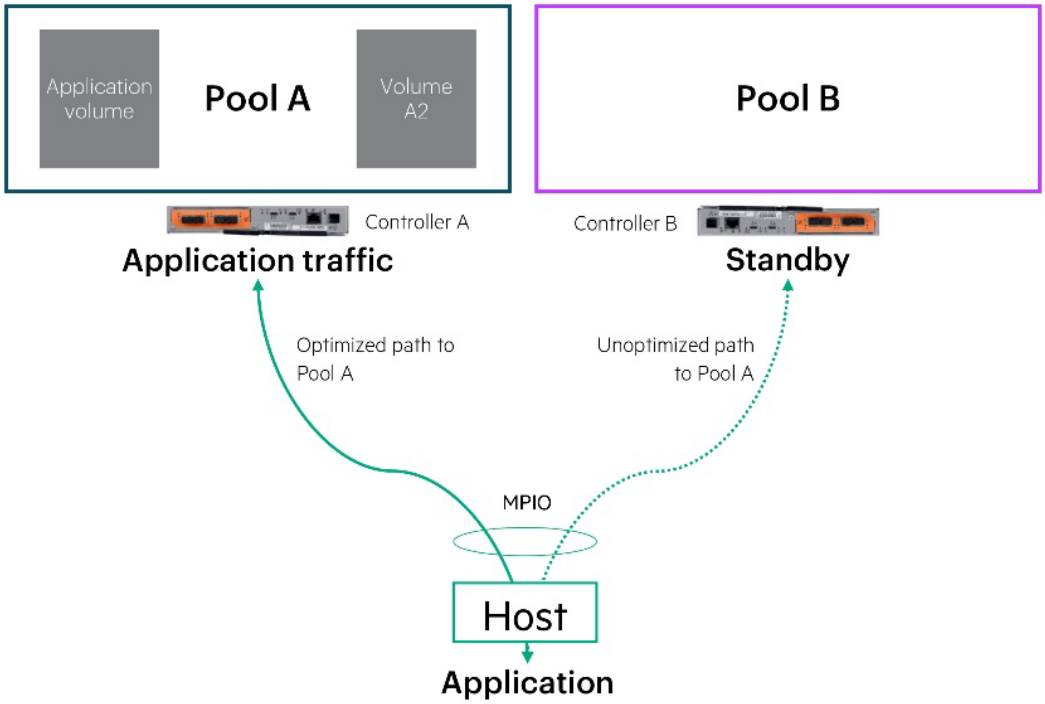


Figure 5. Example of a host with redundant paths to a volume

– Best practice: Do not attach a volume to more than one host unless all hosts support the same cluster-aware file system or are cooperatively managed such that nonshared file systems can be used.

- System default: N/A
- Detail: Sharing volumes between multiple hosts can lead to data loss and corruption if the file system or operating system cannot cooperate and temporarily lock disk regions at a granular level. Refer to operating system documentation for guidance on its file system capabilities and requirements.

Important
Some file systems such as Microsoft Cluster Shared Volumes (CSV) require that hosts be in a cluster, whereas others such as VMware vSphere® VMFS do not. Take appropriate action to prepare hosts before mapping volumes.

- Example: The following are examples of file systems that can be shared under the correct circumstances:
  - VMFS is a cluster-aware file system that can be shared without placing hosts in a cluster
  - CSV is a cluster-aware file system that requires hosts to participate in a failover cluster

Summary

These best practices help administrators achieve the best possible performance and availability of their HPE MSA storage arrays. Use this guide and the documentation listed to deliver the best configuration for your application needs.

Resources

HPE MSA Gen7 virtual storage technical reference guide

HPE MSA Health Check

Sign up for HPE updates

HPE Ninja Online for MSA

Learn more at

HPE.com/storage/msa

Visit HPE.com

Chat now

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, Windows Server, and PowerShell are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware, VMware ESXi, and VMware vSphere VMFS are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.
All third-party marks are property of their respective owners.

a50011790ENW, Rev. 2

HEWLETT PACKARD ENTERPRISE

hpe.com