Technical white paper

HP P2000 G3 MSA

Best practices

Table of contents

About this document

This white paper highlights the best practices for optimizing the HP P2000 G3 MSA, and should be used in conjunction with other HP Storage Modular Smart Array manuals. Modular Smart Array (MSA) technical user documentations can be found at hp.com/go/P2000

Intended audience

This paper is intended for entry-level and mid-range HP P2000 G3 MSA administrators and requires previous SAN knowledge. This document offers Modular Storage Array facts that can contribute to an MSA best customer experience.

This paper outlines a best practice approach to performance and configuration.

Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Fibre Channel, Serial Attached SCSI (SAS), Internet SCSI (iSCSI), and Ethernet protocols

Related documentation

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

- HP P2000 G3 MSA System Racking Instructions
- HP P2000 G3 MSA System Installation Instructions
- HP P2000 G3 MSA System Cable Configuration Guide
- HP P2000 G3 MSA System FC User's Guide HP P2000 G3 MSA System FC/iSCSI User's Guide
- HP P2000 G3 MSA System SAS User's Guide
- HP P2000 G3 MSA System iSCSI User's Guide

You can find these documents from the Manuals page of the HP Business Support Center: hp.com/support/manuals

The P2000 G3 MSA

Topics covered

- This section examines the following:
- What's new
- Hardware overview
- Unified LUN Presentation (ULP)
- Choosing single or dual controllers
- Choosing DAS or SAN attach
- Dealing with controller failures
- Virtual disks
- Volume mapping
- RAID levels
- Cache configuration
- World Wide Name (WWN) naming conventions

- Fastest throughput optimization
- Highest fault-tolerance optimization
- Disk Scrubbing, SMART, and Drive Spin Down
- Cascading Array Enclosures
- 8 Gb Switches and SFP transceivers
- MSA 70 Considerations
- Administering with HP SMU
- P2000 G3 iSCSI MSA Considerations
- Using Optic Cables
- Software

What's new in the P2000 G3 MSA

- Four port 1GbE iSCSI controller or
- Two port 10GbE iSCSI controller or
- Four port 6 Gb SAS controller or
- Two new 8 Gb Fibre Channel controllers:
  - Standard model with two 8 Gb FC host ports each
  - Combo model with two 8 Gb FC host ports and two 1GbE iSCSI ports each

- Controllers come with 2 GB cache memory each
- Increased support to seven P2000 LFF disk enclosures (96 LFF drives)
- Increased support to five D2700 SFF disk enclosures (149 SFF drives)
- 6 Gb SAS back end and HDD support
- 64 Snaps and clone capability come standard on G3 models
- Optional 512 snapshots max (double the MSA2000 G2)
- Optional controller-based replication (Remote Snap) with the FC or the FC/iSCSI Combo controllers only
- 512 max LUN support
- Higher performance with an upgraded controller with increased I/O performance
- Improved Storage Management Utility (SMU) user interface
- Full support for G1/G2 to G3 upgrade, including cross-protocol upgrades

Hardware overview

HP P2000 G3 MSA Modular Smart Array

HP 2000 Family of storage arrays features P2000 G3 MSA arrays with the latest 8 Gb Fibre Channel, 6 Gb SAS, 10GbE iSCSI, and four port 1GbE iSCSI connected models. The arrays are designed for entry-level customers and feature the latest in functionality and host-connect technology while offering excellent price/performance. The P2000 G3 MSA is ideal for companies with small budgets or limited IT expertise, and also for larger companies with departmental or remote requirements. Each solution is designed to be easy to deploy, to be secure, along with low management costs, while driving rapid return on investment.

The P2000 G3 FC is an 8 Gb Fibre Channel, while the P2000 G3 SAS is a 6 Gb SAS connected 2U storage area network (SAN) or direct connect solution (OS dependent) designed for small to medium-sized departments or remote locations.

The P2000 G3 MSA offers a choice of four controllers:

- A high-performance, Fibre Channel 8 Gb dual port model
- A unique dual-purpose Combo controller with two 8 Gb Fibre Channel ports plus two 1GbE iSCSI ports
- A 6 Gb SAS controller with four ports per controller
- A 10GbE iSCSI controller with two ports per controller
- A 1GbE iSCSI controller with four ports per controller

The P2000 G3 arrays are 2U storage area network or direct connect solutions (OS and protocol dependent) offering a choice of five controllers—two FC, one SAS, one 10GbE iSCSI and the newest model features a four port 1 Gb iSCSI. The first Fibre Channel controller is a high-performance, 8 Gb dual port model. The second offering is a unique dual-purpose Combo controller with two 8 Gb Fibre Channel ports with the addition of two 1GbE iSCSI ports. The third controller choice is 6 Gb SAS with four ports per controller. There is also the recent addition of a two port 10 Gb iSCSI, while the latest controller introduction features four 1 Gb iSCSI ports per controller.

The dual-protocol P2000 G3 MSA FC/iSCSI Combo Controller gives exceptional flexibility. The 8 Gb FC ports support a full FC SAN while the two 1GbE iSCSI ports can serve two purposes. With this combination you can economically share the array storage resource with a smaller department accessing it over iSCSI or enable the new optional Remote Snap functionality over iSCSI protocol (also available over FC).

The P2000 G3 SAS is the follow-on product to the MSA2000sa G2, adding the latest 6 Gb SAS technology to the four host ports per controller. The P2000 G3 SAS array is designed for directly attaching up to four dual-path or eight single path rack servers. SAS array support for BladeSystems will continue to come from the MSA2000sa G2.

The P2000 G3 10GbE iSCSI, brings the very latest in high-performance host connection with technology generally found only in higher priced arrays. The bandwidth it provides in conjunction with server consolidation is highly advantageous in shared storage configurations. Array connection to 10GbE switches that are in turn connected to 1GbE NICs is commonplace. Directly attached server support requires the server units to have 10GbE NICs.

The most recent addition is the P2000 G3 iSCSI controller featuring four 1 Gb iSCSI Ethernet ports, double the number of the G2 model. This allows an array that keeps the price of the components, particularly the interconnects, low while markedly increasing the performance capabilities.

All P2000 G3 models can be equipped with single or dual controllers, feature the same scalability, and offer 6 Gb SAS back-end transmission speed to drives and JBODs. Significant data protection advances are delivered by the all P2000 G3 arrays. First, all G3 units come standard with 64 snapshot capability at no extra cost and there is an option for the G3 series of 512 snapshots. Volume Copy (clone) also comes standard. In a further move to protect the user's data, optional Remote Snap (replication) capability is offered on the FC or FC/iSCSI versions (only).

The controller-less P2000 chassis is offered in two models—one comes standard with twelve Large Form Factor (LFF) 3.5-inch drive bays, the other can accommodate twenty-four Small Form Factor (SFF) 2.5-inch drives. Both are able to simultaneously support enterprise-class SAS drives, SAS Midline, and archival-class SATA Midline drives. Either chassis can have one or two matching P2000 G3 controllers (same protocol) and are available with AC or DC power supplies.

Additional capacity can easily be added when needed by attaching either the MSA2000 12 bay drive enclosure, the MSA70 drive enclosure, or the D2700 drive enclosure. Configurations utilizing the SFF drive chassis can grow to a total of 149 SFF drives. The LFF drive chassis can grow up to a total of 96 drives.

Note
Larger hard drives are always in test; refer to hp.com/go/p2000 to get the latest hard drive capacity limits.

In addition to support for Windows® and Linux on x86 and x64 server platforms, the P2000 G3 continues support for HP-UX, OpenVMS, Windows, and Linux on powerful Integrity servers depending on the protocol.

Unified LUN Presentation

The P2000 G3 uses the concept of Unified LUN Presentation (ULP). ULP can expose all LUNs through all host ports on both controllers. The interconnect information is managed in the controller firmware. ULP appears to the host as an active-active storage system where the host can choose any available path to access a LUN regardless of Vdisk ownership.

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

Overview
ULP presents all LUNs to all host ports

• Removes the need for controller interconnect path

• Presents the same World Wide Node Name (WWNN) for both controllers

Shared LUN number between controllers with a maximum of 512 LUNs

• No duplicate LUNs allowed between controllers
• Either controller can use any unused logical unit number

ULP recognizes which paths are "preferred"

• The preferred path indicates which is the owning controller per ALUA specifications
• "Report Target Port Groups" identifies preferred path
• Performance is slightly better on preferred path

Write I/O Processing with ULP

• Write command to controller A for LUN 1 owned by Controller B
• The data is written to Controller A cache and broadcast to Controller A mirror
• Controller A acknowledges I/O completion back to host
• Data written back to LUN 1 by Controller B from Controller A mirror

Figure 1. Write I/O Processing with ULP
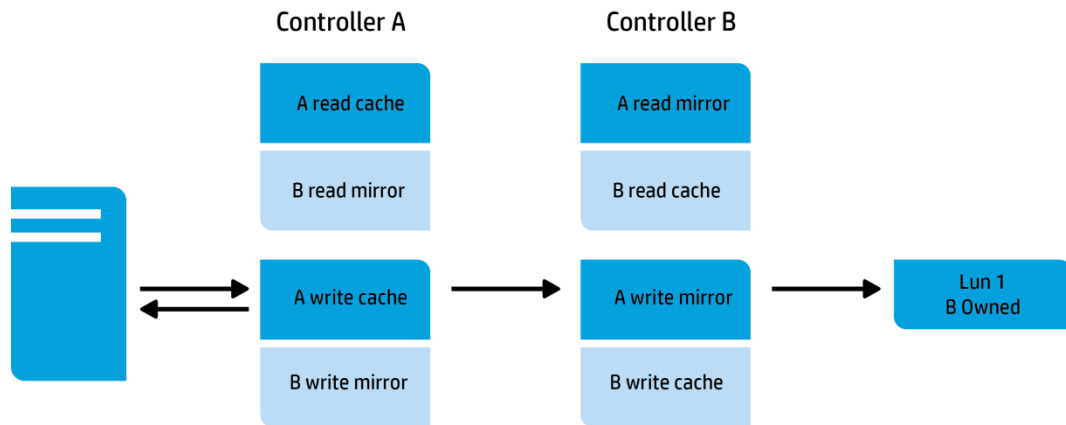
Read I/O Processing with ULP

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

Figure 2. Read I/O Processing with ULP
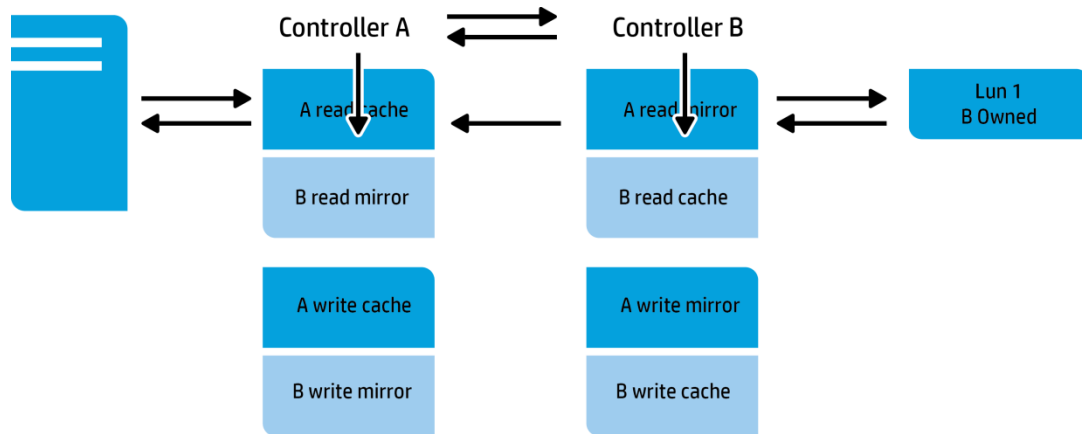
Choosing single or dual controllers

Although you can purchase a single-controller configuration, it is best practice to use the dual-controller configuration to enable high availability and better performance. However, under certain circumstances, a single-controller configuration can be used as an overall redundant solution.

Dual controller

A dual-controller configuration improves application availability because in the unlikely event of a controller failure, the affected controller fails over to the surviving controller with little interruption to the flow of data. The failed controller can be replaced without shutting down the storage system, thereby providing further increased data availability. An additional benefit of dual controllers is increased performance as storage resources can be divided between the two controllers, enabling them to share the task of processing I/O operations. Controller failure results in the surviving controller doing the following:

• Taking ownership of all RAID sets
• Managing the failed controller's cache data
• Restarting data protection services
• Assuming the host port characteristics of both controllers

The dual-controller configuration takes advantage of mirrored cache. By automatically "broadcasting" one controller's write data to the other controller's cache, the primary latency overhead is removed and bandwidth requirements are reduced on the primary cache. Any power loss situation will result in the immediate writing of cache data into both controllers' compact flash devices, reducing any data loss concern. The broadcast write implementation provides the advantage of enhanced data protection options without sacrificing application performance or end-user responsiveness.

Note

When using dual controllers, it is highly recommended that dual-ported hard drives be used for redundancy. If you use single-ported drives in a dual controller system and the connecting path is lost, the data on the drives would remain unaffected, but connection to the drives would be lost until the path to them is restored.

Single controller

A single-controller configuration provides no redundancy in the event that the controller fails; therefore, the single controller is a potential Single Point of Failure (SPOF). Multiple hosts can be supported in this configuration (up to two for direct attach). In this configuration, each host can have access to the storage resources. If the controller fails, the host loses access to the storage.

The single-controller configuration is less expensive than the dual-controller configuration. It is a suitable solution in cases where high availability is not required and loss of access to the data can be tolerated until failure recovery actions are complete. A single-controller configuration is also an appropriate choice in storage systems where redundancy is achieved at a higher level, such as a two-node cluster. For example, a two-node cluster where each node is attached to a P2000 G3 FC with a single controller and the nodes do not depend upon shared storage. In this case, the failure of a controller is equivalent to the failure of the node to which it is attached.

Another suitable example of a high-availability storage system using a single controller configuration is where a host uses a volume manager to mirror the data on two independent single-controller P2000 G3 storage systems. If one P2000 G3 storage system fails, the other P2000 G3 storage system can continue to serve the I/O operations. Once the failed controller is replaced, the data from the survivor can be used to rebuild the failed system.

Note
When using a single-controller system, the controller must be installed in the slot A of the array.

Choosing DAS or SAN attach

There are two basic methods for connecting storage to data hosts: Direct Attached Storage (DAS) and Storage Area Network (SAN). The option you select depends on the number of hosts you plan to connect and how rapidly you need your storage solution to expand.

Direct attach
DAS uses a direct connection between a data host and its storage system. The DAS solution of connecting each data host to a dedicated storage system is straightforward and the absence of storage switches can reduce cost. Like a SAN, a DAS solution can also share a storage system, but it is limited by the number of ports on the storage system.

A powerful feature of the storage system is its ability to support four direct attach single-port data hosts, or two direct attach dual-port data hosts without requiring storage switches. The P2000 G3 FC can also support two single-connected hosts and one dual connected host for a total of three hosts.

If the number of connected hosts is not going to change or increase beyond four then the DAS solution is appropriate. However, if the number of connected hosts is going to expand beyond the limit imposed by the use of DAS, it is best to implement a SAN.

Tip
It is a best practice to use a dual-port connection to data hosts when implementing a DAS solution. This includes using dual-ported hard drives for redundancy.

Switch attach
A switch attach solution, or SAN, places a switch between the servers and storage systems. This strategy tends to use storage resources more effectively and is commonly referred to as storage consolidation. A SAN solution shares a storage system among multiple servers using switches and reduces the total number of storage systems required for a particular environment, at the cost of additional element management (switches), and path complexity.

Using switches increases the number of servers that can be connected. Essentially, the maximum number of data hosts that can be connected to the SAN becomes equal to the number of available switch ports.

Note
The HP P2000 G3 FC supports 64 hosts.

Tip
It is a best practice to use a switched SAN environment anytime more than four hosts or when growth in required or storage or number of hosts is expected.

Dealing with controller failovers

Since the P2000 G3 uses Unified LUN Presentation, all host ports see all LUNs; thus failovers are dealt with differently than with the MSA2000 G1 (excluding the MSA2000sa G1).

FC direct-attach configurations
In a dual-controller system, both controllers share a unique node WWN so they appear as a single device to hosts. The controllers also share one set of LUNs to use for mapping volumes to hosts.

A host can use any available data path to access a volume owned by either controller. The preferred path, which offers slightly better performance, is through target ports on a volume's owning controller.

Note 1
Ownership of volumes is not visible to hosts. However, in SMU you can view volume ownership and change the owner of a virtual disk and its volumes.

Note 2
Changing the ownership of a virtual disk should never be done with I/O in progress. I/O should be quiesced prior to changing ownership.

In the following configuration, both hosts have redundant connections to all mapped volumes.

Figure 3. FC storage presentation during normal operation (high-availability, dual-controller, and direct attach with two hosts)
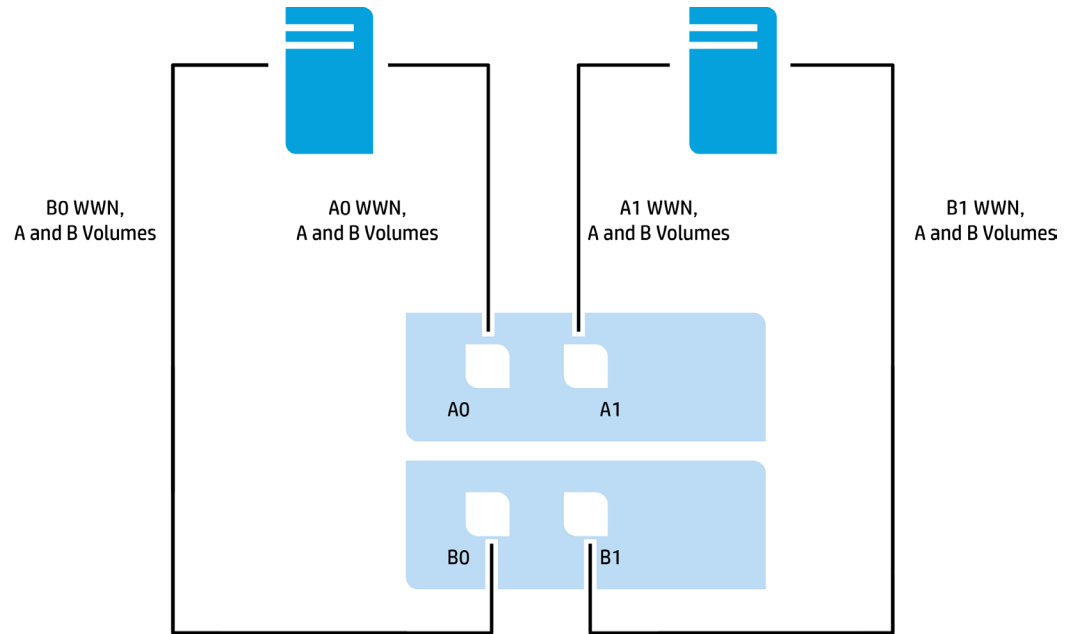
If a controller fails, the hosts maintain access to all of the volumes through the host ports on the surviving controller, as shown in the Figure 4.

Figure 4. FC storage presentation during failover (high-availability, dual-controller, and direct attach with two hosts)
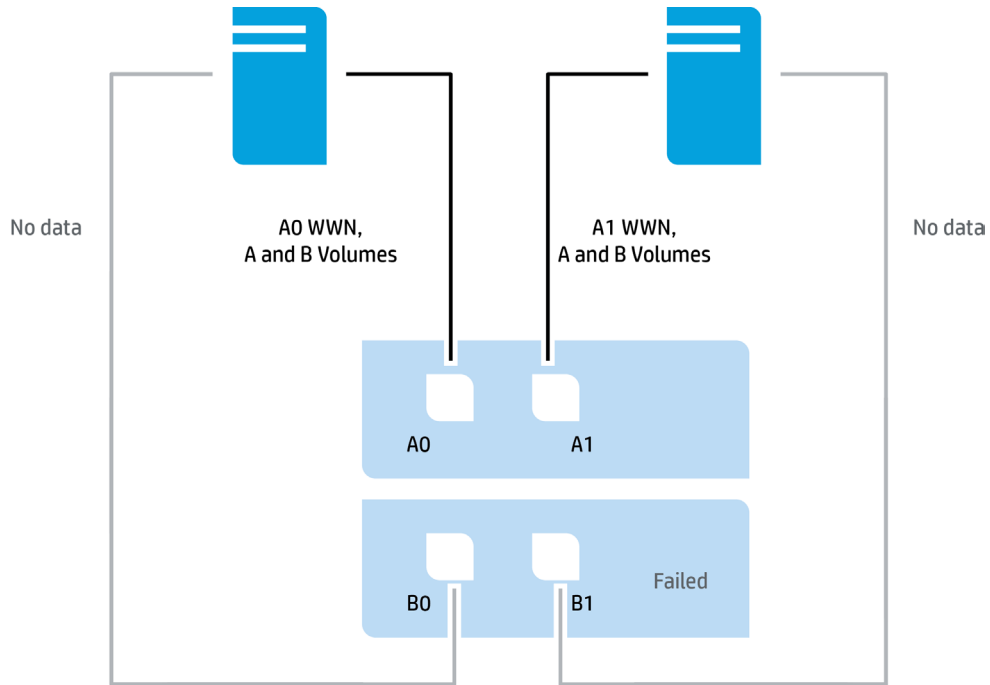
No data

A0 WWN,
A and B Volumes

A1 WWN,
A and B Volumes

No data

A0   A1

B0   B1   Failed

In the following configuration, each host has a non-redundant connection to all mapped volumes. If a controller fails, the hosts connected to the surviving controller maintain access to all volumes owned by that controller. The hosts connected to the failed controller will lose access to volumes owned by the failed controller.

Figure 5. FC storage presentation during normal operation (high-availability, dual-controller, direct attach with four hosts)

B0 WWN,
A and B Volumes

A0 WWN,
A and B Volumes

A1 WWN,
A and B Volumes

B1 WWN,
A and B Volumes

A0   A1

B0   B1

FC switch-attach configuration

When using a switch configuration, it is important to have at least one port connected from each switch to each controller for redundancy. See Figure 6.

Figure 6. FC storage presentation during normal operation (high-availability, dual-controller, and switch attach with four hosts)
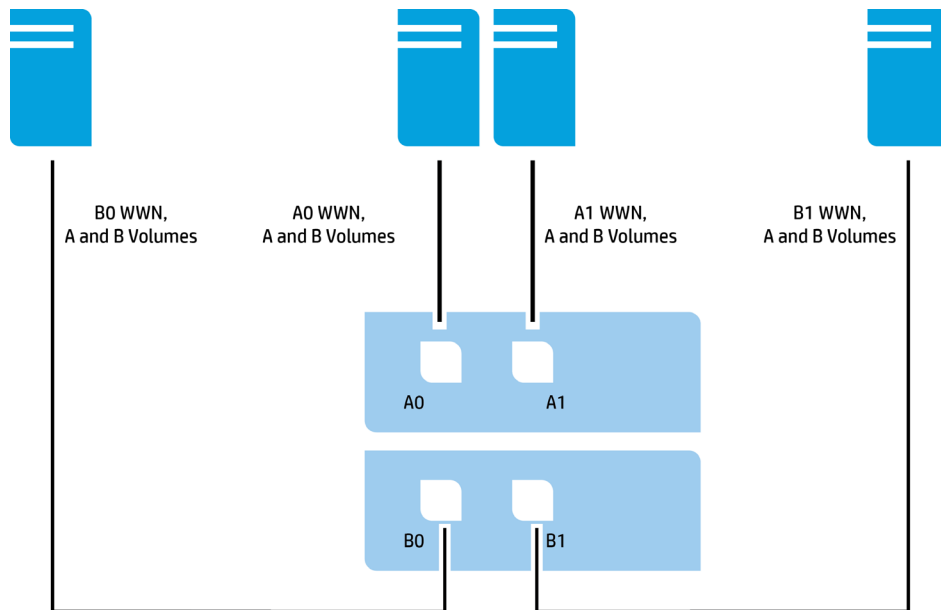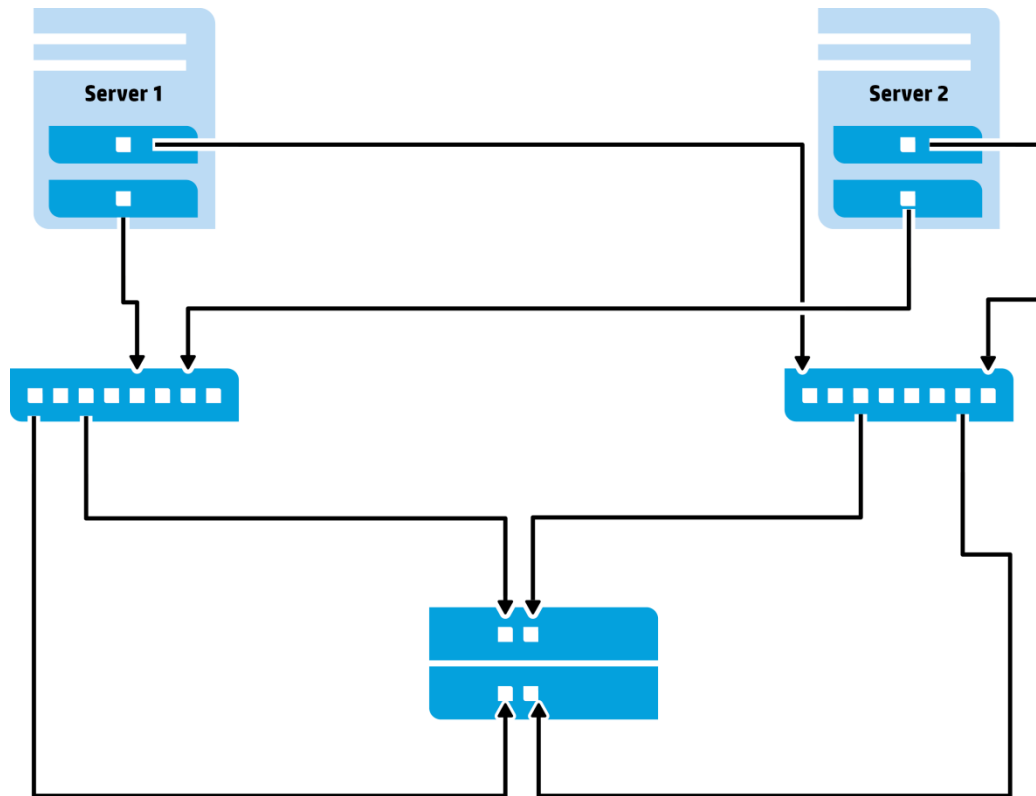

If controller B fails in this setup, the preferred path will shift to controller A and all volumes will be still accessible to both servers as in Figure 6. Each switch has a redundant connection to all mapped volumes; therefore, the hosts connected to the surviving controller maintain access to all volumes.

Virtual disks

A Vdisk is a group of disk drives configured with a RAID level. Each virtual disk can be configured with a different RAID level. A virtual disk can contain SATA drives or SAS drives, but not both. The controller safeguards against improperly combining SAS and SATA drives in a virtual disk. The system displays an error message if you choose drives that are not of the same type.

The HP P2000 G3 system can have a maximum of 16 virtual disks per controller in a dual controller system for a total of 32 virtual disks. In a single controller system, the maximum is 32 virtual disks.

For storage configurations with many drives, it is recommended to consider creating a few virtual disks each containing many drives, as opposed to many virtual disks each containing a few drives. Having many virtual disks is not very efficient in terms of drive usage when using RAID 3 or RAID 5. For example, one 12-drive RAID 5 virtual disk has one parity drive and 11 data drives, whereas four 3- drive RAID 5 virtual disks each have one parity drive (four total) and two data drives (only eight total).

A virtual disk can be larger than 2 TB. This can increase the usable storage capacity of configurations by reducing the total number of parity disks required when using parity-protected RAID levels. However, this differs from using volumes larger than 2 TB, which requires specific operating system, HBA driver, and application-program support.

Note

The P2000 G3 can support a maximum volume size of 16 TB.

Supporting large storage capacities requires advanced planning because it requires using large virtual disks with several volumes each or many virtual disks. To increase capacity and drive usage (but not performance), you can create virtual disks larger than 2 TB and divide them into multiple volumes with a capacity of 2 TB or less.

The largest supported Vdisk is the number of drives allowed in a RAID set multiplied by the largest drive size.

- RAID 0, 3, 5, 6, 10 can support up to 16 drives.
- RAID 50 can support up to 32 drives.

Tip 1
The best practice for creating virtual disks is to add them evenly across both controllers. With at least one virtual disk assigned to each controller, both controllers are active. This active-active controller configuration allows maximum use of a dual-controller configuration's resources.

Tip 2
Another best practice is to stripe virtual disks across shelf enclosures to enable data integrity in the event of an enclosure failure. A virtual disk created with RAID 1, 10, 3, 5, 50, or 6 can sustain an enclosure failure without loss of data depending on the number of shelf enclosures attached. The design should take into account whether spares are being used and whether the use of a spare can break the original design. A plan for evaluation and possible reconfiguration after a failure and recovery should be addressed. Non-fault tolerant Vdisks do not need to be dealt with in this context because a shelf enclosure failure with any part of a non-fault tolerant Vdisk can cause the Vdisk to fail.

Chunk size
When you create a virtual disk, you can use the default chunk size or one that better suits your application. The chunk (also referred to as stripe unit) size is the amount of contiguous data that is written to a virtual disk member before moving to the next member of the virtual disk. This size is fixed throughout the life of the virtual disk and cannot be changed. A stripe is a set of stripe units that are written to the same logical locations on each drive in the virtual disk. The size of the stripe is determined by the number of drives in the virtual disk. The stripe size can be increased by adding one or more drives to the virtual disk.

Available chunk sizes include:

- 16 KB
- 32 KB
- 64 KB (default)

If the host is writing data in 16 KB transfers, for example, then that size would be a good choice for random transfers because one host read would generate the read of exactly one drive in the volume. That means if the requests are random-like, then the requests would be spread evenly over all of the drives, which is good for performance.

If you have 16-KB accesses from the host and a 64 KB block size, then some of the host's accesses would hit the same drive; each stripe unit contains four possible 16-KB groups of data that the host might want to read.

Alternatively, if the host accesses were 128 KB in size, then each host read would have to access two drives in the virtual disk. For random patterns, that ties up twice as many drives.

Note
On RAID 50 drives, the chunk size is displayed as: <requested Chunk size> * (Num drives in sub Vdisk—1)

For example: A requested chunk size of 32 KB with 4 drives in a sub array. The chunk size is reported as 96 KB.

Using the formula: 32 K byte* (4-1) = 96 KB.

Tip
The best practice for setting the chuck size is to match the transfer block size of the application.

Vdisk initialization

During the creation of a Vdisk, the manage user has the option to create a Vdisk in online mode (default) or offline mode, only after the manage user has the advanced user type. By default, the manage user has the standard user type.

If the "online initialization" option is enabled, you can use the Vdisk while it is initializing, but because the verify method is used to initialize the Vdisk, initialization takes more time. Online initialization is fault tolerant.

If the "online initialization" option is unchecked ("offline initialization"), you must wait for initialization to complete before using the Vdisk, but the initialization takes less time.

To assign the advanced user type to the manage user, log into the HP Storage Management Utility (SMU) and make sure the P2000 G3 on the left frame is highlighted and then click the Configuration drop-down box. Then click Users → Modify User.
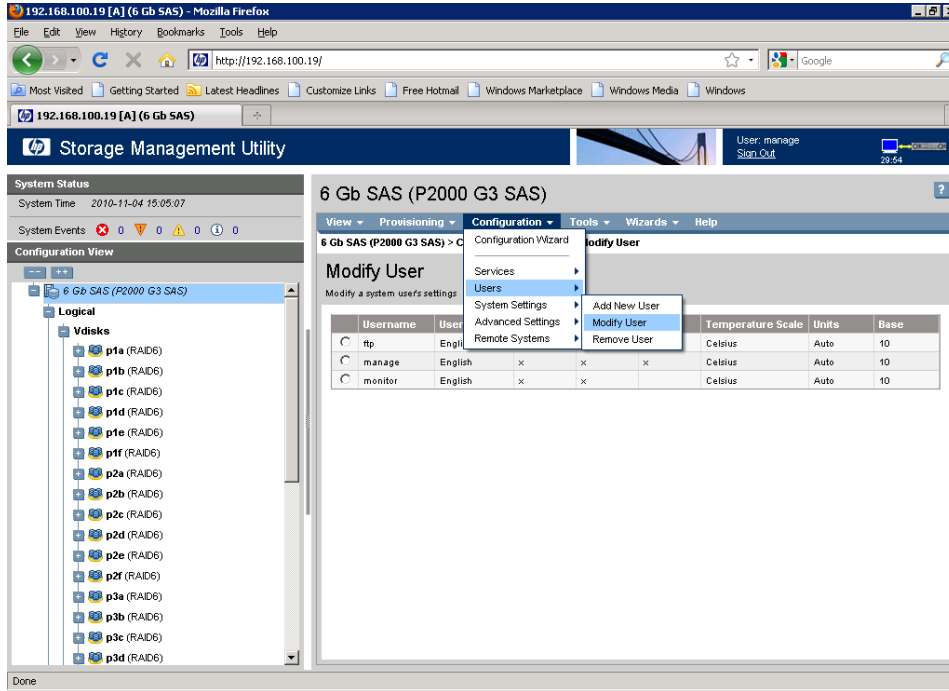
Click the radio button next to the manage user and type in the manage user password. From User Type, select "Advanced" and then to save the change, click the Modify User button → then OK.
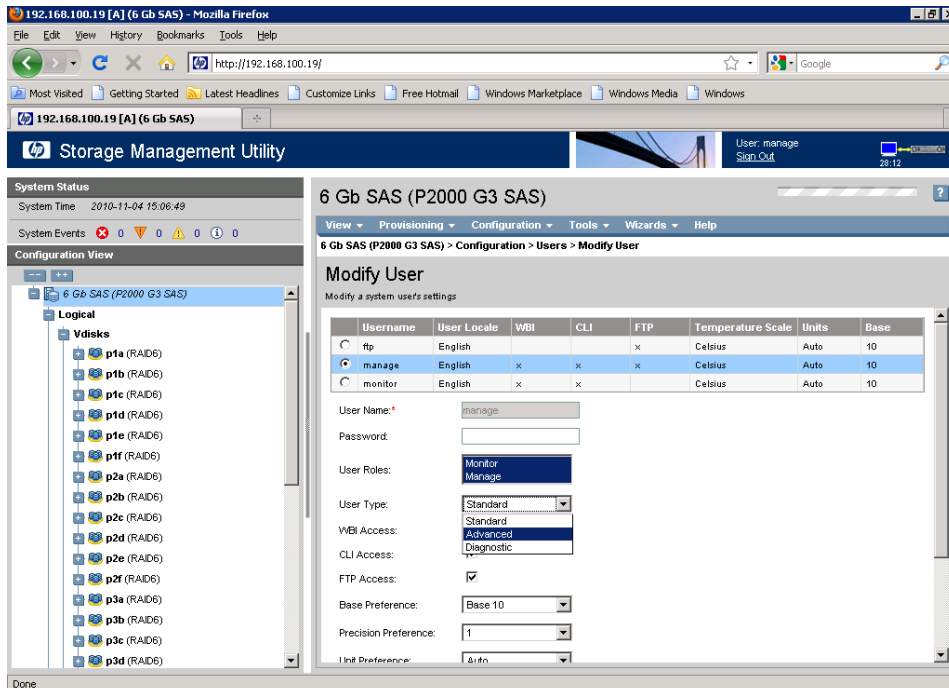
Virtual Disk Expansion Best Practices

With the ever changing storage needs seen in the world today, there comes a time when storage space gets exhausted quickly. The P2000 G3 MSA gives you the option to grow the size of a LUN to keep up with your dynamic storage needs.

A Virtual Disk (Vdisk) expansion allows you to grow the size of a Vdisk in order to expand an existing volume or create volumes from the newly available space on the Vdisk. Depending on several factors, Vdisk expansion can take a significant amount of time to complete. For faster alternatives, see the "Vdisk Expansion Recommendations" section on page 15.

These factors include but are not limited to:

Physical Disk size, Number of disks to expand (1-4), and I/O activity during Vdisk expansion.

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

Vdisk Expansion Capability for Supported RAID Levels

The chart below gives information on the expansion capability for the P2000 G3 MSA supported RAID levels.

Expansion capability for each RAID level

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

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

If an expanding Vdisk becomes DEGRADED (example: RAID 6 with a single drive failure) the storage administrator should determine the level of risk of continuing to allow the expansion to complete versus the time required to backup, re-create the Vdisk (see below under "Vdisk Expansion Recommendations") and restore the data to the volumes on the Vdisk.

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

Vdisk Expansion

Vdisk expansion is the standard process of increasing the available capacity in a Vdisk. This process can be very time consuming. There is no way to reliably determine when the expansion will be complete and when other disk utilities will be available.

Follow the procedure below
Procedure:

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

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

Vdisk Expansion Recommendations

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

Allocate "quiet" period(s) to help optimize Vdisk expansion

This method of expansion utilizes the expand capability of the system and requires manual intervention from the administrator to get the best possible performance from the expansion.

In this context, a "quiet" period indicates a length of time when there is no Host I/O or data access to the system.

Follow the procedure below
Procedure:

1. Backup the current data from the existing Vdisk.
2. Shutdown all hosts connected to the P2000 G3 system.
3. Label and disconnect host side cables from the P2000 G3 system.
4. Using the WBI or CLI, start the Vdisk expansion.
5. Monitor the Vdisk expansion percentage complete.

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

1. Re-connect host side cables to the P2000 G3 system.
2. Restart hosts connected to the P2000 G3 system.

If additional "quiet" periods are required:

1. Shutdown all hosts connected to the P2000 G3 system.
2. Label and disconnect host side cables from the P2000 G3 system.
3. Monitor the Vdisk expansion percentage complete.

Vdisk expansion can take a few hours with no data access for smaller capacity hard drives (72 GB) and may take several days to complete with larger capacity hard drives (2 TB). This procedure can be paused by reconnecting the host side cables and restarting hosts. Priority is given to Host I/O or data access over the expansion process during normal array operation. While the system is responding to Host I/O or data access requests, it may seem as if the expansion process has stopped. When expanding during "quiet" periods, expansion time is minimized and will allow quicker restoration of other disk utilities.

Re-create the Vdisk with additional capacity and restore data

This method is the easiest and fastest method for adding additional capacity to a Vdisk. The online Vdisk initialization allows a user to access the Vdisk almost immediately and will complete quicker than the expansion process on a Vdisk that is also servicing data requests.

Follow the procedure below
Procedure:

1. Backup the current data from the existing Vdisk.
2. Delete the current Vdisk.
3. Using the WBI or CLI, create a new Vdisk with the available hard drives.
4. Create a volume.
5. Restore data to the new volume.

Consult the "Virtual disks" section for Best Practices when creating Virtual Disks.

Volume mapping

It is a best practice to map volumes to the preferred path. The preferred path is both ports on the controller that owns the Vdisk.

If a controller fails, the surviving controller will report it is now the preferred path for all Vdisks. When the failed controller is back online, the Vdisks and preferred paths switch back.

Best Practice
For fault tolerance, HP recommends mapping the volumes to all available ports on the controller.

For performance, HP recommends mapping the volumes to the ports on the controller that owns the Vdisk. Mapping to the non-preferred path results in a slight performance degradation.

Optimum performance with MPIO can be achieved with volumes mapped on all paths. 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.

Note
By default, a new volume will have the "all other hosts read-write access" mapping, so the manage user must go in and explicitly assign the correct volume mapping access.

RAID levels

Choosing the correct RAID level is important whether your configuration is for fault tolerance or performance. Table 1 gives an overview of supported RAID implementations highlighting performance and protection levels.

Note
Non-RAID is supported for use when the data redundancy or performance benefits of RAID are not needed; no fault tolerance.

Table 1. An overview of supported RAID implementations

RAID level	Cost	Performance	Protection level
RAID 0 Striping	N/A	Highest	No data protection
RAID 1 Mirroring	High cost 2x drives	High	Protects against individual drive failure
RAID 3 Block striping with dedicated parity drive	1 drive	Good	Protects against individual drive failure
RAID 5 Block striping with striped parity drive	1 drive	Good	Protects against any individual drive failure; medium level of fault tolerance
RAID 6 Block striping with multiple striped parity	2 drives	Good	Protects against multiple (2) drive failures; high level of fault tolerance
RAID 10 Mirrored striped array	High cost 2x drives	High	Protects against certain multiple drive failures; high level of fault tolerance
RAID 50 Data striped across RAID 5	At least 2 drives	Good	Protects against certain multiple drive failures; high level of fault tolerance

Spares

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

At least one Vdisk must exist before you can add a global spare. A spare must have sufficient capacity to replace the smallest disk in an existing Vdisk. If a drive in the virtual disk fails, the controller automatically uses the Vdisk spare for reconstruction of the critical virtual disk to which it belongs. A spare drive must be the same type (SAS or SATA) as other drives in the virtual disk. You cannot add a spare that has insufficient capacity to replace the largest drive in the virtual disk. If two drives fail in a RAID 6 virtual disk, two properly sized spare drives must be available before reconstruction can begin. For RAID 50 virtual disks, if more than one sub-disk becomes critical, reconstruction and use of Vdisk spares occur in the order sub-Vdisks are numbered.

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

Working with Failed Drives and Global Spares

When a failed drive rebuilds to a spare, the spare drive now becomes the new drive in the virtual disk. At this point, the original drive slot position that failed is no longer part of the virtual disk. The original drive now becomes a "Leftover" drive.

In order to get the original drive slot position to become part of the virtual disk again, do the following:

1. Replace the failed drive with a new drive.
2. If the drive slot is still marked as "Leftover", use the "Clear Disk Metadata" option found in the "Tools" submenu.
3. When the new drive is online and marked as "Available", configure the drive as a global spare drive.
4. Fail the drive in the original global spare location by removing it from the enclosure. The RAID engine will rebuild to the new global spare which will then become an active drive in the RAID set again.
5. Replace the drive you manually removed from the enclosure.
6. If the drive is marked as "Leftover", clear the metadata as in step 2 above.
7. Re-configure the drive as the new global spare.

Tip

A best practice is to designate a spare disk drive for use if a drive fails. Although using a dedicated Vdisk spare is the best way to provide spares for your virtual disks, it is also expensive to keep a spare assigned to each virtual disk. An alternative method is to enable dynamic spares or to assign one or more unused drives as global spares.

Cache configuration

Controller cache options can be set for individual volumes to improve a volume's fault tolerance and I/O performance.

Note

To change the following cache settings, the user—who logs into the HP SMU—must have the "advanced" user credential. The manage user has the "standard" user credential by default. This credential can be changed using the HP SMU and click on "Configuration," then "Users," then "Modify Users."

Write-back cache settings

Write back is a cache-writing strategy in which the controller receives the data to be written to disk, 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 drive. Write-back cache mirrors all of the data from one controller module cache to the other. Write-back cache improves the performance of write operations and the throughput of the controller.

When write-back cache is disabled, write-through becomes the cache-writing strategy. Using write-through cache, the controller writes the data to the disk before signaling the host operating system that the process is complete. Write-through cache has lower throughput and write operation performance than write back, but it is the safer strategy, with low risk of data loss on power failure. However, write-through cache does not mirror the write data because the data is written to the disk before posting command completion and mirroring is not required. You can set conditions that cause the controller to switch from write-back caching to write-through caching as described in "Auto Write Through Trigger and Behavior Settings" later in this paper.

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

You can enable and disable the write-back cache for each volume. By default, volume write-back cache is enabled. Data is not lost if the system loses power because controller cache is backed by super capacitor technology.

For most applications this is the correct setting, but because backend bandwidth is used to mirror cache, if you are writing large chunks of sequential data (as would be done in video editing, telemetry acquisition, or data logging) write-through cache has much better performance. Therefore, you might want to experiment with disabling the write-back cache. You might see large performance gains (as much as 70 percent) if you are writing data under the following circumstances:

• Sequential writes
• Large I/Os in relation to the chunk size
• Deep queue depth

If you are doing any type of random access to this volume, leave the write-back cache enabled.

Caution
Write-back cache should only be disabled if you fully understand how your operating system, application, and HBA (SAS) move data. You might hinder your storage system's performance if used incorrectly.

Auto-write through trigger and behavior settings
You can set the trigger conditions that cause the controller to change the cache policy from write-back to write-through. While in write-through mode, system performance might be decreased.

A default setting makes the system revert to write-back mode when the trigger condition clears.

To make sure that this occurs and that the system doesn't operate in write-through mode longer than necessary, make sure you check the setting in HP SMU or the CLI.

You can specify actions for the system to take when write-through caching is triggered:

• Revert when Trigger Condition Clears: Switches back to write-back caching after the trigger condition is cleared. The default and best practice is Enabled.
• Notify Other Controller: In a dual-controller configuration, the partner controller is notified that the trigger condition is met. The default is Disabled.

Cache configuration summary
The following guidelines list the general best practices. When configuring cache:

• For a fault-tolerant configuration, use the write-back cache policy, instead of the write-through cache policy.
• For applications that access both sequential and random data, use the standard optimization mode, which sets the cache block size to 32 KB. For example, use this mode for transaction-based and database update applications that write small files in random order.
• For applications that access only sequential data and that require extremely low latency, use the super-sequential optimization mode, which sets the cache block size to 128 KB. For example, use this mode for video playback and multimedia post-production video- and audio-editing applications that read and write large files in sequential order.

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

Table 2. Optimizing performance for your application

Application	RAID level	Read ahead cache size	Cache optimization
Default	5 or 6	Default	Standard
High-Performance Computing (HPC)	5 or 6	Maximum	Standard
Mail Spooling	1	Default	Standard
NFS_Mirror	1	Default	Standard
Oracle_DSS	5 or 6	Maximum	Standard
Oracle_OLTP	5 or 6	Maximum	Standard
Oracle_OLTP_HA	10	Maximum	Standard
Random 1	1	Default	Standard
Random 5	5 or 6	Default	Standard
Sequential	5 or 6	Maximum	Super-Sequential
Sybase_DSS	5 or 6	Maximum	Standard
Sybase_OLTP	5 or 6	Maximum	Standard
Sybase_OLTP_HA	10	Maximum	Standard
Video Streaming	1 or 5 or 6	Maximum	Super-Sequential
Exchange Database	5 for data; 10 for logs	Default	Standard
SAP	10	Default	Standard
SQL	5 for data; 10 for logs	Default	Standard

For SQL server and Microsoft® Exchange Best Practices and Performance on the P2000 G3 MSA, consult ActiveAnswers at: h71019.www7.hp.com/ActiveAnswers/us/en/aa-categories.html

World Wide Name naming conventions

A best practice for acquiring and renaming World Wide Names (WWN) for the P2000 G3 MSA is to plug-in one cable connection at a time and then rename the WWN to an identifiable name.

The procedure below outlines the steps using the P2000 G3 SAS model as an example.

Procedure
1. Login to the Storage Management Utility (SMU). The Status Summary page will be displayed.
2. Click "+" next to "Hosts" from the left Windows frame. This will expand the list to show all connected hosts.
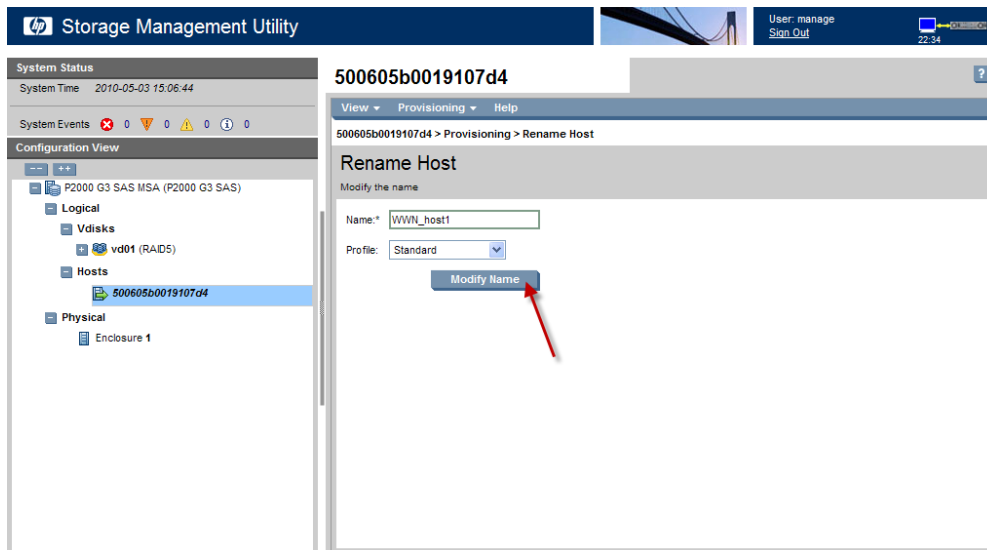
3. Highlight the host in the list that you want to rename by clicking the WWN name.
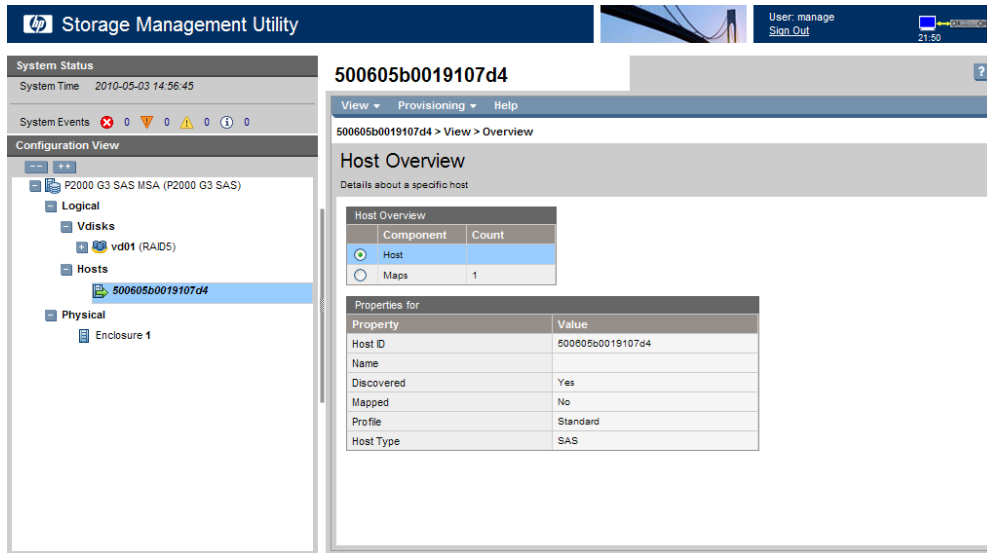
4. On the right window frame, click Provisioning -> Rename Host.
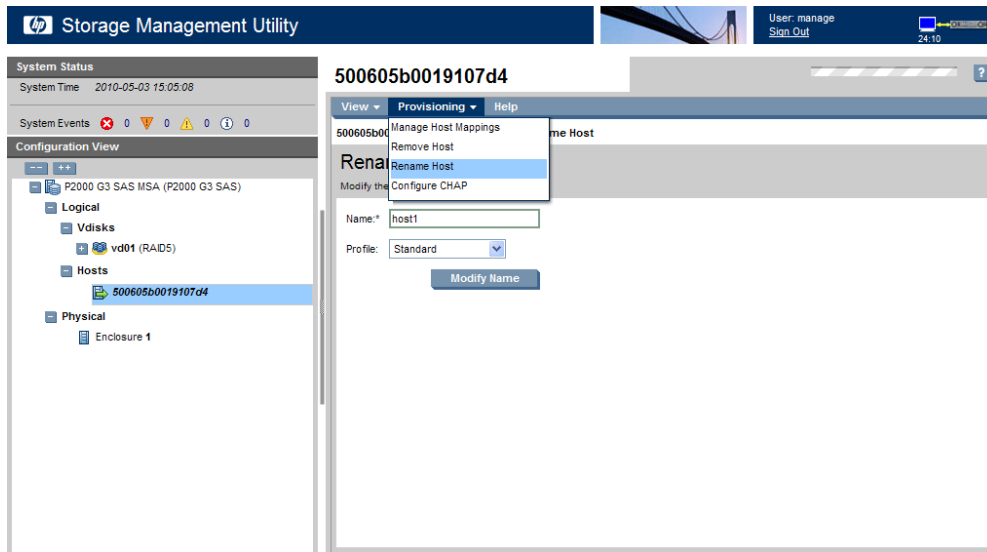
5. Type in the host nickname and choose the Profile and then click Modify Name. Click OK on the pop-up window.

6. Plug in the SAS port of the HBA on the second server into the P2000 G3 SAS MSA controller port. Make sure the server is powered on.

7. Repeat steps 3–5 for the remaining servers.

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 for 8 Gb/sec on the P2000 G3 FC MSA.
- Host ports should be configured for 1 Gb/sec on the P2000 G3 iSCSI MSA.
- Virtual disks should be balanced between the two controllers.
- Disk drives should be balanced between the two controllers.
- Cache settings should be set to match (Optimizing performance for your application) for the application.
- In order to get the maximum sequential performance from a Vdisk, you should only create one volume per Vdisk. Otherwise you will introduce randomness into the workload when multiple volumes on the Vdisk are being exercised.
- Distribute the load across as many drives as possible.

Highest fault tolerance optimization

The following guidelines list the general best practices to follow when configuring your storage system for highest fault tolerance:

- Use dual controllers
- Use two cable connections from each host Use
- Multipath Input/Output (MPIO) software

Boot from storage considerations

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

Disk Background Scrub, Drive Spin Down, and SMART

The P2000 G3 MSA also uses the disk background scrubbing feature. You can scrub disk drives that are in a Vdisk or have not yet been assigned to a Vdisk.

The P2000 G3 MSA has now added the power saving feature called drive spin down (DSD). The drive spin down is a cost and power saving tool. The drive spin down feature will stop virtual disks, available disk drives and global spares disk from spinning.

Self-Monitoring Analysis and Reporting Technology (SMART) can alert the controller of impending disk failure. When SMART is enabled, the system checks for SMART events one minute after a restart and every five minutes thereafter. SMART events are recorded in the event log.

Configuring background scrub for Vdisks

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

Vdisk-level errors reported include:

Hard errors, medium errors, and bad block replacements (BBRs).

Disk-level errors reported include:

Metadata read errors, SMART events during scrub, bad blocks during scrub, and new disk defects during scrub.

For RAID 3, 5, 6, and 50, the utility checks all parity blocks to find data-parity mismatches.

For RAID 1 and 10, the utility compares the primary and secondary disks to find data inconsistencies. For NRAID (Non-RAID, non-striped) and RAID 0, the utility checks for media errors.

You can use a Vdisk while it is being scrubbed. Background Vdisk scrub runs at background utility priority, which reduces to no activity if CPU usage is above a certain percentage or if I/O is occurring on the Vdisk being scrubbed. A Vdisk scrub may be in process on multiple Vdisks at once.

A new Vdisk will first be scrubbed 20 minutes after creation. After a Vdisk is scrubbed, scrub will start again after the interval specified by the Vdisk Scrub Interval (hours) option. When a scrub is complete, the number of errors found is reported with event code 207 in the event log.

Note
If you choose to disable background Vdisk scrub, you can still scrub a selected Vdisk by using Media Scrub Vdisk.

Utility priority
- High: Use when your highest priority is to get the system back to a fully fault-tolerant state. This causes heavy I/O with the host to be slower than normal. This value is the default.
- Medium: Use when you want to balance data streaming with data redundancy.
- Low: Use when streaming data without interruption, such as for a Web server, is more important than data redundancy. This enables a utility such as Reconstruct to run at a slower rate with minimal effect on host I/O.
- Background: Utilities run only when the processor has idle cycles.

Best Practice: Leave the default setting of Background Scrub ON in the background priority for both Vdisks and available disks.

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

Note
DSD affects disk operations as follows:

Spun-down disks are not polled for SMART events.

Operations requiring access to disks may be delayed while the disks are spinning back up.

Best Practice: Set DSD for unconfigured disks, spares, and configured Vdisks that do not perform a read/write operation at least once every 24 hours.

SMART settings
- Don't Modify: Allows current disks to retain their individual SMART settings and does not change the setting for new disks added to the system.
- Enabled: Enables SMART for all current disks after the next rescan and automatically enables SMART for new disks added to the system. This option is the default.
- Disabled: Disables SMART for all current disks after the next rescan and automatically disables SMART for new disks added to the system.

Best Practice: HP recommends using the default value "Enabled."

Cascading Array Enclosures

Since the P2000 G3 MSA can be upgraded from any MSA2000 G1/G2 model, the existing MSA70s and MSA2000 expansion array enclosures might hold data that you need. The MSA70 and MSA2000 expansion enclosures operate at 3 Gb, see the Figures 7 and 8.

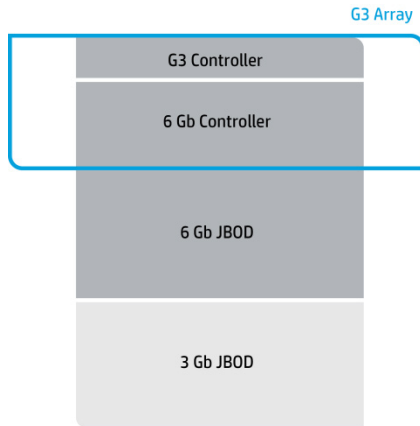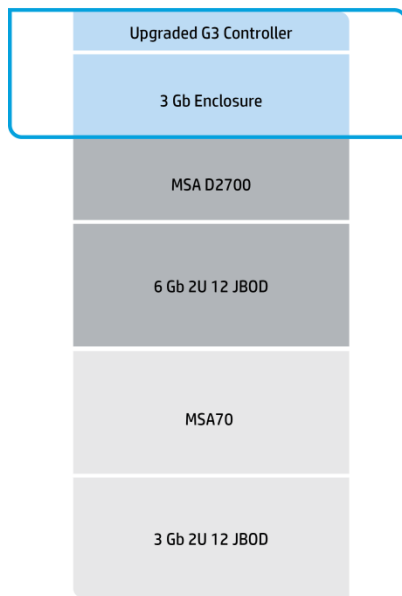**Figure 7.** Mixed 3 Gb and 6 Gb JBODs behind a P2000 G3 FC MSA array

Figure 8. Mixed 3 Gb and 6 Gb JBODs behind a P2000 G3 FC MSA array

Note

The 6 Gb MSA D2700 must come before the new 6 Gb MSA2000 2U 12 Expansion Enclosure in the cascade chain.

If MSA2000 2U 12 Expansion Enclosures are used in conjunction with D2700's, the MSA2000 2U 12 Expansion Enclosures MUST come at the end of the cascade chain, and ONLY straight through cabling is allowed. Otherwise, either straight through cabling or reverse cabling is allowed. Note that access to 6 Gb enclosures following a 3 Gb enclosure is restricted to 3 Gb; therefore, if reverse cabling is used, to avoid the drop, place the 3 Gb enclosure in the middle of the cascade chain, and arrange Virtual Disks such that they do not span 6 Gb enclosures at the beginning and end of the cascade chain, and are owned by the controller closest to them in the cascade chain.

Table 3 shows a list of supported enclosures and the SAS rate of the enclosures.

Table 3. SAS rates of supported enclosures for the P2000 G3 MSA

HP Storage MSA System Model No.	Disk Form	Disk Quantity	SAS Rate
P2000 G3 iSCSI MSA SFF (controller enclosure)	2.5"	24-drive	6 Gb
P2000 G3 iSCSI MSA LFF (controller enclosure)	3.5"	12-drive	6 Gb
P2000 G3 MSA 10GbE iSCSI SFF (controller enclosure)	2.5"	24-drive	6 Gb
P2000 G3 MSA 10GbE iSCSI LFF (controller enclosure)	3.5"	12-drive	6 Gb
P2000 G3 MSA SAS SFF (controller enclosure)	2.5"	24-drive	6 Gb
P2000 G3 MSA SAS LFF (controller enclosure)	3.5"	12-drive	6 Gb
P2000 G3 MSA FC/iSCSI SFF (controller enclosure)	2.5"	24-drive	6 Gb
P2000 G3 MSA FC/iSCSI LFF (controller enclosure)	3.5"	12-drive	6 Gb
P2000 G3 MSA FC SFF (controller enclosure)	2.5"	24-drive	6 Gb
P2000 G3 MSA FC LFF (controller enclosure)	3.5"	12-drive	6 Gb
P2000 G3 6 Gb 3.5" 12-drive enclosure	3.5"	12-drive	6 Gb
D2700 6 Gb drive enclosure	2.5"	25-drive	6 Gb
MSA2000 3.5" 12-drive enclosure	3.5"	12-drive	3 Gb
MSA70 drive enclosure	2.5"	25-drive	3 Gb

For the P2000 G3 MSA cabling, consult the document titled "HP P2000 G3 MSA System Cable Configuration Guide" found at bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c02254377/c02254377.pdf

8 Gb Switches and SFP transceivers

The 8 Gb switches that HP offers use differing models of Small Form-Factor Pluggable (SFP) transceivers. The correct SFPs must be loaded into the correct supported 8 Gb switches when connecting to the P2000 G3 FC MSA. If the wrong SFPs are used in an unsupported 8 Gb switch, the storage on the P2000 G3 FC MSA will not be visible.

Here is a list of 2 SFPs and 1 supported switch for each:

• SFP part number AJ71 8A will work with the HP 8/20q Fibre Channel Switch (HP P\N: AM868A Revision: 0A) (20 ports)
• SFP part number AJ71 6A will work with the Brocade 8 Gb SAN Switch (HP P/N: AK242-63001 Revision: 0C) (24 ports)
• Consult the QuickSpecs for other supported switches

MSA70 considerations

Dual-domains
When using the MSA70 with dual-domains, dual I/O modules, make sure the following procedure is followed.

MSA70 systems with firmware earlier than 1.50
If your MSA70 has installed firmware earlier than version 1.50, you must replace the chassis backplane before installing a second I/O module in the chassis. To determine your installed firmware version, use a server-based tool such as HP Systems Insight Manager or your Management Agents.

If installed firmware is earlier than 1.50, do the following:

• Contact HP Support and order a replacement backplane:
  MSA70:   430149-001

Caution
Be sure to order the part number indicated in this notice, not the spare part number printed on your existing backplanes.

Be sure to order a quantity of two replacement kits.

- Install the replacement backplane using instructions shipped with the backplane.
- Install the additional I/O module using instructions shipped with the I/O module.

Firmware versions
If there are MSA70 enclosures connected to the P2000 G3 MSA, make sure that the firmware on the enclosure is 2.18 or greater. If the MSA70 has a firmware version prior to 2.18, the MSA70 will be in a degraded state and virtual disks cannot be created or accessed from the MSA70.

Administering with HP SMU

If you choose to use the HP Storage Management Utility (SMU) for administration, it is best to use either the Firefox 3.0 or later or Internet Explorer 7 or later Web browsers.

P2000 G3 iSCSI MSA considerations

When using the P2000 FC/iSCSI combo G3 MSA (for iSCSI traffic), P2000 G3 10GbE iSCSI MSA, or P2000 G3 iSCSI MSA, 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 makes sure that the storage network is isolated from the other networks.

The private LAN is the network that goes from the server to the P2000 FC/iSCSI combo G3 MSA (for iSCSI traffic), P2000 G3 10GbE iSCSI MSA, or the P2000 G3 iSCSI MSA. This private LAN is the storage network. The storage network should be isolated from the Public network to improve performance. See Figure 9.

Figure 9. P2000 FC/iSCSI combo G3 MSA (for iSCSI traffic) Network
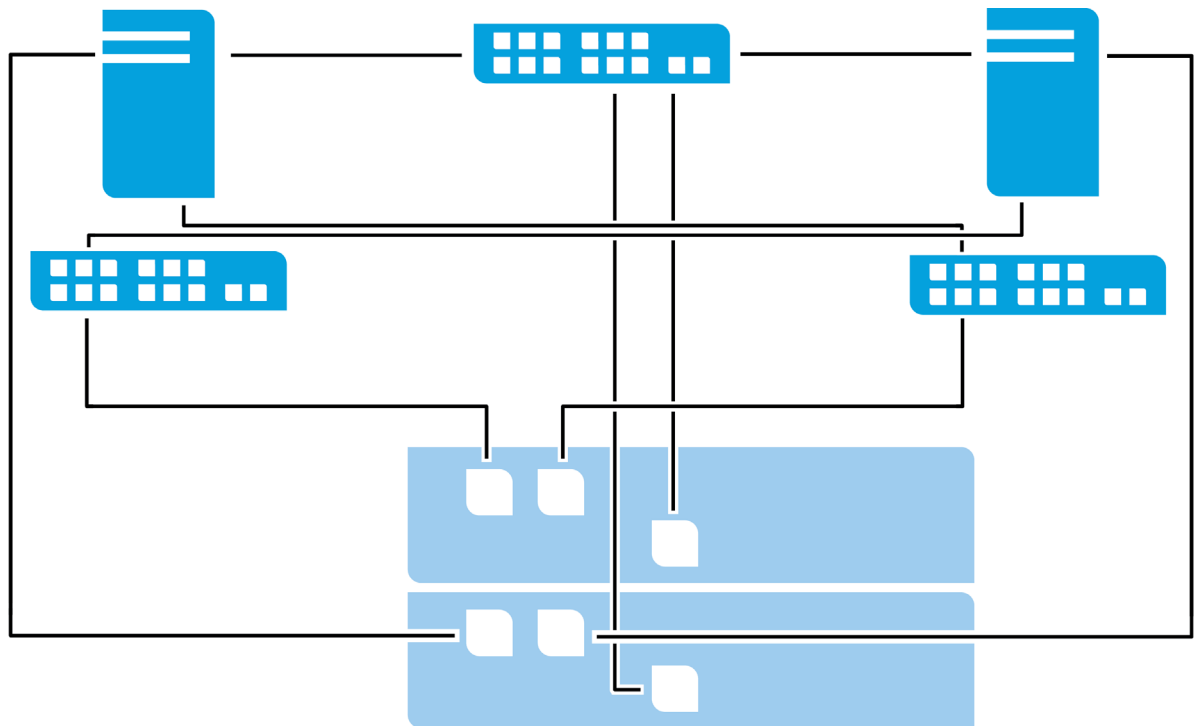
IP Address scheme for the controller pair
The P2000 FC/iSCSI combo G3 MSA (for iSCSI traffic) and P2000 G3 10GbE iSCSI MSA use port 0 of each controller as one failover pair, and port 1 of each controller as a second failover pair. Therefore, port 0 of each controller must be in the same subnet, and port 1 of each controller should be in a second subnet.

For example (with a netmask of 255.255.255.0):

- Controller A port 0: 10.10.10.100
- Controller A port 1: 10.11.10.120
- Controller B port 0: 10.10.10.110
- Controller B port 1: 10.11.10.1 30

Note
In the case of the P2000 G3 iSCSI MSA, set up the scheme similar to the following example (with a netmask of 255.255.255.0):

- Controller A port 0: 10.10.10.100
- Controller A port 1: 10.11.10.120
- Controller A port 2: 10.10.10.110
- Controller A port 3: 10.11.10.130
- Controller B port 0: 10.10.10.140
- Controller B port 1: 10.11.10.150
- Controller B port 2: 10.10.10.160
- Controller B port 3: 10.11.10.170

Using optic cables

The P2000 G3 10GbE iSCSI MSA supports fiber optic cables with Short Range (SR) and Long Range Multi-mode (LRM) SFP+ transceivers only.

Consult the QuickSpecs for more information.

Software

The section below introduces the HP P2000 Modular Smart Array Software Support/Documentation CD.

Versions

3.30—found in the shipping software kit

3.35—Web launch version

Description
The HP P2000 Modular Smart Array Software Support/Documentation CD provides deployment and maintenance support for the HP P2000 G3 Modular Smart Array Family products which includes the P2000 G3 FC MSA, the P2000 G3 FC/iSCSI MSA, P2000 G3 SAS MSA, P2000 G3 10GbE MSA, and the P2000 G3 iSCSI MSA.

HP P2000 Modular Smart Array Software Support/Documentation CD

HP P2000 Modular Smart Array Software Support/Documentation CD contains a common Windows/Linux navigation HTML framework to provide customers with a common installation experience.

This CD contains end-user documents, host server software deliverables, and deployment and installation tools to simplify the setup and maintenance of your HP P2000 G3 Modular Smart Array Family product. The CD contains tabulated groups for documents, software, firmware, setup, tools, and service. This CD also contains the latest software drivers and user documents along with search links to secure the latest version from hp.com/us/en/support-drivers.html

Here are some of the significant features of the HP P2000 Modular Smart Array Software Support/Documentation CD:

- Provides step-by-step install instructions for each supported product with links to user documents included on the CD.
- Contains host software and drivers in various forms. Available only for the Web launch CD (3.25).
  - OS specific Host Software bundles.
  - Additional packages (Microsoft hot fixes and other drivers) that cannot be installed through the bundles.

- Contains the listing of all current P2000 G3 Modular Smart Array Family Firmware.
  - Provides search links to get the latest firmware from <u>hp.com</u>.

- Contains all the P2000 G3 Modular Smart Array user documents in one place.
  - User documents included on the CD with Internationalized versions where available.
  - Links to additional documentation resources like white papers are also available on the CD.

- Provides additional tools that assist customers with various management tasks.
  - MSA Device Discovery Tool
    - Assists in discovering HP 2000 Modular Smart Array Family (MSA2000) and HP P2000 Modular Smart Array Systems that are direct attach or reachable over the network.
    - Allows users to launch any of the management interfaces like SMU/Telnet/CLI/FTP for the selected device.
    - Provides an option to schedule log collection from the selected MSA device to pull the storage debug logs onto the host server at specified intervals.
    - Generates XML/Text output reports with inventory details of the local host and the discovered HP 2000 Modular Smart Array Family (MSA2000) and HP P2000 Modular Smart Array devices.
  - SNMP MIBs—The MSA2000 SNMP MIBs provide MIB browsers and other SNMP-aware software with the necessary information to query, update, and properly display SNMP variables on supported hardware.
  - Links to HP SMI-S documentation on the Web.

- Provides various HP Support links for services like Product registration, Warranty, Service Care Packs, Learning Center, and so on.

Host Server Software

HP P2000 Modular Smart Array Software Support/Documentation CD provides various Host Software products corresponding to HP P2000 G3 Modular Smart Array Family products. It also contains links to <u>hp.com</u> where newer versions of these software products may be available.

The following are the key Host Software products contained in HP P2000 Modular Smart Array Software Support/Documentation CD:

Host Software bundles
Separate Host Software bundles are provided for each of the supported OS platforms.

Windows bundles
Windows bundles are based on Windows ProLiant Support Packs. Separate bundles are available for Windows 2003 and 2008, and for each of the supported hardware architectures (x86, x64, and IA64). Bundles include installable Smart Components and the HP SUM engine. The HP SUM engine will have "pull from Web updates," so users can get the latest from Web automatically.

Linux bundles
Linux bundles are based on Linux ProLiant Support Packs. Separate bundles are available for RHEL4, RHEL5, SLES10, SLES11. Bundles include installable RPM packages and the HP SUM engine.

Individual Host Software Smart Components
Individual Smart Components are available for each of the drivers contained in the bundles so customers can choose to install or update a specific driver without going through the bundle installation. Here again, the individual drivers are available for each of the supported OS platforms and hardware architectures. In addition to the drivers locally hosted on the CD, links are provided to <u>hp.com</u> where newer versions of the drivers may be available.

Microsoft Hot Fixes
Also available are the Microsoft Hot Fixes for various Microsoft dependent products like the Storport storage driver, VDS, and VSS. The Setup page provides detailed instructions on the sequence of steps required to install these hot fixes.

Best Practices for Firmware Updates

The sections below details common firmware updates best practices for all generations of the MSA2000/P2000. This includes the MSA2000 G1, MSA2000 G2, and the P2000 G3 MSA.

General P2000/MSA2000 Device Firmware Update Best Practices

- As with any other firmware upgrade, it is a recommended best practice to ensure that you have a full backup prior to the upgrade.
- Before upgrading the firmware, make sure that the storage system configuration is stable and is not being reconfigured or changed in any way. If any configurations changes are in progress, monitor them using the SMU or CLI and wait until they are completed before proceeding with the upgrade.
- Do not power cycle or restart devices during a firmware update. If the update is interrupted or there is a power failure, the module could become inoperative. Should this happen, contact HP customer support.
- After the device firmware update process is completed, confirm the new firmware version is displayed correctly via one of the MSA management interfaces—SMU GUI, MSA CLI, and so on.

P2000/MSA2000 Array Controller or I/O Module Firmware Update Best Practices

- The array controller (or I/O module) firmware can be updated in an online mode only in redundant controller systems.
- When planning for a firmware upgrade, schedule an appropriate time to perform an online upgrade.
  - For single domain systems, I/O must be halted.
  - For dual domain systems, because the online firmware upgrade is performed while host I/Os are being processed, I/O load can impact the upgrade process. Select a period of low I/O activity to ensure the upgrade completes as quickly as possible and avoid disruptions to hosts and applications due to timeouts.
- When planning for a firmware upgrade, allow sufficient time for the update.
  - In single-controller systems, it takes approximately 10 minutes for the firmware to load and for the automatic controller restart to complete.
  - In dual-controller systems, the second controller usually takes an additional 20 minutes, but may take as long as one hour.
- When reverting to a previous version of the firmware, ensure that the Management Controller (MC) Ethernet connection of each storage controller is available and accessible before starting the downgrade.
  - When using a Smart Component firmware package, the Smart Component process will automatically first disable Partner Firmware Update (PFU) and then perform downgrade on each of the controllers separately (one after the other) through the Ethernet ports.
  - When using a Binary firmware package, first disable the PFU option and then downgrade the firmware on each of the controller separately (one after the other).
- When performing firmware updates to MSA70 drive enclosures, each enclosure will need to have a power cycle performed.

P2000/MSA2000 Disk Drive Firmware Update Best Practices

- Disk drive upgrades on the HP P2000/MSA2000 storage systems is an off line process. All host and array I/O must be stopped prior to the upgrade.
- If the drive is in a virtual disk, verify that it is not being initialized, expanded, reconstructed, verified, or scrubbed. If any of these tasks is in progress, before performing the update wait for the task to complete or terminate it. Also verify that background scrub is disabled so that it doesn't start. You can determine this using SMU or CLI interfaces. If using a firmware smart component, it would fail and report if any of the above pre-requisites are not being met.
- Disk Drives of the same model in the storage system must have the same firmware revision. If using a firmware smart component, the installer would ensure all the drives are updated.

Summary

The HP Storage MSA administrators should determine the appropriate levels of fault tolerance and performance that best suits their needs. Following the configuration options listed in this paper can help you make sure that the HP Storage MSA family enclosure is optimized accordingly.

Learn more at
hp.com/go/P2000

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document