



HPE 3PAR Storage best practices guide

A reference and best practices guide for HPE 3PAR Storage

Contents

Executive summary.....	3
Units of measure.....	3
Related documentation.....	3
HPE 3PAR block storage concepts and terminology.....	3
What’s new in HPE 3PAR OS 3.3.1 MU3.....	6
Licensing changes.....	6
Performance enhancements.....	6
Data reduction.....	6
Ease of use.....	6
Increased volume size.....	7
Getting started with HPE 3PAR Storage best practices.....	7
Set up HPE 3PAR system ports.....	8
Fibre Channel host zoning.....	8
Persistent Ports.....	10
High availability.....	11
Naming conventions considerations for hosts and host sets.....	11
Hosts and host sets.....	12
Provisioning block storage from an HPE 3PAR system.....	13
Common provisioning groups.....	13
Adaptive Flash Cache.....	14
Provisioning VVs.....	15
Priority Optimization.....	16
Remote Copy.....	17
Asynchronous streaming Remote Copy.....	18
Adaptive Optimization.....	18
Security.....	20
HPE System Reporter and HPE InfoSight in the SSMC.....	20
Threshold alerts.....	20
Storage analytics on the web.....	21
Autonomic rebalance.....	21
Appendix A: Supported host personas.....	22
Summary.....	23



Executive summary

This reference and best practices guide for HPE 3PAR Storage highlights techniques to ensure optimal use of HPE 3PAR OS 3.3.1 MU3 on HPE 3PAR arrays. It covers configuring for high availability as well as for the HPE 3PAR OS software suite. Following recommended best practices provides the most robust and efficient implementation and paves the way for easy upgrades.

This best practices guide is for system and storage administrators of all levels. Anyone who plans storage policies, configures storage resources, or monitors the storage usage of HPE 3PAR Storage arrays should read this guide.

Units of measure

All units of storage (capacity) are calculated in base 2 (x 1,024). Therefore:

- 1 KiB = 1,024 bytes
- 1 MiB = 2²⁰ bytes = 1,048,576 bytes
- 1 GiB = 2³⁰ bytes = 1,024 MiB = 1,073,741,824 bytes
- 1 TiB = 2⁴⁰ bytes = 1,024 GiB = 1,099,511,627,776 bytes

All units of performance (speed) are calculated in base 10 (x 1000). Therefore:

- 1 KB = 1,000 bytes
- 1 MB = 10⁶ bytes = 1,000,000 bytes
- 1 GB = 10⁹ bytes = 1000 MiB = 1,000,000,000 bytes
- 1 TB = 10¹² bytes = 1000 GiB = 1,000,000,000,000 bytes

Related documentation

Table 1. Places to look for more information

Description	Document location
Complete description of CLI commands	HPE 3PAR Command Line Interface Reference
Overview and explanation of HPE 3PAR technology	HPE 3PAR StoreServ Storage Concepts guide
Using the SSMC to configure and administer the system	HPE 3PAR StoreServ Management Console 3.4 User Guide
Using HPE 3PAR Remote Copy software	HPE 3PAR Remote Copy Software User Guide
Using the HPE 3PAR Common Information Model (CIM)	HPE 3PAR CIM API Programming Reference
Using the HPE 3PAR 3PARinfo tool	HPE 3PAR 3PARinfo for HPE 3PAR StoreServ Storage
Using HPE 3PAR Host Explorer software	HPE 3PAR Host Explorer Software User Guide

HPE 3PAR block storage concepts and terminology

The HPE 3PAR array comprises the following logical data layers:

- Physical disks (PDs)
- Chunklets
- Logical disks (LDs)
- Common provisioning groups (CPGs)
- Virtual volumes (VVs)

The relationship between system data layers is illustrated in Figure 1. Each layer is created from elements of the layer above. Chunklets are drawn from physical disks. Logical disks are created from groups of chunklets. CPGs are groups of logical disks. Virtual volumes use storage space provided by CPGs. The VVs are exported to hosts and are the only data layer visible to hosts.



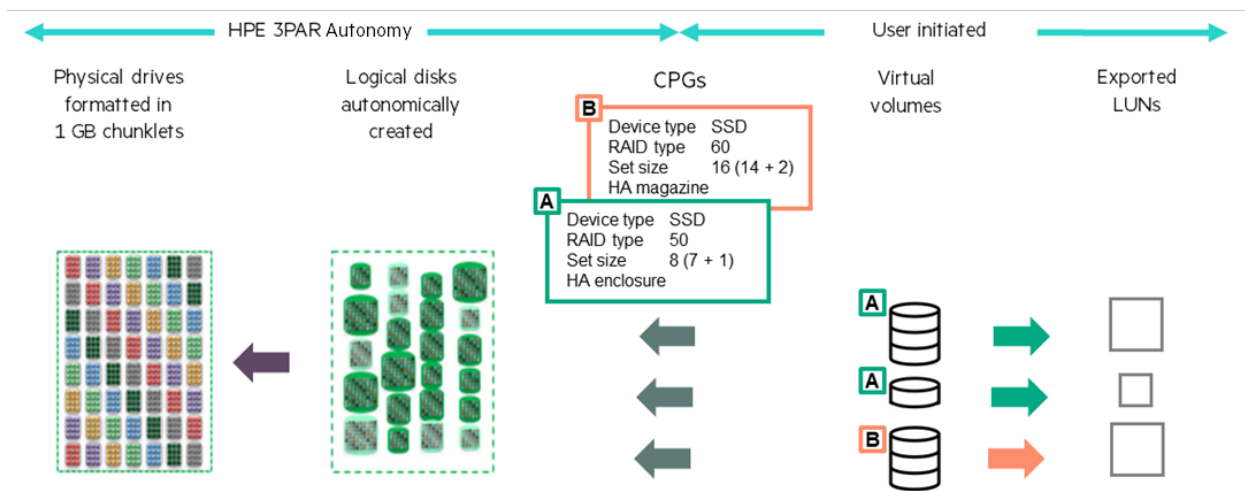


Figure 1. HPE 3PAR StoreServ system data layers

Terminology used when discussing HPE 3PAR block storage concepts include:

- **Cage:** “Cage” is a legacy HPE 3PAR term and is interchangeable with “drive enclosure,” “enclosure,” and “drive shelf.”
- **Chunklet:** Physical disks are divided into chunklets. Each chunklet occupies physically contiguous space on a Fast Class, nearline (NL), or solid-state drive (SSD) disk. On all current HPE 3PAR arrays, all chunklets are 1 GB. Chunklets are automatically created by the HPE 3PAR OS, and they are used to create logical disks. A chunklet is assigned to only one logical disk.
- **Physical disk:** A physical disk is a hard drive (spinning media or SSD) located in an HPE 3PAR drive enclosure.
- **Logical disk:** An LD is a collection of chunklets arranged as rows of RAID sets. Each RAID set is made up of chunklets from different physical disks. LDs are pooled together in CPGs, which allocate space to VVs.

The underlying logical disks are automatically created by the HPE 3PAR OS when you create VVs. The RAID type, space allocation, growth increments, and other logical disk parameters are specified when you create a CPG, or they can be modified later. HPE 3PAR arrays support the following RAID types:

- RAID 1 is also known as RAID 1+0.
- RAID Multi-Parity (MP) is also known as RAID 6 or double parity.

- **Virtual copy:** “Virtual copy” is a legacy HPE 3PAR term and is interchangeable with “snapshot.”
- **Common provisioning group:** A CPG is a template for the creation of logical disks that allocate space to VVs on demand. A CPG allows up to 65,536 virtual volumes to share resources assigned to a CPG. You can create fully provisioned virtual volumes (FPVVs), and compressed VVs that draw space from the logical disks in a CPG. An individual VV may be up to 64 TiB in size; however, deduplicated VVs have a limit of 16 TiB. It is important to note that if no volumes of any type have been created in a CPG, it consumes no space.
- **Virtual volume:** VVs draw their resources from the LDs in CPGs and are exported as logical unit numbers (LUNs) to hosts. VVs are the only data layer visible to the hosts. You can create clones (previously known as “full copies”) or snapshots (previously known as “virtual copies”) of VVs. Clones remain available if the original base volume becomes unavailable. You can create VVs by using the CPGs created at installation time or the CPGs defined by the user.
- **Exporting virtual volumes:** For a host to see a VV, the volume must be exported as a LUN. Volumes are exported by creating VV-LUN pairings (VLUNs) on the system. When you create VLUNs, the system produces both VLUN templates that establish export rules, and active VLUNs that the host sees as LUNs and attached disk devices. A VLUN is created for each path available to the host for each VV exported.



- **Fully provisioned virtual volume (FPVV):** An FPVV is a volume that uses logical disks that belong to a CPG. Unlike thinly provisioned virtual volume (TPVVs), FPVVs have a set amount of user space that is allocated for user data. The fully provisioned volume size is allocated and consumed at the time of provisioning. Size limits range from 256 MB to 64 TB (compressed and deduplicated VVs have a maximum size of 16 TiB). The volume size can be increased at any time (if free space is available) up to the maximum 64 TiB without any downtime. However, the VV size cannot be decreased below the initial allocation.

Note

In previous versions of the HPE 3PAR OS, there was a provisioning type termed “copy-provisioned virtual volumes (CPVV),” which simply meant that the provisioned VV had associated snapshot space capacity available. As of HPE 3PAR OS 3.2.2, all volumes created are associated with snapshot space in the same CPG. By using the additional menu options during VV creation in the SSMC or the `-snap_cpg` option of the `createvv CLI` command, you can choose a different CPG for snapshot space. Because snapshots are reservationless, if no changes are generated, no space is consumed.

- **Thinly provisioned virtual volume (TPVV):** A TPVV is a volume that uses logical disks that belong to a CPG. TPVVs and deduped and compressed (DECO) VVs associated with the same CPG draw space from the logical disks belonging to those CPGs as needed, allocating space on demand in 16 KiB increments for each VV. As the volumes that draw space from the LDs (including compressed and deduped) require additional storage, the HPE 3PAR OS automatically creates additional LDs or expands the size of existing LDs and adds them to the CPG until the CPG reaches a user-defined growth limit, if one has been set, or the system runs out of space.
- **Clone:** A clone (also known as a “physical copy”) duplicates all the data from a base volume to a destination volume. The base volume is the original volume that is copied to the destination volume. The clone on the destination volume remains available if the original base volume becomes unavailable.

A clone requires the destination volume have usable capacity equal to or greater than the usable capacity of the base volume being cloned. As of HPE 3PAR OS 3.2.1, the clone can be exported immediately after creation while the data copy continues in the background.

- **Snapshot:** Unlike a clone, which is a block-for-block duplicate of an entire volume, a snapshot (also known as a “virtual copy”) preserves a bitmap of a VV at a particular point in time. Updates to VVs are written to snap data (SD) space and the bitmap (snap admin space) of the VV.

Snapshots for FPVVs, TPVVs, and other snapshots are created using copy-on-write techniques available (Virtual Copy). Hundreds of snapshots of each VV can be created if there is enough storage space available. It is worth noting that snapshots do not consume any space unless data on the base volume has been updated and the original data copied to the SD space.

Changed data is copied only once regardless of the number of snapshots taken. Snapshots are particularly useful for test/dev environments because they can be created in seconds and exported without affecting production data. Also, testers and developers can be granted the ability to create snapshots without having any other administrative privileges, lowering administration requirements. (Consult the [HPE 3PAR CLI Reference](#) guide for correct procedure for granting the `updatevv` privilege).

Snapshots can now be updated without the requirement to unexport and re-export the VV. For more information on snapshot/Virtual Copy technologies, refer to the [HPE 3PAR Software Products QuickSpecs](#).



What's new in HPE 3PAR OS 3.3.1 MU3

HPE 3PAR OS 3.3.1 MU3 introduces new features and functions that make it easier to use and result in better performance.

Licensing changes

Every HPE 3PAR system now ships with every feature you need to manage a single array.

- Data-at-rest encryption is licensed as a separate option. This option is available retroactive to HPE 3PAR OS 3.2.2.
- For multiple systems, add a single Multi-system Software option for Remote Copy, Federation, Peer Persistence, and CLX.
- The all-inclusive license is supported on HPE 3PAR 8000, 9000, 10000, and 20000 platforms running at least HPE 3PAR OS 3.2.2.

Table 2. Licensing in HPE 3PAR OS 3.3.1 MU3

Software	Licensing for 8000, 9000, 20000, and future 3PAR models	Delivered as:
All-inclusive single system software	OS Suite, Virtual Copy, Dynamic Optimization, Adaptive Optimization, Priority Optimization, Virtual Domains, Virtual Lock, Online Import, Recovery Manager Central App Suite, File Persona, Smart SAN	Software bundled into base controller
All-inclusive multi-system software Software license for titles requiring more than a single array	Peer Motion, Remote Copy, Peer Persistence, and CLX Frame LTU	One frame license per array participating in federation or replication group
Data at Rest Encryption	Frame LTU	One frame license per array

Performance enhancements

The following performance enhancements were introduced with HPE 3PAR OS 3.3.1 MU3:

- Support of Memory Driven Flash using Storage Class Memory (SCM) devices (with T05)
- Express writes for iSCSI and 16 Gb/s Fibre Channel
- Read and write performance optimizations
- Faster VM clones
- Smoother I/O transition during code upgrades
- Adaptive Flash Cache enhancements

Data reduction

HPE 3PAR OS 3.3.1 MU3 introduced:

- Improvements in Adaptive Data Reduction
- Inline compression
- Enhanced deduplication engine
- Data packing

Ease of use

HPE 3PAR OS 3.3.1 MU3 is easier to use because of the following enhancements:

- SSMC 3.4 and later appliances supporting VMware® hypervisors and Microsoft® Hyper-V
 - HPE InfoSight integration
 - Configuration migration
 - Replication of File Persona
 - Performance analytics
- Improved capacity and data reduction reporting
- Smart SAN 2.0
- Bigger virtual volume sizes
- VVols support for iSCSI and replication



Increased volume size

In HPE 3PAR 3.2.2 and earlier, the maximum VV size was 16 TiB on all hardware platforms. HPE 3PAR OS 3.3.1 MU3 introduces support for greater maximum values for VV sizes.

Table 3. Maximum VV sizes

HPE 3PAR array	TPVV/FPVV	TPVV with compression
7000 Series	64 TiB	16 TiB
8000/9000 Series	64 TiB	16 TiB
10000 Series	64 TiB	16 TiB
20000 Series	64 TiB	16 TiB

Refer to the latest feature availability matrix for the most current information from [SPOCK](#).

Getting started with HPE 3PAR Storage best practices

To begin, prepare the environment:

- Get the environment ready
 - Create an [HPE Passport Account](#) to gain access to product documentation, compatibility (SPOCK) as well as drivers and host software
 - [Install the SSMC appliance on either VMware ESXi™ or Microsoft Hyper-V](#)
 - Configure the SSMC for [InfoSight](#) and Advanced Analytics
- Get the array ready
 - Follow the Physical Planning Installation Guide available on [SPOCK](#)
 - Set up the system ports
 - Set up the HPE 3PAR and host zones
 - Review naming conventions
 - Create hosts
 - Provision storage
 - Configure security
 - Set up automated reporting
 - Configure threshold reporting

In addition to following the recommendations in the [HPE 3PAR physical planning guides](#), it is important to keep the number and type of physical disks as well as the number of drive enclosures as evenly distributed as possible behind node pairs to facilitate maximum performance and load distribution.

When tuning CPGs/VVs as a result of a hardware upgrade or changing requirements, use the default tunesys concurrent task level of two in order to limit processing overhead in a production environment.



Set up HPE 3PAR system ports

Follow HPE best practices when setting up HPE 3PAR system ports.

Port locations and nomenclature

The HPE 3PAR CLI and SSMC display the controller node, Fibre Channel, iSCSI, one Gigabit, and 10 Gigabit Ethernet port locations in the following format: <Node>:<Slot>:<Port>. For example: 2:4:1.

- Node
 - Valid node numbers are 0–7 depending on the number of nodes installed in the system when viewing a system from the service side (rear) of a cabinet.
 - In HPE 3PAR 8000/9000 series arrays, nodes are numbered 0–3 from bottom left to the top right when facing the rear of the nodes.
 - In HPE 3PAR 20000 series arrays, nodes are numbered 0–7 from bottom left to top right when facing the rear of the nodes.
- Slot
 - The HPE 3PAR 8000/9000 series arrays have a single onboard slot in each node; numbering starts at 0.
 - The HPE 3PAR 20000 arrays' slots are numbered left to right/top to bottom, starting at 0 from left to right, bottom to top in a node in the lower enclosure. In the upper enclosure, slots are numbered 0–9 from left to right, top to bottom.
- Port
 - Valid node port numbers are 1–4 for all add-in adapters, counting from the bottom up.
 - HPE 3PAR 8000/9000 ports are horizontal and labeled beginning with 1 on the HBA or iSCSI adapter.
 - HPE 3PAR 20000 ports are numbered from bottom to top in a node in the lower enclosure. In the upper enclosure, ports are numbered from top to bottom.

Front-end port cabling

Keep in mind the following best practices when attaching front-end cables:

- Each HPE 3PAR controller node should be connected to two fabrics to protect against fabric failures.
- Ports of the same pair of nodes with the same ID should be connected to the same fabric. Example:
 - 0:2:3 and 1:2:3 on fabric 1
 - 0:2:4 and 1:2:4 on fabric 2
- Connect odd ports to fabric 1 and even ports to fabric 2 and so forth.

Table 4. Example with a four-node HPE 3PAR 8400 array with eight host ports

Fabric 1	Fabric 2
0:2:3, 1:2:3, 2:2:3, 3:2:3	0:2:4, 1:2:4, 2:2:4, 3:2:4

Table 5. Example with a four-node HPE 3PAR 20400 array with 32 host ports

Fabric 1	Fabric 2
0:2:1, 0:2:3, 0:5:1, 0:5:3, 1:2:1, 1:2:3, 1:5:1, 1:5:3, 2:2:1, 2:2:3, 2:5:1, 2:5:3, 3:2:1, 3:2:3, 3:5:1, 3:5:3	0:2:2, 0:2:4, 0:5:2, 0:5:4, 1:2:2, 1:2:4, 1:5:2, 1:5:4, 2:2:2, 2:2:4, 2:5:2, 2:5:4, 3:2:2, 3:2:4, 3:5:2, 3:5:4

Fibre Channel host zoning

Keep the following considerations in mind when configuring Fibre Channel host zoning:

- Zoning options (in order of preference)
 - TDPZ – HPE Smart SAN with B-series switches
 - Smart Zoning – MDS switches
 - Single initiator-single target (optionally single initiator-target port pair)



- Best practices
 - In a B-series fabric, always make sure Smart SAN is enabled, and configure zones either using the HPE 3PAR SSMC or CLI.
 - Smart SAN helps reduce SAN complexity with automation and target orchestration.
 - Target driven peer zoning (TDPZ) allows a single target port to be zoned with multiple initiators automatically while provisioning hosts to 3PAR ports.
 - Although it is not a requirement, HPE recommends that the initiator port be provisioned/zoned with both native and guest ports of the persistent port pair.
 - Zoning is completely automated and avoids preconfiguration of zones manually using switch management tools.
 - TDPZ provides all of the benefits of initiator-based zoning, also known as *single initiator-single target zoning* (the most commonly used type), with fewer zones and optimal use of switch resources.
- Smart Zoning
 - Use Cisco Smart Zoning in an MDS fabric.
 - Smart Zoning allows multiple initiators and multiple target ports to be included in a single zone by uniquely identifying initiator ports and target ports while configuring zones.
 - This is still manual operation and zoning must be preconfigured, unlike in TDPZ.
 - Smart Zoning also has the advantage of single initiator single target benefits with much smaller number of zones.
- Single initiator-single target zoning
 - When the preceding options are not viable for any reason, configure single initiator, single target zones using switch management tools.
 - In case of HPE 3PAR persistent ports, HPE recommends that you zone the initiator port with both native and guest ports of the persistent port pair. This is to reduce the complexity of total number of zones in single initiator, single target approach and to some extent is in line with the overall objective.
 - Optionally, if required, single initiator and a pair of target ports can also be included in the same zone.
 - Always use port WWNs while configuring these zones.

The following diagram is a representation of a typical configuration where a host is connected to an HPE 3PAR port pair on two different nodes with persistent ports enabled. Depending on the context, any of the three zoning options can be used to configure the required zones.

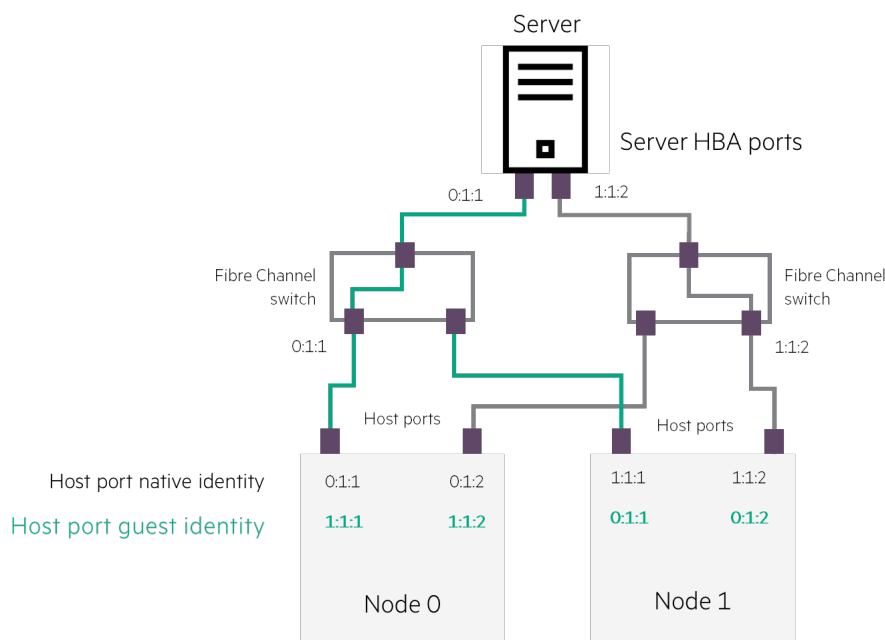


Figure 2. HPE 3PAR controller nodes



Use at least two separate SAN switches to enable availability in the event of a switch failure. Additional considerations include:

- Zoning should be performed using the Worldwide Port Name (WWPN), which is the WWN of each individual port on the HPE 3PAR array. Port Persistence is not compatible with destination address (DID) zoning.
- Hosts should be mirrored to node pairs. For example, a host should be zoned to nodes 0 and 1, or nodes 2 and 3. Hosts should not be zoned to nonmirrored nodes, such as 0 and 3.
- Nonhypervisor hosts
 - A single nonhypervisor host port should be zoned with a minimum of two ports from the two nodes of the same pair. In addition, the ports from a host’s zoning should be mirrored across nodes. In the case of hosts attaching with multiple host HBA ports attached to dual switches, each port should be zoned to at least the two-mirrored nodes.
 - Nonhypervisor hosts do not need to be connected to all nodes because of the way the volumes are spread on all the nodes.
- Hypervisor hosts
 - Each hypervisor HBA port should be zoned to a minimum of two ports from each node pair on two-node systems (At least one or more ports from at least four nodes on four-, six-, and eight-node systems support throughput across multiple-node buses.)

Table 6. Examples of valid zoning

Host type	Host HBA port	HPE 3PAR ports
Nonhypervisor, single HBA, single port	HBA1:port 1	0:1:1, 1:1:1
Nonhypervisor, single HBA, two HBA ports to separate switches	HBA1:port 1	0:1:1, 1:1:1
	HBA1:port 2	0:1:2, 1:1:2
Hypervisor, two HBAs, two HBA ports each, connected to two separate switches in a four-node StoreServ array	HBA1:port 1	0:1:1, 1:1:1
	HBA1:port 2	0:1:2, 1:1:2
	HBA 2:port 1	2:1:1, 3:1:1
	HBA 2:port 2	2:2:2, 3:2:2

- The maximum number of initiators supported on each HPE 3PAR system depends on the model and configuration. Regarding this maximum, one initiator = one path from one host.
 - A single HBA zoned with two Fibre Channel ports is counted as two initiators.
 - A host with two HBAs, each zoned with two ports, counts as four initiators.
- No more than 256 connections (128 if using Persistent Ports) are supported per front-end/host port.

Persistent Ports

Persistent Ports functionality is supported for HPE 3PAR OS 3.1.2 and later only (with functionality restrictions). HPE 3PAR OS 3.1.3 introduced support for Fibre Channel over Ethernet (FCoE) connected hosts and iSCSI connected hosts and the ability to detect an array node suffering “loss_sync” (a physical layer problem occurring between the HPE 3PAR controller node and the switch it is connected to). HPE 3PAR OS versions earlier than 3.1 do not support Persistent Ports.

For HPE 3PAR Fibre Channel host ports, the following requirements must be met:

- The same host port on host-facing HBAs in the nodes in a node pair must be connected to the same Fibre Channel fabric and preferably different Fibre Channel switches on the fabric (for example, 0:1:1 and 1:1:1).
- The host-facing HBAs must be set to “target” mode.
- The host-facing HBAs must be configured for point-to-point connection. (There is no support for loop mode.)
- The Fibre Channel fabric being used must support NPIV and have NPIV enabled.

For HPE 3PAR FCoE host ports, the following requirements must be met:

- The same converged network adapter (CNA) port on host-facing HBAs in the nodes in a node pair must be connected to the same FCoE network and preferably different FCoE switches on the network (for example, 0:1:1 and 1:1:1).
- The FCoE network being used must support N-Port ID Virtualization (NPIV) and have NPIV enabled.



For HPE 3PAR iSCSI host ports, the same host port on host-facing converged network adapters (CNAs) in the nodes in a node pair must be connected to the same IP network and preferably different IP switches on the fabric (for example, 0:1:1 and 1:1:1).

Persistent Ports configuration considerations

Persistent Ports requires that corresponding native and guest host ports on a node pair be connected to the same Fibre Channel fabric or IP network. In addition, the switches they are connected to must support and be configured for NPIV in the case of Fibre Channel and FCoE. This means that for a minimum configuration to provide Persistent Ports functionality, where the node pair is connected to redundant Fibre Channel SAN fabrics, each node in a node pair must have at least two Fibre Channel host ports cabled with one port connected to each fabric.

Ensure that the same <slot>:<port> on each node in the node pair are connected to the same Fibre Channel fabric. This is a requirement to enable the Persistent Port functionality.

High availability

Best practices to consider when configuring storage for high availability include:

- Use Cage Level availability, if the physical configuration allows, when configuring CPGs. Cage level availability ensures that no two members of the same RAID set reside in the same physical drive enclosure.
- Edit the properties of the system to ensure that 3PAR Differential Integrity Field (DIF) or Standard SCSI DIF is enabled to error free Fibre Channel data transmission. The DIF ensures end-to-end, host-to-disk data integrity.
- HPE encourages all HPE 3PAR Storage customers to upgrade to the latest recommended HPE 3PAR OS. Upgrading to the most current HPE 3PAR OS allows the storage system to benefit from the ongoing design improvements and enhancements. For customers participating in the Get 6-Nines Guarantee Program, the program identifies the latest HPE 3PAR OS version that is covered under the Guarantee program.
- Size the system appropriately so that all workloads and applications dependent on the HPE 3PAR system can perform as needed under the conditions of a node being down. This might occur during an unplanned controller node failure or planned maintenance of a controller node. In no situation should the maximum limits of the system as defined in this document and product specifications be exceeded.
- In systems with four or more nodes, a resilience feature called *Persistent Cache* is automatically enabled. The Persistent Cache feature helps to ensure that no storage controller node is placed into performance-limiting “cache write thru” mode as a result of a losing its partner in the node pair. Any node that loses its adjacent node can dynamically form a mirrored cache relationship with another storage controller node. This limits the performance impact of unplanned downtime or controller node maintenance.

Naming conventions considerations for hosts and host sets

A concise naming convention is important to manage an HPE 3PAR system effectively. A naming convention should be consistent and descriptive. HPE 3PAR systems are case sensitive. Preferably, use CamelCase to ensure ease of readability and a descriptive name. When using the bundled Call Home feature, use of a clear naming convention allows support personnel to handle support calls more efficiently and use of descriptive comments enables further support case resolution.

Use hierarchical names that allow multiple objects of the same classification group to begin with the same characters.

Examples:

```
prd.UnixDatavg01.vv
```

```
prd.UnixAppvg02.vv
```

Use a naming convention and a suffix that allows all objects of the same type to be grouped together when sorting on the name field and that allows effective search of all objects when using patterns (question mark [?] or asterisk [*]).

Example: `showcpg *.cpg.*`



Naming conventions

Define the naming convention early in the implementation. Examples are as follows:

- **Hosts:** Host names support up to 31 characters. Host name are of the form <TYPE>.<OS>.<HOST>.<OBJECT TYPE>, where:
 - <TYPE> can be `prd` for production servers, `dev` for development servers, `tst` for test servers, and so forth.
 - <OS> can be `win` for Microsoft Windows® hosts, `vmw` for VMware hosts, `lin` for Linux® hosts, `sol` for Solaris hosts, `aix` for IBM AIX® hosts, `hpux` for HP-UX hosts, and so forth.
 - <HOST> is the host name.
 - <OBJECT TYPE> is one of `VV`, `cpg`, `VVset`, and so forth.

Examples:

```
- prdWinServer1.vv
- prdHpuxServer2.vv
- dev.lin.server3.vv
```

- **Host sets:** Host set names support up to 31 characters. Host set name will be of the form <TYPE>.<OS>.<CLUSTER NAME>, where:
 - <TYPE> can be `prd` for production servers, `dev` for development servers, `tst` for test servers, and so forth.
 - <OS> can be `win` for Windows hosts, `vmw` for VMware hosts, `lin` for Linux hosts, `sol` for Solaris hosts, `aix` for AIX hosts, `hpux` for HP-UX hosts, and so forth.
 - <CLUSTER NAME> is the name of the cluster.

Examples:

```
- prd.win.sqlcluster1.vvset
- prd.vmw.esxcluster2.vvset
- dev.lin.cluster3.vvset
```

Hosts and host sets

When creating hosts, follow the implementation guide for each platform. Selecting the correct host persona (specifying host operating system) for each host is important. [Implementation guides are available for download.](#)

Each physical server should have a different host defined, containing the WWN or iSCSI Qualified Name for this host.

Best practice for creating a new host

Follow these best practices on systems running Windows, Red Hat® Linux, Solaris, SUSE, HP-UX, VMware ESXi™, or HPE Helion OpenStack®.

1. Install the (optional) [HPE 3PAR Host Explorer for HPE 3PAR StoreServ Storage software](#) in the host if available for the host platform. The Host Explorer software communicates with the storage array over SCSI ID 254 and transmits host name, operating system version, and WWNs as well as other information to the array.
2. Zone in all the ports according to the zoning best practices.
3. From the host CLI, execute the `tpdhostagent-start` and then `tpdhostagent-push` commands.
4. This automatically creates the host on the HPE 3PAR system.

Best practice for creating a new host manually:

1. Zone in the host ports to the HPE 3PAR array using the zoning best practices, one host at a time, or use [HPE SmartSAN for 3PAR](#) for automated management control.
2. For each host, select the host and then create the new host.
3. In the WWN selection screen, select the WWNs associated with the new host.
4. Zone in a single host and then create the host on the HPE 3PAR array to reduce the possibility of assigning incorrect WWNs to a host. Repeat until all hosts are zoned in.



Add FC
?

Select existing unassigned FC WWNs or enter a new name if the host is not connected to the array.

Existing WWN
 New WWN

1 selected
Total 18

Name	System Ports (N:S:P)	System	Reported Host Name	Operating System
1000C4346B20F5 0A	0:0:1, 1:0:1	m2-tme8200-02	m2-dl360g9-40.mip.sto rage.hpecorp.net	VMware (ESXi)
1000C4346B20F5 0B	0:0:2, 1:0:2	m2-tme8200-02	m2-dl360g9-40.mip.sto rage.hpecorp.net	VMware (ESXi)
1000C4346B20F5 7A	0:0:1, 1:0:1	m2-tme8200-02	m2-dl360g9-41.mip.sto rage.hpecorp.net	VMware (ESXi)
1000C4346B20F5 7B	0:0:2, 1:0:2	m2-tme8200-02	m2-dl360g9-41.mip.sto rage.hpecorp.net	VMware (ESXi)

Add
Add +
Cancel

Figure 3. Creating a host in the SSMC

For clusters, create a host set containing all the hosts used by the cluster. This allows you to export shared VVs to all hosts of the cluster in a single operation.

Provisioning block storage from an HPE 3PAR system

Provisioning block storage from an HPE 3PAR array includes provisioning block storage for use by File Persona.

Common provisioning groups

CPGs are primarily used as templates for the creation of LDs but have other characteristics as well, such as capacity limits. CPGs define:

- RAID level for the LDs to be created
- Availability level (HA CAGE, HA PORT, or HA MAG)
- Step size in addition to other characteristics, such as drive geometry

CPGs are created only across drives of the same type (SSD, Fast Class, or NL) and include drives of different relative performance metrics (RPMs) unless otherwise specified by using a disk filter.

Keep the following best practices in mind when configuring CPGs:

- For all drive types, use RAID 6 for maximum availability.
- When creating CPGs, accept defaults according to performance/capacity requirements.
- The number of CPGs should be kept to a minimum.
- You might need to create more CPGs than the default in certain cases:
 - When using HPE 3PAR Virtual Domain software, because a given CPG can only be in one domain.
 - When capacity reporting is required per customer or application CPGs to ease capacity reporting.
 - When snapshots are used heavily and the snapshot data is kept in a tier different from the source data. In this case, use a CPG that matches the production CPG performance characteristics as close as possible to improve performance.
- Do not set growth limits on CPGs. If a warning threshold is required, set a growth warning (warning in terms of capacity), not an allocation warning (warning in percentage).
- HPE recommends that the largest set size used be smaller than the number of physical drives in a 2N configuration. For example, an eight-drive, 2N system should not have a CPG configured to use a set size of RAID 6 (6+2).



Solid-state drive CPGs

Consider the following best practices when creating SSD CPGs:

- SSD CPGs should be configured with RAID 6 with the default set size. This will yield a superior performance/capacity ratio and 6+2 on systems that have enough SSDs to support 6+2, for larger systems R6 14+2 gives the best capacity / performance mix.
- The growth increment should be set to the minimum value, which is 8 GB per node pair.
 - On two-node systems, set the value to 8 GB.
 - On four-node systems, set the value to 16 GB.
 - On six-node systems, set the value to 24 GB.
 - On eight-node systems, set the value to 32 GB.

To set the CPG growth increment to a lower value than the default, check the **Show advanced option** box.

- Availability should be left to “cage level” (the default option) if the system’s configuration allows for it. If not, it should be set to “magazine level” availability. You can change this by checking the **Advanced options** box in the SSMC.

Fast Class CPGs

Follow these best practices when creating Fast Class CPGs:

- Fast Class (SAS) CPGs should be RAID 6 type by default. This configuration yields the highest availability for modern high-capacity drives. You can change the set size (data to parity ratio) from the default value of 6+2 if the system configuration supports it. If usable capacity is the primary concern, use a larger set size.
- For applications that have a very high write ratio (more than 50% of the access rate), create a CPG using RAID 1 if performance (as opposed to usable capacity) is the primary concern.
- The growth increment should be left to the default value (32 GB per node pair).
- Availability should be left to **cage level** (the default option) if the system’s configuration allows for it. If not, it should be set to **magazine level** availability. You can change this by checking the **Advanced options** box in the SSMC.

NL CPGs

Follow these best practices when creating NL CPGs:

- NL CPGs should be RAID 6, which is the default.
- You can change the set size (data to parity ratio) from the default value of 8 (6+2) if the system configuration supports it.
- Do not use RAID 5 with NL disks.
- The growth increment should be left to the default value (32 GB per node pair).
- Availability should be left to **cage level** (the default option) if the system’s configuration allows for it. If not, it should be set to **magazine level** availability. You can change this by checking the **Additional settings** box in the SSMC during VV creation.

Adaptive Flash Cache

HPE 3PAR Adaptive Flash Cache is included as part of the HPE 3PAR OS Suite 3.2.1 and later. It is supported in all HPE 3PAR arrays that have a blend of SSDs and hard disk drives (HDDs). Adaptive Flash Cache uses CPGs to tier data according to access requirements, for example, heavily accessed, routinely accessed, and rarely accessed; that is, hot, medium, and cold.

Benefits

Benefits of Adaptive Flash Cache include:

- Reduces latency for random read-intensive workloads
- Responds quickly, providing smart and flexible second stage data caching based on application and workload demands
- Enables HPE 3PAR Adaptive Flash Cache across the entire system or for particular workloads to accelerate



Requirements

Adaptive Flash Cache requires the following components:

- HPE 3PAR OS 3.2.2 or later
- Two SSDs in HPE 3PAR 8000/9000 series or four SSDs in the HPE 3PAR 20000 series per node pair

Refer to the [HPE Adaptive Flash Cache](#) guide for information regarding the limits for Adaptive Flash Cache limits.

- If SSD drives are not available, use the CLI command `createflashcache sim`. This command simulates how much data could possibly move into the flash cache tier and the possible increase in performance.

HPE 3PAR Memory Driven Flash is a persistent memory technology that takes advantage of DRAM and flash technologies to accelerate I/O on the memory bus. Keep the following best practices in mind when configuring SCM devices:

- With HPE 3PAR OS 3.3.1 MU3, SCM devices are used only for Memory Driven Flash, making setup of SCM easy.
- There are no sizing guidelines for the creation of the cache; the HPE 3PAR array uses the full capacity of the SCM devices.
- SCM devices should be the same size for each node to make the most efficient use of space. System checks will inform of a size difference.
- SCM is used only for read cache; therefore, the workload on the VV will affect performance.

- On systems with four or more SSD drives per node pair available, enable flash cache system-wide by issuing the following commands:

```
createflashcache
```

If there is insufficient space available for the maximum amount, the command returns the maximum amount available.

```
setflashcache enable sys:all
```

This command provides an alternative to enabling flash cache for the entire system. You can also select VVsets.

For more information, refer to the [HPE 3PAR Adaptive Flash Cache](#) white paper.

Provisioning VVs

Follow these best practices for provisioning VVs:

- Because the granularity of deduplication is 16 KiB, the efficiency is greatest when I/Os are aligned to this granularity. For hosts that use file systems with tunable allocation units, consider setting the allocation unit to a multiple of 16 KiB.
- Deduplication is performed on the data contained within the VVs of a CPG. For maximum deduplication, store data with duplicate affinity on VVs within the same CPG.
- Deduplication is ideal for data that has a high level of repetition. Data that has been previously deduplicated, compressed, or encrypted is not a good candidate for deduplication and should be stored on thinly provisioned volumes.
- When using an HPE 3PAR array as external storage to a third-party array, deduplication might not function optimally.
- Use TPVVs with dedup enabled when there is a high level of redundant data and the primary goal is capacity efficiency.
- Use TPVVs with dedup enabled with the appropriate dataset.

Suitable workloads that work well with TPVVs, DECO VVs, and other types of VVs include:

- Virtual Server Infrastructure
- Virtual Desktop Infrastructure (VDI)
- File shares and web servers
- SAP HANA® Tailored Datacenter Integration (TDI)
- Microsoft SharePoint with AvePoint® DocAve®

Unsuitable workloads include:

- Databases
- Compressed video/images/files
- Host or application-encrypted/compressed data



Follow these best practices with file systems:

- Format file systems on hosts and VMs with 16 Kb (or multiples of 16 Kb) allocations.
- Deploy VM operating system disks on dedicated datastores on a deduped VV.
- Clone VMs from a small number of templates.
- If dedup savings at the CPG level are 1.1:1 or less, consider converting the VVs to TPVVs.

For more information, refer to the [HPE 3PAR Thin Technologies](#) white paper.

Priority Optimization

As of HPE 3PAR OS 3.1.2 MU2, [Priority Optimization](#) Quality of Service (QoS) enables you to set a maximum limit of IOPS or bandwidth available to a VVset.

With HPE 3PAR OS 3.1.3 and later, this feature set was extended to add priority levels (high, normal, and low), latency goals, and minimum goals along with the ability to set QoS rules against virtual domains.

Take care when setting QoS rules. System performance must be understood, and historical performance data must be taken into account when setting maximum limits, minimum goals, and latency goals. A good understanding of the applications that reside on the VVs is also equally important.

Table 7. Enhancements to QoS over previous versions

Control type	Minimum HPE 3PAR OS version	Description	Dependencies and best practices
Max limit	3.1.2 MU2 and later	Maximum threshold for IOPS and bandwidth for a QoS object.	Maximum limit has no dependencies on the other control types.
Min goal	3.1.3+	Minimum floor for IOPS and bandwidth below which HPE 3PAR Priority Optimization will not throttle a QoS object.	When a minimum goal is set on an object, the user must also configure a maximum limit on the same object within the same rule. Minimum goal is ignored if the system has no rules with latency goal set.
Latency goal	3.1.3+	The Service Time target the system will try to achieve for a given workload in a QoS object.	This control type requires other rules in the system with minimum goal to be set because the latency goal algorithm needs direction on which workload to target and throttle. The order in which these will be throttled is provided by the priority levels.
Priority level	3.1.3+	Precedence order for QoS subsystem to throttle workloads to meet performance levels.	High priority should be used against critical applications, lower priority on less critical applications.

Best practices for Priority Optimization:

- When implementing maximum limits (the maximum amount of IOPS/bandwidth a given VVset or domain is allowed to achieve), the best practice is to use System Reporter data in order to quantify volume performance and set maximum limits rules accordingly.
- When implementing the minimum goal (the minimum amount of IOPS/bandwidth below which the system will not throttle a given VVset or domain in order to meet the latency goal of a higher priority workload), you should set the minimum goal by looking at the historical performance data and understanding the minimum amount of performance that should be granted to the applications that reside in that VVset. The volumes in the VVset might use more IOPS/bandwidth than what is set by the minimum goal, but they will be throttled to the given limit as the system gets busier. The performance might also go below the minimum goal. This can happen if the application is not pushing enough IOPS or if the sum of all minimum goals defined is more than the I/O capability of the system or a given tier of storage.
- Latency goal (the service time [svctime] the system attempts to fulfill for a given QoS rule) requires rules with a minimum goal specification to exist so the system can throttle those workloads. A reasonable latency goal should be set. You can do this by looking at historical performance data. The latency goal is also influenced by the tier on which the volume resides.



Table 8. Latency goal guidelines

Tier	Latency goal guidelines
All NL	>= 20 ms
All Fast Class	>= 10 ms
All SSD	>= 1 ms
Subtiered volumes (Adaptive Optimization)	Guideline of the middle tier in a three-tier Adaptive Optimization configuration or of the lowest tier in a two-tier Adaptive Optimization configuration

When a latency goal is set on a QoS object in a given tier, you should also create QoS rules with a minimum goal and lower priority on other objects that reside in the same tier. This will allow QoS to throttle VVsets that share the same tier towards their minimum goal if there is contention of resource on a that tier of storage in order to meet the latency goal of the higher priority workloads.

- Use VVsets for setting QoS rules as well as exporting VVs.

Adaptive Optimization and QoS interoperate, but the Adaptive Optimization data migration process might impact latency goal settings. For example, when a workload is partially migrated from Fast Class drives to NL drives after the collection and analysis of the I/O access pattern by Adaptive Optimization, depending on the host access I/O pattern, this might have a negative impact on latency and trigger workloads to be throttled towards their minimum goal. This is an expected behavior if the QoS rule was modeled after drive performance.

- Three priority levels exist:
 - High
 - Normal
 - Low

As the system gets busier, it will start targeting lower-priority workloads and throttling their performance to meet the latency goals of higher-priority workloads. High priority level should be used against critical applications, lower priority on less critical applications.

Remote Copy

Follow these best practices when configuring Remote Copy. Review the [Remote Copy Users Guide](#) for more detailed information.

- Follow the [HPE 3PAR Remote Copy Software User Guide](#) strictly when configuring Remote Copy. Pay particular attention to latency tolerance.
- Remote Copy Fibre Channel ports should only be in one zone with one Remote Copy Fibre Channel port from another HPE 3PAR system.
- Do not zone hosts or other devices with Remote Copy Fibre Channel other than the destination Remote Copy Fibre Channel port.
- HPE 3PAR OS 3.1.3 substantially improved topology flexibility.
- All VVs belonging to the same application should be added to the same Remote Copy group. A Remote Copy group can contain up to 300 virtual volumes. VVs that have logical links include:
 - VVs used by the same application (data and log volumes, for example)
 - VVs used by a Logical Volume Manager (LVM) volume group
 - VVs that contain virtual disks of the same VM
- Do not add VVs that have no logical link (host or application) to the same Remote Copy group. This will give the best granularity of Remote Copy failover by allowing a failover of only one host or application.
- Source and destination VVs must have the same size.
- Although source and destination VVs do not need to be of the same type of disk or RAID level, consider performance before mixing different types of disks/RAID levels on the source and destination systems.



- When Remote Copy is used in synchronous mode, it is not advised to replicate Fast Class VVs to NL virtual volumes because the reduced performance of the NL disks might impact the primary VV performance as the source system waits for the destination system to acknowledge the write.
- When Remote Copy is used in combination with virtual domains, the source and destination virtual volumes must be in domains of the same name on the source and destination systems. However, they do not need to be of the same type (type of provisioning, RAID, or disk).
- In case of complete communication failure between two HPE 3PAR systems, wait for a period of low activity before restarting the Remote Copy groups.
- If the hosts on the secondary sites are powered on, do not export the secondary (destination) virtual volumes unless using a geographically distributed cluster with automatic failover, such as CLX, MetroCluster, or Geocluster, where this is specifically required.

Asynchronous streaming Remote Copy

Asynchronous streaming Remote Copy (also known as *replication*) is perfect for environments where very small recovery point objectives (RPOs) are required or for environments where synchronous replication is needed but the link latencies are so large it will result in unacceptable write I/O response times for the data to be replicated synchronously.

Asynchronous streaming Remote Copy solutions generally require replication link speed to be sized within 95% to 99% of the maximum data generation rate.

Refer to [HPE SPOCK](#) for the most current information.

Note

Asynchronous streaming Remote Copy is supported on Fibre Channel and FCIP transports only.

Adaptive Optimization

Best practices to follow when configuring Adaptive Optimization include:

- The following combinations are acceptable within the same [Adaptive Optimization](#) configuration (policy):
 - SSD, Fast Class/SAS, and NL
 - Fast Class/SSD and Fast Class
 - Fast Class/SAS and NL
- Using different RAID levels within the same policy is acceptable.
- When configuring two- or three-tier solutions containing SSDs, if region density data is not available for sizing the SSD tier, assume that the SSD tier will only provide the greatest number of IOPS per drive. The increase in the estimated number of IOPS on larger drives is not because of a difference in technology, but rather the increased probability that “hot” data regions are on the larger SSDs and not on the smaller SSDs.
- Always size the solution assuming the NL tier will contribute 0% of the IOPS required from the solution.
- Configurations that only contain SSD and NL are not recommended unless this is for a well-known application with a very small ratio of active capacity compared to the total usable capacity (1%–2%).

For more information on Adaptive Optimization, refer to the [Adaptive Optimization for HPE 3PAR StoreServ Storage](#) white paper.

- When using thin provisioning volumes along with Adaptive Optimization, select a CPG using Fast Class disks for the user CPG of the thin provisioning volumes. This means that when new data is written, it will be on a mid-performance tier by default, which can then be distributed appropriately.
- Ensure that the default tier (Fast Class) has enough capacity and performance to accommodate the requirement of new applications until data is migrated to other tiers. When new data is created (new virtual volumes or new user space for a thin volume), it will be created in the Fast Class tier, and Adaptive Optimization will not migrate regions of data to other tiers until the next time the Adaptive Optimization configuration is executed. It is therefore important that the Fast Class disks have enough performance and capacity to accommodate the performance or capacity requirements of new applications (or applications that are in the process of being migrated to HPE 3PAR arrays) until the moment when the regions of data will be migrated to the other tiers.



- If SSDs are used in Adaptive Optimization configurations, no thin provisioning volumes should be directly associated with SSD CPGs. The thin provisioning volumes should only be associated with Fast Class CPGs. This ensures that SSD capacity is consumed by Adaptive Optimization and allows this capacity to be safely used to 95% or even 100%. To help ensure that no TPVV is associated with an SSD or NL CPG, run the `showcpg` command and confirm that only the Fast Class CPG reports a TPVV value greater than 0.

In the following example, only the Fast Class CPG has TPVVs associated with it.

Id Name	Domain	Warn%	VVs	TPVVs	Usr	Snp	Total	Used	Total	Used	Total
4 AO_01_SSD	-	-	0	0	0	0	1201792	1201792	11648	0	16384
8 AO_01_FC	-	-	469	469	469	0	11665920	11665920	8560128	0	120832
12 AO_01_NL	-	-	0	0	0	0	29161088	29161088	729472	0	325632

- All CPGs used in an Adaptive Optimization configuration should have the same level of availability. Using a CPG with magazine level availability in an Adaptive Optimization configuration with CPGs with cage level availability will mean that all virtual volumes will have an effective availability equivalent to magazine level.

Refer to the [“Common provisioning groups”](#) section of this white paper for details of CPG best practices.

- Schedule the various Adaptive Optimization configurations to run at the same time, preferably at night. Adaptive Optimization will execute each policy in a serial manner but will calculate what needs to be moved at the same time.
- It is preferable not to set any capacity limit on the Adaptive Optimization configuration level or on the CPG (no allocation warning or limit). This allows the Adaptive Optimization software to make excellent use of the different tiers available in the system. In HPE 3PAR OS 3.1.2 and later, use the SSMC to set the CPG capacities.
- If a capacity limit is required for a given tier, set a capacity-warning threshold (not limit) on the CPG itself through the SSMC or the CLI. Adaptive Optimization will not attempt to move more data to a CPG than the capacity warning set on the CPG.
- Always ensure that at least one of the CPGs used by the Adaptive Optimization configuration does not have any growth warning or limit. Ideally, this should be the CPG to which the virtual volumes are linked—the CPG in which new user space will be created if needed.
- Use a simple Adaptive Optimization configuration model as often as possible. For most applications, use generic Adaptive Optimization configurations that:

- Use all the tiers available in the system
- Run during the days of the workweek only (for example, Monday–Friday)
- Execute once a day, preferably at night
- Use a measurement/hours of 24
- Use a mode of “balanced”

For well-known applications that require a high level of performance, use tailored Adaptive Optimization configurations that:

- Preferably use all the tiers available in the system
- Execute immediately at the end of the high-activity period
- Use a measurement/hours that only covers the length of the high-activity period
- Use a mode of “performance.”

For test environments where performance is not a requirement, use an Adaptive Optimization configuration that:

- Uses only Fast Class and NL tiers.
- Run during the days of the work week only (for example, Monday through Friday).
- Executes once a day, preferably at night.
- Uses a measurement/hours of 24.

- Do not mix Adaptive Optimization with any other application or process that moves data on a LUN or between LUNs. Any application or process that moves data on a LUN or between LUNs, such as VMware Storage DRS, should be disabled because they might conflict with each other. Only Adaptive Optimization, being the lowest level (storage level), should be used.



Security

Follow these security best practices:

- Do not change the password or remove the accounts for the 3parsvc, 3parservice, 3paredit, or 3parbrowse users. These are randomly generated at the time of system initialization and are required for communication between the service processor and the HPE 3PAR system.
- Create a different user for each system administrator that will use the system. Alternatively, configure the system to use Active Directory and make sure all users use their own accounts to log in to the system.
- When scripting, use the lowest privilege level required. If a script requires only read access to the system, use a browse account. If a script does not need to remove objects, use a create account.

HPE System Reporter and HPE InfoSight in the SSMC

Two object selections are available. By default, **All** is selected, which enables you to chart all objects for the category selected. For example, if you want to chart real-time exported volume performance, the choices for object selection are All or Filter by objects. Using the default of All then charts performance of all volumes created on the array.

Selecting to filter by objects enables you to choose which object is charted. You can choose to monitor an object by the user of a virtual volume, host, or port number. You may highlight several objects and select the **Add** radio button at the bottom of the screen. If you need to select multiples items that are not all displayed, you can highlight one object, click the **Add +** radio button, scroll down the list of menu items, and add more objects by either clicking **Add +** for multiple objects or **Add** if only one additional object is added.

You may also search for objects to chart by using the search bar identified by the magnifying glass icon. Using Filter by objects, you must also select the plotting style. The two choices are:

- Individual series for each selected object
- Aggregated series for all selected objects

Individual series for each selected object is applicable only when **Filter by objects** is selected. This option is not valid when selecting for all objects.

- Use the search bar to identify objects to enter rather than scrolling through menu items.
- Combining rules does not allow you to differentiate the outcome of the object.

Follow these best practices when using System Reporter and InfoSight:

- For all-flash arrays (AFAs), enable Advanced Analytics from the Settings menu of the main SSMC drop-down menu. This enables the population of the dashboard panels including performance outliers.
- Use the SSMC System Reporter function to schedule daily reports grouped by application sets. If application sets are not available, group VVs/CPGs together to form application performance profiles.

For more information, download the [SSMC documentation](#).

Threshold alerts

Threshold alerting allows real-time regex-based alerting for performance issues. Best practices include:

- Configure an SMTP relay from the main settings menu available from the SSMC drop-down menu.
- Use **Exported Volumes Read IOPs and Service Time** to create an alert to detect potential performance issues. Specific conditions will vary based on PD type. Examples are listed in Table 9.

Table 9. Suggested threshold alerts

Tier	Latency goal guidelines	Suggested threshold alert
All NL	>= 20 ms	>= 30 ms
All Fast Class	>= 10 ms	>= 20 ms
All SSD	>= 1 ms	>= 5 ms



- Create a second threshold alert for Host port performance time to detect potential congestion on a host port. Host port congestion might not be reflected in VV service time because VVs may be serviced quickly but face congestion at the port, resulting in delayed responses to hosts.

Storage analytics on the web

HPE [InfoSight](#) is built on the HPE 3PAR Call Home infrastructure.

Best practices for using InfoSight:

- Configure HPE 3PAR arrays to send configuration information to HPE Support by using the 3PAR Service Processor. This configuration enables HPE to offer meaningful analytics about storage devices, advisories to improve storage infrastructures, and proactive support.
- Register HPE 3PAR arrays in InfoSight by logging in using an HPE Passport ID and then following the steps listed in the portal. You can view devices in InfoSight within a few hours of registration.
- Use System Access Groups in SF Remote to grant access rights to other users.
- Use the column filter options and the status filters to create a combination of filters and save it as a custom filter for easy retrieval and reuse.
- Use the notes feature to tag devices with comments and share them with others as needed.

Autonomic rebalance

Starting with HPE 3PAR OS 3.1.3, the rebalancing of data after hardware upgrades is an integrated feature that does not require an HPE Dynamic Optimization license.

After purchasing and installing new hardware, you can start rebalancing by selecting the HPE 3PAR array in the SSMC and selecting **Tune** from the system menu in the main drop-down menu. The process is invoked automatically after the `adm.t.hw` command is issued on HPE 3PAR platforms.

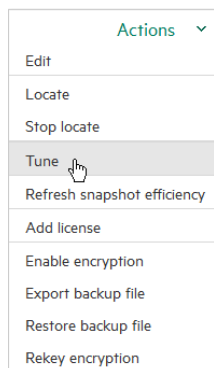


Figure 4. Autonomic rebalancing (tunesys) selection

You can also start autonomic rebalancing (tune system) from the CLI by entering `tunesys`.

Depending on the amount of hardware added to a system and the degree to which the system might have been unbalanced before, the re-leveling process can take several hours to several days or longer to complete. It is also possible that if the system was near capacity before the hardware upgrade, it might take multiple runs of `tunesys` to level the system. Other variables also affect how long `tunesys` takes to complete, including the speed of the drives affected by the re-leveling, how close to capacity the previously installed drives are at upgrade, how busy the system is, and so forth. The autonomic rebalancing process gives priority to servicing host I/O.

Best practices for autonomic rebalancing include:

- Execute an autonomic rebalance (tune) on a monthly basis
- Schedule an autonomic rebalance to run during off peak hours



Appendix A: Supported host personas

Persona_Id	Persona_Name	Persona_Caps
1	Generic	UARepLun, SESLun
2	Generic-ALUA	UAPreLun, SESLun, RTPG, ALUA
7	HPE-UX-legacy	VolSetAddr, Lun0SCC
8	AIX-legacy	NACA
9	Egenera	Softinq
10	ONTAP-Legacy	Softinq
11	VMware	SubLun, ALUA
12	OpenVMS	UARepLun, RTPG, SESLun, Lun0SCC
13	HPUX	UARepLun, VolSetAddr, SESLun, ALUA, Lun0SCC
15	Windows Server	UAPeLun, SESLun, ALUA, WSC



Summary

HPE 3PAR Storage is the only storage architecture needed whether your organization is a small- or medium-sized business or a large global enterprise. With a range of models to meet the needs of small to large data centers running key business applications up through enterprise-wide deployments of mission-critical applications and beyond, HPE 3PAR Storage meets a range of business and technical requirements. It offers storage with the effortless performance and flexibility you need to accelerate new application deployment and support server virtualization, the cloud, IT as a service (ITaaS), or whatever else the future might hold. This advanced and versatile storage platform helps you overcome unpredictability—effortlessly, without exception, and without compromise.

Resources

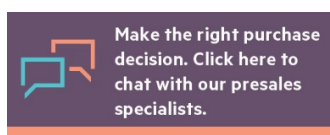
HPE Support Center
support.hpe.com

HPE 3PAR Host Explorer for HPE 3PAR StoreServ Storage
h20392.www2.hpe.com/portal/swdepot/displayProductInfo.do?productNumber=HP3PARHE

HPE 3PAR Remote Copy Software User Guide
support.hpe.com/hpsc/doc/public/display?docId=emr_na-c03618143

HPE Single Point of Connectivity Knowledge (SPOCK)
hpe.com/storage/spock

Learn more at
hpe.com/3par



 **Share now**

 **Get updates**

© Copyright 2018-2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for 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, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. SAP HANA is the trademark or registered trademark of SAP SE in Germany and in several other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. SD is a trademark or registered trademark of SD-3C in the United States and other countries or both. VMware ESXi is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. IBM and AIX are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. AvePoint and DocAve are registered trademarks of AvePoint, Inc. with the United States Patent and Trademark Office. All other third-party marks are property of their respective owners.

