Technical white paper

HP 3PAR StoreServ Storage best practices guide

A reference and best practices guide for HP 3PAR StoreServ Storage

Table of contents

Typographical conventions

This guide uses the following typographical conventions:

Table 1. Typographical conventions

Typeface	Meaning	Example
ABCDabcd	Used for dialog elements such as titles, button labels, and other screen elements	When prompted, click finish to complete the installation
ABCDabcd	Used for user input, filenames, commands, paths, and screen output	`Start the \OS\Windows\setup.exe`
<ABCDabcd>	Used for variables in user input, filenames, paths, and screen output	To add a user to a domain issue the set user—add domain <domainname>:role <username> command
Best practice	Used to highlight best practices for a particular topic section	**Best practice:** use RAID 5

Advisories

To avoid injury to people or damage to data and equipment, be sure to observe the cautions and warnings in this guide.

Always be careful when handling any electrical equipment.

WARNING!
Warnings alert you to actions that can cause injury to people or irreversible damage to data or the OS.

CAUTION!
Cautions alert you to actions that can cause damage to equipment, software, or data.

Note
Notes are reminders, tips, or suggestions that supplement the procedures included in this guide.

Note
The HP InServ Storage Server has been rebranded as the HP 3PAR StoreServ Storage system. There are instances in this document where menu items and command output refer to the HP 3PAR StoreServ Storage system as InServ or InServ Storage Server.

Introduction

Audience

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

User interfaces

Two user interfaces are available for the administration of HP 3PAR StoreServ: the HP 3PAR OS CLI software and the HP 3PAR Management Console software. Unless otherwise stated, all tasks can be performed with both the CLI and the Management Console software. Refer to the *HP 3PAR OS CLI administrator's manual* and the *HP 3PAR OS Management Console online help* for instructions on how to perform the tasks described at a conceptual level in this guide.

Units of measure

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

- 1 KiB = 1,024 bytes
- 1 MiB = 2^{20} bytes = 1,048,576 bytes
- 1 GiB = 2^{30} bytes = 1,024 MB = 1,073,741,824 bytes
- 1 TiB = 2^{40} bytes = 1,024 GB = 1,099,511,627,776 bytes

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

- 1 KB = 1,000 bytes
- 1 MB = 10^6 bytes = 1,000,000 bytes
- 1 GB = 10^9 bytes = 1000 MB = 1,000,000,000 bytes
- 1 TB = 10^{12} bytes = 1000 GB = 1,000,000,000,000 bytes

Related documentation

Complete description of CLI commands	HP 3PAR CLI Reference
Overview and explanation of HP 3PAR technology	HP 3PAR concept guide
Using the Management Console to configure and administer the system	HP 3PAR Management Console online help
Using the CLI to configure and administer the system	HP 3PAR CLI Administrator's manual
Identifying storage system components and detailed alert information	HP Guided Troubleshooting
Using HP 3PAR Remote Copy software	HP 3PAR Remote Copy software user guide
Using the HP 3PAR CIM	HP 3PAR CIM API Programming Reference
Using HP 3PAR Host Explorer software	HP 3PAR Host Explorer user guide

For identifying storage system configuration specifications and compatibility information, go to the Single Point of Connectivity Knowledge (SPOCK) website at hp.com/storage/spock.

Overview

HP 3PAR StoreServ Storage concepts and terminology

HP 3PAR StoreServ is comprised of 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. Common provisioning groups are groups of logical disks. And virtual volumes use storage space provided by CPGs. The virtual volumes are exported to hosts and are the only data layer visible to hosts.

Figure 1. HP 3PAR StoreServ system data layers

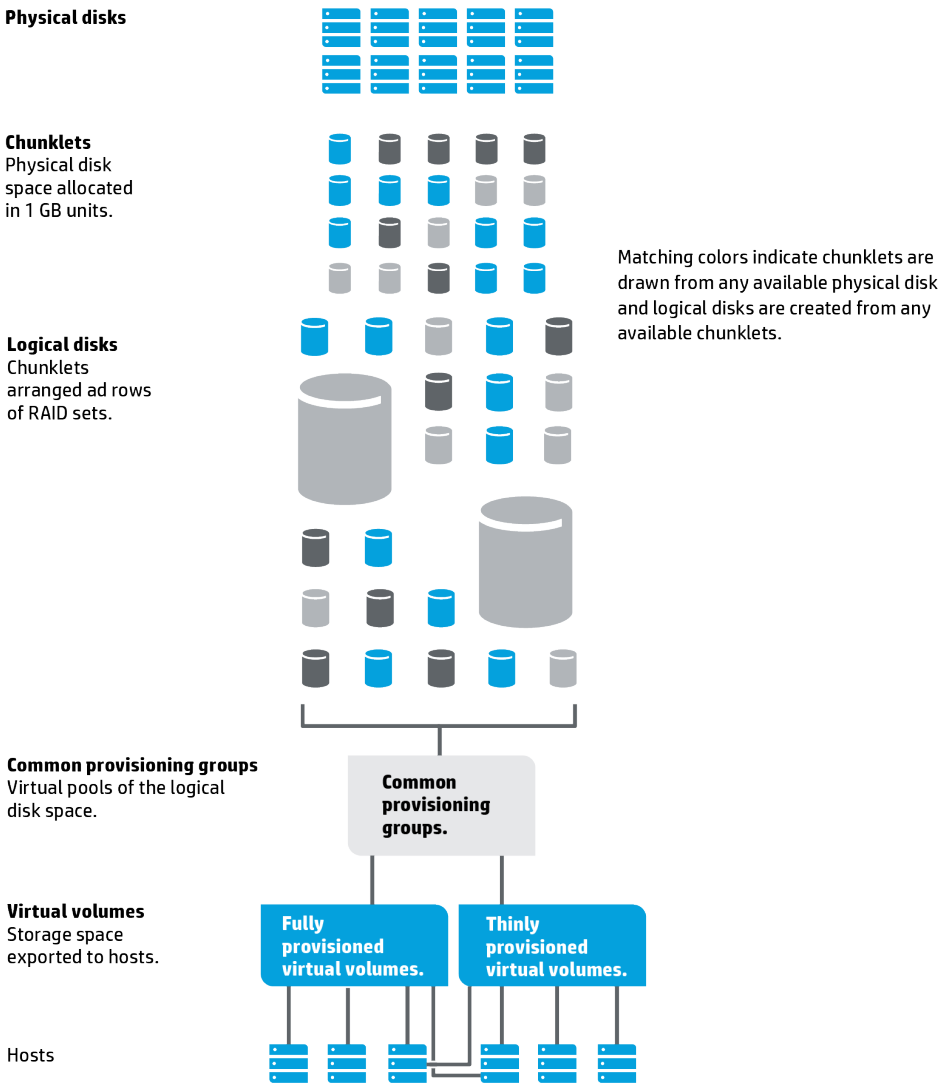

Physical disks

A *physical disk* is a hard drive mounted on a drive magazine located in an HP 3PAR StoreServ drive enclosure.

Chunklets

Physical disks are divided into *chunklets*. Each chunklet occupies contiguous space on a physical disk. On F-Class and T-Class systems, all chunklets are 256 MB. On HP 3PAR StoreServ 10000 and 7000 Storage, all chunklets are 1 GB.

Chunklets are automatically created by the HP 3PAR OS, and they are used to create logical disks. A chunklet is assigned to only one logical disk.

Logical disks

A *logical disk* is a collection of chunklets arranged as rows of RAID sets. Each RAID set is made up of chunklets from different physical disks. Logical disks are pooled together in common provisioning groups, which allocate space to virtual volumes. The underlying logical disks are automatically created by the HP 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 can be modified later. HP 3PAR StoreServ support the following RAID types:

- RAID 10 (RAID 1)
- RAID 50 (RAID 5)
- RAID Multi-Parity (MP) or RAID 6

Common provisioning groups

A *CPG* is a virtual pool of logical disks that allocates space to virtual volumes on demand. A CPG allows up to 65,536 virtual volumes to share the CPG's resources. You can create fully provisioned virtual volumes and thinly provisioned virtual volumes (TPVVs) that draw space from a CPG's logical disk pool. It is important to note that if no volumes (thick or thin) have been created in a CPG, it consumes no space.

Virtual volumes

Virtual volumes draw their resources from CPGs and are exported as logical unit numbers (LUNs) to hosts. Virtual volumes are the only data layer visible to the hosts. You can create physical copies or HP 3PAR Virtual Copy software snapshots of virtual volumes. Full copies remain available if the original base volume becomes unavailable. Before creating virtual volumes, you must first create CPGs to allocate space to the virtual volumes.

Fully provisioned virtual volumes

A *fully provisioned* virtual volume is a volume that uses logical disks that belong to a CPG. Unlike TPVVs, fully provisioned virtual volumes have a set amount of user space that is allocated for user data. The fully provisioned volume size is fixed, and the size limits range from 256 MB to 16 TB. Note that the volume size can be increased at any time (provided space is available) up to the maximum 16 TB size without any downtime however, the VV size cannot be decreased below the initial allocation.

Thinly provisioned virtual volumes

A TPVV is a volume that uses logical disks that belong to a CPG. TPVVs associated with the same CPG draw space from that pool as needed, allocating space on demand in small increments for each TPVV. As the volumes that draw space from the CPG require additional storage, the HP 3PAR OS automatically creates additional logical disks and adds them to the pool until the CPG reaches the user-defined growth limit which restricts the CPG's maximum size.

Physical copies

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 physical copy on the destination volume remains available if the original base volume becomes unavailable. Unlike a virtual copy or snapshot, a physical copy can maintain the performance of the base VV provided the physical copy has the same disk characteristics (type, speed, RAID level, etc.).

In addition, the destination volume must have a user space size at least as large as the user space of the base volume being copied. In addition, the HP 3PAR OS now allows the export of the physical copy immediately after the creation of the copy, while the data copy continues to completion in the background.

Virtual copy snapshots

A snapshot is a virtual copy of a base volume. The base volume is the original volume that is copied. Unlike a physical copy, which is a duplicate of an entire volume, a virtual copy only records changes to the base volume. This allows an earlier state of the original virtual volume to be recreated by starting with its current state and rolling back all the changes that have been made since the virtual copy was created.

You can make snapshots of fully provisioned virtual volumes, TPVVs, physical copies, or another virtual copy snapshot. Snapshots are created using copy-on-write techniques available only with the HP 3PAR Virtual Copy software license. Thousands of snapshots of each virtual volume can be created assuming that there is sufficient 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 snapshot data space. Changed data is copied only once regardless of the number of snapshots taken.

Note
Creating virtual copies requires an HP 3PAR Virtual Copy software license.

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 or attached disk devices.

What's new in HP 3PAR OS 3.1.3

Performance

- Multiple changes to the StoreServ OS to increase the efficiency of processor cycles, snapshot updates, and HP 3PAR Adaptive Optimization software region moves

Scalability

- Increased Fibre Channel (FC) initiators per port to 128 and maximum system initiators from 2,048 to 8,192

Model	3.1.3
7200	1,024
7400/7450	2,048
10400	4,096
10800	8,192

- Increased maximum number of volumes (base and snapshot) from 12,288 to 65,536
- Increased maximum base volumes from 4,096 to 32,768
- Increased maximum raw capacity limits

Model	Maximum raw capacity
7200	400 TB
7400	1.1 PB
7450	220 TB
10400	1.6 PB
10800	3.2 PB

- Increased maximum VLUN limits

Model	Maximum VLUNs
F/T	32,768
7000 series	65,536
10400	131,072
10800	131,072

Functionality

- HP 3PAR Priority Optimization software
  - 3.1.2 MU2 introduced IOP and throughput limits
  - 3.1.3 introduces latency goals
- HP 3PAR System Reporter software
  - Expanded reports and alert reporting (currently CLI only)
  - Improved performance of Adaptive Optimization during region moves
- HP 3PAR StoreServ Persistent Ports
  - Port failover now supports FCoE and iSCSI ports
  - Improved failover speed
  - Support for loss of synch for FC
- HP 3PAR Peer Motion software
  - Ability to migrate a single volume between arrays
  - Removed the requirement to unzone the hosts from the legacy array
  - Support for online migration of Windows® clusters

Remote Copy scalability enhancements

- Increased maximum number of replicated volumes from 2,400 to 6,000.
- Increased maximum number of Remote Copy groups.

Model	3.1.3 asynchronous periodic
F/T Series	800
7000 Series	2400
10400	6000
10800	6000

- Remote Copy link scalability enhancements

	3.1.3 (7000/10000)
Max RCFC ports/node	4
Max RCFC ports/InServ	32
Max RCIP ports/node	1
Max RCIP ports/InServ	8
Host HBA sharing	RCFC

- In the matrix below, multiple links/node can be from same host bus adapter (HBA) or different HBAs or a combination without any restrictions subject to maximum allowed.
- The limits below are software limits; however, some platforms (e.g., 7200) might have lower maximum limits because of slot limitations. On these platforms, their maximum possible hardware limits will apply.

	3.1.3 (2 nodes)	**3.1.3 (> 2 nodes)**
Max synch targets	4	8
Max periodic targets	4	4
Max mixed targets	4	8

Remote Copy topology enhancements

- For the 7000 and 10000 series InServ arrays, many of the topology restrictions have been removed and the supported topologies greatly expanded.
- Remote copy support for MxN replication.
- 4 x 4 bidirectional fan-in/fan-out replication.
- Flexible, cost-effective disaster recovery (DR) without dedicated replication pairs.
- Enables DR-as-a-service.
- Mix and match RCFC, RCIP, and FCIP.
- Mix and match replication modes (sync, periodic).
- Up to four RCFC links per node and one RCIP link per node.

Figure 2. Example of a newly supported configuration

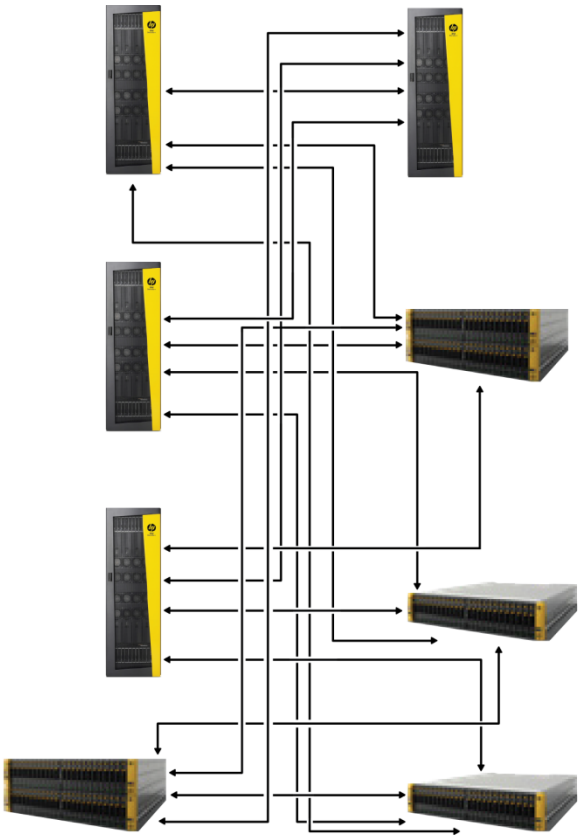

Getting started with HP 3PAR StoreServ Storage

Setting up your HP 3PAR StoreServ system ports

Port locations and nomenclature

The HP 3PAR CLI and/or Management Console displays the controller node, FC, iSCSI, and 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 your system. When viewing a system from the rear of a cabinet:
  - F-Class nodes are numbered 0–3 from top to bottom.
  - T-Class nodes are numbered 0–7 from left to right, top to bottom.
  - 7000 nodes are numbered 0–3 from bottom to top.
  - 10000 nodes are numbered 0–7 from left to right, bottom to top.
- Slot: Valid node slot numbers are 0–9 depending on the class of nodes installed in your system.
  - F-Class slots are numbered 0–5 from left to right.
  - T-Class slots are numbered 0–7 from left to right.
  - 7000 has a single slot in each node, numbered starting at 0.
  - 10000 slots are numbered 0–9 from left to right, bottom to top in a node in the lower chassis. In the upper chassis, slots are numbered 0–9 from left to right, top to bottom.
- Port: Valid node port numbers are 1–4 for all HBAs.
  - F-Class ports are numbered from top to bottom.
  - T-Class ports are numbered from top to bottom.
  - 7000 ports are horizontal and labeled beginning with 1 on the HBA or iSCSI adapter.
  - 10000 ports are numbered from bottom to top in a node in the lower chassis. In the upper chassis, ports are numbered from top to bottom.

Front-end port cabling

Best practice: Each HP 3PAR StoreServ controller node should be connected to two fabrics. This is to protect against fabric failures.

Best practice: 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

Best practice: Odd ports should connect to fabric 1 and even ports to fabric 2. Example with a 4 -node F400 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

Example with a 4 node HP 3PAR 10400 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

FC hosts zoning

Best practice: One initiator to multiple targets per zone (zoning by HBA). This zoning configuration is recommended for the HP 3PAR StoreServ Storage. Zoning by HBA is required for coexistence with other HP Storage arrays.

Best practice: Zoning should be done using Worldwide Port Names (WWPN, the WWN of each individual port on HP 3PAR StoreServ).

Hardware zoning should not be used (switch port/DID).

Best practice: A host 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. This is to support HP 3PAR Persistent Ports.

- Example from a single host: 0:2:1 and 1:21

Hosts do not need to be connected to all nodes because of the way the volumes are spread on all the nodes. A connection to two nodes is enough.

Best practice: Hosts need to be mirrored to node pairs. For example: zoned to nodes 0 and 1, or nodes 2 and 3. Hosts should **not** be zoned to non-mirrored nodes, such as 0 and 3.

Best practice: No more than 128 initiators are supported per front-end/host port.

Best practice: Each HP 3PAR StoreServ system has a maximum number of initiators supported that depends on the model and configuration. In regard to this maximum, *1 initiator = 1 path from a host*.

A single HBA zoned with two FC ports will be counted as two initiators.

A host with two HBAs, each zoned with two ports, will count as four initiators.

In order to keep the number of initiators below the maximum supported value, use the following recommendation:

- Hypervisors: four paths maximum, preferably to four different nodes.
- Other hosts (not hypervisors): two paths to two different nodes on the same port pairs.

Hosts and host sets

Best practice: When creating hosts, follow the implementation guide for each platform.

Selecting the correct host persona for each host is important. Implementation guides are available for download at the following address: hp.com/go/storage

Each physical server should have a different host defined, containing the WWNs or iQN for this host.

Best practice for creating a new host (Windows, SPARC Solaris, VMware®, and Red Hat® Linux®)

1. Install the Host Explorer software in the host if available for the host platform.
2. Zone in all the ports according to the zoning best practices.
3. From the host CLI, execute tpdhostagent—start; then tpdhostagent—push.
4. This will automatically create the host on the HP 3PAR StoreServ Storage system.

Best practice for creating a new host manually

1. Zone in host ports to HP 3PAR StoreServ using the zoning best practices, one host at a time.
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 HP 3PAR StoreServ to reduce the possibility of assigning incorrect WWNs to a host. Repeat until all hosts are zoned in.

Figure 3. Creating a host in the Management Console

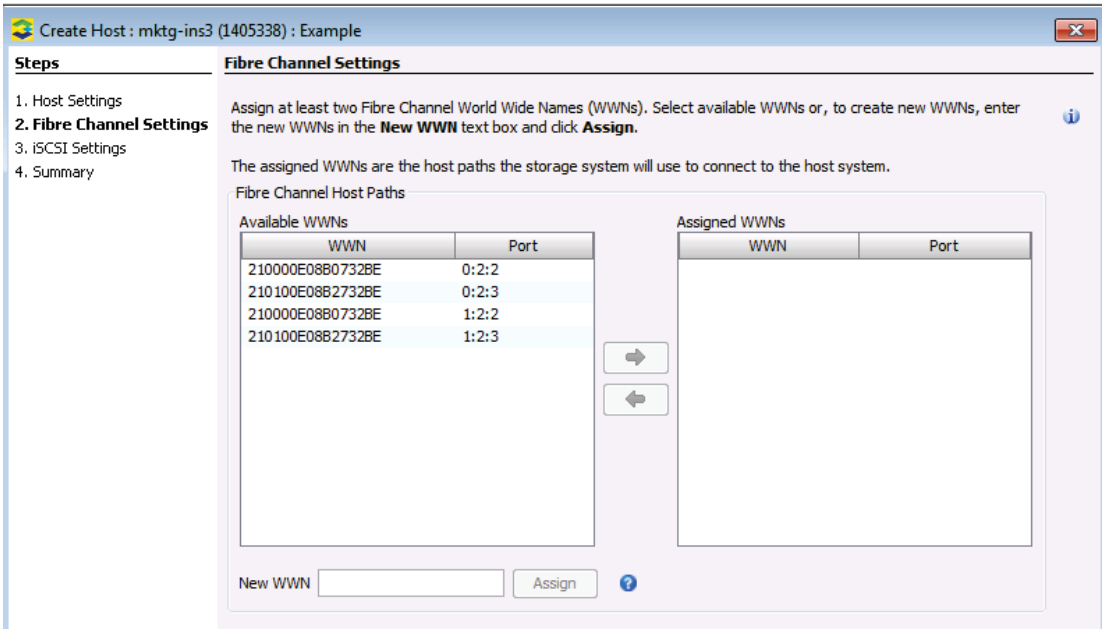

Best practice: For clusters, create a host set containing all the hosts used by the cluster. This will allow export of shared VVs to all hosts of the cluster in a single operation.

Provisioning from HP 3PAR StoreServ Storage

Host-based volume managers

The HP 3PAR StoreServ is, at its core, a very sophisticated volume manager—a volume manager that extends its control all the way down to the raw electrical device. As such, use of external (host-based) volume managers for device concatenation or external striping is **strongly discouraged**. Testing instances have been created where the external volume manager decreased performance when attempting to stripe on top of the HP 3PAR striping.

Common provisioning groups

Notes regarding CPGs

- CPGs primarily act as templates for the creation of LDs.
- If there are no volumes created "inside" a CPG, it will consume no space.
- CPGs define:
  - The RAID level for the LDs to be created
  - Availability level (HA CAGE, HA PORT, or HA MAG)
- CPGs will only be created across drives of the same type.

Best practice: When creating CPGs, always accept defaults if possible.

High availability

Best practice: HP encourages all HP 3PAR StoreServ Storage customers to upgrade to the latest recommended 3PAR OS. Upgrading to the most current 3PAR OS ensures that the storage system benefits from the ongoing design improvements and enhancements. For customers participating in the Get 6-Nines Guarantee Program, the Program will identify the latest 3PAR OS version that is covered under the Guarantee program.

Best practice: Size the system appropriately so that all workloads and applications dependent on the 3PAR system can perform as needed under the conditions of a node being down, which may happen during an unplanned controller 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 4 or more nodes, a resilience feature called Persistent Cache is automatically enabled. The Persistent Cache feature helps 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 helps ensure availability of the data in the event of unplanned downtime or controller node maintenance.

Persistent Ports

Persistent Ports functionality is supported for HP 3PAR OS 3.1.2 and later only (with functionality restrictions on HP 3PAR OS 3.1.2). Starting with HP 3PAR OS 3.1.3 support for FCoE connected hosts and iSCSI connected hosts has been added, and the ability to detect an array node suffering "loss_sync" (a physical layer problem occurring between the HP 3PAR node and the switch it is connected to) has been added. There is no Persistent Ports support on pre-HP 3PAR OS 3.1.2 releases.

For HP 3PAR StoreServ FC 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 FC fabric and preferably different FC 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 (no support for "loop").
- The FC fabric being used must support NPIV and have NPIV enabled.

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

- The same CNA (Converged Network Adapter) 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 host facing CNAs must be set to "target" mode.
- The FCoE network being used must support NPIV and have NPIV enabled.

For HP 3PAR StoreServ iSCSI host ports, the following requirements must be met:

- The same host port on host facing 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).

  There are no special requirements placed on the server HBAs and CNAs other than they must be supported as per Single Point of Connectivity Knowledge (SPOCK). See hp.com/storage/spock

Persistent Ports configuration considerations

Persistent Ports requires that corresponding "Native" and "Guest" host ports on a node pair be connected to the same FC fabric or IP network and the switches they are connected to must support and be configured for NPIV in the case of FC and FCoE. This means that for a minimum configuration to provide Persistent Ports functionality, where the node pair is connected to redundant FC SAN fabrics, each node in a node pair must have at least two FC host ports cabled with one port connected to each fabric. Figure 1 shows a minimum two-fabric redundant configuration that supports Persistent Ports for FC and figure 2 shows a configuration that will not support Persistent Ports functionality.

Best practice: Ensure that the same slot:port on each node in the node pair are connected to the same FC fabric.

Drive enclosure availability

Best practice: When creating CPGs, always select HA cage availability if possible. There are three levels of availability that can be selected with HP 3PAR StoreServ.

First, it is important to understand that drive magazines consist of four drives for HP 3PAR StoreServ 10000. Drive magazines consist of only a single drive in the F-Class and HP 3PAR StoreServ 7000 Storage systems.

- HA CAGE means that no two members of the same RAID set can be in the same drive enclosure. For example, to support RAID 5 3+1 (set size 4), four drive chassis connected to the same node pair are required. This helps ensure that data is still available in the event that access to an entire drive cage is lost. This applies to drive chassis that are point-to-point connected to the nodes (no daisy chain).
- HA MAG means that no two members of the same RAID set are in the same drive magazine. This allows a wider stripe with fewer drive chassis; for example, a RAID 5 stripe size of 7+1 (set size 8) would be possible with only four drive chassis, provided each chassis had at least two drive magazines.
- HA PORT applies only to daisy-chained drive chassis. When this level of availability is selected, no two members of the same RAID set can be in drive chassis that are dependent on one another for node connectivity. For example, in a system in which there are eight drive chassis and four of the drive chassis are connected to another drive chassis for node access, HA PORT would only allow RAID 5 3+1 (set size 4) in order to prevent the loss of one drive chassis causing a loss of data access.

Best practice: The number of CPGs should be kept to a minimum.

Model	Max. base VVs	Max. VVs + snapshots
F/T	4,096	8,192
7000	16,384	32,768
10000	32,768	65,536

Best practice: There are cases in which having more CPGs than the minimum will be required.

- When using HP 3PAR Virtual Domain software: a given CPG can only be in one domain.
- When using Adaptive Optimization software: a given CPG can only be in one Adaptive Optimization policy.
- When requiring easy capacity reporting by customer or application: creating a different CPG or set of CPGs per customer/application will make capacity planning easier.
- When Virtual Copy snapshots are heavily used: if the snapshots are to be kept in a different tier than the source data, use a different CPG for the user CPG and the copy CPG of a VV.
- When Virtual Copy snapshots are heavily used: if reporting of the snapshots' overall capacity is required, use a different CPG for the user CPG and the copy CPG of a VV.

Best practice: 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).

Best practice: Never create RAID 0 CPGs or VVs. RAID 0 does not offer any protection against drive failure.

Best practice: When using thin provisioning, schedule regular CPG compactions during low activity periods.

CPG compaction will allow capacity that is allocated to a CPG but no longer being used to hold data to be returned to the pool of free chunklets. This is not necessary if only one CPG per tier is used as the system will automatically reuse the reserved space for new volume allocations.

This should be scheduled during periods of low activity to reduce the potential performance impact of chunklet initialization (zeroing), which happens automatically when chunklets are freed.

In order to schedule a nightly CPG compaction at midnight, execute the following CLI command:

```
cli% createsched "compactcpg -f -pat *" "0 0 * * *" compactcpg
```

Note
If using Adaptive Optimization on all the CPGs, scheduling a CPG compaction is not required. CPG compaction is part of the Adaptive Optimization process.

Solid-state drive CPGs

Best practice: Solid-state drive (SSD) CPGs should be of the RAID 5 type with a "set size" of 3+1 by default.

This will bring superior performance/capacity ratio. If maximum performance is required, use RAID 1.

Best practice: 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 to 16 GB, on six-node systems to 24 GB, and on eight-node systems to 32 GB.

In order to set the CPG growth increment to a lower value than the default, the "show advanced option" box must be checked.

Best practice: Availability should be left to "cage level" availability (the default option) if the system's configuration allows for it. If not, it should be set to "magazine level" availability. This can be changed using the "advanced options" checkbox of the Management Console.

Other advanced settings such as "preferred chunklets" and "step size" should not be changed from their default values. No disk filtering should be used either.

Fast class (FC/SAS) CPGs

Best practice: FC CPGs should be of the RAID 5 type by default. This will bring superior performance/capacity ratio.

The "set size" (data to parity ratio) can be changed from the default value of 3+1 if the system configuration supports it. If unknown, use 3+1.

Best practice: For applications that have a very high write ratio (over 50 percent of the access rate), create a CPG using RAID 1 if maximum performance is required.

Best practice: The growth increment should be left to the default value (32 GB per node pair).

Best practice: 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. This can be changed using the "advanced options" checkbox of the Management Console.

Other advanced settings such as "preferred chunklets" and "step size" should not be changed from their default values. No disk filtering should be used either.

NL CPGs

Best practice: NL CPGs should be of the RAID 6 type by default.

The "set size" (data to parity ratio) can be changed from the default value of 8 (6+2) if the system configuration supports it. RAID 5 is not recommended with NL disks.

Best practice: The growth increment should be left to the default value (32 GB per node pair).

Best practice: 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. This can be changed using the "advanced options" checkbox of the Management Console.

Other advanced settings such as "preferred chunklets" and "step size" should not be changed from their default values. No disk filtering should be used either.

Priority Optimization

In HP 3PAR OS 3.1.2 MU2, Priority Optimization (quality of service [QoS]) offered the ability to set a maximum limit of IOPS or bandwidth available to a VVset.

HP 3PAR OS 3.1.3 introduces priority levels (high, normal, and low), latency goals, and minimum goals along with the ability to set QoS rules against Virtual Domains.

Care must be taken when setting QoS rules, system performance has to be understood, and historical performance data has to be taken into account when setting maximum limits, minimum goals, and latency goals.

A good understanding of the applications that reside on the volumes is also equally important.

Control type	Minimum HP 3PAR OS version	Description	Dependencies and best practices
Max. limit	3.1.2 MU2 and later	Maximum threshold for IOPS and/or 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/or bandwidth below which HP 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 will be 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 as 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 practice: The *HP 3PAR Priority Optimization* white paper should be followed when configuring QoSs rules. It is available for download at hp.com/go/support.

Best practice: Maximum limit, this is the maximum amount of IOPS or bandwidth, or both, which a given VVset or domain is allowed to achieve. Best practice is to use the System Reporter data in order to quantify volume performance and set maximum limits rules accordingly.

Best practice: Minimum goal, this is the minimum amount of IOPS or bandwidth, or both, below which the system will not throttle a given VVset or domain in order to meet the Latency Goal of a higher priority workload. The minimum goal should be set by looking at the historical performance data and by understanding what is the minimum amount of performance that should be granted to the applications that reside in that VVset. The volumes in the VVset may use more IOPS/bandwidth than what is set by the minimum goal, but will be throttled to the given limit as the system gets busier. The performance may also go below the minimum goal; this may happen if the application is not pushing enough IOPS or if the sum of all minimum goals defined is higher than the I/O capability of the system or a given tier of storage.

Best practice: Latency goal, this is the svctime the system will goal to fulfill for a given QoS rule. In order for the goal to work, rules with a minimum goal must exist so the system can throttle those workloads. Reasonable latency goal should be set; this can be done by looking at historical performance data. The latency goal will also be influenced by the tier the volume resides on, below are some guidelines:

Tier	Latency goal guidelines
All NL	>= 20 ms
All FC	>= 10 ms
All SSD	>= 1 ms
Sub-tiered volumes (Adaptive Optimization)	Guideline of the middle tier in a 3-tier Adaptive Optimization configuration or of the lowest tier in a 2-tier Adaptive Optimization configuration

When a latency goal is set on a QoS object in a given tier, users 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.

HP 3PAR Adaptive Optimization and QoS can cooperate, but HP 3PAR Adaptive Optimization data migration may impact a latency goal settings. For example, when a workload is partially migrated from FC drives into NL ones after the collection and analysis of the I/O access pattern by HP 3PAR Adaptive Optimization, depending on the host access I/O pattern this may have a negative impact on latency and trigger workload to be throttled towards their minimum goal. This is an expected behavior if the QoS rule was modeled after the FC performance.

Best practice: Three priority levels exist: high, normal, and low. As the system gets busier it will start targeting lower priority workloads and throttling their performance in order to meet higher priority workloads latency goals. High priority level should be used against critical applications, lower priority on less critical applications.

Virtual volumes

Best practice: Zero detection should be enabled on TPVVs that are periodically "zeroed out." Zero detect is enabled by default in HP 3PAR OS version 3.1.2 and later.

Best practice: Thinly provisioned virtual volumes can have an allocation warning, but should never have an allocation limit. Do not set an allocation limit, not even 100 percent.

Best practice: Virtual volumes should have both a "user CPG" and "copy CPG" selected. A "copy CPG" is required to use any of the following features:

- Full copy (clones)
- Virtual copy (snapshots)
- Remote copy (remote replication)

The user CPG and copy CPG can be different CPGs. This can be useful if snapshots are to be created in a different class of service from the source data.

Unless this is specifically required, use the same CPG for user CPG and copy CPG.
Virtual LUNs (exports) and volume sets

Best practice: Virtual volumes should be exported to host objects, not to ports for all hosts.

Best practice: "Matched sets" (a combination of export to hosts and exports to ports) should not be used.

Limiting the number of paths for a given host should only be done by zoning the host with only specific front-end ports.

Best practice: Boot from SAN virtual volumes should be exported with LUN 0.

Best practice: Use volume sets when exporting multiple virtual volumes to a host or host set.

Volume sets allow multiple volumes to be exported to the same host or host set in one operation. When exporting a volume set to a host or host set, the user selects which LUN ID will be used for the first virtual volume of the volume set, and the LUN IDs are incremented for every virtual volume in the volume set.

Best practice: Individual volumes (belonging only to one host) should be exported with LUN IDs starting from 1. Volume sets should be exported with LUN IDs starting from 10.

Best practice: Always leave a gap of LUN IDs between two volume sets.

When adding a new virtual volume to a volume set that is already exported, the next LUN ID needs to be free or the operation will fail.

Therefore, it is important that if multiple volume sets are exported to the same host/host set, there are gaps left for adding later virtual volumes to the volume set.

Example: If exporting two volume sets containing 10 virtual volumes each to a host set:

- Export the first volume set with LUN ID 10.
  - LUN IDs 10 to 19 will be used.
- Export the second volume set with LUN ID 30.
  - LUN IDs 30 to 39 will be used.

Remote Copy

Best practice: The *Remote Copy user guide* should be followed strictly when configuring Remote Copy. It is available for download at hp.com/go/support.

Best practice: Remote Copy FC ports should only be in one zone with one Remote Copy FC port from another HP 3PAR StoreServ system.

Do not zone hosts or other devices with the Remote Copy FC other than the destination Remote Copy FC port. For synchronous long distance remote copy SLD RC deployments, a given Remote Copy group will support both transport methods.

HP 3PAR OS version 3.1.3 has substantially improved topology flexibility.

Best practice: All virtual volumes belonging to the same application should be added to the same Remote Copy group. Virtual volumes that have logical links include:

- Virtual volumes used by the same application (data and log volumes, for example)
- Virtual volumes used by a Logical Volume Manager (LVM) volume group
- Virtual volumes that contain virtual disks of the same VM

A Remote Copy group can contain up to 300 virtual volumes.

Best practice: Do not add virtual volumes 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.

Best practice: Source and destination virtual volumes must have the same size.

Best practice: While source and destination virtual volumes do not need to be of the same type of disk or RAID level, performance must be considered before mixing different types of disks/RAID levels on the source and destination systems.

When using Remote Copy in synchronous mode, it is not advised to replicate FC virtual volumes to NL virtual volumes, as the reduced performance of the NL disks might impact the primary virtual volume performance.

Best practice: When using Remote Copy 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).

Best practice: In case of complete communication failure between two HP 3PAR StoreServ systems, wait for a period of low activity before restarting the Remote Copy groups.

Remote Copy initial synchronization is modified to throttle bandwidth utilization in order to not impact system performance.

Best practice: 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.

Adaptive Optimization

Best practice: The following combinations are acceptable within the same Adaptive Optimization configuration (policy):

- SSD, FC/SAS, and NL
- FC/SSD and FC
- FC/SAS and NL

Using different RAID levels within the same policy is acceptable.

Table 2. Recommended Adaptive Optimization configurations

Configuration type	SSD tier	SAS 10K/15K tier	NL tier
2 tiers SSD-FC/SAS	At least 5% of the capacity*	95% of the capacity	N/A
2 tiers FC/SAS-NL	N/A	Minimum 60% of the capacity	Maximum 40% of the capacity
		100% of the IOPS target	0% of the IOPS target
3 tiers SSD-FC/SAS-NL	At least 5% of the capacity*	Minimum 55% of the capacity	Maximum 40% of the capacity
			0% of the IOPS target

Note: Configurations with SSD and NL tiers only are not recommended; four controller configurations require the same type of drives behind controller pairs.

* Or minimum disk requirement for SSD.

Best practice: When configuring two or three tier solutions containing SSDs, if region density data is not available for sizing the SSD tier assume the SSD tier will only provide the following IOPS per SSD drive:

- 200 GB SSD ~= 550 IOPS
- 400 GB SSD ~= 1000 IOPS[1]
- 480 GB SSD ~= 1100 IOPS[1]
- 920 GB SSD ~= 2150 IOPS[1]

With Specific regard to HP 3PAR AO, 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 vs. the smaller SSDs.

Best practice: Always size the solution assuming the NL tier will contribute 0 percent of the IOPS required from the solution.

Best practice: 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 percent).

[1] Estimated IOPS for SSDs used in an Adaptive Optimization configuration. This is not reflective of how many IOPS and SSD can deliver.

Best practice: When using thin provisioning volumes along with Adaptive Optimization, select a CPG using FC 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.

Best practice: Ensure that the default tier (FC) 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 FC 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 FC 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 HP 3PAR StoreServ) until the moment when the regions of data will be migrated to the other tiers.

Best practice: 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 FC CPGs.

This will help ensure that SSD capacity is consumed by Adaptive Optimization and will allow this capacity to be safely used to 95 percent or even 100 percent.

To help ensure that no TPVV is associated with an SSD or NL CPG, run the "showcpg" command and confirm that only the FC CPG reports a TPVV value greater than 0.

In the following example, only the FC CPG has TPVVs associated with it:

```
                                                    ---------------------(MB)--------------------
                                 --Volumes--   ---Usage---  ---- Usr ----  ---- Snp ----  ---- Adm ----
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
```

Best practice: 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."

Best practice: CPGs for SSDs should use RAID 5, with the minimum growth increment supported (8 GB per node pair). CPGs for FC disks should use RAID 5, with the default growth increment.

CPGs for NL disks should use RAID 6, with the default growth increment.

Refer to the "Common provisioning groups" section for details of CPG best practices.

Best practice: Schedule the different 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.

Best practice: It is preferable to not set any capacity limit on the Adaptive Optimization configuration level, or on the CPG (no allocation warning or limit).

This will allow the Adaptive Optimization software to make excellent use of the different tiers available in the system. In HP 3PAR OS version 3.1.2 and later, use the MC 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 Management Console or the CLI. Adaptive Optimization will not attempt to move more data to a CPG than the capacity warning set on the CPG.

Best practice: 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.

Best practice: 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 FC and NL tiers.
- Run during the days of the workweek only (for example, Monday–Friday).
- Executes once a day, preferably at night.
- Uses a measurement/hours of 24.

Best practice: 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, as they might conflict with each other.

Only Adaptive Optimization, being the lowest level (storage level) should be used.

Security

Best practice: Change the password for the "3paradm" user.

The password of the "3parcim" user can also be changed if the CIM/SMI-S service is enabled.

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 the initialization of the system and are required for communication between the service processor and the HP 3PAR StoreServ system.

Best practice: 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.

Best practice: 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 doesn't need to remove objects, use a create account.

Best practice: Do not use a super user (such as "3paradm") for the System Reporter user.

In HP 3PAR OS 3.1.1 and earlier versions, use a browse user if Adaptive Optimization is not required, or a 3PAR_AO user if Adaptive Optimization is required. In 3.1.2 and later versions of the HP 3PAR OS, user selection is no longer necessary.

Naming convention

A good naming convention is important to effectively manage an HP 3PAR system. A naming convention should be consistent and descriptive. HP 3PAR systems are case sensitive. Therefore, use either only lowercase or only uppercase names.

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

prd.unix.datavg01.vv prd.unix.appvg02.vv

Use a naming convention as 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 [*]).

Examples: showcpg *.cpg.*

Define the naming convention early in the implementation. (See the following for an example.)

Naming convention example

Hosts

Host names support up to 31 characters.

Host name will be of the form <TYPE>.<OS>.<HOST>.<OBJECT TYPE>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <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 HPUX hosts, etc.
- <HOST> is the host name.
- <OBJECT TYPE> is one of VV, cpg, VVset, etc.

Examples:

  - prd.win.server1.vv
  - prd.hpux.server2.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, etc.
- <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 HPUX hosts, etc.
- <CLUSTER NAME> is the name of the cluster.

Examples:

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

Virtual Volumes

VV names support up to 31 characters.

For standalone servers, VV name will be of the form <TYPE>.<OS>.<HOST>.<FS NAME>.<ID>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <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 HPUX hosts, etc.
- <HOST> is the host name.
- <FS NAME> is the file system name or drive letter.
- <ID> is an optional ID for volume groups containing multiple LUNs.

Examples:

  - prd.win.server1.e.vv
  - prd.hpux.server2.datavg1.01.vv
  - dev.lin.server3.data1.vv

For clusters, VV name will be of the form <TYPE>.<OS>.<CLUSTER>.<RES NAME>.<ID>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <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 HPUX hosts, etc.
- <CLUSTER> (optional) contains the cluster name.
- <RES NAME> is the resource name of the cluster object mapped to this virtual volume.
- <ID> is an optional ID for volume groups containing multiple LUNs.

Examples:

  - prd.win.cluster1.sql1.m.vv
  - prd.vmw.datastore1.vv
  - prd.aix.datavg1.01.vv

Virtual Volume sets

VV set names support up to 31 characters.

VV set name will be of the form <TYPE>.<OS>.<HOST/CLUSTER>.<FS NAME>, where:

- <TYPE> can be prd for production servers, dev for development servers, tst for test servers, etc.
- <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 HPUX hosts, etc.
- <HOST/CLUSTER> is the host name or cluster name.
- <FS NAME> is the OS name for the group of LUNs, volume group, and datastore group.

Examples:

  - prd.win.hyperv.vmdata1.vvset
  - prd.vmw.esx.datastoregrp2.vvset
  - dev.lin.server3.datavg.vvset

System Reporter

HP 3PAR System Reporter has been traditionally run from an external system prior to HP 3PAR version 3.1.2. The external system reporter (ESR) can still be run from an external server to capture data for reporting but this data is totally independent from the on node System Reporter and the two databases cannot be combined for reporting purposes.

On node System Reporter data can be queried via CLI or via HP 3PAR Management Console. Reporting capabilities have been greatly expanded in version 4.5 of the Management Console.

Best practice: The System Reporter white paper should be followed when configuring and using System Reporter. It is available for download at hp.com/go/support.

Best practice: For users that don't have needs to control how long their historical performance data is retained, it's recommended to use the reporting functionalities from HP 3PAR Management Console. This provides historical reporting by using the data stored on the internal database that resides in the HP 3PAR OS.

Best practice: If you are running ESR, it's recommended to use MySQL database.

Best practice: If you are running ESR, it's recommended to skip LD performance data collection as Adaptive Optimization now runs as part of the HP 3PAR OS. If the System Reporter server and database are running on the same host, double the values found by the sizing tool. The sizing tool is an excel sheet provided with the HP 3PAR System Reporter Installer. Running System Reporter in virtual machines (VMs) is acceptable as long as the VMs are sized correctly.

Best practice: Do not increase the default retention times for high-resolution and hourly data.

If a longer history is required than the default allows, schedule automatic reports to be created every day. This will allow the data to be viewed for a long time without requiring a lot of CPU, memory, and disk resources.

Best practice: To secure access to the System Reporter Web server: the easiest way to secure a System Reporter server is to enable the local firewall and prevent remote access to the HTTP port 80.

This will help ensure that only users who can connect locally to the System Reporter server will be able to access the System Reporter Web server.

It is also possible to configure Apache to require password authentication. See the following link for a quick tutorial on how to do so: sniptools.com/vault/windows-apache-and-htaccess-authentication.

Best practice: Do not use a super user (such as "3paradm") for the System Reporter user. Use a browse user.

Best practice: Schedule high-resolution reports every day for VLUN, PD, CPU, and port performance.

This will make detailed performance analysis possible long after the high-resolution data has been purged by System Reporter. For every HP 3PAR StoreServ system managed by System Reporter, schedule high-resolution reports (through the Policy Settings page) to be executed every day at the same time and kept (Retention) for the length of the support period of the system.

These scheduled reports should cover: VLUNs, ports (hosts, disks, and Remote Copy), PDs (individual reports for SSD, FC, and NL), CPU, as well as Adaptive Optimization (region density, space moved, and VV tiers).

Ongoing management and growth

Autonomic rebalance

With HP 3PAR StoreServ 3.1.3, the rebalancing of data after hardware upgrades is now an integrated feature, not requiring a Dynamic Optimization license. After purchasing and installing new hardware, the rebalancing can be started by selecting the HP 3PAR StoreServ in the GUI and selecting “tune system.”

Figure 4. Autonomic rebalancing selection

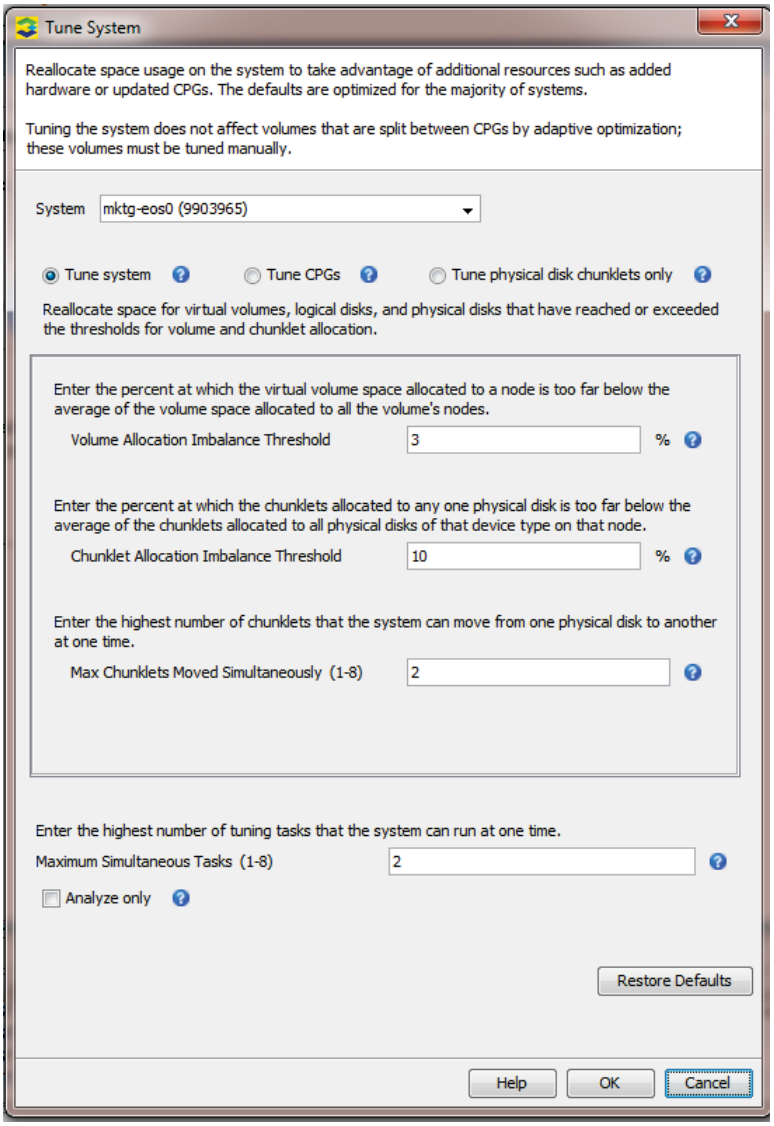

Autonomic rebalancing (tune system) can also be started from the CLI by typing tunesys.

Depending on the amount of hardware added to a system and the degree to which the system may 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 may take multiple runs of tunesys to level the system out. Other variables will 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, etc. The autonomic rebalancing process will give priority to servicing host I/O.

Best practice: Execute tunesys during off peak hours. Leave the leveling tolerance at the default values.

Appendix

Chart of supported host personas

Persona_Id	Persona_Name	Persona_Caps
1	Generic	UARepLun, SESLun
2	Generic-ALUA	UAPreLun, SESLun, RTPG, ALUA
7	HP-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®	UAPelLun, SESLun, ALUA, WSC

Summary

HP 3PAR StoreServ Storage is the last storage architecture you will ever need—regardless of 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, HP 3PAR StoreServ storage has you covered. It's storage that offers 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 your future may hold. It's one of the most advanced storage platforms, which is here to help you master unpredictability—effortlessly, without exception, and without compromise.

Learn more at
hp.com/3par

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA4-4524ENW, September 2014, Rev. 5