

# HPE 3PAR StoreServ Fast Track HODH4S

The Managing HPE 3PAR StoreServ Fast Track course is designed for the HPE 3PAR administrators from entry to experienced level. The goal of the class is to acquaint the 3PAR administrator with the most common day-to-day tasks and best practices associated with administration with additional topics and advanced features of the 3PAR array. The levels of provisioning storage are emphasized.

HPE 3PAR Fast Track course also provides a comprehensive understanding of everyday administration within an HPE B-series SAN solution covering technologies and concepts.

This training reflects the newest release of the HPE 3PAR OS: 3.3.1. NOTE: Hosts used in the lab environment are MS Windows.

<b>HPE course number</b>	HODH4S
<b>Course length</b>	5 days
<b>Delivery mode</b>	ILT, vILT
<b>View schedule, local pricing, and register</b>	<a href="#">View now</a>
<b>View related courses</b>	<a href="#">View now</a>

## Why HPE Education Services?

- IDC MarketScape leader 4 years running for IT education and training\*
- Recognized by IDC for leading with global coverage, unmatched technical expertise, and targeted education consulting services\*
- Key partnerships with industry leaders OpenStack®, VMware®, Linux®, Microsoft®, ITIL, PMI, CSA, and (ISC)2
- Complete continuum of training delivery options—self-paced eLearning, custom education consulting, traditional classroom, video on-demand instruction, live virtual instructor-led with hands-on lab, dedicated onsite training
- Simplified purchase option with HPE Training Credits

## Audience

HPE 3PAR administrators who desire training from the basic concepts and best practices needed to administer the array and keen to get additional training on the advanced features of the HPE 3PAR array.

Technical professionals seeking a learning path that includes both conceptual knowledge of Fibre Channel SAN technologies and experience in HPE B-series SAN environments.

## Prerequisites

At the conclusion of this course, the student should be able to:

- An understanding of general storage concepts including fibre channel technology and RAID
- Operator level functionality in a Windows environment
- Basic technical understanding of networking and storage, concepts and terminology
- Experience managing Windows® or UNIX® systems
- Recommended free, Web-based training: SAN Fundamentals (U5527AAE)

## Course Objectives

At the conclusion of this course, the student should be able to:

- Describe features and functionality of the switch hardware
- Discuss ICL technology
- Describe blade, HB/CAN, SFP, and cable options
- Discuss FC theory
- Identify switch configuration parameters
- List the steps for port initialization
- List Fabric services and features
- Perform an out-of-box initial configuration
- Perform common administrative tasks
- List basic security features
- Work with ports
- Identify routing capabilities
- Discuss trucking
- Manage Inter Switch Link (ISL)
- Zoning theory
- New zoning features
- Implement zoning using the CLI syntax
- Activate and deactivate a default zone
- List zoning tools
- Identify when a long distance license is required
- List SAN extension technologies
- Configure long distance connection
- Describe buffer concepts
- Administer using different interfaces
- List SNMP capabilities
- Back up and manage configuration files
- Upgrade firmware
- Create a diagram of a fabric using collected data
- Perform basic troubleshooting and performance monitoring
- Explain the numbering schemes for the HPE 3PAR hardware components: controllers, ports, and physical disks
- Use the StoreServ Management Console (SSMC) GUI and the CLI to perform administrative tasks
- Create and work with a Common Provisioning Group (CPG)
- Administer Virtual Volumes using the SSMC and the CLI
- Understand the advantages of Thin Provisioning and create a Thin Provisioned Virtual Volume (TPVV)
- Understand the advantages of Dedup and Compression for storage allocated from SSDs
- Export and un-export virtual volumes from hosts
- Use HPE 3PARinfo to analyze luns presented to hosts
- Use Host Explorer to simplify addition of hosts
- Use Host Sets and Volume Sets to simplify provisioning storage
- Create a Snapshot and promote (restore) from a Snapshot
- Create a Clone and promote a Clone
- Convert a Virtual Volume (i.e. from fully provisioned to thin provisioned or thin provisioned to thin dedup or vice versa)
- Use the SSMC reporting feature to monitor capacity and performance
- Use the CLI stat commands to troubleshoot a 3PAR array solution
- Manage a private cloud infrastructure by using HPE CloudSystem Operations Console (OpsConsole)
- Create provider and tenant networks
- Check available resources by using an OpenStack (Horizon) portal
- Create, manage, and publish compute, network, and storage services for users to deploy
- Explore and manage the HPE Helion CloudSystem Enterprise Marketplace Portal (MPP)
- Work with Dynamic Optimization to change volumes characteristics (media type, raid level, etc.)
- Work with Adaptive Optimization to implement virtual volume tiering to save storage costs
- Work with Priority Optimization to prioritize higher priority workloads over lower priority workloads
- Use Remote Copy to replicate data between arrays
- Understand the Peer Persistence high availability feature when used in conjunction with Remote Copy
- Migrate data between arrays using Peer Motion in a Storage Federation

## Detailed course outline

---

### Brocade switches

- Brocade Fibre Channel switch family overview
  - Features and functionality of the switch hardware
  - Switch blades and compatibility
  - ICL connectivity
  - Different Brocade FC HBAs
  - Brocade CNAs
  - Different types of fiber optic cable and SFPs
- 

### Fibre Channel theory

- SA theory and introduction
  - FC layers
  - Fibre Channel networking model
  - Different class of service
  - Frame structure
  - WWN format
  - Node and port types
  - Basic switch parameters
  - Steps for port initialization
  - FC addressing
  - Fabric services well-known addresses
  - Fabric and N-port login sequence
  - RSCN
  - Name server and other fabric services
  - NPIV implementation and support
  - Fabric and port name
  - Interoperability
- 

### Installation and configuration

- Out-of-box initial configuration
  - Checking switch status
  - Basic management tools
  - Important fabric parameters
  - Initial configuration
  - IP settings
  - Time settings
  - Login banner
  - Activating licenses
  - Setting switch and chassis name
  - Working with syslog server settings
  - Working with security
  - Switch and port status
  - Booting a switch
-

---

**Brocade FCP routing**

- Terminology review
- Principle switch role and up/downstream links
- Fabric initialization
- ISL oversubscription
- Virtual channels
- Port-based routing
- Exchange-based routing
- Dynamic Load Sharing (DLS)
- In-Order Delivery (IOD)
- Inter Switch Link (ISL)
- Working and checking routing
- Trunking theory, benefits, and configuration

---

**Zoning**

- Basic concepts associated with zoning
- Zoning elements
- Zoning hierarchy
- New zoning features
- Zoning using the CLI syntax
- How to check the maximum size of a zoning database
- Displaying zoning information
- Activate and deactivate a default zone
- Zoning configuration example
- Hardware and session enforcement differences
- Zoning tools
- Issues when adding a switch to an existing fabric with zoning enabled
- Best practices when implementing a zone
- Fabric segmentation

---

**Long distance connectivity**

- Cabling
- Extension options
- Use and effects of buffer credits on distance and speed
- Long distance modes, settings, and supported distances for Brocade switches
- Configuration
- When a long distance license is required
- Limitations of long distance connection when using trunking
- Long distance commands
- Buffer monitoring

---

**Administration—Firmware**

- Interfaces and management tools to manage Brocade switches and fabrics
  - Using BNA, Web Tools, and Telnet/SSH/HTTP/SSL/SNMP
  - Host Connectivity Manager (HCM) to manage Brocade HBAs
  - Brocade SAN health
  - Fabric watch
  - SNMP management
  - SNMP commands
  - Back up and manage configuration files
  - Firmware upgrade steps
-

---

**Basic troubleshooting**

- Troubleshooting techniques
- Data-gathering process
- Supportsave and supportshow
- How to create a diagram of a fabric using collected data
- Documentation tools
- Brocade SAN health
- HPE SAP® visibility tool
- Common SAN problems
- Switch and Field Replaceable Units (FRU) status
- Performance monitoring and troubleshooting
- Diagnostic tools

---

\* The switch part is only theory, no labs will be provided for this part of the session

## Detailed Course Outline

<b>HPE 3PAR Solution Introduction and Review</b>	<ul style="list-style-type: none"> <li>• HPE 3PAR Hardware Current models</li> <li>• Hardware overview for the 7000, 8000, and 20000 models: controllers, drive cages, and disks</li> <li>• Controller port number naming conventions for the 7000, 8000, and 20000 models</li> <li>• Disk number naming convention for disks in the HPE 3PAR arrays</li> </ul>	<ul style="list-style-type: none"> <li>• ASIC chip functionality</li> <li>• HPE 3PAR virtual storage architecture benefits</li> <li>• Big Picture HPE 3PAR component connectivity</li> <li>• SSMC GUI introduction and functionality overview</li> <li>• HPE 3PAR CLI overview and introduction</li> </ul>
<b>Storage Concepts and Terminology</b>	<ul style="list-style-type: none"> <li>• HPE 3PAR provisioning terminology</li> <li>• HPE 3PAR concept of a disk chunklet and Logical Disk (LD)</li> <li>• HPE 3PAR concept of a Common Provisioning Group (CPG)</li> <li>• HPE 3PAR Virtual Volumes (VV) types: full provisioning and thin provisioning</li> <li>• HPE 3PAR provisioning terminology</li> </ul>	<ul style="list-style-type: none"> <li>• HPE 3PAR concept of a disk chunklet and Logical Disk (LD)</li> <li>• HPE 3PAR concept of a Common Provisioning Group (CPG)</li> <li>• HPE 3PAR Virtual Volumes (VV) types: full provisioning and thin provisioning</li> </ul>
<b>Storage Configuration</b>	<ul style="list-style-type: none"> <li>• CPGs using SSMC and the CLI</li> <li>• Working with virtual volume templates</li> </ul>	<ul style="list-style-type: none"> <li>• Working with fully provisioned and thin provisioned volumes using SSMC and the CLI</li> </ul>
<b>Host Connectivity and Storage Allocation</b>	<ul style="list-style-type: none"> <li>• How to prepare a host to access an HPE 3PAR storage array</li> <li>• HPE 3PAR Port Persistence</li> <li>• Adding hosts in an HPE 3PAR storage array</li> <li>• Adding FC ports to a host</li> <li>• Export virtual volumes to hosts as VLUNs</li> </ul>	<ul style="list-style-type: none"> <li>• Unexport VVs/VLUNs from a host</li> <li>• Using SSMC and CLI to work with hosts and storage</li> <li>• Working with Smart SAN</li> <li>• Use Host Explorer to add hosts and explore host details</li> <li>• Use HPE 3PARInfo to gather storage information</li> </ul>
<b>Host Sets and Virtual Volume Sets</b>	<ul style="list-style-type: none"> <li>• Host and virtual volume sets advantages</li> <li>• Host and virtual volume sets use cases</li> </ul>	<ul style="list-style-type: none"> <li>• Creating and maintaining host and volume sets using SSMC and the CLI</li> <li>• Host and volume sets guidelines and rules</li> </ul>
<b>Thin Features: A Deep Dive</b>	<ul style="list-style-type: none"> <li>• Benefits of the Zero Detection/Thin Persistence</li> <li>• Zero Detection use cases</li> <li>• Space reclamation: UNMAP and other utilities</li> <li>• Compaction ratio explanation and thin overprovisioning</li> </ul>	<ul style="list-style-type: none"> <li>• Thin provisioning: seeing what is allocated vs. used by a host</li> <li>• Online Virtual Volume conversion: converting volumes from full to thin or thin to full using SSMC and the CLI</li> <li>• Compacting a CPG to free up space</li> </ul>
<b>Adaptive Data Reduction for Flash: Dedup and Compression</b>	<ul style="list-style-type: none"> <li>• Thin Dedup: saving space using deduplication and dedup detail</li> <li>• Thin Dedup: what to dedup: good candidates and bad candidates</li> <li>• Compression and how it works</li> <li>• Compression: what to compress--good candidates and bad candidates</li> <li>• Data Packing: packing data for space and cost savings</li> <li>• DECO: Dedup and Compression together</li> </ul>	<ul style="list-style-type: none"> <li>• Working with Dedup and Compression using the SSMC GUI and the CLI</li> <li>• Performing a Dedup Estimate and a Compression Estimate</li> <li>• Seeing if Dedup and Compression are working</li> <li>• Analyzing Adaptive Data Reduction ratios: Dedup ratio, Compression ratio and overall Data Reduction ratio</li> <li>• Online Virtual Volume conversion: converting volumes to thin-dedup, compressed or DECO</li> </ul>
<b>Snapshots and Clones</b>	<ul style="list-style-type: none"> <li>• Snapshots and Clones: benefits</li> <li>• Creating, exporting, unexporting, and deleting a snapshot</li> <li>• Rules and relationships regarding snapshots</li> <li>• Restore from a snapshot</li> <li>• Scheduling the creation of a snapshot</li> </ul>	<ul style="list-style-type: none"> <li>• Using the Virtual Lock feature to delete outdated snapshots</li> <li>• Resynchronize a clone to a base volume</li> <li>• Promote a clone to a base volume</li> <li>• Use SSMC and the CLI to manage snapshots and clones</li> </ul>
<b>Reporting: Performance and Capacity</b>	<ul style="list-style-type: none"> <li>• The on-node database: the .srd data database volume</li> <li>• Using SSMC reports to analyze capacity needs and trends</li> <li>• Using SSMC reports to analyze performance</li> <li>• Using SSMC default reports and creating custom reports</li> <li>• Scheduling and emailing reports</li> </ul>	<ul style="list-style-type: none"> <li>• Using the Excel client</li> <li>• Using the CLI stat commands to analyze performance and capacity</li> <li>• Using the CLI sr* commands to analyze performance and capacity</li> </ul>

## Detailed Course Outline

<b>Dynamic Optimization</b>	<ul style="list-style-type: none"> <li>• Introduction to the Dynamic Optimization feature and its benefits: changing raid levels, media types and set sizes using DO</li> <li>• Dynamic Optimization use cases: cost, availability, and performance</li> <li>• Using SSMC and the CLI to tune a volume's user space and copy space</li> </ul>	<ul style="list-style-type: none"> <li>• Performing a DO tune and converting a volume simultaneously</li> <li>• Dynamic Optimization troubleshooting: performing a Restart and Rollback</li> </ul>
<b>Adaptive Optimization</b>	<ul style="list-style-type: none"> <li>• Introduction to the Adaptive Optimization feature and using AO to balance performance and cost</li> <li>• Dynamic Optimization vs. Adaptive Optimization</li> <li>• Components of AO: the 128MG region, region movement, region analysis</li> <li>• AO sizing goals and the AO algorithm</li> </ul>	<ul style="list-style-type: none"> <li>• Configuring AO and managing AO using the SSMC GUI and the CLI</li> <li>• Monitoring AO and looking at AO effectiveness</li> <li>• AO latency thresholds</li> <li>• AO best practices</li> <li>• AO reports in SSMC</li> </ul>
<b>Priority Optimization</b>	<ul style="list-style-type: none"> <li>• Introduction to Priority Optimization and Quality of Service implementation</li> <li>• Priority Optimization use cases</li> <li>• Priority Optimization: how it works</li> <li>• Priority Optimization configuration using SSMC and the CLI</li> </ul>	<ul style="list-style-type: none"> <li>• Priority Optimization using Priorities and the System Busy level</li> <li>• Using Latency Goals with Priority Optimization</li> <li>• Performance considerations and Priority Optimization</li> <li>• Priority Optimization monitoring using SSMC and CLI commands</li> </ul>
<b>Remote Copy</b>	<ul style="list-style-type: none"> <li>• Introduction to replication principles and Remote Copy</li> <li>• Replication using RCIP and RCFC protocols</li> <li>• Replication and thin provisioning, dedup, and compression</li> <li>• Remote copy groups and data integrity</li> <li>• Remote copy modes: Synchronous, Periodic Asynchronous, and Async Streaming</li> <li>• Replication using Remote Copy between three arrays using Synchronous Long Distance modes</li> </ul>	<ul style="list-style-type: none"> <li>• Remote Copy configurations: One-to-Many, Many-to-One, and M-to-N</li> <li>• Remote Copy failure scenarios</li> <li>• Failing over a remote copy group</li> <li>• Remote copy configuration and administration using SSMC and the CLI</li> </ul>
<b>Peer Persistence</b>	<ul style="list-style-type: none"> <li>• Introduction to Peer Persistence and using Peer Persistence with Remote Copy</li> <li>• Using Peer Persistence for disaster tolerant data and load balancing</li> <li>• Performing an Automatic Transparent failover using Quorum Witness or a Manual Transparent failover</li> <li>• Quorum Witness details and requirements</li> </ul>	<ul style="list-style-type: none"> <li>• Failover scenarios with Peer Persistence</li> <li>• Peer Persistence path management</li> <li>• 3DC (3 Data Center) Peer Persistence</li> <li>• Peer Persistence vs. Synchronous Long Distance</li> </ul>
<b>Storage Federation with Peer Motion</b>	<ul style="list-style-type: none"> <li>• Data Migration concepts using Peer Motion</li> <li>• Bi-directional migration using Storage Federation</li> <li>• Migration of data using Peer Motion using SSMC</li> </ul>	<ul style="list-style-type: none"> <li>• Storage Federation use cases and features</li> <li>• Storage Federation supported configurations</li> <li>• Migration of data using the PMU CLI</li> </ul>

Learn more at  
[hpe.com/au/education](http://hpe.com/au/education)

Follow us:



© Copyright 2017 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 is either a registered trademark or trademark 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. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

c04623212 June 2017, H0DH4S