

Managing HPE 3PAR StoreServ I: Management and local replication HK902S

HPE course number	HK902S
Course length	3 days
Delivery mode	ILT, vILT
View schedule, local pricing, and register	View now
View related courses	View now

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

The Managing HPE 3PAR StoreServ I course is designed for the new or entry-level HPE 3PAR administrator. 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 of the 3PAR array. The levels of provisioning storage are emphasized.

This training reflects the newest release of the HPE 3PAR OS: 3.3.1. The course is approximately 50% lecture and 50% hands-on labs using HPE 3PAR arrays. NOTE: Hosts used in the lab environment are MS Windows.

Audience

HPE 3PAR administrators who desire training on basic concepts and best practices needed to administer the array.

Prerequisites

- An understanding of general storage concepts including fibre channel technology and RAID
- Operator level functionality in a Windows environment
- Please register [here](#) to access to the H6LH4AAE Pre Study for Managing HPE 3PAR StoreServ I

Course Objectives

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

- 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

Detailed Course Outline

H6LH4AAE Pre Study for Managing HPE 3PAR StoreServ I	<ul style="list-style-type: none"> Register here to access to the Pre-Study WBT course 	
HPE 3PAR Solution Introduction and Review	<ul style="list-style-type: none"> HPE 3PAR Hardware Current models Hardware overview for the 7000, 8000, and 20000 models: controllers, drive cages, and disks Controller port number naming conventions for the 7000, 8000, and 20000 models Disk number naming convention for disks in the HPE 3PAR arrays 	<ul style="list-style-type: none"> ASIC chip functionality HPE 3PAR virtual storage architecture benefits Big Picture HPE 3PAR component connectivity SSMC GUI introduction and functionality overview HPE 3PAR CLI overview and introduction
Storage Concepts and Terminology	<ul style="list-style-type: none"> HPE 3PAR provisioning terminology HPE 3PAR concept of a disk chunklet and Logical Disk (LD) HPE 3PAR concept of a Common Provisioning Group (CPG) HPE 3PAR Virtual Volumes (VV) types: full provisioning and thin provisioning HPE 3PAR provisioning terminology 	<ul style="list-style-type: none"> HPE 3PAR concept of a disk chunklet and Logical Disk (LD) HPE 3PAR concept of a Common Provisioning Group (CPG) HPE 3PAR Virtual Volumes (VV) types: full provisioning and thin provisioning
Storage Configuration	<ul style="list-style-type: none"> CPGs using SSMC and the CLI Working with virtual volume templates 	<ul style="list-style-type: none"> Working with fully provisioned and thin provisioned volumes using SSMC and the CLI
Host Connectivity and Storage Allocation	<ul style="list-style-type: none"> How to prepare a host to access an HPE 3PAR storage array HPE 3PAR Port Persistence Adding hosts in an HPE 3PAR storage array Adding FC ports to a host Export virtual volumes to hosts as VLUNs 	<ul style="list-style-type: none"> Unexport VVs/VLUNs from a host Using SSMC and CLI to work with hosts and storage Working with Smart SAN Use Host Explorer to add hosts and explore host details Use HPE 3PARInfo to gather storage information
Host Sets and Virtual Volume Sets	<ul style="list-style-type: none"> Host and virtual volume sets advantages Host and virtual volume sets use cases 	<ul style="list-style-type: none"> Creating and maintaining host and volume sets using SSMC and the CLI Host and volume sets guidelines and rules
Thin Features: A Deep Dive	<ul style="list-style-type: none"> Benefits of the Zero Detection/Thin Persistence Zero Detection use cases Space reclamation: UNMAP and other utilities Compaction ratio explanation and thin overprovisioning 	<ul style="list-style-type: none"> Thin provisioning: seeing what is allocated vs. used by a host Online Virtual Volume conversion: converting volumes from full to thin or thin to full using SSMC and the CLI Compacting a CPG to free up space
Adaptive Data Reduction for Flash: Dedup and Compression	<ul style="list-style-type: none"> Thin Dedup: saving space using deduplication and dedup detail Thin Dedup: what to dedup: good candidates and bad candidates Compression and how it works Compression: what to compress--good candidates and bad candidates Data Packing: packing data for space and cost savings DECO: Dedup and Compression together 	<ul style="list-style-type: none"> Working with Dedup and Compression using the SSMC GUI and the CLI Performing a Dedup Estimate and a Compression Estimate Seeing if Dedup and Compression are working Analyzing Adaptive Data Reduction ratios: Dedup ratio, Compression ratio and overall Data Reduction ratio Online Virtual Volume conversion: converting volumes to thin-dedup, compressed or DECO
Snapshots and Clones	<ul style="list-style-type: none"> Snapshots and Clones: benefits Creating, exporting, unexporting, and deleting a snapshot Rules and relationships regarding snapshots Restore from a snapshot Scheduling the creation of a snapshot 	<ul style="list-style-type: none"> Using the Virtual Lock feature to delete outdated snapshots Resynchronize a clone to a base volume Promote a clone to a base volume Use SSMC and the CLI to manage snapshots and clones
Reporting: Performance and Capacity	<ul style="list-style-type: none"> The on-node database: the .srdata database volume Using SSMC reports to analyze capacity needs and trends Using SSMC reports to analyze performance Using SSMC default reports and creating custom reports Scheduling and emailing reports 	<ul style="list-style-type: none"> Using the Excel client Using the CLI stat commands to analyze performance and capacity Using the CLI sr* commands to analyze performance and capacity

Learn more at
hpe.com/ww/learnstorage

Follow us:



© Copyright 2018 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.

HK902S G.00, June 2018