

Managing HPE 3PAR StoreServ II HK904S

The Managing HPE 3PAR StoreServ II course is designed for the experienced HPE 3PAR administrator. The goal of the class is to acquaint the 3PAR administrator with additional topics and features of the HPE 3PAR array. The class is intended to be a follow-up to the Managing HPE 3PAR StoreServ I course.

This training reflects the newest release of the HPE 3PAR OS: 3.3.1. The course is approximately 60% lecture and 40% hands-on labs using HPE 3PAR arrays.

HPE course number	HK904S
Course length	2 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

Audience

HPE 3PAR administrators who desire additional training on the advanced features of the HPE 3PAR array.

Prerequisites

- An understanding of general storage concepts
- Successful completion of the Managing HPE 3PAR StoreServ I course
- Operator level functionality in a Windows environment

Course Objectives

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

- 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

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

Learn more at
hpe.com/ww/learnstorage

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.

c04599841, June 2017, HK904S G.00