

Szkozenie: HPE Managing HPE Alletra 6000



Cel szkolenia:

The Managing HPE Alletra 6000 course describes the HPE Alletra 6000 and HPE Alletra 5000 portfolio hardware building blocks, theories of operation and features. Using hands-on labs, students learn to perform common day-to-day management tasks, including how to create hosts, volumes, and collections, as well as how to monitor the product. This course also provides knowledge of more advanced features, including local and remote replication, disaster recovery, scaling-out, QoS, and maintenance. You gain a practical understanding of HPE Alletra 6000 and HPE Alletra 5000 array capabilities using extensive hands-on lab exercises, performed on all applicable user interfaces.

Course objectives

After completing this course, you should be able to:

- Describe HPE Alletra 6000 and HPE Alletra 5000 hardware, architecture, and software features
- List and compare HPE Alletra 6000 and HPE Alletra 5000 management options
- Explain HPE Alletra 6000 and HPE Alletra 5000 provisioning terminology, features, and read/write operations flow
- Describe HPE Alletra 6000 and HPE Alletra 5000 array initialization and Data Services Cloud Console onboarding
- Prepare and create hosts and initiator groups for HPE Alletra 6000 and HPE Alletra 5000 storage arrays
- Work with volumes, volume collections, and quality of service (QoS)
- Describe and manage snapshots, clones, and protection templates
- Describe the concept, use, and benefits of the HPE Alletra 6000 and HPE Alletra 5000 storage scale-out architectures
- Explain asynchronous remote replication concepts, implementation, and failure scenarios
- Describe Peer Persistence concepts, architecture, requirements, and the Automatic Switchover (ASO) process
- Describe HPE Alletra 6000 and HPE Alletra 5000 monitoring options and tools
- Describe the HPE Alletra 6000 and HPE Alletra 5000 OS update procedure

Audience

Customers, administrators, and channel partner sales or technical sales

Plan szkolenia:

- Hardware and Features
 - HPE Alletra 6000 and HPE Alletra 5000 overview
 - The architecture behind HPE Alletra 6000 and HPE Alletra 5000 systems
 - PCIe expansion options
 - Power supply unit details
 - HPE Alletra 6000 and HPE Alletra 5000 series head shelf drive layout
 - RAID options
 - HPE Alletra 2140 4u all flash expansion shelves
 - HPE Alletra 2120 expansion shelves
 - HPE Alletra 6000 and HPE Alletra 5000 controller upgrades
 - Leading enterprise features
- Management Options
 - HPE Alletra 6000 and HPE Alletra 5000 manageability
 - HPE GreenLake, Data Ops Manager, and Data Services Cloud Console
 - HPE GreenLake edge-to-cloud platform services
 - HPE GreenLake dashboard
 - Data Services Cloud Console overview
 - Provisioning options
 - Data Ops Manager dashboard
 - Block Storage dashboard
 - Intent-based provisioning
 - Data Services Cloud Console REST API
 - HPE Alletra 6000 and HPE Alletra 5000 web user interface (UI)
 - HPE Alletra 6000 and HPE Alletra 5000 command line interface (CLI)
 - HPE Alletra 6000 and HPE Alletra 5000 Rest API
 - HPE Alletra 6000 and HPE Alletra 5000 WSAPI
- Provisioning Terminology
 - Industry-leading data protection and efficiency
 - Read and write operations
 - What is an NVDIMM?
 - NVDIMM—how it works
 - Resiliency and data integrity
 - RAID options review
 - Intra-drive parity (aka chunk parity)

- HPE Alletra 6000 and HPE Alletra 5000 Checksums and self-ids
- Integrated spare: rebuild operation
 - Integrated spare RAID: single drive failure scenario
 - Integrated spare RAID: replacement drive rebuild process
- Quick RAID rebuild
- HPE Alletra 6000 and HPE Alletra 5000 data reduction
- Thin provisioning overview
- Space reclamation
- Inline dedupe in action
- What to deduplicate and compress
- Log structured file system design
- HPE Alletra 6000 and HPE Alletra 5000 storage file system—CASL
- SmartSecure encryption
- Initialization and Array Status
 - Installation
 - Documentation and resources
 - End-to-end onboarding
 - HPE Alletra 6000 and HPE Alletra 5000—preinstallation and software subscription
 - HPE Storage Toolkit for Windows
 - Setup Manager for HPE Alletra 6000, HPE Alletra 5000, and HPE Nimble arrays
 - Discover the array
 - Connect the array to the HPE GreenLake edge-to-cloud platform
 - Onboarding device to Data Service Cloud Console
 - Data Service Cloud Console setup service
 - Array initialization using the CLI
 - Diagnostics setup and test
- Hosts and Initiator Groups
 - Networking concepts and terminology
 - Understanding Internet Protocol addresses (IPs)
 - Discovery IP addresses
 - Typical subnets
 - Optimizing switch links: IP address zones
 - Fibre channel components
 - Initiator groups and access control
 - Protocol connectivity options
 - Obtaining host Host Bus Adapter (HBA) World Wide Names (WWNs)

- Fibre channel zoning overview
- HPE Smart SAN for HPE Alletra 6000 and HPE Alletra 5000
- Fibre channel configuration using Target Driven Zoning (TDZ)
- Working with initiators from Data Ops Manager and Data Services Cloud Console
- Working with initiators from Data Services Cloud Console
- Working with initiators from HPE Alletra 6000 and HPE Alletra 5000 user interfaces (UIs)
 - Initiator groups
 - Access control
- Working with initiators from the command line interface (CLI)
- HPE Storage Toolkit for Windows and HPE Connection Manager for Windows
- Volumes, Volume Collections, and QoS
 - Volumes overview
 - Thin provisioning
 - Volume reserves
 - Volume quotas
 - Performance policies
 - Workload-specific requirements
 - Custom performance policies
 - Application-aware storage
 - Protection templates and volume collections
 - Quality of Service (QoS)
 - Monitoring volume limits
 - Provisioning options review
 - Intent-based provisioning
 - Basic provisioning
 - Provisioning from Data Services Cloud Console
 - Provisioning from HPE Alletra 6000 and HPE Alletra 5000 UIs
 - Provisioning from the command line interface (CLI)
 - Volumes and QoS—some CLI commands
- Protection Policies, Snapshots, and Clones
 - Snapshots
 - Understanding snapshots
 - Anatomy of a snapshot: backup
 - Understanding clones
 - Managing from Data Ops Manager and Data Services Cloud Console
 - Protection policy examples

- Managing from HPE Alletra 6000 and HPE Alletra 5000 UIs
 - A manual snapshot
 - Scheduled snapshot
 - Schedules
- Managing from the command line interface (CLI)
 - Protection templates
 - Snapshots and clones
- Asynchronous Remote Replication
 - Replication Introduction
 - SmartReplicate (asynchronous replication)
 - Replication use cases
 - Replication components
 - How replication works—the basics
 - Volume collection and replication schedules
 - Replication monitoring
 - Replication QOS—bandwidth limit
 - Replication considerations
 - SmartReplicate disaster recovery (DR)
 - General DR operations
 - Replication concepts
 - Promote
 - Handover
 - Demote
 - Recovery scenarios
 - Managing from Data Ops Manager and Data Services Cloud Console
 - Replication using Data Services Cloud Console—Block
 - Creating and protecting volumes in block storage
 - Protection policies available in Data Service Cloud Console
 - Managing from HPE Alletra 6000 and HPE Alletra 5000 UIs
 - Managing from the command line interface (CLI)
- Scale-Out
 - Scale-to-fit: flexible and nondisruptive scalability
 - Scale-out introduction
 - Overview of group/pool operations with new arrays
 - Scale-out technology
 - Nondisruptive data migrations

- Spanned pools
 - Volume creation
 - Node addition
 - Node evacuation
- Peer Persistence
 - Peer Persistence: sync replication with auto-failover
 - Peer Persistence overview
 - Peer Persistence builds on scale-out group
 - Peer Persistence Quorum Witness
 - Peer Persistence customer benefits
 - Peer Persistence architecture
 - Basic terms
 - Connectivity
 - Performance
 - Peer Persistence volume collection
 - Peer Persistence snapshots
 - Sync operations
 - Automatic switchover (ASO)
 - Peer Persistence scenarios
- Monitoring, Alerts, and Events
 - HPE InfoSight
 - HPE InfoSight Cross-Stack Analytics for VMware® environments
 - Options available in Data Ops Manager and Data Services Cloud Console
 - Options available in HPE Alletra 6000 and HPE Alletra 5000 UIs
 - About SNMP
 - SNMP support
 - Syslog support
 - Options available in the command line interface (CLI)
- HPE Alletra 6000 and HPE Alletra 5000 OS Update
 - Data Ops Manager OS upgrade
 - Downloading updates, running readiness checks, and installing updates
 - After readiness check completes
 - Installing update

Wymagania:

- An understanding of general storage concepts including fiber channel, iSCSI technology and RAID
- Operator level functionality in a Windows environment

Poziom trudności



Certyfikaty:

The participants will obtain certificates signed by HPE (course completion).

Prowadzący:

Authorized HPE Trainer.