Hewlett Packard Enterprise

Managing HPE Alletra 6000 H61N2S

HPE course number	H61N2S
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?

- Comprehensive worldwide HPE technical, IT industry and personal development training

- Training and certification preparation for ITIL®, Security, VMware®, Linux, Microsoft and more

- Innovative training options that match individual learning styles

- Anytime, anywhere remote learning via HPE Digital Learner subscriptions

- Verifiable digital badges for proof of training, skill recognition and career development

- Simplified purchase options with HPE Training Credits

The Managing HPE Alletra 6000 course describes the HPE Alletra 6000 portfolio hardware building blocks, theory of operation and features. Using hands-on labs (HOL), 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, and QoS, as well as maintenance. You gain a practical understanding of HPE Alletra 6000 array capabilities using extensive hands-on lab exercises, performed on all applicable user interfaces.

Audience

Customers, administrators, and channel partner sales or technical sales

Prerequisites

- An understanding of general storage concepts including fiber channel, iSCSI technology and RAID

- Operator level functionality in a Windows environment

Course objectives

After completing this course, you should be able to:

- Describe HPE Alletra 6000 hardware, architecture, and software features

- List and compare HPE Alletra 6000 management options

- Explain HPE Alletra 6000 provisioning terminology, features, and read/write operations flow

- Describe HPE Alletra 6000 array initialization and Data Services Cloud Console onboarding

- Prepare and create hosts and initiator groups for an HPE Alletra 6000 storage array

- 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 storage scale-out architecture

- 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 monitoring options and tools

- Describe the HPE Alletra 6000 OS update procedure

Detailed course outline

Module 0: Course Introduction

Module 1: Hardware and Features	
• HPE Alletra 6000 overview	• RAID options
• The architecture behind HPE Alletra 6000 systems	• HPE Alletra 2140 4u all flash expansion shelves
• PCIe expansion options	• HPE Alletra 6000 controller upgrades
• Power supply unit details	• Leading enterprise features
• HPE Alletra 6000 series head shelf drive layout	

Module 2: Management Options	
• HPE Alletra 6000 manageability	• Block Storage dashboard
• HPE GreenLake, Data Ops Manager, and Data Services Cloud Console	• Intent-based provisioning
• HPE GreenLake edge-to-cloud platform services	• Data Services Cloud Console REST API
• HPE GreenLake dashboard	• HPE Alletra 6000 web user interface (UI)
• Data Services Cloud Console overview	• HPE Alletra 6000 command line interface—CLI
• Provisioning options	• HPE Alletra 6000 Rest API
• Data Ops Manager dashboard	• HPE Alletra 6000 WSAPI

Module 3: Provisioning Terminology	
• Industry-leading data protection and efficiency	• Quick RAID rebuild
• Read and write operations	• HPE Alletra 6000 data reduction
• What is an NVDIMM?	• Thin provisioning overview
• NVDIMM—how it works	• Space reclamation
• Resiliency and data integrity	• Inline dedupe in action
• RAID options review	• What to deduplicate and compress
• Intra-drive parity (aka chunk parity)	• Log structured file system design
• HPE Alletra 6000 Checksums and self-ids	• HPE Alletra 6000 storage file system—CASL
• Integrated spare: rebuild operation	• SmartSecure encryption
– Integrated spare RAID: single drive failure scenario	
– Integrated spare RAID: replacement drive rebuild process	

Module 4: Initialization and Array Status	
• Installation	• Connect the array to the HPE GreenLake edge-to-cloud platform
• Documentation and resources	• Onboarding device to Data Service Cloud Console
• End-to-end onboarding	• Data Service Cloud Console setup service
• HPE Alletra 6000: preinstallation and software subscription	• Array initialization using the CLI
• HPE Storage Toolkit for Windows	• Diagnostics setup and test
• Setup Manager for HPE Alletra 6000 and HPE Nimble arrays	
• Discover the array	

Module 5: Hosts and Initiator Groups	• Networking concepts and terminology	• Fibre channel zoning overview
	• Understanding Internet Protocol addresses (IPs)	• HPE Smart SAN for HPE Alletra 6000
	• Discovery IP addresses	• Fibre channel configuration using Target Driven Zoning (TDZ)
	• Typical subnets	• Working with initiators from Data Ops Manager and Data Services Cloud Console
	• Optimizing switch links: IP address zones	• Working with initiators from Data Services Cloud Console
	• Fibre channel components	• Working with initiators from HPE Alletra 6000 User Interface (UI)
	• Initiator groups and access control	– Initiator groups
	• Protocol connectivity options	– Access control
	• Obtaining host Host Bus Adapter (HBA) World Wide Names (WWNs)	• Working with initiators from the command line interface (CLI)
		• HPE Storage Toolkit for Windows and HPE Connection Manager for Windows
Module 6: Volumes, Volume Collections, and QoS	• Volumes overview	• Quality of Service (QoS)
	– Thin provisioning	• Monitoring volume limits
	– Volume reserves	• Provisioning options review
	– Volume quotas	• Intent-based provisioning
	• Performance policies	• Basic provisioning
	• Workload-specific requirements	• Provisioning from Data Services Cloud Console
	• Custom performance policies	• Provisioning from HPE Alletra 6000 UI
	• Application-aware storage	• Provisioning from the command line interface (CLI)
	• Protection templates and volume collections	• Volumes and QoS—some CLI commands
Module 7: Protection Policies, Snapshots, and Clones	• Snapshots	• Managing from HPE Alletra 6000 UI
	• Understanding snapshots	– A manual snapshot
	• Anatomy of a snapshot: backup	– Scheduled snapshot
	• Understanding clones	– Schedules
	• Managing from Data Ops Manager and Data Services Cloud Console	• Managing from the command line interface (CLI)
	• Protection policy examples	– Protection templates
		– Snapshots and clones
Module 8: Scale-Out	• Scale-to-fit: flexible and nondisruptive scalability	• Nondisruptive data migrations
	• Scale-out introduction	• Spanned pools
	• Overview of group/pool operations with new arrays	– Volume creation
	• Scale-out technology	– Node addition
		– Node evacuation

Module 9: Asynchronous Remote Replication	• Replication Introduction	• Replication concepts
	• SmartReplicate (asynchronous replication)	– Promote
	• Replication use cases	– Handover
	• Replication components	– Demote
	• How replication works—the basics	• Recovery scenarios
	• Volume collection and replication schedules	• Managing from Data Ops Manager and Data Services Cloud Console
	• Replication monitoring	• Replication using Data Services Cloud Console—Block
	• Replication QOS—bandwidth limit	• Creating and protecting volumes in block storage
	• Replication considerations	• Protection policies available in Data Service Cloud Console
	• SmartReplicate disaster recovery (DR)	• Managing from HPE Alletra 6000 UI
	• General DR operations	• Managing from the command line interface (CLI)
Module 10: Peer Persistence	• Peer Persistence: sync replication with auto-failover	• Connectivity
	• Peer Persistence overview	• Performance
	• Peer Persistence builds on scale-out group	• Peer Persistence volume collection
	• Peer Persistence Quorum Witness	• Peer Persistence snapshots
	• Peer Persistence customer benefits	• Sync operations
	• Peer Persistence architecture	• Automatic switchover (ASO)
	• Basic terms	• Peer Persistence scenarios
Module 11: Monitoring, Alerts, and Events	• HPE InfoSight	• About SNMP
	• HPE InfoSight Cross-Stack Analytics for VMware® environments	• SNMP support
	• Options available in Data Ops Manager and Data Services Cloud Console	• Syslog support
		• Options available in the command line interface (CLI)
	• Options available in HPE Alletra 6000 UI	
Module 12: HPE Alletra 6000 OS Update	• Data Ops Manager OS upgrade	• HPE Alletra 6000 UI
	– Downloading updates, running readiness checks, and installing updates	– HPE Alletra 6000 OS upgrade
	– After readiness check completes	• Command line interface (CLI)
	– Installing update	

Detailed lab outline

Lab 0: HPE vLabs Access

Lab 1: Getting Started with UIs	• Exercise 1: Accessing and browsing hardware information in Data Services Cloud Console Data Ops Manager	• Exercise 2: Accessing and browsing hardware information in HPE Alletra 6000 UI
		• Exercise 3: Accessing and browsing hardware information in HPE Alletra 6000 CLI

Lab 4: Initialization and Array Status	• Exercise 1: Welcome Center	• Exercise 2: Cabling tools

---

Lab 5: Hosts Configuration and Management	• Exercise 1: Initial preparation of host	• Exercise 2b: Working with hosts in HPE Alletra 6000 UI
	• Exercise 2a: Working with hosts in Data Services Cloud Console Data Ops Manager	• Exercise 3: Working with hosts in HPE Alletra 6000 CLI

Lab 6: Volumes Configuration and Management	• Exercise 1: Working with volumes and collections in Data Services Cloud Console Block Storage	• Exercise 3: Working with volumes and collections in HPE Alletra 6000 CLI
	• Exercise 2: Working with volumes and collections in HPE Alletra 6000 UI	• Exercise 4: Disks preparation

---

Lab 7: Working with Protection Policies, Snapshots and Clones	• Exercise 1: Data preparation	• Exercise 3: Working with snapshot schedules and snapshots in HPE Alletra 6000 UI
	• Exercise 2: Working with protection policies and snapshots in Data Services Cloud Block Storage	• Exercise 4: Working with snapshots and clones using HPE Alletra 6000 CLI

Lab 8: Working with QoS	• Exercise 1: Workloads preparation	• Exercise 3: Working with QoS in HPE Alletra 6000 UI
	• Exercise 2: Working with QoS in Data Services Cloud Block Storage	• Exercise 4: Working with QoS using HPE Alletra 6000 CLI

Lab 9: Working with Replication	• Exercise 1: Initial setup	• Exercise 4: Clean-up
	• Exercise 2: Creating and managing replication collections	• Exercise 5: Configuring replication using CLI
	• Exercise 3: Failover to the remote site and failback	• Exercise 6: Final clean-up

---

Lab 10: Working with Scale-Out technology	• Exercise 1: Initial setup	• Exercise 3: Configuring Scale-Out Pool
	• Exercise 2: Configuring Scale-Out Group	• Exercise 4: Performing data migration

---

Lab 11: Working with Peer Persistence	• Exercise 1: Initial setup	• Exercise 3: Configuring Peer Persistence
	• Exercise 2: Working with the Quorum Witness	• Exercise 4: Performing failover and failback

---

Lab 12: Performing Monitoring, Managing, Alerts and Events	• Exercise 1: Initial setup	• Exercise 3: Managing alerts and events
	• Exercise 2: Monitoring performance	

Lab 13: Working with HPE Alletra 6000 Maintenance and OS Update	• Exercise 1: Browsing maintenance options	• Exercise 2: Firmware update preparation

Learn more at
hpe.com/ww/learnstorage

Follow us:

in

**Hewlett Packard
Enterprise**

H61N2S A.00, August 2022