Hewlett Packard Enterprise



# HPE Ezmeral Data Fabric Cluster Administration

## HG7E0S

| | |
|---|---|
| **Course ID** | HG7E0S |
| **Duration** | 3 days |
| **Format** | ILT, VILT |
| **View schedule, local pricing, and register** | View now |
| **Browse related courses** | View now |

In this course, students gain the knowledge and skills to plan, install, maintain, and manage a secure **HPE Ezmeral Data Fabric File and Object Store cluster**. Using lecture and labs, you learn how to design and install a cluster, and perform pre- and post-installation testing. You configure users, groups, and work with key features of an HPE Ezmeral Data Fabric cluster, including volumes, snapshots, and mirrors—learn how to use remote mirrors for disaster recovery. This course also covers monitoring and maintaining disks, nodes, and how to troubleshoot basic cluster problems.

**5 reasons to choose HPE as your training partner**

1. Learn HPE and in-demand IT industry technologies from expert instructors.
2. Build career-advancing power skills.
3. Enjoy personalized learning journeys aligned to your company's needs.
4. Choose how you learn: in-person, virtually, or online—anytime, anywhere.
5. Sharpen your skills with access to real environments in virtual labs.

Explore our simplified purchase options, including HPE Education Learning Credits.

## Audience

System administrators who will be installing, configuring, and maintaining an HPE Ezmeral Data Fabric File and Object Store cluster environment.

## Prerequisites

Participants in this course should have:

- Basic Hadoop knowledge and intermediate Linux knowledge
- Experience using a Linux text editor such as vi
- Familiarity with the Linux command line options such as mv, cp, ssh, grep, and useradd

## Course objectives

By the end of the course, the participant should be able to:

- Demonstrate basic cluster administration skills
- Audit and prepare cluster hardware prior to installation
- Run pre-installation tests to verify performance
- Plan a service layout according to cluster configuration and business needs
- Describe the primary architectural components of an HPE Data Fabric installation (nodes, storage pools, volumes, containers, chunks, blocks)
- Use the UI installer or the manual method to install the HPE Ezmeral Data Fabric File and Object Store cluster
- Define and implement an appropriate node topology
- Define and implement an appropriate volume topology
- Set permissions and quotas for users and groups
- Set up email and alerts
- Locate and review configuration files used by the cluster
- Start and stop services
- Use Hadoop commands to perform basic functions
- Use “maprcli” commands to perform basic functions
- Use the MCS
- Assist with data ingestion
- Configure, monitor, and respond to alerts
- Detect and replace failed disks
- Detect and replace failed nodes
- Create and delete snapshots using both maprcli and the MCS
- Create and delete mirrors using both maprcli and the MCS
- Use mirrors and snapshots for data protection
- Add, remove, and upgrade ecosystem components
- Monitor and tune job performance
- Set up NFS access to the cluster

## Detailed course outline

| | | |
|---|---|---|
| **Module 1 - Introduction to HPE Ezmeral Data Fabric File and Object Store** | • Why data fabric?<br>• Key features | • Architecture basics |
| **Module 2 - Prepare for Installation** | • Security modes<br>• Configuring data fabric security and managing impersonations<br>• Plan the cluster and service layout | • Prepare and verify cluster hardware<br>• Test nodes |
| **Module 3 - Install a Data Fabric Cluster** | • Installer options<br>• Perform a manual installation | • License the cluster |
| **Module 4 - Verify and Test Cluster** | • Verify cluster status<br>• Run post install benchmark tests | • Explore the cluster structure |
| **Module 5 - Work with Volumes** | • Data fabric architecture<br>• Object store architecture and operations<br>• Node/volume topology<br>• Attributes for standard volumes | • Design a volume plan<br>• Name container sizing<br>• Data tiering |
| **Module 6 - Work with Snapshots** | • Describe how snapshots work | • Use and maintain snapshots |
| **Module 7 - Work with Mirrors** | • How mirrors work<br>• Mirroring types | • Using mirrors for disaster recovery |
| **Module 8 - Configure Users and Cluster Parameters** | • Manage users and groups<br>• Use access control expressions<br>• Accountable entities and quotas | • Configuring policy-based security<br>• Configuring object store access policies |
| **Module 9 - Configure Cluster Access** | • Access cluster data<br>• Configure virtual IP addresses | • Set up client access |
| **Module 10 - Monitor and Manage the Cluster** | • Monitor using the control system and CLI<br>• Using HPE Ezmeral Data Fabric Monitoring (Spyglass Initiative) | • Configure and respond to alarms |
| **Module 11 - Disk and Node Maintenance** | • Managing disks<br>• Managing nodes | • Configuring balancer settings |
| **Module 12 - Troubleshoot Cluster Problems** | • Basic troubleshooting | • Tools and utilities |

## Detailed lab outline

| | | |
|---|---|---|
| **Lab 0: Prepare for Lab access** | • Task 1: Prepare/access lab environment | |
| **Lab 1: Prepare for Lab access** | • No Labs | |
| **Lab 2: Prepare for Installation** | • Task 1: Plan a service layout<br>• Task 2: Audit the cluster | • Task 3: Run pre-install tests |
| **Lab 3: Install the Data Fabric** | • Task 1: Install and license a secure cluster | |
| **Lab 4: Verify and Test the Cluster** | • Task 1: Run RWSpeedTest | • Task 2: Explore the cluster |
| **Lab 5: Work with Volumes** | • Task 1: Configure node topology | • Task 2: Create volumes and set quotas |
| **Lab 6: Snapshots** | • Task 1: Work with snapshots | • Task 2: Restore data from a snapshot |
| **Lab 7: Work with Mirrors** | • Task 1: Configure local mirrors | |
| **Lab 8: Work with Object Store** | • Task 1: Creating account, user, group | • Task 2:  Creating bucket, uploading object |
| **Lab 9: Configure Cluster Settings** | • Task 1: Set up users and groups | • Task 2: Control access to the cluster |
| **Lab 10: Configure Cluster Access** | • Task 1: Configure VIPs | |
| **Lab 11: Monitor and Manage the Cluster** | • Task 1: Configure alerts | |
| **Lab 12: Disk and Node Maintenance** | • Task 1: Replace a failed disk | |
| **Lab 13: Troubleshoot Cluster Problems** | • Task 1: Troubleshooting | • Task 2: Collect logs for support |

## Learn more at

hpe.com/ww/learnbigdata

**Follow us:**



HG7E0S D.00, October 2023

Hewlett Packard Enterprise