
Tytuł szkolenia: HPE Ezmeral Data Fabric Cluster Administration

Kod szkolenia: HG7E0X

Wprowadzenie

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

Adresaci szkolenia

The audience for the basic cluster administration course is system administrators who will be creating and maintaining an HPE Ezmeral Data Fabric cluster environment.

Prerequisites

- 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
- Laptop with internet access and an installed browser

Cel szkolenia

At the end of the course, the participant is expected to acquire basic cluster administration skills, including the following.

- 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 a HPE Data Fabric installation (nodes, storage pools, volumes, containers, chunks, blocks)
- Use the UI installer to install the HPE Ezmeral Data Fabric distribution
- 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
- Create and implement a disaster recovery plan
- Add, remove and upgrade ecosystem components
- Monitor and tune job performance
- Set up NFS access to the cluster

Czas i forma szkolenia

- 21 godzin (3 dni x 7 godzin), w tym wykłady i warsztaty praktyczne.

Plan szkolenia

Lesson 0: Getting Started

- Course introduction
- Prepare HPE MyRoom
- Prepare/access your lab environment

Lesson 1: Introduction to the HPE Ezmeral Data Fabric

- Key components of HDFS
- Key components of HPE Ezmeral Data Fabric File System
- Compare and contrast HPE Ezmeral Data Fabric File System and HDFS

Lesson 2: Prepare for Installation

- Select a security mode
- Plan a service layout
- Prepare and verify cluster hardware
- Test nodes

Lesson 3: Install a HPE Ezmeral Data Fabric Cluster

- Use of Installer
- Performing a manual installation
- License the cluster

Lesson 4: Verify and Test the Cluster

- Verify cluster status
- Run post-Install benchmark tests
- Explore the cluster

Lesson 5: Work with Volumes

- Overview of volumes
- Volume placement (topology)
- Attributes for standard volumes
- Designing a volume plan
- Creating and configuring volumes
- Data tiering

Lesson 6: Work with Snapshots

- How snapshots work
- Create and schedule Snapshot
- Use and maintaining snapshots

Lesson 7: Work with Mirrors

- How mirrors work
- Configure and use local mirrors
- Remote mirrors
- Remote mirrors and disaster recovery

Lesson 8: Configure Users and Cluster Parameters

- Manage users and groups
- Access Control Expressions (ACEs)
- User and group quotas
- Configure policy based security
- Configuring topology and email

Lesson 9: Configure Cluster Access

- Accessing cluster data
- Configure virtual IP addresses
- Setup client access

Lesson 10: Monitor and Manage the Cluster

- Using the MCS and CLI
- Monitoring using data fabric features
- Configure and respond to alarms

Lesson 11: Disk and Node Maintenance

- Adding disks
- Replacing failed disks
- Remove or maintain nodes
- Adding nodes

Lesson 12: Troubleshoot Cluster Problems

- Basic troubleshooting
- Tools and utilities

Lesson 0: Welcome to Class

- Lab 0.1: Prepare HPE MyRoom
- Lab 0.2 : Prepare/access lab environment

Lesson 1: Introduction to the HPE Data Fabric

- No labs

Lesson 2: Prepare for Installation

- Lab 2.2: Plan a service layout
- Lab 2.3: Audit the cluster
- Lab 2.4: Run pre-install tests

Lesson 3: Install the Data Fabric

- Lab 3.3: Install and license a secure cluster

Lesson 4: Verify and Test the Cluster

- Lab 4.2: Run RWSpeedTest
- Lab 4.3: Explore the cluster

Lesson 5: Work with Volumes

- Lab 5.2: Configure node topology
- Lab 5.5: Create volumes and set quotas

Lesson 6: Snapshots

- Lab 6.2: Work with snapshots
- Lab 6.3: Restore data from a snapshot

Lesson 7: Work with Mirrors

- Lab 7.2: Configure local mirrors

Lesson 8: Configure Cluster Settings

- Lab 8.1: Set up users and groups
- Lab 8.2: Control access to the cluster

Lesson 9: Configure Cluster Access

- Lab 9.2: Configure VIPs

Lesson 10: Monitor and Manage the Cluster

- Lab 10.3: Configure alerts

Lesson 11: Disk and Node Maintenance

- Lab 11.2: Replace a failed disk

Lesson 12: Troubleshoot Cluster Problems

- Lab 12.1: Troubleshooting
- Lab 12.2: Collect logs for support