# Red Hat High Availability Clustering

Course code: RH436

Deploy reliable, available critical production services in a high availability clusterIn the Red Hat High Availability Clustering (RH436) course, you will learn how to provide highly available network services to a mission-critical enterprise environment through the deployment and management of shared storage and server clusters. Created for senior Linux system administrators, this 4-day course strongly emphasizes lab-based activities.You will set up a cluster of systems running the Pacemaker component of the Red Hat Enterprise Linux High-Availability Add-On, and deploy Linux-based services such as web servers and databases on that cluster. Cluster storage components from the Resilient Storage Add-On are also covered; installations and applications that require multiple cluster nodes can access the same storage simultaneously. This includes Logical Volume Manager (LVM) Shared Volume Groups, Red Hat Global File System 2 (GFS2), and Device-Mapper Multipath.This course is based on Red Hat Enterprise Linux 8.3.

## Who is the course for

Senior Linux system administrators who use high-availability clustering and fault-tolerant shared storage technologies

to maximize resiliency of production services.

## What we teach you
- Install and configure a Pacemaker-based high availability cluster.
- Create and manage highly available services.
- Troubleshoot common cluster issues.
- Work with shared storage (iSCSI) and configure multipathing.
- Implement Logical Volume Manager (LVM) in cluster-aware configurations.
- Configure GFS2 file systems on storage shared by multiple nodes.

## Required skills
- Red Hat Certified System Administrator (RHCSA) exam (EX200) and associated courses.
- Red Hat Cerfitied Engineer (RHCE) exam (EX294) and associated courses.

## Course Outline

### Creating high availability clusters

Create a basic high availability cluster.

### Managing cluster nodes and quorum

Manage node membership in the cluster and describe how it impacts cluster operation.

### Isolating malfunctioning cluster nodes

Isolate unresponsive cluster nodes to protect data and recover services and resources after a failure.

### Creating and configuring resources

Create basic resources and resource groups to provide highly available services.

### Troubleshooting high availability clusters

Identify, diagnose, and fix cluster issues.

### Automating cluster and resource deployment

Deploy a new high availability cluster and cluster resources using Ansible automation.

### Managing two-node clusters

Operate two-node clusters while identifying and avoiding issues specific to a two-node cluster configuration.

### Accessing iSCSI storage

Configure iSCSI initiators on your servers to access block-based storage devices provided by network storage arrays or

Ceph storage clusters.

**GOPAS**

### Accessing storage devices resiliently

Configure resilient access to storage devices that have multiple access paths.

### Configuring LVM in clusters

Select, configure, and manage the correct LVM configuration for use in your cluster.

### Providing storage with the GFS2 cluster file system

Use the GFS2 cluster file system to simultaneously pProvide tightly coupled shared storage that can be accessed by multiple nodes.

### Eliminating single points of failure

Identify and eliminate single points of failure in your cluster to decrease risk and increase average service availability.
Note: Course outline is subject to change with technology advances and as the nature of the underlying job evolves.

## What you need to know

### Impact on the organization

High availability clustering can improve reliability, availability, and resiliency of your mission-critical services, resulting in reduced downtime and easier hardware maintenance.
Red Hat has created this course in a way intended to benefit our customers, but each company and infrastructure is unique, and actual results or benefits may vary.

### Impact on the individual

As a result of attending this course, you will be able to create, manage, and troubleshoot highly available network services and tightly-coupled cluster storage for business-critical applications.
Students should be able to demonstrate the following skills:
Improve application uptime by using high availability clustering.

Manage storage in an high availability environment using iSCSI initiators, HA-LVM or LVM Shared Volume Groups as appropriate, and GFS2 cluster file systems.

Implement strategies to identify and eliminate single points of failure in high availability clusters.