

DO370 Enterprise Kubernetes Storage with Red Hat OpenShift Data Foundation

Teaches the essential skills required to design, implement, and manage a Red Hat OpenShift Data Foundation cluster and perform day-to-day Kubernetes storage management tasks.

Traditional storage options available to Kubernetes administrators are limited and lack flexibility and/or versatility. Red Hat OpenShift Data Foundation provides real advantages, even when it is backed by cloud storage such as AWS EBS and sophisticated on-prem legacy storage like SAN arrays. Many companies rely on third-party solutions to manage backup and disaster recovery in production. However, proper planning to implement these solutions requires knowledge of the Kubernetes CSI and OAPD APIs. This course walks the student through the recommended steps of configuring and managing storage services for container and Kubernetes services.

Course content summary

- Deploy Red Hat OpenShift Data Foundation in internal and external mode.
- Provision non-shareable block storage to applications like databases.
- Provision shareable block storage to applications like virtual machines.
- Provision shareable file storage to such applications as CI/CD pipelines and AI/ML.
- Provision shareable object storage to applications, such as AI/ML and media streaming.
- Provision storage for Red Hat OpenShift cluster services, such as monitoring and registry.
- Monitor and expand storage capacity and performance
- Attach and detach storage from an application for backup and archiving.
- Create and access volume snapshots and clones.
- Troubleshoot internal Ceph components of Red Hat OpenShift Data Foundation.
- Perform backup and restore operations using the OADP API.

Target Audience

The intended audience for this course includes:

- Cluster administrators (systems administrators, cloud administrators, cloud engineers)
- Cluster engineers (systems administrators, cloud administrators, cloud engineers)
- Site reliability engineers (SREs)

Outline for this course

- Describing Red Hat OpenShift Data Foundation deployment architectures
- Deploying OpenShift Data Foundation on Red Hat OpenShift using Internal, Converged Mode
- Configuring Red hat OpenShift Cluster Services to use OpenShift Data Foundation
- Configuring application workloads to use OpenShift Data Foundation block and file storage
- Monitoring and expanding OpenShift Data Foundation block and file storage capacity
- Troubleshooting Ceph components from OpenShift Data Foundation
- Expanding OpenShift Data Foundation block and file storage volumes
- Performing backup and restore of OpenShift Data Foundation block and file volumes
- Configuring application workloads to use OpenShift Data Foundation object storage
- Monitoring and expanding OpenShift Data Foundation object storage capacity
- Performing backup and restore of OpenShift Data Foundation object buckets
- Deploying OpenShift Data Foundation on Red Hat OpenShift using external mode

Impact on the individual

Enterprise Kubernetes Storage with Red Hat OpenShift Data Foundation teaches the essential skills required to provision and manage storage that fits the availability and performance requirements of applications, such as:

- Deploying Red Hat OpenShift Data Foundation on a Red Hat OpenShift cluster using local or cloud storage.
- Selecting and configuring storage classes based on workload requirements.
- Monitoring and proactively expanding storage capacity.
- Creating and attaching snapshots and clones of persistent volumes.