



Red Hat Storage Server Administration

In association with :  **redhat.**



Red Hat Storage Server Administration

Duration: 3 hours

Prerequisites for this exam

- Be a Red Hat Certified System Administrator (RHCSA) or have comparable work experience and skills (Red Hat Certified Engineer (RHCE) would be even better)
- Have taken Red Hat Ceph Storage Architecture and Administration (CEPH125) or have comparable work experience
- Review the Red Hat Certified Specialist in Ceph Storage Administration exam objectives
- Install Red Hat Ceph Storage Server
 - Install Red Hat Ceph Storage Server on both physical and virtual systems
 - Be familiar with the Ansible installation files for Ceph
 - Be able to install a Ceph Storage Server using Ansible
- Work with existing Red Hat Ceph Storage Server appliances
 - Be able to change a Ceph Storage Server configuration
 - Add monitor (MON) nodes and object storage device (OSD) nodes
- Configure Red Hat Ceph Storage Server
 - Configure a replicated storage pool
 - Store objects in storage pool
 - Store objects within a namespace within a storage pool
 - Create and configure erasure-coded pools
 - Create an erasure-coded pool profile with specified parameters
 - Upload a file to an erasure-coded pool
 - Change default settings in the Ceph configuration files
 - Manage Ceph authentication
 - Create a Ceph client with restricted read or write access to MONs, OSDs, pools, and namespaces
- Provide block storage with RBD
 - Create a RADOS block device image
 - Obtain information about a RADOS block device image
 - Map a RADOS block device image on a server
 - Use a RADOS block device image
 - Create an RBD snapshot
 - Create an RBD clone
 - Configure RBD mirrors
 - Deploy a RBD mirror agent

- Configure one-way RBD mirroring in pool mode
 - Configure one-way RBD mirroring in image mode
 - Check the status of the mirroring process
 - Import and export RBD images
 - Export a RADOS block device to an image file
 - Create an incremental RBD image file
 - Import a full RBD image file
 - Import a full RBD image file updated with an incremental RBD image file
- Provide object storage with RADOSGW
 - Deploy a RADOS gateway
 - Deploy a multisite RADOS gateway
 - Provide object storage using the Amazon S3 API
 - Be able to create a RADOSGW user that will use the S3 clients commands
 - Be able to upload and download objects to a RADOSGW using the S3 client commands
 - Export S3 objects using NFS
 - Provide object storage for Swift
 - Be able to create a RADOSGW user that will use the Swift interface
 - Be able to upload or download objects to a RADOSGW using Swift commands
- Provide file storage with CephFS
 - Create a Ceph file system
 - Mount a Ceph file system on a client node, persistently
 - Configure CephFS quotas
 - Create a CephFS snapshot
- Configure a CRUSH map
 - Be able to create a bucket hierarchy in a CRUSH map that can be used in an erasure profile or a replicant rule
 - Be able to remap a PG
 - Be able to remap all PGs in a pool for an optimal redistribution
- Manage and update cluster maps
 - Manage MON and OSD maps
 - Be able to monitor and change OSD storage limits for monitoring available space on an OSD
- Manage a Red Hat Ceph Storage cluster
 - Determine the general status of a Ceph cluster
 - Troubleshoot problems with OSDs and MONs
- Tune Red Hat Ceph Storage

- Specify and tune key network tuning parameters for a Ceph cluster
 - Control and manage scrubbing and deep scrubbing
 - Control and manage recovery and rebalancing processes
 - Control and manage RAM utilization against I/O performance
- Troubleshoot Red Hat Ceph Storage server problems
 - Troubleshoot client issues
 - Enable debugging mode on RADOS gateway
 - Optimize RBD client access using key tuning parameters
- Integrate Red Hat Ceph Storage with Red Hat OpenStack
 - Integrate Ceph using both Glance and Cinder
 - Identify key Glance configuration files
 - Configure Glance to use Ceph as a backend to store images in the Ceph cluster
 - Identify key Cinder configuration files
 - Configure Cinder to use Ceph RBDs for block storage backing volumes

Thank You