

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : EX267

**Title : Red Hat Certified Specialist
in OpenShift AI**

Version : DEMO

1.What is the primary role of the odh-install Operator in OpenShift AI?

- A. Manages the installation of Open Data Hub components
- B. Installs Red Hat Enterprise Linux (RHEL) nodes
- C. Deploys load balancer services
- D. Configures Persistent Volumes for AI workloads

Answer: A

Explanation:

The odh-install Operator simplifies the deployment of Open Data Hub components by automating installation processes and managing configurations necessary for OpenShift AI environments.

2.Which command validates a successful OpenShift AI installation?

- A. `oc get odh -n ai`
- B. `oc logs -f openshift-ai-installer`
- C. `oc status --ai`
- D. `oc get pods -n openshift-ai`

Answer: D

Explanation:

Running `oc get pods -n openshift-ai` displays the status of pods in the OpenShift AI namespace, helping ensure all required components are running successfully.

3.What is the purpose of kustomize during OpenShift AI installation?

- A. Managing cluster storage configurations
- B. Customizing Kubernetes resources declaratively
- C. Automating network policies
- D. Enabling high availability

Answer: B

Explanation:

Kustomize enables users to customize Kubernetes manifests declaratively without modifying the original files, making the configuration process for OpenShift AI flexible and reusable.

4.Which two components are essential in OpenShift AI's Machine Learning stack?

- A. TensorFlow Serving
- B. JupyterHub
- C. OpenShift Pipelines
- D. OpenShift Service Mesh

Answer: A, B

Explanation:

TensorFlow Serving manages model serving, ensuring efficient deployment of ML models, while JupyterHub provides interactive development environments for data scientists.

5.What prerequisite must be met for Persistent Storage in OpenShift AI?

- A. Only NFS storage is allowed
- B. A StorageClass must be configured
- C. Dynamic provisioning must be disabled

D. Static IP allocation is required

Answer: B

Explanation:

A StorageClass must be configured in OpenShift AI to dynamically provision Persistent Volumes, allowing seamless storage management for AI workloads.