

KCNA^{Q&As}

Kubernetes and Cloud Native Associate (KCNA)

Pass Linux Foundation KCNA Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.certbus.com/kcna.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Linux Foundation Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

What is the command to list all the available objects in your Kubernetes cluster?

- A. kubectl get all
- B. kubectl get api-resources
- C. kubectl api-resources
- D. kubectl get pods

Correct Answer: C

Explanation: <https://kubernetes.io/docs/reference/kubectl/cheatsheet/>

Resource types

List all supported resource types along with their shortnames, **API group**, whether they are **namespaced**, and **Kind**:

```
kubectl api-resources
```

QUESTION 2

Which kubernetes object do deployments use behind the scenes when they need to scale pods?

- A. Horizontal pod autoscaler
- B. ReplicaSets
- C. kubectl
- D. Replication controller

Correct Answer: B

Explanation: <https://kubernetes.io/docs/concepts/workloads/controllers/replicaset/>

ReplicaSet

A ReplicaSet's purpose is to maintain a stable set of replica Pods running at any given time. As such, it is often used to guarantee the availability of a specified number of identical Pods.

QUESTION 3

What command to view the kube config?

- A. kubectl view config
- B. kubectl config view
- C. kubectl get kubeconfig

Correct Answer: B

Explanation: [https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#-em-view-em-](https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#-em-view-em-view)

view

Display merged kubeconfig settings or a specified kubeconfig file.

You can use `--output jsonpath={...}` to extract specific values using a jsonpath expression.

Usage

```
$ kubectl config view
```



QUESTION 4

Which of the following is an advantage a cloud-native microservices application has over monolithic applications?

- A. Cloud-native microservices applications tend to be faster and more responsive than monolithic applications.
- B. Cloud-native microservice applications tend to be easier to troubleshoot.
- C. Cloud-native microservice applications tend to be easier to scale and perform updates on.

Correct Answer: C

Explanation: Cloud-native applications tend to be microservice base, they have individual services that can be independently scaled, updated and rolled back. This makes scaling and update operations simpler and less risky.

QUESTION 5

What is etcd used for in Kubernetes?

- A. Integration with cloud platforms
- B. Network routing for the cluster
- C. Kubernetes API security
- D. Backend object storage for the Kubernetes API

Correct Answer: D

Explanation: etcd serves as a distributed object store that backs the Kubernetes API.

QUESTION 6

What does the 'kops' acronym mean?

- A. Kubernetes Open Platform Specification
- B. Kubernetes Operations
- C. Kubernetes Operators
- D. Kubernetes Operation Policy Specification

Correct Answer: B

Explanation: <https://github.com/kubernetes/kops>

☰ README.md

kOps - Kubernetes Operations

go report A+  reference

The easiest way to get a production grade Kubernetes cluster up and running.

What is kOps?

We like to think of it as `kubect1` for clusters.

`kops` will not only help you create, destroy, upgrade and maintain production-grade, highly available, Kubernetes cluster, but it will also provision the necessary cloud infrastructure.

AWS (Amazon Web Services) is currently officially supported, with DigitalOcean, GCE, and OpenStack in beta support, and Azure and AliCloud in alpha.

Graphical user interface, text, application, email

QUESTION 7

What is the command used to scale the application?

- A. `kubect1 run`
- B. `kubect1 explain`
- C. `kubect1 scale`

Correct Answer: C

Explanation: <https://kubernetes.io/docs/reference/generated/kubect1/kubect1- commands#scale>

scale

Set a new size for a deployment, replica set, replication controller, or stateful set.

Scale also allows users to specify one or more preconditions for the scale action.

If `--current-replicas` or `--resource-version` is specified, it is validated before the scale is attempted, and it is guaranteed that the precondition holds true when the scale is sent to the server.

Usage

```
$ kubectl scale [--resource-version=version] [--current-replicas=count] --replicas=COUNT (-f FILENAME | TYPE NAME)
```

example

Scale a replica set named 'foo' to 3

```
kubectl scale --replicas=3 rs/foo
```

Scale a resource identified by type and name specified in "foo.yaml" to 3

```
kubectl scale --replicas=3 -f foo.yaml
```

If the deployment named mysql's current size is 2, scale mysql to 3

```
kubectl scale --current-replicas=2 --replicas=3 deployment/mysql
```

Scale multiple replication controllers

```
kubectl scale --replicas=5 rc/foo rc/bar rc/c
```

QUESTION 8

What is a commonly used package manager for kubernetes applications?

- A. npm
- B. apt
- C. helm
- D. kubernetes manifest

Correct Answer: C

Explanation: <https://helm.sh/>

QUESTION 9

What standard does kubelet use to communicate with the container runtime?

- A. Service Mesh Interface (SMI)
- B. CRI-O
- C. ContainerD

D. Container Runtime Interface (CRI)

Correct Answer: D

Explanation: kubelet can communicate with any runtime that supports the CRI standard.

QUESTION 10

Which part of a Kubernetes cluster is responsible for running container workloads?

- A. Worker Node
- B. kube-proxy
- C. Control plane
- D. etcd

Correct Answer: A

Explanation: Worker Nodes are responsible for executing containerized workloads.

QUESTION 11

What is the name for a service that has no clusterIp address?

- A. Headless
- B. NodePort
- C. ClusterIP
- D. LoadBalancer

Correct Answer: A

Explanation: <https://kubernetes.io/docs/concepts/services-networking/service/#headless-services>

Headless Services

Sometimes you don't need load-balancing and a single Service IP. In this case, you can create what are termed "headless" Services, by explicitly specifying "None" for the cluster IP (`.spec.clusterIP`).

You can use a headless Service to interface with other service discovery mechanisms, without being tied to Kubernetes' implementation.

For headless Services , a cluster IP is not allocated, kube-proxy does not handle these Services, and there is no load balancing or proxying done by the platform for them. How DNS is automatically configured depends on whether the Service has selectors defined:

QUESTION 12

Which of the following provides cloud-native storage orchestration?

- A. Cloud Provider Specific storage (EBS, EFS, Cloud Storage)
- B. Cloud Storage
- C. Storage IO

Correct Answer: A

Explanation: <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#types-of-persistent-volumes>

Types of Persistent Volumes

PersistentVolume types are implemented as plugins. Kubernetes currently supports the following plugins:

- `awsElasticBlockStore` - AWS Elastic Block Store (EBS)
- `azureDisk` - Azure Disk
- `azureFile` - Azure File
- `cephfs` - CephFS volume
- `csi` - Container Storage Interface (CSI)
- `fc` - Fibre Channel (FC) storage
- `gcePersistentDisk` - GCE Persistent Disk
- `glusterfs` - Glusterfs volume
- `hostPath` - HostPath volume (for single node testing only; WILL NOT WORK in a multi-node cluster; consider using `local` volume instead)
- `iscsi` - iSCSI (SCSI over IP) storage
- `local` - local storage devices mounted on nodes.
- `nfs` - Network File System (NFS) storage
- `portworxVolume` - Portworx volume
- `rbd` - Rados Block Device (RBD) volume
- `vsphereVolume` - vSphere VMDK volume

The following types of PersistentVolume are deprecated. This means that support is still available but will be removed in a future Kubernetes release.

- `cinder` - Cinder (OpenStack block storage) (**deprecated** in v1.18)

QUESTION 13

What is scheduling in Kubernetes

- A. Determining when to execute a cron-job
- B. Assigning pods to nodes
- C. Joining a new nodes to the clusters

D. Setting a time for automated tasks

Correct Answer: B

Explanation: <https://kubernetes.io/docs/concepts/scheduling-eviction/>

Scheduling

- [Kubernetes Scheduler](#)
 - [Assigning Pods to Nodes](#)
 - [Pod Overhead](#)
 - [Taints and Tolerations](#)
 - [Scheduling Framework](#)
 - [Scheduler Performance Tuning](#)
 - [Resource Bin Packing for Extended Resources](#)
-

QUESTION 14

What kind of limitation cgroups allows?

- A. Prioritization
- B. Resource limiting
- C. Accounting
- D. None of the options
- E. Control
- F. Server cpu and memory

Correct Answer: ABCE

QUESTION 15

Which of the following best describes a cloud-native app?

- A. An application where all logic is coded into a single large binary.

- B. An application that publishes an HTTPS web front-end.
- C. An application that takes advantages of cloud computing frameworks and their loosely coupled cloud services.
- D. An application that leverages services that are native to public cloud platforms such as Azure, GCP, and/or AWS.

Correct Answer: C

Explanation: Cloud-native apps leverage cloud computing frameworks and tend to be microservices based, where individual components of the app are coded as individual.

[KCNA Practice Test](#)

[KCNA Study Guide](#)

[KCNA Braindumps](#)