



*Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ---- Guaranteed.*



CKA Dumps  
CKA Braindumps  
CKA Real Questions  
CKA Practice Test  
CKA Actual Questions



[killexams.com](https://killexams.com)

**CNCF**

**CKA**

*Certified Kubernetes Administrator 2025*

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/CKA>



**Question: 476**

You need to monitor the health of your Kubernetes cluster continuously. Which tool would you use to set up monitoring and alerts?

- A. Grafana
- B. Prometheus
- C. Fluentd
- D. Calico

Answer: B

Explanation: Prometheus is widely used for monitoring and alerting in Kubernetes environments, collecting metrics and providing a robust query language.

**Question: 477**

The ----- service provides a way to route external traffic to services within a Kubernetes cluster.

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

Answer: D

Explanation: An Ingress resource is used to manage external access to services, typically HTTP, within a Kubernetes cluster.

**Question: 478**

You need to define a Role that allows a user to create and delete ConfigMaps in the development namespace. What YAML snippet would you use?

- A. `apiVersion: rbac.authorization.k8s.io/v1kind: Rolemetadata:name: configmap-editornamespace: developmentrules:apiGroups: ["*"]resources: ["configmaps"]verbs: ["create", "delete"]`
- B. `apiVersion: rbac.authorization.k8s.io/v1kind: Rolemetadata:name: configmap-editornamespace: developmentrules:apiGroups: [""]resources: ["configmaps"]verbs: ["create", "delete"]`
- C. `apiVersion: rbac.authorization.k8s.io/v1`  
`kind: Role`

```
metadata:  
name: configmap-editor  
namespace: development  
rules:  
  apiGroups: [""]  
  resources: ["*"]  
  verbs: ["create", "delete"]  
D. apiVersion: rbac.authorization.k8s.io/v1  
kind: Role  
metadata:  
name: configmap-editor  
namespace: development  
rules:  
  apiGroups: ["configmaps"]  
  resources: ["configmaps"]  
  verbs: ["create", "delete"]
```

Answer: B

Explanation: This YAML configuration correctly defines a Role that allows creating and deleting ConfigMaps in the development namespace.

### Question: 479

You want to ensure that a Pod can only be scheduled on nodes with a specific resource available. What feature should you use?

- A. Resource Limits
- B. Resource Requests
- C. Node Affinity
- D. Node Selector

Answer: B

Explanation: Resource Requests specify the minimum amount of resources required for the Pod, influencing the scheduler to place it on a suitable node.

### Question: 480

You are troubleshooting a secret that is not being correctly injected into a pod. What command would you use to check the secret's details?

- A. Both A and B
- B. `kubectl describe secret <secret-name>`
- C. `kubectl logs <pod-name>`
- D. `kubectl get secret <secret-name> -o yaml`

Answer: A

Explanation: Both commands provide detailed information about the secret, helping to diagnose issues with its injection into the pod.

**Question: 481**

You have a multi-node Kubernetes cluster and need to ensure that all kubelet services are running correctly. What command would you use to check the status of the kubelet service on a node?

- A. `kubectl get nodes`
- B. `systemctl status kubelet`
- C. `journalctl -u kubelet`
- D. `kubectl describe node &lt;node-name>`

Answer: B

Explanation: The command `systemctl status kubelet` checks the status of the kubelet service directly on the node, providing insights into whether it is active and any errors that may have occurred.

**Question: 482**

You are configuring a dynamic provisioning for a storage class. Which parameter would you set to specify the disk type for an AWS EBS volume?

- A. `type`
- B. `volumeType`
- C. `diskType`
- D. `storageType`

Answer: B

Explanation: The `volumeType` parameter specifies the type of AWS EBS volume to be provisioned (e.g., `gp2`, `io1`).

**Question: 483**

You want to ensure that your control plane nodes are highly available. Which configuration should you implement?

- A. Deploy a single control plane node with a static IP
- B. Deploy multiple control plane nodes with an etcd cluster
- C. Use an external etcd cluster
- D. Use a single-node etcd instance

Answer: B

Explanation: For high availability, you should deploy multiple control plane nodes along with a

distributed etcd cluster to prevent a single point of failure.

**Question: 484**

To achieve high availability for the controller manager, you decide to run multiple instances. What must you ensure about the API server in this configuration?

- A. It must be run on a single node only.
- B. It must support leader election among the controller managers.
- C. No changes are needed for the API server.
- D. It must be configured to replicate its state across nodes.

Answer: B

Explanation: When running multiple instances of the controller manager, it is crucial to enable leader election to ensure that only one instance actively manages the cluster at any given time.

**Question: 485**

You need to ensure that your Kubernetes nodes can be automatically updated with the latest security patches. Which feature should you consider?

- A. Node Auto-Scaling
- B. Cluster Autoscaler
- C. Manual Node Management
- D. Managed Kubernetes Services

Answer: D

Explanation: Managed Kubernetes services often include automated updates for security patches, ensuring nodes are kept up to date without manual intervention.

**Question: 486**

You need to upgrade a Kubernetes cluster from version 1.20 to 1.22. What is the recommended upgrade path?

- A. Upgrade to 1.21 first, then to 1.22.
- B. Upgrade directly to 1.22.
- C. Upgrade to the latest patch of 1.20, then to 1.22.
- D. Upgrade to 1.19 first, then to 1.21.

Answer: A

Explanation: Kubernetes requires upgrades to be performed sequentially between minor versions. Thus, you must upgrade to 1.21 before proceeding to 1.22.

**Question: 487**

In a scenario where you need to configure Kubernetes with multiple etcd clusters for disaster recovery,

which of the following configurations would be ideal?

- A. All etcd clusters in the same data center.
- B. etcd clusters running on the same node as the API server.
- C. etcd clusters located in different geographical regions.
- D. A single etcd cluster with read replicas in different zones.

Answer: C

Explanation: Placing etcd clusters in different geographical regions provides redundancy and ensures data availability even if one region faces a failure.

### Question: 488

You need to set up a cluster with a specific API server advertising address. What parameter should you use when initializing the cluster with kubeadm?

- A. `--advertise-address=<tip>`
- B. `--apiserver-override=<tip>`
- C. `--advertise-ip=<tip>`
- D. `--bind-address=<tip>`

Answer: A

Explanation: The `--advertise-address` parameter specifies the IP address that the kube-apiserver will advertise to the other nodes in the cluster.

### Question: 489

You want to ensure that a pod cannot be scheduled on nodes with the label `env=production`. Which configuration should you use?

- A. `nodeSelector:`  
`env: production`
- B. `tolerations:- key: envoperator: Equalvalue: production`
- C. `taints:`  
`- key: env`  
`value: production`  
`effect: NoSchedule`
- D. `nodeAffinity:requiredDuringSchedulingIgnoredDuringExecution:nodeSelectorTerms:- matchExpressions:- key: envoperator: NotInvalues:- production`

Answer: D

Explanation: Using NotIn in matchExpressions for node affinity prevents the pod from being scheduled on nodes labeled with env=production.

**Question: 490**

What is the significance of the restartPolicy field in a Pod specification?

- A. It determines how many replicas to run.
- B. It defines the resource limits for the Pod.
- C. It specifies the node on which the Pod should run.
- D. It controls the behavior of the Pod when it fails.

Answer: D

Explanation: The restartPolicy field controls the behavior of the Pod when it fails, determining whether it should be restarted or not.

**Question: 491**

You are configuring a Kubernetes cluster with kubeadm and need to ensure that the API server can handle high availability. What is the minimum number of API server instances you must deploy?

- A. 1
- B. 2
- C. 3
- D. 5

Answer: C

Explanation: To achieve high availability, you should deploy at least three instances of the API server, ensuring quorum can be maintained in etcd.

**Question: 492**

Which command would you use to view the permissions associated with a specific Role in a namespace?

- A. `kubectl describe role &lt;role-name> -n &lt;namespace>`
- B. `kubectl get role &lt;role-name> -n &lt;namespace>`
- C. Both A and B
- D. `kubectl auth can-i --list --as &lt;user> -n &lt;namespace>`

Answer: C

Explanation: Both commands provide insights into the permissions defined in a Role and how they relate to a specific user.

**Question: 493**

You suspect a resource leak in your cluster. What command can you use to assess resource usage across all namespaces?

- A. `kubectl top pods --all-namespaces`
- B. `kubectl get all --all-namespaces`
- C. `kubectl describe nodes`
- D. `kubectl get resourcequotas --all-namespaces`

Answer: A

Explanation: The `kubectl top pods --all-namespaces` command provides metrics about resource usage for pods, helping to identify potential leaks.

#### Question: 494

You are investigating a memory leak in a pod. What command would you use to check the resource usage of all containers in that pod?

- A. `kubectl get pod <pod-name> -o jsonpath='{.status.containerStatuses[*].usage}'`
- B. `kubectl top pod <pod-name>`
- C. `kubectl describe pod <pod-name>`
- D. `kubectl logs <pod-name>`

Answer: B

Explanation: The `kubectl top pod` command provides real-time metrics on CPU and memory usage for the specified pod's containers.

#### Question: 495

A pod is expected to communicate with a service that uses a headless configuration. Which command can you use to verify the endpoints of the service named `my-headless-service` in the default namespace?

- A. `kubectl get endpoints my-headless-service -n default`
- B. `kubectl get service my-headless-service -n default`
- C. `kubectl describe service my-headless-service -n default`
- D. `kubectl logs my-headless-service -n default`

Answer: A

Explanation: The `kubectl get endpoints` command retrieves the endpoints associated with the headless service, showing the actual pod IPs that are part of the service.

#### Question: 496



What command would you use to check the status of etcd in your Kubernetes cluster?

- A. etcdctl endpoint health
- B. kubectl get pods -n kube-system
- C. kubectl logs -n kube-system etcd-&lt;node-name>
- D. Both A and C

Answer: D

Explanation: Both commands provide information about the health and status of the etcd cluster, which is crucial for Kubernetes functionality.

### Question: 497

The \_\_\_\_\_ is responsible for managing the lifecycle of pods and ensuring the desired state is achieved in a Kubernetes cluster.

- A. kubelet
- B. kube-proxy
- C. API server
- D. etcd

Answer: A

Explanation: The kubelet manages the lifecycle of pods on a node and communicates with the API server to report status.

### Question: 498

You need to expose a service internally within the cluster using a ClusterIP service. What YAML snippet correctly defines this service for a deployment named my-app?

- A. 

```
apiVersion: v1
kind: Service
metadata:
  name: my-app
spec:
  type: ClusterIP
selector:
  app: my-app
ports:
  - port: 80
    targetPort: 8080
```
- B. 

```
apiVersion: v1
kind: Service
metadata:
  name: my-app
spec:
  type: NodePort
selector:
  app: my-app
ports:
  - port: 80
    targetPort: 8080
```
- C. 

```
apiVersion: v1
```

```
kind: Service
metadata:
name: my-app
spec:
type: LoadBalancer
selector:
app: my-app
ports:
port: 80
targetPort: 8080
D. apiVersion: v1kind: Servicemetadata:name: my-appspec:type: ClusterIPselector:app: my-
appports:port: 80targetPort: 8080
```

Answer: D

Explanation: This YAML snippet correctly defines a ClusterIP service that routes traffic from port 80 to the target port 8080 of the pods labeled app=my-app.

### Question: 499

Fill in the blank: The \_\_\_\_\_ component is responsible for ensuring that the state of the cluster matches the desired state defined in the deployment specifications.

Answer: Controller Manager

Explanation: The Controller Manager continuously monitors and adjusts the state of the cluster to match the desired state defined in deployments.

### Question: 500

Fill in the blank: To ensure that your etcd cluster can recover from failures, you should regularly back up the etcd data using the command \_\_\_\_\_.

Answer: etcdctl snapshot save

Explanation: The etcdctl snapshot save command is used to create backups of the etcd data, which can be restored in case of failures.

## Question: 501

A Kubernetes worker node, named wk8s-node-0 is in state NotReady. Investigate why this is the case, and perform any appropriate steps to bring the node to a Ready state, ensuring that any changes are made permanent.


You can ssh to the failed node using:

```
[[email protected]] $ | ssh Wk8s-node-0
```

You can assume elevated privileges on the node with the following command:

```
[[email protected]] $ | sudo Ci
```

**Answer:** solution



The screenshot shows a web terminal interface with a dark background. At the top, there is a blue header bar with "Readme" and "Web Terminal" on the left, and "THE LINUX FOUNDATION" on the right. The terminal content shows the following commands and output:

```
root@node-1:~# kubectl config use-context wk8s
Switched to context "wk8s".
root@node-1:~# k get nodes
NAME           STATUS    ROLES    AGE   VERSION
wk8s-master-0  Ready    master   77d   v1.18.2
wk8s-node-0    NotReady <none>   77d   v1.18.2
wk8s-node-1    Ready    <none>   77d   v1.18.2
root@node-1:~# ssh wk8s-node-0
```

The terminal also features a large, semi-transparent watermark in the center that reads "KILL EXAMS" with a stylized "KE" logo above it.

```
wk8s-node-0    NotReady    <none>    77d    v1.18.2
wk8s-node-1    Ready       <none>    77d    v1.18.2
root@node-1:~# ssh wk8s-node-0
Welcome to Ubuntu 16.04.6 LTS (GNU/Linux 4.4.0-1109-aws x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

 * Are you ready for Kubernetes 1.19? It's nearly here! Try RC3 with
   sudo snap install microk8s --channel=1.19/candidate --classic

   https://microk8s.io/ has docs and details.

4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@wk8s-node-0:~$ sudo -i
root@wk8s-node-0:~# systemctl restart kubelet
root@wk8s-node-0:~# systemctl enable kubelet
```

```
https://microk8s.io/ has docs and details.

4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@wk8s-node-0:~$ sudo -i
root@wk8s-node-0:~# systemctl restart kubelet
root@wk8s-node-0:~# systemctl enable kubelet
Created symlink from /etc/systemd/system/multi-user.target.wants/kubelet.service to /lib/systemd/system/kubelet.service.
root@wk8s-node-0:~# exit
logout
student@wk8s-node-0:~$ exit
logout
Connection to 10.250.5.34 closed.
root@node-1:~# k get nodes
NAME           STATUS    ROLES    AGE   VERSION
wk8s-master-0  Ready    master   77d   v1.18.2
wk8s-node-0    Ready    <none>   77d   v1.18.2
wk8s-node-1    Ready    <none>   77d   v1.18.2
root@node-1:~#
```

Question: 502

CORRECT TEXT

Create a pod as follows:

Name: non-persistent-redis


container Image: redis

Volume with name: cache-control

Mount path: /data/redis

The pod should launch in the staging namespace and the volume must not be persistent.

**Answer:** solution



The screenshot shows a web terminal window with a dark background. At the top, there is a blue header bar with the text "THE LINUX FOUNDATION" on the right and "Readme" and "Web Terminal" on the left. The terminal content shows a root user at a node-1 prompt, followed by three lines of text: "root@node-1:~#", "root@node-1:~#", and "root@node-1:~# vim volume.yaml". The terminal is overlaid with a large, semi-transparent watermark that reads "KILL EXAMS" in a stylized font. A vertical scrollbar is visible on the right side of the terminal window.

```
apiVersion: v1
kind: Pod
metadata:
  name: non-persistent-redis
  namespace: staging
spec:
  containers:
  - name: redis
    image: redis
    volumeMounts:
    - name: cache-control
      mountPath: /data/redis
  volumes:
  - name: cache-control
    emptyDir: {}
```

```
~
~
~
~
~
~
~
~
~
~
:w
```

```
root@node-1:~#
root@node-1:~#
root@node-1:~# vim volume.yaml
root@node-1:~# k create -f volume.yaml
pod/non-persistent-redis created
root@node-1:~# k get po -n staging
NAME                READY   STATUS    RESTARTS   AGE
non-persistent-redis 1/1     Running   0           6s
root@node-1:~#
```



# KILLEXAMS.COM

*Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:*



**Actual Exam Questions:** *Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

**Exam Dumps:** *Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

**Practice Tests:** *Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

**Guaranteed Success:** *Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

**Updated Content:** *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

**Technical Support:** *Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*