

# CKS Exam Dumps 100% Guarantee You Get CKS Exam - Getcertkey



2026 Latest Getcertkey CKS PDF Dumps and CKS Exam Engine Free Share: [https://drive.google.com/open?id=1raTQgUf9X6L\\_XyMhsjNoguRzHmXkipX](https://drive.google.com/open?id=1raTQgUf9X6L_XyMhsjNoguRzHmXkipX)

In order to help customers, who are willing to buy our CKS test torrent, make good use of time and accumulate the knowledge, Our company have been trying our best to reform and update our Certified Kubernetes Security Specialist (CKS) exam tool. “Quality First, Credibility First, and Service First” is our company’s purpose, we deeply hope our CKS Study Materials can bring benefits and profits for our customers. So we have been persisting in updating our CKS test torrent and trying our best to provide customers with the latest study materials.

If you really want a learning product to help you, our CKS study materials are definitely your best choice, you can't find a product more perfect than it. And according to the data, our CKS exam questions have really helped a lot of people pass the exam and get their dreaming CKS Certification. As the quality of our CKS practice questions is high, the pass rate of our worthy customers is also high as 98% to 100%. It is hard to find in the market.

>> **CKS Reliable Exam Braindumps** <<

**TOP CKS Reliable Exam Braindumps 100% Pass | Valid Linux Foundation Latest Certified Kubernetes Security Specialist (CKS) Dumps Sheet Pass for**

**sure**

With the CKS exam, you will harvest many points of theories that others ignore and can offer strong prove for managers. So the CKS exam is a great beginning. However, since there was lots of competition in this industry, the smartest way to win the battle is improving the quality of our CKS Learning Materials, which we did a great job. With passing rate up to 98 to 100 percent, you will get through the CKS exam with ease.

## **Linux Foundation Certified Kubernetes Security Specialist (CKS) Sample Questions (Q24-Q29):**

### **NEW QUESTION # 24**

You are working on a Kubernetes cluster that is deployed on a Cloud provider. You need to ensure that the Kubernetes nodes are hardened according to security best practices. Implement a solution that automatically scans the nodes for vulnerabilities and applies necessary security updates.

**Answer:**

Explanation:

Solution (Step by Step):

1. Choose a vulnerability scanning tool. There are many open-source and commercial tools available, such as Trivy, Anchore, and Clair.
2. Deploy the scanning tool in your cluster- This can be done by deploying the tool as a Daemonset, so that it runs on every node.
3. Configure the scanning tool to scan the nodes regularly. This can be done using a CronJob or by configuring the tool to run on a schedule.
4. Integrate the scanning tool with a security information and event management (SIEM) system. This will allow you to centralize security logs and alerts.
5. Configure automatic updates for your nodes. This can be done using your Cloud providers tools or by using a tool like Kured. Important Considerations: False Positives: Tune the scanning tool to minimize false positives. Remediation: Have a process in place for remediating vulnerabilities that are discovered. Node Updates: Ensure that node updates do not disrupt your applications.

### **NEW QUESTION # 25**

**SIMULATION**

Fix all issues via configuration and restart the affected components to ensure the new setting takes effect.

Fix all of the following violations that were found against the API server:- a. Ensure the --authorization-mode argument includes RBAC b. Ensure the --authorization-mode argument includes Node c. Ensure that the --profiling argument is set to false Fix all of the following violations that were found against the Kubelet:- a. Ensure the --anonymous-auth argument is set to false.

b. Ensure that the --authorization-mode argument is set to Webhook.

Fix all of the following violations that were found against the ETCD:-

a. Ensure that the --auto-tls argument is not set to true

Hint: Take the use of Tool Kube-Bench

**Answer:**

Explanation:

API server:

Ensure the --authorization-mode argument includes RBAC

Turn on Role Based Access Control.

Role Based Access Control (RBAC) allows fine-grained control over the operations that different entities can perform on different objects in the cluster. It is recommended to use the RBAC authorization mode.

Fix - Buildtime

Kubernetes

apiVersion: v1

kind: Pod

metadata:

creationTimestamp: null

labels:

component: kube-apiserver

tier: control-plane

name: kube-apiserver

```
namespace: kube-system
spec:
  containers:
  - command:
  + - kube-apiserver
  + - --authorization-mode=RBAC,Node
  image: gcr.io/google_containers/kube-apiserver-amd64:v1.6.0
  livenessProbe:
  failureThreshold: 8
  httpGet:
  host: 127.0.0.1
  path: /healthz
  port: 6443
  scheme: HTTPS
  initialDelaySeconds: 15
  timeoutSeconds: 15
  name: kube-apiserver-should-pass
  resources:
  requests:
  cpu: 250m
  volumeMounts:
  - mountPath: /etc/kubernetes/
  name: k8s
  readOnly: true
  - mountPath: /etc/ssl/certs
  name: certs
  - mountPath: /etc/pki
  name: pki
  hostNetwork: true
  volumes:
  - hostPath:
  path: /etc/kubernetes
  name: k8s
  - hostPath:
  path: /etc/ssl/certs
  name: certs
  - hostPath:
  path: /etc/pki
  name: pki
```

Ensure the --authorization-mode argument includes Node

Remediation: Edit the API server pod specification file /etc/kubernetes/manifests/kube-apiserver.yaml on the master node and set the --authorization-mode parameter to a value that includes Node.

```
--authorization-mode=Node,RBAC
```

Audit:

```
/bin/ps -ef | grep kube-apiserver | grep -v grep
```

Expected result:

```
'Node,RBAC' has 'Node'
```

Ensure that the --profiling argument is set to false

Remediation: Edit the API server pod specification file /etc/kubernetes/manifests/kube-apiserver.yaml on the master node and set the below parameter.

```
--profiling=false
```

Audit:

```
/bin/ps -ef | grep kube-apiserver | grep -v grep
```

Expected result:

```
'false' is equal to 'false'
```

Fix all of the following violations that were found against the Kubelet:- Ensure the --anonymous-auth argument is set to false.

Remediation: If using a Kubelet config file, edit the file to set authentication: anonymous: enabled to false. If using executable arguments, edit the kubelet service file /etc/systemd/system/kubelet.service.d/10-kubeadm.conf on each worker node and set the below parameter in KUBELET\_SYSTEM\_PODS\_ARGS variable.

```
--anonymous-auth=false
```

Based on your system, restart the kubelet service. For example:



## NEW QUESTION # 26

Create a User named john, create the CSR Request, fetch the certificate of the user after approving it.

Create a Role name john-role to list secrets, pods in namespace john

Finally, Create a RoleBinding named john-role-binding to attach the newly created role john-role to the user john in the namespace john. To Verify: Use the kubectl auth CLI command to verify the permissions.

### Answer:

Explanation:

se kubectl to create a CSR and approve it.

Get the list of CSRs:

```
kubectl get csr
```

Approve the CSR:

```
kubectl certificate approve myuser
```

Get the certificate

Retrieve the certificate from the CSR:

```
kubectl get csr/myuser -o yaml
```

here are the role and role-binding to give john permission to create NEW\_CRD resource:

```
kubectl apply -f roleBindingJohn.yaml --as=john
```

```
rolebinding.rbac.authorization.k8s.io/john_external-rosource-rb created kind: RoleBinding apiVersion: rbac.authorization.k8s.io/v1 metadata:
```

```
name: john_crd
```

```
namespace: development-john
```

```
subjects:
```

```
- kind: User
```

```
name: john
```

```
apiGroup: rbac.authorization.k8s.io
```

```
roleRef:
```

```
kind: ClusterRole
```

```
name: crd-creation
```

```
kind: ClusterRole
```

```
apiVersion: rbac.authorization.k8s.io/v1
```

```
metadata:
```

```
name: crd-creation
```

```
rules:
```

```
- apiGroups: ["kubernetes-client.io/v1"]
```

```
resources: ["NEW_CRD"]
```

```
verbs: ["create, list, get"]
```

## NEW QUESTION # 27

### SIMULATION

Analyze and edit the given Dockerfile

```
FROM ubuntu:latest
```

```
RUN apt-get update -y
```

```
RUN apt-install nginx -y
```

```
COPY entrypoint.sh /
```

```
ENTRYPOINT ["/entrypoint.sh"]
```

```
USER ROOT
```

Fixing two instructions present in the file being prominent security best practice issues Analyze and edit the deployment manifest file

```
apiVersion: v1 kind: Pod metadata:
```

```
name: security-context-demo-2
```

```
spec:
```

```
securityContext:
```

```
runAsUser: 1000
```

```
containers:
```

```
- name: sec-ctx-demo-2
```

```
image: gcr.io/google-samples/node-hello:1.0
```

```
securityContext:
```

```
runAsUser: 0
privileged: True
allowPrivilegeEscalation: false
Fixing two fields present in the file being prominent security best practice issues Don't add or remove configuration settings; only
modify the existing configuration settings Whenever you need an unprivileged user for any of the tasks, use user test-user with the
user id 5487
```

**Answer:**

```
Explanation:
FROM debian:latest
MAINTAINER k@bogotobogo.com
# 1 - RUN
RUN apt-get update && DEBIAN_FRONTEND=noninteractive apt-get install -yq apt-utils RUN
DEBIAN_FRONTEND=noninteractive apt-get install -yq htop RUN apt-get clean
# 2 - CMD
#CMD ["htop"]
#CMD ["ls", "-l"]
# 3 - WORKDIR and ENV
WORKDIR /root
ENV DZ version1
$ docker image build -t bogodevops/demo .
Sending build context to Docker daemon 3.072kB
Step 1/7 : FROM debian:latest
---> be2868bebaba
Step 2/7 : MAINTAINER k@bogotobogo.com
---> Using cache
---> e2eef476b3fd
Step 3/7 : RUN apt-get update && DEBIAN_FRONTEND=noninteractive apt-get install -yq apt-utils
---> Using cache
---> 32fd044c1356
Step 4/7 : RUN DEBIAN_FRONTEND=noninteractive apt-get install -yq htop
---> Using cache
---> 0a5b514a209e
Step 5/7 : RUN apt-get clean
---> Using cache
---> 5d1578a47c17
Step 6/7 : WORKDIR /root
---> Using cache
---> 6b1c70e87675
Step 7/7 : ENV DZ version1
---> Using cache
---> cd195168c5c7
Successfully built cd195168c5c7
Successfully tagged bogodevops/demo:latest
```

**NEW QUESTION # 28**

Context: Cluster: prod Master node: master1 Worker node: worker1

You can switch the cluster/configuration context using the following command:

```
[desk@cli] $ kubectl config use-context prod
```

Task: Analyse and edit the given Dockerfile (based on the ubuntu:18:04 image) /home/cert\_masters/Dockerfile fixing two instructions present in the file being prominent security/best-practice issues.

Analyse and edit the given manifest file /home/cert\_masters/mydeployment.yaml fixing two fields present in the file being prominent security/best-practice issues.

Note: Don't add or remove configuration settings; only modify the existing configuration settings, so that two configuration settings each are no longer security/best-practice concerns. Should you need an unprivileged user for any of the tasks, use user nobody with user id 65535

**Answer:**

Explanation:

1. For Dockerfile: Fix the image version & user name in Dockerfile 2. For mydeployment.yaml : Fix security contexts Explanation

```
[desk@cli] $ vim /home/cert_masters/Dockerfile
```

```
FROM ubuntu:latest # Remove this
```

```
FROM ubuntu:18.04 # Add this
```

```
USER root # Remove this
```

```
USER nobody # Add this
```

```
RUN apt get install -y lsof=4.72 wget=1.17.1 nginx=4.2
```

```
ENV ENVIRONMENT=testing
```

```
USER root # Remove this
```

```
USER nobody # Add this
```

```
CMD ["nginx -d"]
```

```
[desk@cli] $ vim /home/cert_masters/mydeployment.yaml
```

```
apiVersion: apps/v1
```

```
kind: Deployment
```

```
metadata:
```

```
creationTimestamp: null
```

```
labels:
```

```
app: kafka
```

```
name: kafka
```

```
spec:
```

```
replicas: 1
```

```
selector:
```

```
matchLabels:
```

```
app: kafka
```

```
strategy: {}
```

```
template:
```

```
metadata:
```

```
creationTimestamp: null
```

```
labels:
```

```
app: kafka
```

```
spec:
```

```
containers:
```

```
- image: bitnami/kafka
```

```
name: kafka
```

```
volumeMounts:
```

```
- name: kafka-vol
```

```
mountPath: /var/lib/kafka
```

```
securityContext:
```

```
{ "capabilities": { "add": ["NET_ADMIN"], "drop": ["all"] }, "privileged": True, "readOnlyRootFilesystem": False, "runAsUser": 65535 } #
```

```
Delete This
```

```
{ "capabilities": { "add": ["NET_ADMIN"], "drop": ["all"] }, "privileged": False, "readOnlyRootFilesystem": True, "runAsUser": 65535 } #
```

```
Add This resources: {} volumes:
```

```
- name: kafka-vol
```

```
emptyDir: {}
```

```
status: {}
```

```
Pictorial View: [desk@cli] $ vim /home/cert_masters/mydeployment.yaml
```

## NEW QUESTION # 29

.....

In order to facilitate the user's offline reading, the CKS study braindumps can better use the time of debris to learn. Our CKS study braindumps can be very good to meet user demand in this respect, allow the user to read and write in a good environment continuously consolidate what they learned. Our CKS prep guide has high quality. So there is all effective and central practice for you to prepare for your test. With our professional ability, we can accord to the necessary testing points to edit CKS Exam Questions. It points to the exam heart to solve your difficulty. So high quality materials can help you to pass your exam effectively, make you feel easy, to achieve your goal.

**Latest CKS Dumps Sheet:** [https://www.getcertkey.com/CKS\\_braindumps.html](https://www.getcertkey.com/CKS_braindumps.html)

