

Brilliant CKS Guide Materials: Certified Kubernetes Security Specialist (CKS) Display First-class Exam Braindumps - PassReview



BONUS!!! Download part of PassReview CKS dumps for free: https://drive.google.com/open?id=1fqbc0IMQCPwvYfR3Th_b35BU28pxjYGJ

We promise you will pass the exam and obtain the Certified Kubernetes Security Specialist (CKS) certificate successfully with our help of CKS exam questions. According to recent survey of our previous customers, 99% of them can achieve their goals, so believe that we can be the helping hand to help you achieve your ultimate goal. Besides we have high-quality CKS test guide for managing the development of new knowledge, thus ensuring you will grasp every study points in a well-rounded way. On the other hand, if you fail to pass the exam with our CKS Exam Questions unfortunately, you can receive a full refund only by presenting your transcript. At the same time, if you want to continue learning, our CKS test guide will still provide free updates to you and you can have a discount more than one year. Finally our refund process is very simple. If you have any question about Certified Kubernetes Security Specialist (CKS) study question, please contact us immediately.

The CKS certification is an essential step for security professionals who want to deepen their knowledge and skills in the Kubernetes environment. It provides comprehensive coverage of Kubernetes security topics and validates the candidate's ability to secure Kubernetes clusters and containerized applications against cyber threats. Candidates who pass the CKS Certification Exam demonstrate their expertise in securing Kubernetes applications and stand out from their peers in a rapidly evolving Kubernetes ecosystem.

>> Valid Dumps CKS Free <<

CKS Actual Torrent - CKS Pass-King Materials & CKS Actual Exam

The software version of the CKS study materials is very practical. This version has helped a lot of customers pass their exam successfully in a short time. The most important function of the software version is to help all customers simulate the real examination environment. If you choose the software version of the CKS Study Materials from our company as your study tool, you can have the right to feel the real examination environment. In addition, the software version is not limited to the number of the computer.

Linux Foundation Certified Kubernetes Security Specialist (CKS) Sample Questions (Q54-Q59):

NEW QUESTION # 54
SIMULATION

You must connect to the correct host . Failure to do so may

result in a zero score.

```
[candidate@base] $ ssh cks000023
```

Task

Analyze and edit the Dockerfile located at `/home/candidate/subtle-bee/build/Dockerfile`, fixing one instruction present in the file that is a prominent security/best-practice issue.

Do not add or remove instructions; only modify the one existing instruction with a security/best-practice concern.

Do not build the Dockerfile, Failure to do so may result in running out of storage and a zero score.

Analyze and edit the given manifest file `/home/candidate/subtle-bee/deployment.yaml`, fixing one fields present in the file that are a prominent security/best-practice issue.

Do not add or remove fields; only modify the one existing field with a security/best-practice concern.

Should you need an unprivileged user for any of the tasks, use user nobody with user ID 65535.

Answer:

Explanation:

See the Explanation below for complete solution

Explanation:

0) Connect to the correct host

```
ssh cks000023
```

```
sudo -i
```

PART A - Fix ONE prominent Dockerfile security/best-practice issue

1) Open the Dockerfile

```
vi /home/candidate/subtle-bee/build/Dockerfile
```

2) Find the "most obvious" security/best-practice problem and modify ONLY THAT ONE instruction Use `/` search in `vi` to quickly find candidates:

Candidate 1 (very common): USER root (or no USER but a USER 0)

Search:

```
/USER
```

If you see:

```
USER root
```

Change that single instruction to:

```
USER 65535
```

(or USER nobody if that exact word is already used in the file-but the task explicitly allows UID 65535, so USER 65535 is safest.)

This is one-instruction change and is a top-tier best practice.

Candidate 2 (very common): FROM `<image>:latest`

Search:

```
/FROM
```

If you see something like:

```
FROM nginx:latest
```

Change ONLY that line to a pinned tag (example):

```
FROM nginx:1.25.5
```

(Any non-latest pinned version is the point. Don't add a digest line; just modify the existing FROM line.) Candidate 3: ADD `http://...`

(remote URL download) Search:

```
/ADD
```

If you see remote URL usage like:

```
ADD https://example.com/app.tar.gz /app/
```

Change that single instruction to COPY only if it's copying local files.

If it's a remote URL, the more "correct" fix would normally be using curl with verification, but that would require adding instructions (not allowed).

So in this exam constraint, do NOT pick this unless it's actually a local add like:

```
ADD ./app
```

Then change just the word:

```
COPY ./app
```

3) Save and exit

```
:wq
```

Don't run docker build (task forbids building).

PART B - Fix ONE prominent security/best-practice issue in the Deployment manifest

4) Open the manifest

```
vi /home/candidate/subtle-bee/deployment.yaml
```

5) Change ONLY ONE existing field that is a clear security issue

Use `/` search in `vi` for the usual "bad fields":

Option 1 (most common): running as root

Search:

/runAsUser

If you see:

runAsUser: 0

Change that one existing field value to:

runAsUser: 65535

This is a single-field change and matches the prompt hint.

Option 2: privileged container

Search:

/privileged

If you see:

privileged: true

Change only that value to:

privileged: false

Option 3: allow privilege escalation

Search:

/allowPrivilegeEscalation

If you see:

allowPrivilegeEscalation: true

Change only that value to:

allowPrivilegeEscalation: false

Option 4: writable root filesystem

Search:

/readOnlyRootFilesystem

If you see:

readOnlyRootFilesystem: false

Change only that value to:

readOnlyRootFilesystem: true

Option 5: image uses :latest

Search:

/image:

If you see:

image: something:latest

Change only that value to a pinned tag, e.g.:

image: something:1.2.3

6) Save and exit

:wq

What to pick (fast decision rule)

If you see run as root in either file, that's usually the highest scoring / most "prominent" security issue.

Dockerfile: USER root → USER 65535

Deployment: runAsUser: 0 → runAsUser: 65535

Those are perfect because you only modify one line/field and it matches the hint.

NEW QUESTION # 55

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:

```
systemctl daemon-reload
```

```
systemctl restart kubelet.service
```

Audit:

```
/bin/ps -fC kubelet
```

Audit Config:

```
/bin/cat /var/lib/kubelet/config.yaml
```

Expected result:

'false' is equal to 'false'

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

Audit

```
docker inspect kubelet | jq -e '.[0].Args[] | match("--authorization-mode=Webhook").string' Returned Value: --authorization-
```

mode=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 Do not use self-signed certificates for TLS. etcd is a highly-available key value store used by Kubernetes deployments for persistent storage of all of its REST API objects. These objects are sensitive in nature and should not be available to unauthenticated clients. You should enable the client authentication via valid certificates to secure the access to the etcd service.

Fix - Buildtime

Kubernetes

apiVersion: v1

kind: Pod

metadata:

annotations:

scheduler.alpha.kubernetes.io/critical-pod: ""

creationTimestamp: null

labels:

component: etcd

tier: control-plane

name: etcd

namespace: kube-system

spec:

containers:

- command:

+ - etcd

+ - --auto-tls=true

image: k8s.gcr.io/etcd-amd64:3.2.18

imagePullPolicy: IfNotPresent

livenessProbe:

exec:

command:

- /bin/sh

--ec

- ETCDCTL_API=3 etcdctl --endpoints=https://[192.168.22.9]:2379 --cacert=/etc/kubernetes/pki/etcd/ca.crt

--cert=/etc/kubernetes/pki/etcd/healthcheck-client.crt --key=/etc/kubernetes/pki/etcd/healthcheck-client.key get foo

failureThreshold: 8 initialDelaySeconds: 15 timeoutSeconds: 15 name: etcd-should-fail resources: {} volumeMounts:

- mountPath: /var/lib/etcd

name: etcd-data

- mountPath: /etc/kubernetes/pki/etcd

name: etcd-certs

hostNetwork: true

priorityClassName: system-cluster-critical

volumes:

```

- hostPath:
  path: /var/lib/etcd
  type: DirectoryOrCreate
  name: etcd-data
- hostPath:
  path: /etc/kubernetes/pki/etcd
  type: DirectoryOrCreate
  name: etcd-certs
status: {}
Explanation:

```

NEW QUESTION # 56

Create a PSP that will prevent the creation of privileged pods in the namespace.

Create a new PodSecurityPolicy named prevent-privileged-policy which prevents the creation of privileged pods.

Create a new ServiceAccount named psp-sa in the namespace default.

Create a new ClusterRole named prevent-role, which uses the newly created Pod Security Policy prevent-privileged-policy.

Create a new ClusterRoleBinding named prevent-role-binding, which binds the created ClusterRole prevent-role to the created SA psp-sa.

Also, Check the Configuration is working or not by trying to Create a Privileged pod, it should get failed.

Answer:

Explanation:

Create a PSP that will prevent the creation of privileged pods in the namespace.

```
$ cat clusterrole-use-privileged.yaml
```

```

---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: use-privileged-priv
rules:
- apiGroups: ['policy']
  resources: ['podsecuritypolicies']
  verbs: ['use']
resourceNames:
- default-priv

```

```

---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: privileged-role-bind
namespace: psp-test
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: use-privileged-priv
subjects:
- kind: ServiceAccount
  name: privileged-sa

```

```
$ kubectl -n psp-test apply -f clusterrole-use-privileged.yaml
```

After a few moments, the privileged Pod should be created.

Create a new PodSecurityPolicy named prevent-privileged-policy which prevents the creation of privileged pods.

```

apiVersion: policy/v1beta1
kind: PodSecurityPolicy
metadata:
  name: example
spec:
  privileged: false # Don't allow privileged pods!

```

The rest fills in some required fields.

```
seLinux:  
rule: RunAsAny  
supplementalGroups:  
rule: RunAsAny  
runAsUser:  
rule: RunAsAny  
fsGroup:  
rule: RunAsAny  
volumes:  
- '*'
```

And create it with kubectl:

```
kubectl-admin create -f example-psp.yaml
```

Now, as the unprivileged user, try to create a simple pod:

```
kubectl-user create -f <<EOF
```

```
apiVersion: v1  
kind: Pod  
metadata:  
name: pause  
spec:  
containers:  
- name: pause  
image: k8s.gcr.io/pause  
EOF
```

The output is similar to this:

Error from server (Forbidden): error when creating "STDIN": pods "pause" is forbidden: unable to validate against any pod security policy: [] Create a new ServiceAccount named psp-sa in the namespace default.

```
$ cat clusterrole-use-privileged.yaml
```

```
---  
apiVersion: rbac.authorization.k8s.io/v1  
kind: ClusterRole  
metadata:  
name: use-privileged-psp  
rules:  
- apiGroups: ['policy']  
resources: ['podsecuritypolicies']  
verbs: ['use']  
resourceNames:  
- default-psp  
---
```

```
apiVersion: rbac.authorization.k8s.io/v1  
kind: RoleBinding  
metadata:  
name: privileged-role-bind  
namespace: psp-test  
roleRef:  
apiGroup: rbac.authorization.k8s.io  
kind: ClusterRole  
name: use-privileged-psp  
subjects:  
- kind: ServiceAccount  
name: privileged-sa
```

```
$ kubectl -n psp-test apply -f clusterrole-use-privileged.yaml
```

After a few moments, the privileged Pod should be created.

Create a new ClusterRole named prevent-role, which uses the newly created Pod Security Policy prevent-privileged-policy.

```
apiVersion: policy/v1beta1  
kind: PodSecurityPolicy  
metadata:  
name: example  
spec:  
privileged: false # Don't allow privileged pods!
```

The rest fills in some required fields.

```
seLinux:  
rule: RunAsAny  
supplementalGroups:  
rule: RunAsAny  
runAsUser:  
rule: RunAsAny  
fsGroup:  
rule: RunAsAny  
volumes:  
- '*'
```

And create it with kubectl:

```
kubectl-admin create -f example-psp.yaml
```

Now, as the unprivileged user, try to create a simple pod:

```
kubectl-user create -f <<EOF
```

```
apiVersion: v1  
kind: Pod  
metadata:  
name: pause  
spec:  
containers:  
- name: pause  
image: k8s.gcr.io/pause  
EOF
```

The output is similar to this:

Error from server (Forbidden): error when creating "STDIN": pods "pause" is forbidden: unable to validate against any pod security policy: [] Create a new ClusterRoleBinding named prevent-role-binding, which binds the created ClusterRole prevent-role to the created SA psp-sa.

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

This role binding allows "jane" to read pods in the "default" namespace.

You need to already have a Role named "pod-reader" in that namespace.

```
kind: RoleBinding
```

```
metadata:  
name: read-pods  
namespace: default  
subjects:
```

You can specify more than one "subject"

```
- kind: User
```

```
name: jane # "name" is case sensitive
```

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

```
roleRef:
```

"roleRef" specifies the binding to a Role / ClusterRole

```
kind: Role #this must be Role or ClusterRole
```

```
name: pod-reader # this must match the name of the Role or ClusterRole you wish to bind to apiGroup: rbac.authorization.k8s.io
```

```
apiVersion: rbac.authorization.k8s.io/v1 kind: Role metadata:
```

```
namespace: default
```

```
name: pod-reader
```

```
rules:
```

```
- apiGroups: [""] # "" indicates the core API group
```

```
resources: ["pods"]
```

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

NEW QUESTION # 57

You are tasked with hardening a Kubernetes cluster running on a public cloud provider. The cluster currently runs Kubernetes version 1.18 and has been exposed to the internet for several months. A security audit has identified several vulnerabilities in the current Kubernetes version, including CVE-2021-25743, which affects all versions prior to 1.22.

How do you upgrade your cluster to Kubernetes 1.22 and patch the vulnerabilities without disrupting the applications running on the cluster?

Answer:

Explanation:

Solution (Step by Step) :

1. Plan the upgrade:

- Identify the workloads running in the cluster.
- Understand the dependencies and configurations of each workload.
- Check compatibility of workloads with the new Kubernetes version.
- Research the recommended upgrade path for your cloud provider.

2. Prepare the environment:

- Create a backup of the cluster configuration. This includes the cluster manifest, service account configurations, and any custom resources.
- Test the upgrade process on a staging environment. This helps to identify potential issues and avoid downtime in the production cluster.
- Identify and fix any issues discovered in the staging environment. This could involve updating application configurations or deploying new versions of workloads.

3. Perform the upgrade:

- Use the recommended upgrade process for your cloud provider. Most cloud providers provide automated tools for Kubernetes upgrades.
- Monitor the upgrade process closely. Keep an eye on logs and metrics for any issues or errors.
- Rollback to the previous version if necessary. Have a plan to revert the upgrade if any critical issues arise.

4. Validate the upgrade:

- Verify/ that all applications are running as expected. Check application logs, metrics, and functionality to ensure that there are no regressions.
- Confirm that the vulnerabilities have been patched. Use tools like 'kubect audit or 'kubeadm upgrade' to verify the patched version.

Example using Google Kubernetes Engine:

- Create a new cluster with the desired Kubernetes version (1.22) in the Google Cloud Console.
- Use 'kubectl get nodes --all-namespaces to list the nodes in the existing cluster.
- Use 'kubectl drain' to drain the nodes in the existing cluster-
- Use 'kubectl cordon' to cordon the nodes in the existing cluster.
- Once the nodes are drained and cordoned, use 'kubectl delete node to delete the nodes in the existing cluster
- Join the nodes to the new cluster using 'kubect join
- Migrate the applications and configurations from the old cluster to the new cluster
- Delete the old cluster

This process ensures a minimal disruption to the applications during the upgrade, and that the vulnerabilities are patched effectively.

NEW QUESTION # 58

SIMULATION

Documentation Namespace, NetworkPolicy, Pod

You must connect to the correct host . Failure to do so may result in a zero score.

```
[candidate@base] $ ssh cks000031
```

Context

You must implement NetworkPolicies controlling the traffic flow of existing Deployments across namespaces.

Task

First, create a NetworkPolicy named deny-policy in the prod namespace to block all ingress traffic.

The prod namespace is labeled env:prod

Next, create a NetworkPolicy named allow-from-prod in the data namespace to allow ingress traffic only from Pods in the prod namespace.

Use the label of the prod names & Click to copy traffic.

The data namespace is labeled env:data

Do not modify or delete any namespaces or Pods . Only create the required NetworkPolicies.

Answer:

Explanation:

See the Explanation below for complete solution

Explanation:

1) Connect to the correct host

```
ssh cks000031
```

sudo -i

2) Use admin kubeconfig (safe default)

```
export KUBECONFIG=/etc/kubernetes/admin.conf
```

PART A - Deny ALL ingress traffic in prod namespace

Requirement:

NetworkPolicy name: deny-policy

Namespace: prod (namespace is labeled env=prod)

Effect: block all ingress

3) Create deny-policy in prod

Create the policy directly with kubectl (fastest & safest):

```
cat <<EOF | kubectl apply -f-
```

```
apiVersion: networking.k8s.io/v1
```

```
kind: NetworkPolicy
```

```
metadata:
```

```
name: deny-policy
```

```
namespace: prod
```

```
spec:
```

```
podSelector: {}
```

```
policyTypes:
```

```
- Ingress
```

```
EOF
```

□ What this does:

podSelector: {} → selects all Pods in prod

No ingress: rules → deny all ingress traffic

4) Verify

```
kubectl -n prod get networkpolicy deny-policy
```

PART B - Allow ingress to data ONLY from Pods in prod

Requirement:

NetworkPolicy name: allow-from-prod

Namespace: data (namespace is labeled env=data)

Allow ingress only from Pods in prod namespace

Use namespace label (env=prod)

5) Create allow-from-prod policy in data

```
cat <<EOF | kubectl apply -f-
```

```
apiVersion: networking.k8s.io/v1
```

```
kind: NetworkPolicy
```

```
metadata:
```

```
name: allow-from-prod
```

```
namespace: data
```

```
spec:
```

```
podSelector: {}
```

```
policyTypes:
```

```
- Ingress
```

```
ingress:
```

```
- from:
```

```
- namespaceSelector:
```

```
matchLabels:
```

```
env: prod
```

```
EOF
```

□ What this does:

Applies to all Pods in data

Allows ingress only from namespaces labeled env=prod

All other ingress traffic is denied by default

6) Verify

```
kubectl -n data get networkpolicy allow-from-prod
```

FINAL CHECK (What the examiner expects)

```
kubectl get networkpolicy -n prod
```

```
kubectl get networkpolicy -n data
```

You should see:

deny-policy in prod

allow-from-prod in data

