

## CKS Dumps

### Certified Kubernetes Security Specialist (CKS) Exam

<https://www.certleader.com/CKS-dumps.html>



**NEW QUESTION 1**

Create a new NetworkPolicy named deny-all in the namespace testing which denies all traffic of type ingress and egress traffic

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

You can create a "default" isolation policy for a namespace by creating a NetworkPolicy that selects all pods but does not allow any ingress traffic to those pods.

```
--  
apiVersion: networking.k8s.io/v1  
kind: NetworkPolicy  
metadata:  
name: default-deny-ingress  
spec:  
podSelector: {}  
policyTypes:  
- Ingress
```

You can create a "default" egress isolation policy for a namespace by creating a NetworkPolicy that selects all pods but does not allow any egress traffic from those pods.

```
--  
apiVersion: networking.k8s.io/v1  
kind: NetworkPolicy  
metadata:  
name: allow-all-egress  
spec:  
podSelector: {}  
egress:  
- {}  
policyTypes:  
- Egress
```

Default deny all ingress and all egress traffic You can create a "default" policy for a namespace which prevents all ingress AND egress traffic by creating the following NetworkPolicy in that namespace.

```
--  
apiVersion: networking.k8s.io/v1  
kind: NetworkPolicy  
metadata:  
name: default-deny-all  
spec:  
podSelector: {}  
policyTypes:  
- Ingress  
- Egress
```

This ensures that even pods that aren't selected by any other NetworkPolicy will not be allowed ingress or egress traffic.

**NEW QUESTION 2**

Given an existing Pod named nginx-pod running in the namespace test-system, fetch the service-account-name used and put the content in /candidate/KSC00124.txt

Create a new Role named dev-test-role in the namespace test-system, which can perform update operations, on resources of type namespaces.

Create a new RoleBinding named dev-test-role-binding, which binds the newly created Role to the Pod's ServiceAccount ( found in the Nginx pod running in namespace test-system).

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Send us your feedback on it.

**NEW QUESTION 3**

Enable audit logs in the cluster, To Do so, enable the log backend, and ensure that-

- \* 1. logs are stored at /var/log/kubernetes/kubernetes-logs.txt.
- \* 2. Log files are retained for 5 days.
- \* 3. at maximum, a number of 10 old audit logs files are retained.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Edit and extend the basic policy to log:

- \* 1. Cronjobs changes at RequestResponse
- \* 2. Log the request body of deployments changes in the namespace kube-system.
- \* 3. Log all other resources in core and extensions at the Request level.
- \* 4. Don't log watch requests by the "system:kube-proxy" on endpoints or Send us your feedback on it.

**NEW QUESTION 4**

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**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 - BuildtimeKubernetesapiVersion: 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

--anonymous-auth=false

variable.

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 - BuildtimeKubernetesapiVersion: 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: {}

## NEW QUESTION 5

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.

- A. Mastered
- B. Not Mastered

**Answer:** A

**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-priv.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-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 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-priv.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 6

Analyze and edit the given Dockerfile

```
FROM ubuntu:latest
```

```
RUN apt-getupdate -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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Whenever you need an unprivileged user for any of the tasks, use user test-user with the user id 5487 Send us the Feedback on it.

**NEW QUESTION 7**

A container image scanner is set up on the cluster. Given an incomplete configuration in the directory /etc/kubernetes/confs and a functional container image scanner with HTTPS endpoint [https://test-server.local:8081/image\\_policy](https://test-server.local:8081/image_policy)

- \* 1. Enable the admission plugin.
- \* 2. Validate the control configuration and change it to implicit deny.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Finally, test the configuration by deploying the pod having the image tag as latest. Send us your Feedback on this.

**NEW QUESTION 8**

Create a RuntimeClass named untrusted using the prepared runtime handler named runsc.

Create a Pods of image alpine:3.13.2 in the Namespace default to run on the gVisor runtime class. Verify: Exec the pods and run the dmesg, you will see output like this:

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Send us your feedback on it.

**NEW QUESTION 9**

Secrets stored in the etcd is not secure at rest, you can use the etcdctl command utility to find the secret value for e.g:ETCDCTL\_API=3 etcdctl get /registry/secrets/default/cks-secret --cacert="ca.crt" --cert="server.crt" --key="server.key" Output

Using the Encryption Configuration, Create the manifest, which secures the resource secrets using the provider AES-CBC and identity, to encrypt the secret-data at rest and ensure all secrets are encrypted with the new configuration.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Send us your feedback on it.

**NEW QUESTION 10**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your CKS Exam with Our Prep Materials Via below:**

<https://www.certleader.com/CKS-dumps.html>