

100% Pass Quiz Latest CKS - Certified Kubernetes Security Specialist (CKS) Valid Test Vce



P.S. Free & New CKS dumps are available on Google Drive shared by ValidBraindumps: https://drive.google.com/open?id=1BTrgZFsZ9yjq_ALz_n-hjQfQypaGzBu3

Now is not the time to be afraid to take any more difficult certification exams. Our CKS learning quiz can relieve you of the issue within limited time. Our website provides excellent CKS learning guidance, practical questions and answers, and questions for your choice which are your real strength. You can take the CKS Training Materials and pass it without any difficulty. As long as you can practice CKS study guide regularly and persistently your goals of making progress and getting certificates smoothly will be realized just like a piece of cake.

Linux Foundation has recently announced the launch of a new certification exam – the Certified Kubernetes Security Specialist (CKS). CKS Exam is designed to assess and validate the skills and knowledge of IT professionals who specialize in securing Kubernetes clusters.

>> CKS Valid Test Vce <<

2026 Useful CKS Valid Test Vce Help You Pass CKS Easily

It is well known that obtaining such a CKS certificate is very difficult for most people, especially for those who always think that their time is not enough to learn efficiently. With our CKS test prep, you don't have to worry about the complexity and tediousness of the operation. As long as you enter the learning interface of our soft test engine of CKS Quiz guide and start practicing on our Windows software, you will find that there are many small buttons that are designed to better assist you in your learning.

Linux Foundation Certified Kubernetes Security Specialist (CKS) Sample Questions (Q38-Q43):

NEW QUESTION # 38

SIMULATION

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-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 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-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 # 39

SIMULATION

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.

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

Answer:

Explanation:

See the Explanation belowExplanation:

```
candidate@cli:~$ kubectl config use-context KSRS00602
Switched to context "KSRS00602".
candidate@cli:~$ ssh ksr00602-master
Warning: Permanently added '10.240.86.243' (ECDSA) to the list of known hosts.
```

The programs included with the Ubuntu system are free software; the exact distribution terms for each program are described in the individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by applicable law.

```
root@ksrs00602-master:~# cat /etc/kubernetes/logpolicy/sample-policy.yaml
---
apiVersion: audit.k8s.io/v1
kind: Policy
# Don't generate audit events for all requests in RequestReceived stage.
omitStages:
- "RequestReceived"
rules:
# Don't log watch requests by the "system:kube-proxy" on endpoints or services
- level: None
  users: ["system:kube-proxy"]
  verbs: ["watch"]
  resources:
  - group: "" # core API group
    resources: ["endpoints", "services"]

# Don't log authenticated requests to certain non-resource URL paths.
- level: None
  userGroups: ["system:authenticated"]
  nonResourceURLs:
  - "/api*" # Wildcard matching.
  - "/version"

# Edit form here below
root@ksrs00602-master:~# vim /etc/kubernetes/logpolicy/sample-policy.yaml
```

```

- "/api*" # Wildcard matching.
- "/version"
# Edit form here below
- level: RequestResponse
  resources:
  - group: ""
    resource: ["cronjobs"]
  level: Request
  resources:
  - group: "" # core API group
    resources: ["pods"]
    namespaces: ["webapps"]
# Log configmap and secret changes in all other namespaces at the Metadata level.
- level: Metadata
  resources:
  - group: "" # core API group
    resources: ["secrets", "configmaps"]

# A catch-all rule to log all other requests at the Metadata level.
- level: Metadata
  # Long-running requests like watches that fall under this rule will not
  # generate an audit event in RequestReceived.
  omitStages:
  - "RequestReceived"

```

```

- "/version"
# Edit form here below
- level: RequestResponse
  resources:
  - group: ""
    resources: ["cronjobs"]
- level: Request
  resources:
  - group: "" # core API group
    resources: ["pods"]
    namespaces: ["webapps"]
# Log configmap and secret changes in all other namespaces at the Metadata level.
- level: Metadata
  resources:
  - group: "" # core API group
    resources: ["secrets", "configmaps"]

# A catch-all rule to log all other requests at the Metadata level.
- level: Metadata
  # Long-running requests like watches that fall under this rule will not
  # generate an audit event in RequestReceived.
  omitStages:
  - "RequestReceived"
root@ksrs00602-master:~# vim /etc/kubernetes/logpolicy/sample-policy.yaml
root@ksrs00602-master:~# vim /etc/kubernetes/manifests/kube-apiserver.yaml

```

```

labels:
  component: kube-apiserver
  tier: control-plane
name: kube-apiserver
namespace: kube-system
spec:
  containers:
    - command:
      - kube-apiserver
      - --advertise-address=10.240.86.243
      - --allow-privileged=true
      - --audit-policy-file=/etc/kubernetes/logpolicy/sample-policy.yaml
      - --audit-log-path=/var/log/kubernetes/kubernetes-logs.txt
      - --audit-log-maxbackup=1
      - --audit-log-maxage=30
      - --authorization-mode=Node,RBAC
      - --client-ca-file=/etc/kubernetes/pki/ca.crt
      - --enable-admission-plugins=NodeRestriction
      - --enable-bootstrap-token-auth=true
      - --etcd-cafile=/etc/kubernetes/pki/etcd/ca.crt

```

```

# A catch-all rule to log all other requests at the Metadata level.
- level: Metadata
  # Long-running requests like watches that fall under this rule will not
  # generate an audit event in RequestReceived.
  omitStages:
    - "RequestReceived"
root@ksrs00602-master:~# vim /etc/kubernetes/logpolicy/sample-policy.yaml
root@ksrs00602-master:~# vim /etc/kubernetes/manifests/kube-apiserver.yaml
root@ksrs00602-master:~# systemctl daemon-reload
root@ksrs00602-master:~# systemctl restart kubelet.service
root@ksrs00602-master:~# systemctl enable kubelet
root@ksrs00602-master:~# exit
logout
Connection to 10.240.86.243 closed.
candidate@cli:~$

```

NEW QUESTION # 40

Given an existing Pod named test-web-pod running in the namespace test-system Edit the existing Role bound to the Pod's Service Account named sa-backend to only allow performing get operations on endpoints.

Create a new Role named test-system-role-2 in the namespace test-system, which can perform patch operations, on resources of type statefulsets.

- A. Create a new RoleBinding named test-system-role-2-binding binding the newly created Role to the Pod's ServiceAccount sa-backend.

Answer: A

NEW QUESTION # 41

SIMULATION

Create a Pod name Nginx-pod inside the namespace testing. Create a service for the Nginx-pod named nginx-svc, using the ingress of your choice, run the ingress on tls, secure port.

Answer:

Explanation:

See the Explanation belowExplanation:

```
$ kubectl get ing -n <namespace-of-ingress-resource>
```

```

NAME HOSTS ADDRESS PORTS AGE
cafe-ingress cafe.com 10.0.2.15 80 25s
$ kubectl describe ing <ingress-resource-name> -n <namespace-of-ingress-resource> Name: cafe-ingress Namespace: default
Address: 10.0.2.15 Default backend: default-http-backend:80 (172.17.0.5:8080) Rules:
Host Path Backends
-----
cafe.com
/tea tea-svc:80 (<none>)
/coffee coffee-svc:80 (<none>)
Annotations:
kubectl.kubernetes.io/last-applied-configuration: {"apiVersion":"networking.k8s.io/v1","kind":"Ingress","metadata":{"annotations":
{"name":"cafe-ingress","namespace":"default","selfLink":"/apis/networking/v1/namespaces/default/ingresses/cafe-ingress"},"spec":
{"rules":[{"host":"cafe.com","http":{"paths":[{"backend":{"serviceName":"tea-svc","servicePort":80},"path":"/tea"}, {"backend":
{"serviceName":"coffee-svc","servicePort":80},"path":"/coffee"}]}]},"status":{"loadBalancer":{"ingress":
[{"ip":"169.48.142.110"}]}}} Events:
Type Reason Age From Message
-----
Normal CREATE 1m ingress-nginx-controller Ingress default/cafe-ingress
Normal UPDATE 58s ingress-nginx-controller Ingress default/cafe-ingress
$ kubectl get pods -n <namespace-of-ingress-controller>
NAME READY STATUS RESTARTS AGE
ingress-nginx-controller-67956bf89d-fv58j 1/1 Running 0 1m
$ kubectl logs -n <namespace> ingress-nginx-controller-67956bf89d-fv58j
----- NGINX Ingress controller Release: 0.14.0 Build:
git-734361d Repository: https://github.com/kubernetes/ingress-nginx
-----
....

```

NEW QUESTION # 42

Context:

Cluster: gvisor

Master node: master1

Worker node: worker1

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

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

Context: This cluster has been prepared to support runtime handler, runc as well as traditional one.

Task:

Create a RuntimeClass named not-trusted using the prepared runtime handler names runc.

Update all Pods in the namespace server to run on newruntime.

Answer:

Explanation:

Find all the pods/deployment and edit runtimeClassName parameter to not-trusted under spec

```
[desk@cli] $ k edit deploy nginx
```

spec:

```
runtimeClassName: not-trusted. # Add this
```

Explanation

```
[desk@cli] $ vim runtime.yaml
```

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

```
kind: RuntimeClass
```

```
metadata:
```

```
name: not-trusted
```

```
handler: runc
```

```
[desk@cli] $ k apply -fruntime.yaml
```

```
[desk@cli] $ k get pods
```

```
NAME READY STATUS RESTARTS AGE
```

```
nginx-6798fc88e8-chp6r 1/1 Running 0 11m
```

```
nginx-6798fc88e8-fs53n 1/1 Running 0 11m
```

```
nginx-6798fc88e8-ndved 1/1 Running 0 11m
```

```
[desk@cli] $ k get deploy
NAME READY UP-TO-DATE AVAILABLE AGE
nginx 3/3 11 3 5m
[desk@cli] $ k edit deploy nginx
```

```
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app: nginx
  name: nginx
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx
  strategy: {}
  template:
    metadata:
      labels:
        app: nginx
    spec:
      runtimeClassName: not-trusted # Add this
      containers:
      - image: nginx
        name: nginx
        resources: {}
status: {}
```

NEW QUESTION # 43

.....

The product we provide with you is compiled by professionals elaborately and boosts varied versions which aimed to help you learn the CKS study materials by the method which is convenient for you. They check the update every day, and we can guarantee that you can get a free update service from the date of purchase. Once you have any questions and doubts about the CKS Exam Questions we will provide you with our customer service before or after the sale, you can contact us if you have question or doubt about our exam materials and the professional personnel can help you solve your issue about using CKS study materials.

CKS Frenquent Update: <https://www.validbraindumps.com/CKS-exam-prep.html>

- CKS Reliable Exam Tips CKS Valid Dumps Book New CKS Cram Materials Search for ➔ CKS and easily obtain a free download on ▷ www.prepawaypdf.com ◁ CKS Cert Guide
- CKS Guide Examcollection CKS Free Dumps CKS Valid Dumps Book Enter ⇒ www.pdfvce.com ⇐ and search for ➔ CKS to download for free CKS Practice Braindumps
- Clearer CKS Explanation CKS Practice Braindumps Reliable CKS Exam Pattern Search for (CKS) and download exam materials for free through ▷ www.validtorrent.com ◁ CKS Practice Braindumps
- Pass Guaranteed Quiz 2026 High Hit-Rate Linux Foundation CKS: Certified Kubernetes Security Specialist (CKS) Valid Test Vce Search for ▷ CKS ◁ on www.pdfvce.com immediately to obtain a free download CKS Test Dates
- Pass Guaranteed Quiz 2026 High Hit-Rate Linux Foundation CKS: Certified Kubernetes Security Specialist (CKS) Valid Test Vce Open “www.dumpsmaterials.com” and search for [CKS] to download exam materials for free CKS Practice Braindumps
- CKS Valid Dumps Book CKS Free Test Questions Preparation CKS Store Copy URL ⇒ www.pdfvce.com

