

Linux-Foundation

Exam Questions CKS

Certified Kubernetes Security Specialist (CKS) Exam



NEW QUESTION 1

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Send us your feedback on this.

NEW QUESTION 2

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 3

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 4

Create a new ServiceAccount named backend-sa in the existing namespace default, which has the capability to list the pods inside the namespace default. Create a new Pod named backend-pod in the namespace default, mount the newly created sa backend-sa to the pod, and Verify that the pod is able to list pods. Ensure that the Pod is running.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A service account provides an identity for processes that run in a Pod.

When you (a human) access the cluster (for example, using kubectl), you are authenticated by the apiserver as a particular User Account (currently this is usually admin, unless your cluster administrator has customized your cluster). Processes in containers inside pods can also contact the apiserver. When they do, they are authenticated as a particular Service Account (for example, default).

When you create a pod, if you do not specify a service account, it is automatically assigned the default service account in the same namespace. If you get the raw json or yaml for a pod you have created (for

example, `kubectl get pods/<podname> -o yaml`), you can see the `spec.serviceAccountName` field has been automatically set.

You can access the API from inside a pod using automatically mounted service account credentials, as described in Accessing the Cluster. The API permissions of the service account depend on the authorization plugin and policy in use.

In version 1.6+, you can opt out of automounting API credentials for a service account by setting `automountServiceAccountToken: false` on the service account:

```
apiVersion:v1
kind:ServiceAccount
```

```
metadata:
name:build-robot
automountServiceAccountToken:false
```

In version 1.6+, you can also opt out of automounting API credentials for a particular pod:

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

```
metadata:
name:my-pod
```

```
spec:
serviceAccountName:build-robot
automountServiceAccountToken:false
```

The pod spec takes precedence over the service account if both specify a `automountServiceAccountToken` value.

NEW QUESTION 5

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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

use 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-resource-rbcreated
```

```
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 6

Create a PSP that will only allow the persistentvolumeclaim as the volume type in the namespace restricted.

Create a new PodSecurityPolicy named prevent-volume-policy which prevents the pods which is having different volumes mount apart from persistentvolumeclaim.

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

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

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

Hint:

Also, Check the Configuration is working or not by trying to Mount a Secret in the pod manifest, it should get failed.

POD Manifest:

```
* apiVersion: v1
```

```
* kind: Pod
```

```
* metadata:
```

```
* name:
```

```
* spec:
```

```
* containers:
```

```
* - name:
```

```
* image:
```

```
* volumeMounts:
```

```
* - name:
```

```
* mountPath:
```

```
* volumes:
```

```
* - name:
```

```
* secret:
```

```
* secretName:
```

A. Mastered

B. Not Mastered

Answer: A

Explanation:

```
apiVersion: policy/v1beta1
```

```
kind: PodSecurityPolicy
```

```
metadata:
```

```
name: restricted
```

```
annotations:
```

```
seccomp.security.alpha.kubernetes.io/allowedProfileNames: 'docker/default,runtime/default'
```

```
apparmor.security.beta.kubernetes.io/allowedProfileNames: 'runtime/default' seccomp.security.alpha.kubernetes.io/defaultProfileName: 'runtime/default'
```

```
apparmor.security.beta.kubernetes.io/defaultProfileName: 'runtime/default'
```

```
spec:
```

```
privileged: false
```

```
# Required to prevent escalations to root.
```

```
allowPrivilegeEscalation: false
```

```
# This is redundant with non-root + disallow privilege escalation,
```

```
# but we can provide it for defense in depth.
```

```
requiredDropCapabilities:
```

```
- ALL
```

```
# Allow core volume types. volumes:
```

```
- 'configMap'
```

```
- 'emptyDir'
```

```
- 'projected'
```

```
- 'secret'
```

```
- 'downwardAPI'
```

```
# Assume that persistentVolumes set up by the cluster admin are safe to use.
```

```
- 'persistentVolumeClaim'
```

```
hostNetwork: false
```

```
hostIPC: false
hostPID: false
runAsUser:
# Require the container to run without root privileges.
rule: 'MustRunAsNonRoot'
seLinux:
# This policy assumes the nodes are using AppArmor rather than SELinux.
rule: 'RunAsAny'
supplementalGroups:
rule: 'MustRunAs'
ranges:
# Forbid adding the root group.
- min: 1
max: 65535
fsGroup:
rule: 'MustRunAs'
ranges:
# Forbid adding the root group.
- min: 1
max: 65535
readOnlyRootFilesystem: false
```

NEW QUESTION 10

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

CKS Practice Exam Features:

- * CKS Questions and Answers Updated Frequently
- * CKS Practice Questions Verified by Expert Senior Certified Staff
- * CKS Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * CKS Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The CKS Practice Test Here](#)