



Free Questions for CKS by vceexamstest

Shared by Flores on 05-09-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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 theAPI server:-

- a. Ensure that the RotateKubeletServerCertificate argument is set to true.
- b. Ensure that the admission control plugin PodSecurityPolicy is set.
- c. Ensure that the --kubelet-certificate-authority argument is set as appropriate.

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

- 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 theETCD:-

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

Hint: Take the use of Tool Kube-Bench

Options:

A) Explanation:

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

a. Ensure that the RotateKubeletServerCertificate argument is set to true.

apiVersion: v1

kind: Pod

metadata:

creationTimestamp: null

labels:

component: kubelet

tier: control-plane

name: kubelet

namespace: kube-system

spec:

containers:

- command:

- kube-controller-manager

+ - --feature-gates=RotateKubeletServerCertificate=true

image: gcr.io/google_containers/kubelet-amd64:v1.6.0

livenessProbe:

failureThreshold: 8

httpGet:
host: 127.0.0.1
path: /healthz
port: 6443
scheme: HTTPS
initialDelaySeconds: 15
timeoutSeconds: 15
name: kubelet
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

b. Ensure that the admission control plugin PodSecurityPolicy is set.

audit: '/bin/ps -ef | grep \$apiserverbin | grep -v grep'

tests:

test_items:

- flag: '--enable-admission-plugins'

compare:

op: has

value: 'PodSecurityPolicy'

set: true

remediation: |

Follow the documentation and create Pod Security Policy objects as per your environment.

Then, edit the API server pod specification file \$apiserverconf

on the master node and set the --enable-admission-plugins parameter to a

value that includes PodSecurityPolicy :

--enable-admission-plugins=...,PodSecurityPolicy,...

Then restart the API Server.

scored: true

c. Ensure that the --kubelet-certificate-authority argument is set as appropriate.

audit: '/bin/ps -ef | grep \$apiserverbin | grep -v grep'

tests:

test_items:

- flag: '--kubelet-certificate-authority'

set: true

remediation: |

Follow the Kubernetes documentation and setup the TLS connection between the apiserver and kubelets. Then, edit the API server pod specification file

\$apiserverconf on the master node and set the --kubelet-certificate-authority parameter to the path to the cert file for the certificate authority.

--kubelet-certificate-authority=<ca-string>

scored: true

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

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

Edit the etcd pod specification file \$etcdconf on the master node and either remove the --auto-tls parameter or set it to false.

--auto-tls=false

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

Edit the etcd pod specification file \$etcdconf on the master node and either remove the --peer-auto-tls parameter or set it to false.

--peer-auto-tls=false

Answer:

A

Question 2

Question Type: MultipleChoice

SIMULATION

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.

Options:

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-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']
```

Answer:

A

Question 3

Question Type: MultipleChoice

SIMULATION

Using the runtime detection tool Falco, Analyse the container behavior for at least 30 seconds, using filters that detect newly spawning and executing processes

store the incident file at /opt/falco-incident.txt, containing the detected incidents. one per line, in the format

[timestamp],[uid],[user-name],[processName]

Options:

A) Sendyoursuggestiononit

Answer:

A

Question 4

Question Type: MultipleChoice

SIMULATION

use the Trivy to scan the following images,

1. amazonlinux:1
2. k8s.gcr.io/kube-controller-manager:v1.18.6

Look for images with HIGH or CRITICAL severity vulnerabilities and store the output of the same in /opt/trivy-vulnerable.txt

Options:

A) Send us the Feedback on it.

Answer:

A

Question 5

Question Type: MultipleChoice

SIMULATION

a. Retrieve the content of the existing secret named `default-token-xxxxx` in the testing namespace.

Store the value of the token in the `token.txt`

b. Create a new secret named `test-db-secret` in the DB namespace with the following content:

```
username:mysql
```

```
password:password@123
```

Create the Pod name `test-db-pod` of image `nginx` in the namespace `db` that can access `test-db-secret` via a volume at path `/etc/mysql-credentials`

Options:

A) Explanation:

To add a Kubernetes cluster to your project, group, or instance:

Navigate to your:

Project's `Operations > Kubernetes` page, for a project-level cluster.

Group's `Kubernetes` page, for a group-level cluster.

Admin Area `>Kubernetes` page, for an instance-level cluster.

Click `Add Kubernetes cluster`.

Click the `Add existing cluster` tab and fill in the details:

Kubernetes cluster name(required) - The name you wish to give the cluster.

Environment scope(required) - The associated environment to this cluster.

API URL(required) - It's the URL that GitLab uses to access the Kubernetes API. Kubernetes exposes several APIs, we want the "base" URL that is common to all of them. For example, `https://kubernetes.example.com` rather than `https://kubernetes.example.com/api/v1`.

Get the API URL by running this command:

```
kubectl cluster-info | grep -E 'Kubernetes master|Kubernetes control plane' | awk '/http/ {print $NF}'
```

CA certificate(required) - A valid Kubernetes certificate is needed to authenticate to the cluster. We use the certificate created by default.

List the secrets with `kubectl get secrets`, and one should be named similar to `default-token-xxxxx`. Copy that token name for use below.

Get the certificate by running this command:

```
kubectl get secret <secret name> -o jsonpath='{[\'data\'][\'ca\.crt\']}'
```

Answer:

A

Question 6

Question Type: MultipleChoice

SIMULATION

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

Options:

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.

Answer:

A

Question 7

Question Type: MultipleChoice

SIMULATION

On the Cluster worker node, enforce the prepared AppArmor profile

```
#include
```

```
profile nginx-deny flags=(attach_disconnected) {
```

```
#include
```

```
file,
```

```
# Deny all file writes.
```

```
deny /** w,
```

```
}
```

```
EOF'
```

Edit the prepared manifest file to include the AppArmor profile.

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
name: apparmor-pod
```

spec:

containers:

- name: apparmor-pod

image: nginx

Finally, apply the manifests files and create the Pod specified on it.

Verify: Try to make a file inside the directory which is restricted.

Options:

A) Send us the Feedback on it.

Answer:

A

Question 8

Question Type: MultipleChoice

SIMULATION

A container image scanner is set up on the cluster.

Given an incomplete configuration in the directory

/etc/kubernetes/confcontrol and a functional container image scanner with HTTPS endpoint https://test-server.local.8081/image_policy

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

Finally, test the configuration by deploying the pod having the image tag as latest.

Options:

A) Send us the Feedback on it.

Answer:

A

Question 9

Question Type: MultipleChoice

SIMULATION

Create a network policy named restrict-np to restrict to pod nginx-test running in namespace testing.

Only allow the following Pods to connect to Pod nginx-test:-

1. pods in the namespace default
2. pods with label version:v1 in any namespace.

Make sure to apply the network policy.

Options:

A) Send us your Feedback on this.

Answer:

A

Question 10

Question Type: MultipleChoice

SIMULATION

Create a RuntimeClass named gvisor-rc using the prepared runtime handler named runsc.

Create a Pods of image Nginx in the Namespace server to run on the gVisor runtime class

Options:

A) Explanation:

Install the Runtime Class for gVisor

```
{ # Step 1: Install a RuntimeClass
```

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

```
apiVersion: node.k8s.io/v1beta1
```

```
kind: RuntimeClass
```

```
metadata:
```

```
name: gvisor
```

```
handler: runsc
```

```
EOF
```

```
}
```

Create a Pod with the gVisor Runtime Class

```
{ # Step 2: Create a pod
```

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

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
name: nginx-gvisor
spec:
  runtimeClassName: gvisor
  containers:
  - name: nginx
    image: nginx
  EOF
}
Verify that the Pod is running
{ # Step 3: Get the pod
kubectl get pod nginx-gvisor -o wide
}
```

Answer:

A

Question 11

Question Type: MultipleChoice

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

Options:

A) Send us the Feedback on it.

Answer:

A

To Get Premium Files for CKS Visit

<https://www.p2pexams.com/products/cks>

For More Free Questions Visit

<https://www.p2pexams.com/linux-foundation/pdf/cks>

