



Certified Kubernetes Security Specialist (CKS) Exam

# Pass Linux Foundation CKS Exam with 100% Guarantee

Free Download Real Questions & Answers PDF and VCE file from:

https://www.pass4itsure.com/cks.html

# 100% Passing Guarantee 100% Money Back Assurance

Following Questions and Answers are all new published by Linux Foundation Official Exam Center

Instant Download After Purchase

100% Money Back Guarantee

😳 365 Days Free Update

800,000+ Satisfied Customers





## **QUESTION 1**

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.

A. See explanation below.

B. PlaceHolder

Correct Answer: A

ssh-add ~/.ssh/tempprivate eval "\$(ssh-agent -s)" cd contrib/terraform/aws vi terraform.tfvars terraform init terraform apply -var-file=credentials.tfvars ansible-playbook -i ./inventory/hosts ./cluster.yml -e ansible\_ssh\_user=core -e bootstrap\_os=coreos -b --become-user=root --flush-cache -e ansible\_user=core

TASX [kubernete5/master : sets kubeadm api version to vlbeta] Turnshay 03 Skylembkr 2019 07:14:20 +0050 (0:50:55.157) okr [kubernetes dev0215joh10503 master] okr [kubernetes] dev0215joh10503 master] okr [kubernetes] dev0215joh10803-master3]	]]
TLSN [kubernelis/maslim : lubeadm [ Create kukeadm conlig] ** Tuesday 03 September 2019 07:14:21 +0000 (0:00:00.560) charged: [kubernetes dev0210john0503 master0] charged: [kubernetes:dev0210john0503 master1] charged: [kubernetes dev0210john0503 master2]	//************************************
MASE (kubarnalas/maslar : %ackup old carls and kays) ******* Tuesday US September 2019 -0/(14/24/10000 (0:00:03.343)	0;22:02.390 *****
<pre>TMSK [kuberneles/masler : tubeadm   initialize first master] Tweeday 03 Geptember 2018 0/14/25 0000 (0100:00.520) FAILED - RETWYING: kubeadm   initialize first master (3 retri FAILED - RETWYING: kubeadm   initialize first master (1 retri failed - RETWYING)   retrieve - RETWYING   retri failed - RETWYING   retrieve - RETWYING   retwying - RETWYING   retrieve - RETWYING</pre>	<pre>District and a second sec</pre>

# **QUESTION 2**

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. See the below:

B. PlaceHolder

Correct Answer: A

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/ -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.

# **QUESTION 3**



candidate@cli:~\$ kubectl config use-context KSSH00401 Switched to context "KSSH00401". candidate@cli:~\$ ssh kssh00401-worker1 Warning: Permanently added '10.240.86.172' (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@kssh00401-worker1:~# head /etc/apparmor.d/nginx apparmor #include <tunables/global> profile nginx-profile-2 flags=(attach disconnected, mediate deleted) { #include <abstractions/base>
network inet tcp, network inet udp, network inet icmp, deny network raw, root@kssh00401-worker1:~# apparmor\_parser -q /etc/apparmor.d/nginx\_apparmor root@kssh00401-worker1:~# exit logout Connection to 10.240.86.172 closed. candidate@cli:~\$ cat KSSH00401/nginx-pod.yaml apiVersion: vl kind: Pod metadata: name: nginx-pod spec: containers: name: nginx-pod image: nginx:1.19.0 ports: containerPort: 80 candidate@cli:~\$ vim KSSH00401/nginx-pod.yaml ind: Pod name: nginx-pod container.apparmor.security.beta.kubernetes.io/nginx-pod: localhost/nginx-pr name: nginx-pod image: nginx:1.19.0 candidate@cli:~\$ vim KSSH00401/nginx-pod.yaml candidate@cli:~\$ kubectl create -f KSSH00401/nginx-pod.yaml pod/nginx-pod created candidate@cli:~\$ cat KSSH00401/nginx-pod.yaml apiVersion: v1 kind: Pod metadata: name: nginx-pod annotations: container.apparmor.security.beta.kubernetes.io/nginx-pod: localhost/nginx-profile-2 spec: containers: name: nginx-pod image: nginx:1.19.0 ports: - containerPort: 80



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:

1.

Ensure that the RotateKubeletServerCertificate argument is set to true.

2.

Ensure that the admission control plugin PodSecurityPolicy is set.

3.

Ensure that the --kubelet-certificate-authority argument is set as appropriate. Fix all of the following violations that were found against the Kubelet:

1.

Ensure the --anonymous-auth argument is set to false.

2.

Ensure that the --authorization-mode argument is set to Webhook. Fix all of the following violations that were found against the ETCD:

1.

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

2.

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

Hint: Take the use of Tool Kube-Bench

A. See the below.

B. PlaceHolder

Correct Answer: A

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=

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 masternode 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 masternode and either remove the -- peer-auto-tls parameter or set it to false.--peer-auto-tls=false

#### **QUESTION 4**

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. See the below.



### B. PlaceHolder

### Correct Answer: A

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-