

Linux-Foundation

Exam Questions CKS

Certified Kubernetes Security Specialist (CKS) Exam



NEW QUESTION 1

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

- * 1. logs are stored at /var/log/kubernetes-logs.txt.
- * 2. Log files are retained for 12 days.
- * 3. at maximum, a number of 8 old audit logs files are retained.
- * 4. set the maximum size before getting rotated to 200MB

Edit and extend the basic policy to log:

- * 1. namespaces changes at RequestResponse
 - * 2. Log the request body of secrets changes in the namespace kube-system.
 - * 3. Log all other resources in core and extensions at the Request level.
 - * 4. Log "pods/portforward", "services/proxy" at Metadata level.
 - * 5. Omit the Stage RequestReceived
- All other requests at the Metadata level

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Kubernetes auditing provides a security-relevant chronological set of records about a cluster. Kube-apiserver performs auditing. Each request on each stage of its execution generates an event, which is then pre-processed according to a certain policy and written to a backend. The policy determines what's recorded and the backends persist the records.

You might want to configure the audit log as part of compliance with the CIS (Center for Internet Security) Kubernetes Benchmark controls.

The audit log can be enabled by default using the following configuration in cluster.yml:

```
services:
kube-api:
audit_log:
enabled:true
```

When the audit log is enabled, you should be able to see the default values at
/etc/kubernetes/audit-policy.yaml

The log backend writes audit events to a file in JSONlines format. You can configure the log audit backend using the following kube-apiserver flags:

- > --audit-log-path specifies the log file path that log backend uses to write audit events. Not specifying thi flag disables log backend. - means standard out
- > --audit-log-maxbackup defines the maximum number of audit log files to retain
- > --audit-log-maxsize defines the maximum size in megabytes of the audit log file before it gets rotated

If your cluster's control plane runs the kube-apiserver as a Pod, remember to mount the location of the policy file and log file, so that audit records are persisted.

For example:-hostPath-to the

```
--audit-policy-file=/etc/kubernetes/audit-policy.yaml\
--audit-log-path=/var/log/audit.log-
```

NEW 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. 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 servic 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 3

Using the runtime detection tool Falco, Analyse the container behavior for at least 20 seconds, using filters that detect newly spawning and executing processes in

a single container of Nginx.

store the incident file art /opt/falco-incident.txt, containing the detected incidents. one per line, in the format [timestamp],[uid],[processName]

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Send us your feedback on it.

NEW QUESTION 4

On the Cluster worker node, enforce the prepared AppArmor profile

```
#include<tunables/global>
profile docker-nginx flags=(attach_disconnected,mediate_deleted) {
#include<abstractions/base>
network inet tcp,
network inet udp,
network inet icmp,
deny network raw,
deny network packet,
file,
umount,
deny /bin/** wl,
deny /boot/** wl,
deny /dev/** wl,
deny /etc/** wl,
deny /home/** wl,
deny /lib/** wl,
deny /lib64/** wl,
deny /media/** wl,
deny /mnt/** wl,
deny /opt/** wl,
deny /proc/** wl,
deny /root/** wl,
deny /sbin/** wl,
deny /srv/** wl,
deny /tmp/** wl,
deny /sys/** wl,
deny /usr/** wl,
audit /** w,
/var/run/nginx.pid w,
/usr/sbin/nginx ix,
deny /bin/dash mrwxl,
deny /bin/sh mrwxl,
deny /usr/bin/top mrwxl,
capability chown,
capability dac_override,
capability setuid,
capability setgid,
capability net_bind_service,
deny @{PROC}/* w, # deny write for all files directly in /proc (not in a subdir)
# deny write to files not in /proc/<number>/** or /proc/sys/**
deny @{PROC}/{[^1-9],[^1-9][^0-9],[^1-9s][^0-9y][^0-9s],[^1-9][^0-9][^0-9][^0-9]*}/** w,
deny @{PROC}/sys/[^k]** w, # deny /proc/sys except /proc/sys/k* (effectively /proc/sys/kernel)
deny @{PROC}/sys/kernel/{?,??,[^s][^h][^m]**} w, # deny everything except shm* in
/proc/sys/kernel/
deny @{PROC}/sysrq-trigger rwxl,
deny @{PROC}/mem rwxl,
deny @{PROC}/kmem rwxl,
deny @{PROC}/kcore rwxl,
deny mount,
deny /sys/[^f]** wklx,
deny /sys/f[^s]** wklx,
deny /sys/fs/[^c]** wklx,
deny /sys/fs/c[^g]** wklx,
deny /sys/fs/cg[^r]** wklx,
deny /sys/firmware/** rwxl,
deny /sys/kernel/security/** rwxl,
}
```

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 use command ping, top, sh

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Send us your feedback on it.

NEW QUESTION 5

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 art /opt/falco-incident.txt, containing the detected incidents. one per line, in the format [timestamp],[uid],[user-name],[processName]

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Send us your suggestion on it.

NEW QUESTION 6

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

- A. Mastered
- B. Not Mastered

Answer: 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

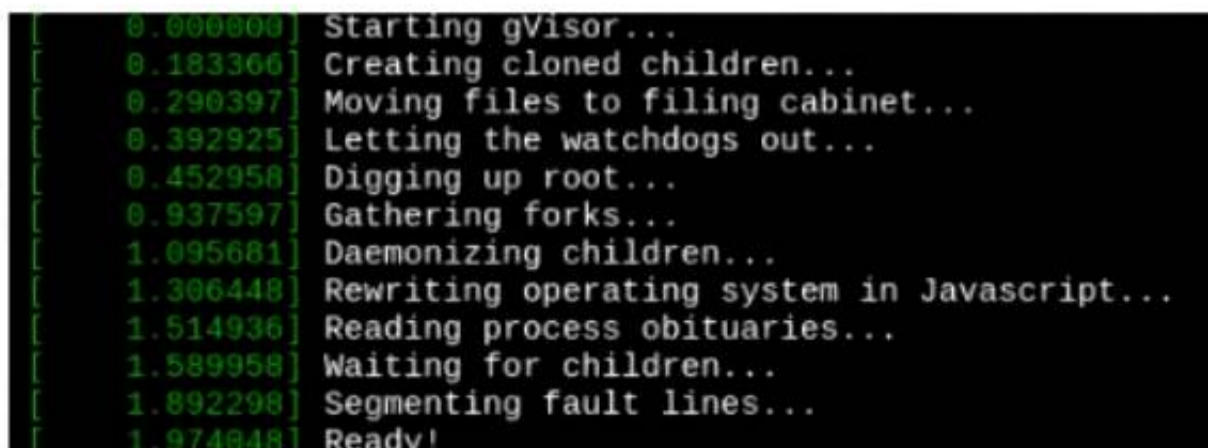
kubectl get pod nginx-gvisor -o wide

}

NEW QUESTION 7

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:



```
[ 0.000000] Starting gVisor...
[ 0.183366] Creating cloned children...
[ 0.290397] Moving files to filing cabinet...
[ 0.392925] Letting the watchdogs out...
[ 0.452958] Digging up root...
[ 0.937597] Gathering forks...
[ 1.095681] Daemonizing children...
[ 1.306448] Rewriting operating system in Javascript...
[ 1.514936] Reading process obituaries...
[ 1.589958] Waiting for children...
[ 1.892298] Segmenting fault lines...
[ 1.974848] Ready!
```

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Send us your feedback on it.

NEW QUESTION 8

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:

se 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 certificateRetrieve 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: kubectlapply-froleBindingJohn.yaml--as=john

rolebinding.rbac.authorization.k8s.io/john_external-rosource-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"]

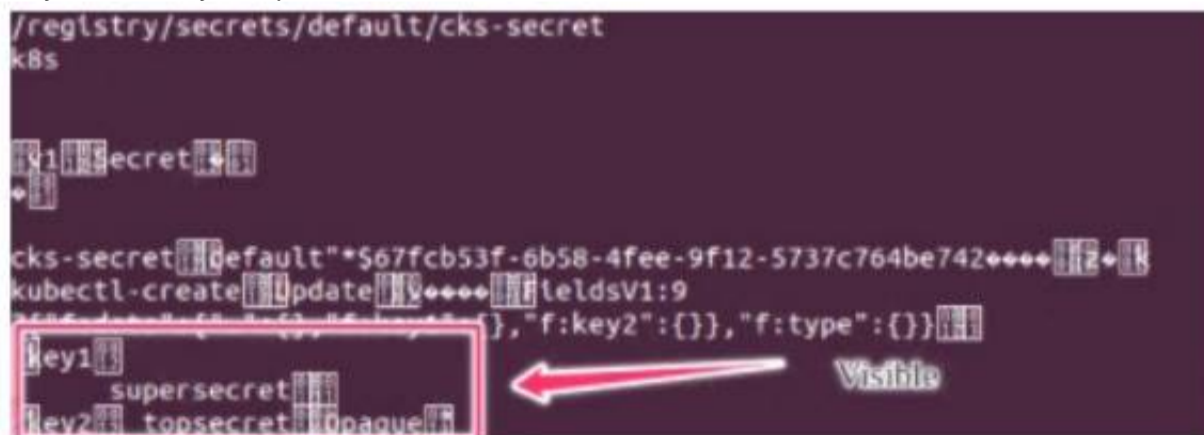
verbs:["create, list, get"]

NEW QUESTION 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

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Send us your suggestion on it.

NEW QUESTION 10

Before Making any changes build the Dockerfile with tag base:v1 Now Analyze and edit the given Dockerfile(based on ubuntu 16:04)

Fixing two instructions present in the file, Check from Security Aspect and Reduce Size point of view.

Dockerfile:

```
FROM ubuntu:latest
RUN apt-getupdate -y
RUN apt install nginx -y
COPY entrypoint.sh /
RUN useradd ubuntu
ENTRYPOINT ["/entrypoint.sh"]
USER ubuntu
entrypoint.sh
#!/bin/bash
echo"Hello from CKS"
```

After fixing the Dockerfile, build the docker-image with the tag base:v2 To Verify: Check the size of the image before and after the build.

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Send us your feedback on it.

NEW QUESTION 11

On the Cluster worker node, enforce the prepared AppArmor profile

```
#include<tunables/global>
profile nginx-deny flags=(attach_disconnected) {
#include<abstractions/base>
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.

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Send us your Feedback on this.

NEW QUESTION 14

.....

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 First Try
- * CKS Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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