



Free Questions for CKS by ebraindumps

Shared by Roberson on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Given an existing Pod named nginx-pod running in the namespace test-system, fetch the service-account-name used and put the content in /candidate/KSC00124.txt

Create a new Role named dev-test-role in the namespace test-system, which can perform update operations, on resources of type namespaces.

Create a new RoleBinding named dev-test-role-binding, which binds the newly created Role to the Pod's ServiceAccount (found in the Nginx pod running in namespace test-system).

Options:

A- Explanation:

```
candidate@cli:~$ kubectl config use-context KSCH00201
Switched to context "KSCH00201".
candidate@cli:~$ kubectl get pods -n security
NAME          READY   STATUS    RESTARTS   AGE
web-pod       1/1     Running   0           6h9m
candidate@cli:~$ kubectl get deployments.apps -n security
No resources found in security namespace.
candidate@cli:~$ kubectl describe rolebindings.rbac.authorization.k8s.io -n security
Name:          dev-role
Labels:        <none>
Annotations:   <none>
Role:
  Kind: Role
  Name: dev-role
Subjects:
  Kind          Name          Namespace
  ----          -
  ServiceAccount sa-dev-1
candidate@cli:~$ kubectl describe role dev-role -n security
Name:          dev-role
Labels:        <none>
Annotations:   <none>
PolicyRule:
  Resources  Non-Resource URLs  Resource Names  Verbs
  -
  *          []                 []              [*]
```

candidate@cli:~\$ kubectl edit role/dev-role -n security █

```
uid: b4c9ddd6-2729-43bd-8fbd-b2d227f4c4cd
rules:
- apiGroups:
  - ""
  resources:
  - services
  verbs:
  - watch
```

```

candidate@cli:~$ kubectl describe role dev-role -n security
Name:          dev-role
Labels:        <none>
Annotations:   <none>
PolicyRule:
  Resources    Non-Resource URLs  Resource Names  Verbs
  -----
  *           []                []              [*]
candidate@cli:~$ kubectl edit role/dev-role -n security
role.rbac.authorization.k8s.io/dev-role edited
candidate@cli:~$ kubectl describe role dev-role -n security
Name:          dev-role
Labels:        <none>
Annotations:   <none>
PolicyRule:
  Resources    Non-Resource URLs  Resource Names  Verbs
  -----
  services    []                []              [watch]
candidate@cli:~$ kubectl get pods -n security
NAME        READY   STATUS    RESTARTS   AGE
web-pod    1/1     Running   0           6h12m
candidate@cli:~$ kubectl get pods/web-pod -n security -o yaml | grep serviceAccount
  serviceAccount: sa-dev-1
  serviceAccountName: sa-dev-1
    - serviceAccountToken:
candidate@cli:~$ kubectl create role role-2 --verb=update --resource=namespaces -n security
role.rbac.authorization.k8s.io/role-2 created
candidate@cli:~$ kubectl create rolebinding role-2-binding --role
--role --role=
candidate@cli:~$ kubectl create rolebinding role-2-binding --role=role-2 --serviceaccount=se
curity:sa-dev-1 -n security
rolebinding.rbac.authorization.k8s.io/role-2-binding created
candidate@cli:~$ █

```

Answer:

A

Question 2

Question Type: MultipleChoice

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:

Options:

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
```

Answer:

A

Question 3

Question Type: MultipleChoice

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 4

Question Type: MultipleChoice

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 5

Question Type: MultipleChoice

Two tools are pre-installed on the cluster's worker node:

Using the tool of your choice (including any non pre-installed tool), analyze the container's behavior for at least 30 seconds, using filters that detect newly spawning and executing processes.

Store an incident file at /opt/KSRS00101/alerts/details, containing the detected incidents, one per line, in the following format:

```
timestamp,uid/username,processName
```

The following example shows a properly formatted incident file:

```
01:40:19.601363716,root,init
01:40:20.606013716,nobody,bash
01:40:21.137163716,1000,tar
```

Keep the tool's original
timestamp-format as-is.



Make sure to store the
incident file on the cluster's
worker node.



Options:

A- Explanation:

```
candidate@cli:~$ kubectl config use-context KSRS00101
Switched to context "KSRS00101".
candidate@cli:~$ ssh ksrs00101-worker1
Warning: Permanently added '10.240.86.96' (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@ksrs00101-worker1:~# falco
falco                falco-driver-loader
root@ksrs00101-worker1:~# ls -l /etc/falco/
total 200
-rw-r--r-- 1 root root  12399 Jan 31 16:06 aws_cloudtrail_rules.yaml
-rw-r--r-- 1 root root  11384 Jan 31 16:06 falco.yaml
-rw-r--r-- 1 root root   1136 Jan 31 16:06 falco_rules.local.yaml
-rw-r--r-- 1 root root 132112 Jan 31 16:06 falco_rules.yaml
-rw-r--r-- 1 root root  27289 Jan 31 16:06 k8s_audit_rules.yaml
drwxr-xr-x 2 root root   4096 Feb 16 01:07 rules.available
drwxr-xr-x 2 root root   4096 Jan 31 16:28 rules.d
root@ksrs00101-worker1:~# vim /etc/falco/falco_rules.local.yaml
```



```
# Copyright (C) 2019 The Falco Authors.
#
#
# Licensed under the Apache License, Version 2.0 (the "License");
# you may not use this file except in compliance with the License.
# You may obtain a copy of the License at
#
#     http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
# See the License for the specific language governing permissions and
# limitations under the License.
#

#####
# Your custom rules!
#####

# Add new rules, like this one
# - rule: The program "sudo" is run in a container
#   desc: An event will trigger every time you run sudo in a container
#   condition: evt.type = execve and evt.dir=< and container.id != host and proc.name = sudo
#   output: "Sudo run in container (user=%user.name %container.info parent=%proc.pname cmdli
ne=%proc.cmdline)"
#   priority: ERROR
#   tags: [users, container]

# Or override/append to any rule, macro, or list from the Default Rules
- rule: Container Drift Detected (chmod)
  desc: New executable created in a container due to chmod
  condition: >
    out.type in (open openat create) and
```

```
root@ksrs00101-worker1:~# vim /etc/falco/falco_rules.local.yaml
root@ksrs00101-worker1:~# systemctl status falco.service
● falco.service - Falco Runtime Security
   Loaded: loaded (/lib/systemd/system/falco.service; disabled; vendor preset: enabled)
   Active: inactive (dead)
root@ksrs00101-worker1:~# systemctl enable falco.service
Created symlink /etc/systemd/system/multi-user.target.wants/falco.service → /lib/systemd/system/falco.service.
root@ksrs00101-worker1:~# systemctl start falco.service
root@ksrs00101-worker1:~# exit
logout
Connection to 10.240.86.96 closed.
candidate@cli:~$ ssh ksrs00101-worker1
Last login: Fri May 20 15:59:48 2022 from 10.240.86.88
root@ksrs00101-worker1:~# vim /etc/falco/falco.yaml █
```



```
# When using json output, whether or not to include the "tags" property
# itself in the json output. If set to true, outputs caused by rules
# with no tags will have a "tags" field set to an empty array. If set to
# false, the "tags" field will not be included in the json output at all.
json_include_tags_property: true

# Send information logs to stderr and/or syslog Note these are *not* security
# notification logs! These are just Falco lifecycle (and possibly error) logs.
log_stderr: true
log_syslog: true
log_file: /opt/KSRS00101/alerts/details

# Minimum log level to include in logs. Note: these levels are
# separate from the priority field of rules. This refers only to the
# log level of falco's internal logging. Can be one of "emergency",
# "alert", "critical", "error", "warning", "notice", "info", "debug".
log_level: info
```

```
root@ksrs00101-worker1:~# vim /etc/falco/falco.yaml
root@ksrs00101-worker1:~# grep log /etc/falco/falco.yaml
# cloudtrail log files.
# If true, the times displayed in log messages and output messages
# Send information logs to stderr and/or syslog Note these are *not* security
# notification logs! These are just Falco lifecycle (and possibly error) logs.
log_stderr: true
log_syslog: true
log_file: /opt/KSRS00101/alerts/details
# Minimum log level to include in logs. Note: these levels are
# log level of falco's internal logging. Can be one of "emergency",
log_level: info
#   - log: log a DEBUG message noting that the buffer was full
# Notice it is not possible to ignore and log/alert messages at the same time.
# The rate at which log/alert messages are emitted is governed by a
#   - log
# The timeout error will be reported to the log according to the above log_* settings.
syslog_output:
#   - logging (alternate method than syslog):
#       program: logger -t falco-test
# this information will be logged, however the main Falco daemon will not be stopped.
root@ksrs00101-worker1:~# systemctl restart falco.service
root@ksrs00101-worker1:~# exit
logout
Connection to 10.240.86.96 closed.
candidate@cli:~$ █
```


Answer:

A

Question 6

Question Type: MultipleChoice

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 your suggestion on it.

B- Send us your suggestion

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>

