



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**

CORRECT TEXT



You **must** complete this task on the following cluster/nodes:

Cluster	Master	Worker
	node	node
KSRS001	ksrs00101	kerc00101
01	-master	KSISUUTUT-
		worker1

You can switch the cluster/configuration context using the following command:

[candidate@cli] \$ kubec
tl config use-context KS
RS00101

You may use your browser (1) to open **one additional tab** to access Falco's documentation.



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

1.

sysdig

2.

falco

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:



The following example shows a properly formatted incident file:

01:40:19.601363716, root, init 01:40:20.606013716, nobody, ba sh 01:40:21.137163716,1000,tar

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



- A. See the explanation below:
- B. PlaceHolder



Correct Answer: A

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 -1 /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





Text



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/sys tem/falco.service. root@ksrs00101-worker1:~# systemct1 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 og file: /opt/KSRS00101/alerts/details g level: info

Text



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", g 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:~\$

**QUESTION 2** 



VCE & PDF Pass4itSure.com

networking.k8s.io/v1 d: NetworkPolicy namespace: dev-team - Ingress matchLabels: environment: dev environment: testing candidate@cli:~\$ vim np.yaml candidate@cli:~\$ cat np.yaml apiVersion: networking.k8s.io/v1 kind: NetworkPolicy metadata: name: pod-access namespace: dev-team pec: podSelector: matchLabels: environment: dev policyTypes: - Ingress ingress: - from: namespaceSelector: matchLabels: matchLabers. environment: dev - podSelector: matchLabels: environment: testing matchLabels: environment: testing candidate@cli:-\$ candidate@cli:-\$ candidate@cli:-\$ candidate@cli:-\$ candidate@cli:-\$ kubectl create -f np.yaml -n dev-team networkpolicy.networking.k8s.io/pod-access created candidate@cli:-\$ kubectl describe netpol -n dev-team Name: pod-access Namespace: dev-team Created on: 2022-05-20 15:35:33 +0000 UTC Labels: <none> Annotations: <none> Spec: Inotations, bec: PodSelector: environment=dev Allowing ingress traffic: To Port: <any> (traffic allowed to all ports) From: NamespaceSelector: environment=dev From: PodSelector: environment=testing From: PodSelector: environment=testing Not affecting egress traffic Policy Types: Ingress Indidate@cli:-& cat KSSH00301/network-policy.yaml apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
 name: ""
 namespace: "" spec: podSelector: {} policyTypes: - Ingress - Ingress ingress: - from: [] - from: [] candidate@cli:~\$ cp np.yaml KSSH00301/network-policy.yaml candidate@cli:~\$ cat KSSH00301/network-policy.yaml andidate@cli:~\$ cat KSSH00301/network-policy.yaml apiVersion: networking.k8s.io/v1 kind: NetworkPolicy netadata: name: pod-access namespace: dev-team spec: podSelector: matchLabels: environment: dev policyTypes: - Ingress ingress: from: - namespaceSelector: matchLabels: environment: dev - podSelector: matchLabels: environment: testing candidate@cli:~\$



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. See the explanation below:

B. PlaceHolder

Correct Answer: A

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 this flag disables log backend. - means standard out --audit-log-maxage defined the maximum number of days to retain old audit log files

--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 hostPath to the location of the policy file and log file, so that audit records are persisted.

For example:

--audit-policy-file=/etc/kubernetes/audit-policy.yaml \

--audit-log-path=/var/log/audit.log

#### **QUESTION 3**

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

/registry/secrets/default/ck k8s	s-secret
∰1 ●	
cks-secret default"*567fcb5 kubectl-create	3f-6b58-4fee-9f12-5737c764be742↔↔↔ (27+) ♦ (21) teldsv1:9
ey1 supersecret	<pre>{};"f:key2":{};"f:type":{}]</pre>

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. See explanation below.

B. PlaceHolder



Correct Answer: A

1.

ETCD secret encryption can be verified with the help of etcdctl command line utility.

2.

ETCD secrets are stored at the path /registry/secrets/\$namespace/\$secret on the master node.

3.

The below command can be used to verify if the particular ETCD secret is encrypted or not.

# ETCDCTL\_API=3 etcdctl get /registry/secrets/default/secret1 [...] | hexdump -C

#### **QUESTION 4**

You can switch the cluster/configuration context using the following command:

[desk@cli] \$ kubectl config use-context prod-account

Context:

A Role bound to a Pod\\'s ServiceAccount grants overly permissive permissions. Complete the following tasks to reduce the set of permissions.

Task:

Given an existing Pod named web-pod running in the namespace database.

1.

Edit the existing Role bound to the Pod\\'s ServiceAccount test-sa to only allow performing get operations, only on resources of type Pods.

2.

Create a new Role named test-role-2 in the namespace database, which only allows performing update operations, only on resources of type statuefulsets.

3.

Create a new RoleBinding named test-role-2-bind binding the newly created Role to the Pod\\'s ServiceAccount. Note: Don\\'t delete the existing RoleBinding.

A. See the explanation below

B. PlaceHolder

Correct Answer: A



candidate@cli:~\$ kubectl config use-context KSCH00201 Switched to context "KSCH00201". candidate@cli:~\$ kubectl get pods -n security READY STATUS RESTARTS NAME 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



VCE & PDF Pass4itSure.com

b4c9ddd6-2729-43bd-8fbd-b2d227f4c4cd services watch candidate@cli:~\$ kubectl describe role dev-role -n security Name: dev-role Labels: <none> Annotations: <none> PolicyRule: Non-Resource URLs Resource Names Resources 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 1/1 Running 0 6h12m web-pod candidate@cli:~\$ kubectl get pods/web-pod -n security -o yaml | grep serviceAccount Account: sa-dev-1 AccountName: sa-dev-1 vice - service ountToken: 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:~\$ 🗍

#### **QUESTION 5**

Create a network policy named allow-np, that allows pod in the namespace staging to connect to port 80 of other pods in the same namespace.

Ensure that Network Policy:

1.

Does not allow access to pod not listening on port 80.

2.



Does not allow access from Pods, not in namespace staging.

A. See the explanation below:

B. PlaceHolder

Correct Answer: A

apiVersion: networking.k8s.io/v1

kind: NetworkPolicy

metadata:

name: network-policy

spec:

podSelector: {} #selects all the pods in the namespace deployed policyTypes:

-Ingress ingress:

-ports: #in input traffic allowed only through 80 port only

-protocol: TCP port: 80

#### **QUESTION 6**

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-