

CKS Pdf Torrent | CKS Paper



P.S. Free & New CKS dumps are available on Google Drive shared by Test4Sure: <https://drive.google.com/open?id=1QN1yCibnVghoGpUhcTFGjgi2XDiCKQhX>

When preparing for the test CKS certification, most clients choose our products because our CKS learning file enjoys high reputation and boost high passing rate. Our products are the masterpiece of our company and designed especially for the certification. Our CKS latest study question has gone through strict analysis and verification by the industry experts and senior published authors. The clients trust our products and treat our products as the first choice. So the total amounts of the clients and the sales volume of our CKS learning file is constantly increasing.

The CKS certification exam is a rigorous and challenging test of the candidate's knowledge and skills in securing Kubernetes platforms. CKS exam consists of 17 questions, which are a combination of multiple-choice and hands-on tasks. The hands-on tasks require the candidate to demonstrate their ability to perform specific security-related tasks in a Kubernetes cluster. CKS Exam is conducted online and is proctored to ensure the integrity of the certification process.

>> **CKS Pdf Torrent** <<

100% Pass-Rate CKS Pdf Torrent Spend Your Little Time and Energy to Pass CKS exam one time

At Test4Sure, we strive hard to offer a comprehensive Certified Kubernetes Security Specialist (CKS) (CKS) exam questions preparation material bundle pack. The product available at Test4Sure includes Certified Kubernetes Security Specialist (CKS) (CKS) real dumps pdf and mock tests (desktop and web-based). Practice exams give an experience of taking the Certified Kubernetes Security Specialist (CKS) (CKS) actual exam.

Linux Foundation Certified Kubernetes Security Specialist (CKS) Sample Questions (Q173-Q178):

NEW QUESTION # 173

You are running a critical application within a Kubernetes cluster. Your application relies on a base image with several unnecessary packages installed. These packages increase the attack surface of your application and make it more vulnerable to exploits. You

want to minimize the base image footprint to enhance the security posture of your application. Explain how you can achieve this in a production environment.

Answer:

Explanation:

Solution (Step by Step) :

1. Identify unnecessary Packages:

- Use tools like 'alpine-pkg-info' or 'dpkg -l' to list installed packages within the base image.
- Analyze the package list to identify packages that are not strictly required for your application's functionality.
- Example: If you are running a Node.js application, you might identify development tools like 'gcc' or 'make' as unnecessary.

2. Create a Custom Base Image:

- Dockerfile: Start by creating a Dockerfile that inherits from a minimal base image like 'alpine:latest' or 'ubuntu:latest' (depending on your application's requirements).

- Install Essential Packages: Include only the absolutely necessary packages for your application in the Dockerfile. Use the 'apt-get install' (for Debian/Ubuntu) or 'apk add' (for Alpine) commands to install these packages.

- Example Dockerfile:

```
FROM alpine:latest
# Install necessary packages
RUN apk add --no-cache bash openssl curl nodejs npm
# Copy your application code
COPY ./app
# Set working directory and execute start script
WORKDIR ./app
CMD ["npm", "start"]
```

3. Test the Custom Image:

- Build the custom image using 'docker build -t custom-base-image'
- Create a container from the custom image and run your application to ensure everything works correctly. This step is critical to catch any compatibility issues before deploying to your Kubernetes cluster.

4. Update Your Deployments:

- Modify your Deployment YAML files to use the custom base image instead of the original image. Update the 'image' field to reference the custom base image tag.

- Example:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-app
spec:
  replicas: 3
  selector:
    matchLabels:
      app: my-app
  template:
    metadata:
      labels:
        app: my-app
    spec:
      containers:
        - name: my-app
          image: custom-base-image
          ports:
            - containerPort: 8080
```

5. Deploy the Updated Application: - Use 'kubectl apply -f deployment.yaml' to update your deployment with the new image. - Monitor the deployment to ensure a successful rollout with your minimal base image. 6. Regular - Periodically review your application's requirements and ensure that the base image still meets your needs. - As you add new features or update dependencies, you might need to add additional packages to the base image. - Keep the image as minimal as possible and use the least-privilege principle when selecting packages.

NEW QUESTION # 174

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

Cluster: apparmor

Master node: master

Worker node: worker1

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

```
[desk@cli] $ kubectl config use-context apparmor
```

Given: AppArmor is enabled on the worker1 node.

Task:

On the worker1 node,

1. Enforce the prepared AppArmor profile located at: /etc/apparmor.d/nginx
2. Edit the prepared manifest file located at /home/cert_masters/nginx.yaml to apply the apparmor profile
3. Create the Pod using this manifest

Answer:

Explanation:

```
[desk@cli] $ ssh worker1
[worker1@cli] $ apparmor_parser -q /etc/apparmor.d/nginx
[worker1@cli] $ aa-status | grep nginx
```

nginx-profile-1

```
[worker1@cli] $ logout
```

```
[desk@cli] $ vim nginx-deploy.yaml
```

Add these lines under metadata:

annotations: # Add this line

```
container.apparmor.security.beta.kubernetes.io/<container-name>: localhost/nginx-profile-1
```

```
[desk@cli] $ kubectl apply -f nginx-deploy.yaml
```

Explanation

```
[desk@cli] $ ssh worker1
```

```
[worker1@cli] $ apparmor_parser -q /etc/apparmor.d/nginx
```

```
[worker1@cli] $ aa-status | grep nginx
```

nginx-profile-1

```
[worker1@cli] $ logout
```

```
[desk@cli] $ vim nginx-deploy.yaml
```

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx-deploy
  annotations:
    container.apparmor.security.beta.kubernetes.io/hello: localhost/nginx-profile-1
  Add this line
    Add this line be sure that container name is hello here, not nginx-deploy
spec:
  containers:
  - name: hello
    image: nginx
```

```
[desk@cli] $ kubectl apply -f nginx-deploy.yaml pod/nginx-deploy created Reference:
```

```
https://kubernetes.io/docs/tutorials/clusters/apparmor/ pod/nginx-deploy created
```

```
[desk@cli] $ kubectl apply -f nginx-deploy.yaml pod/nginx-deploy created Reference:
```

```
https://kubernetes.io/docs/tutorials/clusters/apparmor/
```

NEW QUESTION # 175

Context

A container image scanner is set up on the cluster, but it's not yet fully integrated into the cluster's configuration. When complete, the container image scanner shall scan for and reject the use of vulnerable images.

Task

 To complete the entire task on the cluster's master node, where all services and files have been prepared and placed.

Given an incomplete configuration in directory /etc/kubernetes/epconfig and a functional container image scanner with HTTPS endpoint https://wakanda.local:8081 /image_policy :

1. Enable the necessary plugins to create an image policy
2. Validate the control configuration and change it to an implicit deny
3. Edit the configuration to point to the provided HTTPS endpoint correctly Finally, test if the configuration is working by trying to deploy the vulnerable resource /root/KSSC00202/vulnerable-resource.yml

You can find the container image scanner's log file at
`/var/log/imagepolicy/acme.log`



 THE
LINUX
FOUNDATION

Answer:

Explanation:

```
Switched to context "KSSC00202".
candidate@cli:~$ ssh kssc00202-master
Warning: Permanently added '10.177.80.12' (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@kssc00202-master:~# ls /etc/kubernetes/epconfig/
admission_configuration.json  apiserver-client.key.pem  apiserver-client.pem  kubeconfig.yaml  webhook-key.pem  webhook.pem
root@kssc00202-master:~# vim /etc/kubernetes/epconfig/admission_configuration.json

"imagePolicy": {
  "kubeConfigFile": "/etc/kubernetes/epconfig/kubeconfig.yaml",
  "allowTTL": 50,
  "denyTTL": 50,
  "retryBackoff": 100,
  "defaultAllow": false}

root@kssc00202-master:~# vim /etc/kubernetes/epconfig/admission_configuration.json
root@kssc00202-master:~# vim /etc/kubernetes/epconfig/admission_configuration.json
root@kssc00202-master:~# vim /etc/kubernetes/epconfig/kubeconfig.yaml

apiVersion: v1
clusters:
- cluster:
  certificate-authority: /etc/kubernetes/epconfig/webhook.pem # CA for verifying the remote service.
  server: https://wakanda.local:8081/image_policy
  name: kubernetes
contexts:
- context:
  cluster: kubernetes
  user: kubernetes-admin
  name: kubernetes-admin@kubernetes
current-context: kubernetes-admin@kubernetes
kind: Config
preferences: {}
users:
- name: kubernetes-admin
  user:
    client-certificate: /etc/kubernetes/epconfig/apiserver-client.pem
    client-key: /etc/kubernetes/epconfig/apiserver-client-key.pem

root@kssc00202-master:~# vim /etc/kubernetes/epconfig/admission_configuration.json
root@kssc00202-master:~# vim /etc/kubernetes/epconfig/admission_configuration.json
root@kssc00202-master:~# vim /etc/kubernetes/epconfig/kubeconfig.yaml
root@kssc00202-master:~# vim /etc/kubernetes/manifests/kube-apiserver.yaml p
```



```
apiVersion: v1
kind: Pod
metadata:
  annotations:
    kubeadm.kubernetes.io/kube-apiserver.advertise-address.endpoint: 10.177.80.12:6443
  creationTimestamp: null
  labels:
    component: kube-apiserver
    tier: control-plane
  name: kube-apiserver
  namespace: kube-system
spec:
  containers:
    - command:
        - kube-apiserver
        - --advertise-address=10.177.80.12
        - --allow-privileged=true
        - --authorization-mode=Node, RBAC
        - --client-ca-file=/etc/kubernetes/pki/ca.crt
        - --enable-admission-plugins=NodeRestriction
        - --enable-bootstrap-token-auth=true
        - --etcd-cafile=/etc/kubernetes/pki/etcd/ca.crt
        - --etcd-certfile=/etc/kubernetes/pki/apiserver-etcd-client.crt
        - --etcd-keyfile=/etc/kubernetes/pki/apiserver-etcd-client.key
        - --etcd-servers=https://127.0.0.1:2379
        - --kubelet-client-certificate=/etc/kubernetes/pki/apiserver-kubelet-client.crt
        - --kubelet-client-key=/etc/kubernetes/pki/apiserver-kubelet-client.key
        - --kubelet-preferred-address-types=InternalIP,ExternalIP,Hostname
        - --proxy-client-cert-file=/etc/kubernetes/pki/front-proxy-client.crt
        - --proxy-client-key-file=/etc/kubernetes/pki/front-proxy-client.key
        - --requestheader-allowed-names=front-proxy-client
        - --requestheader-client-ca-file=/etc/kubernetes/pki/front-proxy-ca.crt
        - --requestheader-extra-headers-prefix=X-Remote-Extra-
      "/etc/kubernetes/manifests/kube-apiserver.yaml" 135L, 4626C
```



```
root@kssc00202-master:~# vim /etc/kubernetes/manifests/kube-apiserver.yaml p
2 files to edit
root@kssc00202-master:~# rm -f p
root@kssc00202-master:~# vim /etc/kubernetes/manifests/kube-apiserver.yaml
```

```
apiVersion: v1
kind: Pod
metadata:
  annotations:
    kubeadm.kubernetes.io/kube-apiserver.advertise-address.endpoint: 10.177.80.12:6443
  creationTimestamp: null
  labels:
    component: kube-apiserver
    tier: control-plane
  name: kube-apiserver
  namespace: kube-system
spec:
  containers:
    - command:
        - kube-apiserver
        - --advertise-address=10.177.80.12
        - --allow-privileged=true
        - --authorization-mode=Node, RBAC
        - --client-ca-file=/etc/kubernetes/pki/ca.crt
        - --enable-admission-plugins=NodeRestriction, ImagePolicyWebhook
        - --admission-control-config-file=/etc/kubernetes/epconfig/admin.conf
        - --enable-bootstrap-token-auth=true
        - --etcd-cafile=/etc/kubernetes/pki/etcd/ca.crt
        - --etcd-certfile=/etc/kubernetes/pki/apiserver-etcd-client.crt
        - --etcd-keyfile=/etc/kubernetes/pki/apiserver-etcd-client.key
        - --etcd-servers=https://127.0.0.1:2379
        - --kubelet-client-certificate=/etc/kubernetes/pki/apiserver-kubelet-client.crt
        - --kubelet-client-key=/etc/kubernetes/pki/apiserver-kubelet-client.key
        - --kubelet-preferred-address-types=InternalIP, ExternalIP, Hostname
        - --proxy-client-cert-file=/etc/kubernetes/pki/front-proxy-client.crt
        - --proxy-client-key-file=/etc/kubernetes/pki/front-proxy-client.key
        - --requestheader-allowed-names=front-proxy-client
        - --requestheader-client-ca-file=/etc/kubernetes/pki/front-proxy-ca.crt
```

```

root@kssc00202-master:~# rm -f p
root@kssc00202-master:~# vim /etc/kubernetes/manifests/kube-apiserver.yaml
root@kssc00202-master:~# systemctl daemon-reload
root@kssc00202-master:~#
root@kssc00202-master:~# systemctl restart kubelet.service
root@kssc00202-master:~# systemctl enable kubelet.service
root@kssc00202-master:~#
root@kssc00202-master:~# systemctl status kubelet
root@kssc00202-master:~# ls
KSSC00202 snap
root@kssc00202-master:~# cat KSSC00202/vulnerable-resource.yml
root@kssc00202-master:~# cat KSSC00202/vulnerable-resource.yml
apiVersion: v1
kind: ReplicationController
metadata:
  name: nginx-latest
spec:
  replicas: 1
  selector:
    app: nginx-latest
  template:
    metadata:
      name: nginx-latest
      labels:
        app: nginx-latest
    spec:
      containers:
        - name: nginx-latest
          image: nginx
          ports:
            - containerPort: 80
root@kssc00202-master:~# kubectl create -f KSSC00202/vulnerable-resource.yml
root@kssc00202-master:~# kubectl create -f KSSC00202/vulnerable-resource.yml
The connection to the server 10.177.80.12:6443 was refused - did you specify the right host or port?
root@kssc00202-master:~# kubectl get pods
The connection to the server 10.177.80.12:6443 was refused - did you specify the right host or port?
root@kssc00202-master:~# ls -al .kube/
total 20
drwxr-xr-x  3 root root 4096 Aug  3 04:07 .
drwx-----  9 root root 4096 Oct 11 15:36 ..
drwxr-x---  4 root root 4096 Aug  3 04:07 cache
-rw-r--r--  1 root root 5636 Aug  3 04:07 config
root@kssc00202-master:~# crictl ps -a
[...]
012ea8587130e      a634548d10b03      2 months ago      Exited      kube-proxy      0      1460a9f
a0f1e0      kube-proxy-cm1b5      405227dfa49d0      aeb0758cef4cd      2 months ago      Exited      etcd      0      cfb6522
e720fb      etcd.kssc00202-master
root@kssc00202-master:~# ls -al .kube/ | grep kube-api
root@kssc00202-master:~# crictl ps -a | grep kube-api
WARN[0000] runtime connect using default endpoints: [unix:///var/run/docker-shim.sock unix:///run/containerd/containerd.sock unix:///run/crio/crio.sock unix:///var/run/cri-dockerd.sock]. As the default settings are now deprecated, you should set the endpoint instead.
ERROR[0000] unable to determine runtime API version: rpc error: code =Unavailable desc = connection error: desc = "transport: Error while dialing dial unix /var/run/docker-shim.sock: connect: no such file or directory"
WARN[0000] image connect using default endpoints: [unix:///var/run/docker-shim.sock unix:///run/containerd/containerd.sock unix:///run/crio/crio.sock unix:///var/run/cri-dockerd.sock]. As the default settings are now deprecated, you should set the endpoint instead.
ERROR[0000] unable to determine image API version: rpc error: code =Unavailable desc = connection error: desc = "transport: Error while dialing dial unix /var/run/docker-shim.sock: connect: no such file or directory"
a003b3fdfb61c      d3377ffb7177c      30 seconds ago      Exited      kube-apiserver      3      2dadab4e
984a91      kube-apiserver-kssc00202-master
5e70b9a70f9ed      d3377ffb7177c      7 hours ago      Exited      kube-apiserver      0      68a9f31
6c2559      kube-apiserver-kssc00202-master
root@kssc00202-master:~#
root@kssc00202-master:~#
root@kssc00202-master:~#
root@kssc00202-master:~# exit
logout
Connection to 10.177.80.12 closed.
candidate@cli:~$ 
```

NEW QUESTION # 176

You are running a Kubernetes cluster with a deployment named "my-app" that has been experiencing unexpected crashes. The crash logs indicate that the container's memory consumption is exceeding the resource limits defined in the deployment YAML. Explain how you can utilize the Kubernetes resource quotas and admission controller to prevent this from happening again.

Answer:

Explanation:

Solution (Step by Step) :

1. Create a ResourceQuota:

- Define a ResourceQuota that limits the resources that can be consumed by pods in a specific namespace.
- Specify the limits for CPU, memory, storage, and other resources.
- For example, to limit memory usage to 2Gi per pod in the "my-app" namespace:

```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: memory-limit
  namespace: my-app
spec:
  limits:
    memory: "2Gi"
```

2. Enable the ResourceQuota Admission Controller: - Ensure that the "ResourceQuota" admission controller is enabled in your Kubernetes cluster. This can usually be done by setting the 'admissioncontroller' flag in the 'kube-apiserver' configuration. 3. Apply the ResourceQuota: - Apply the ResourceQuota to the "my-app" namespace using 'kubectl apply -f resource-quota.yaml'. 4. Update the Deployment - Modify the deployment's YAML file to specify the resource requests and limits for the container, ensuring they are within the defined ResourceQuota limits. For example:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-app
spec:
  template:
    # ...
    spec:
      containers:
        - name: my-app-container
          resources:
            requests:
              memory: "1Gi"
            limits:
              memory: "2Gi"
```

5. Apply the updated deployment - Apply the updated deployment using 'kubectl apply -f deployment.yaml'. 6. Monitor and Evaluate: - Monitor the resource consumption of pods in the "my-app" namespace and adjust the ResourceQuota limits as needed to ensure that your cluster remains stable.

NEW QUESTION # 177

You are responsible for securing the software supply chain for your organization, which uses a GitLab CI/CD pipeline to build and deploy containerized applications. You want to implement a robust mechanism to prevent unauthorized code changes from being introduced into the production environment. How would you utilize GitLab's built-in features and best practices to achieve this goal?

Answer:

Explanation:

Solution (Step by Step) :

1. Implement Code Review:

- Configure GitLab to enforce mandatory code reviews for all changes to production branches.
- Configure a minimum number of reviewers required for each merge request.
- Use GitLab's built-in code review features to facilitate discussion and feedback.

2. Enable Branch Protection:

- Protect the production branch by configuring the following:
 - Allow only specific users or groups to merge: Restrict merge rights to authorized personnel.
 - Require merge requests: Prevent direct pushes to the production branch.
- Enforce minimum approval count: Mandate a specific number of approvals for each merge request

3. Utilize GitLab CI/CD Security Features:

- Dependency Scanning: Integrate GitLab's dependency scanning feature to analyze your code for known vulnerabilities in external libraries.
- Container Scanning: use GitLab's container scanning feature to check for vulnerabilities in your Docker images before deployment.
- SAST (Static Application Security Testing): Integrate a SAST tool into your CI/CD pipeline to identify potential vulnerabilities in your code.

- DAST (Dynamic Application Security Testing): Utilize a DAST tool to test your application for security flaws during runtime.

4. Enforce Access Control:

- Implement role-based access control (RBAC) within GitLab.

- Assign roles with specific permissions to users and groups based on their responsibilities.

- Audit user activity regularly to identify any suspicious behavior.

5. Utilize GitLab's Security Integrations:

- Integrate with Vulnerability Databases: Connect your GitLab instance to vulnerability databases such as NIST NVD to receive alerts about newly discovered vulnerabilities.

- Integrate with Security Tools: Connect GitLab with security tools like Security Information and Event Management (SIEM) systems to automate vulnerability reporting and incident response.

6. Develop a Secure Coding Culture:

- Promote secure coding practices within your development team.

- Provide training and resources on secure coding principles.

- Conduct regular code reviews to catch potential vulnerabilities.

```

# Example GitLab CI/CD YAML file for secure deployment
stages:
  - build
  - test
  - scan
  - deploy

variables:
  # Configure your security tool integration
  # Example: Snyk for dependency scanning
  SNYK_TOKEN: "$SNYK_TOKEN"

build:
  stage: build
  image: docker:latest
  script:
    - docker build -t my-app .
    - docker push my-app

test:
  stage: test
  image: node:latest
  script:
    - npm install
    - npm test

scan:
  stage: scan
  image: snyk/snyk
  script:
    - snyk test
  # Configure snyk to fail the pipeline if vulnerabilities are found
  # Use Snyk's API to integrate with GitLab for vulnerability reports
  environment:
    SNYK_TOKEN: $SNYK_TOKEN

deploy:
  stage: deploy
  image: docker:latest
  script:
    - docker login -u $DOCKER_USER -p $DOCKER_PASSWORD $DOCKER_REGISTRY
    - docker push my-app:latest
    - kubectl apply -f deployment.yaml

```



NEW QUESTION # 178

.....

Being anxious for the exam ahead of you? Have a look of our CKS practice materials please. Presiding over the line of CKS practice materials over ten years, our experts are proficient as elites who made our CKS practice materials, and it is their job to officiate the routines of offering help for you. All points are predominantly related with the exam ahead of you. Every page is full of well-turned words for your reference related wholly with the real exam.

CKS Paper: <https://www.test4sure.com/CKS-pass4sure-vce.html>

- CKS Study Questions - Certified Kubernetes Security Specialist (CKS) Guide Torrent -amp; CKS Exam Torrent □

Immediately open [www.vce4dumps.com] and search for (CKS) to obtain a free download CKS Pass Guide

DOWNLOAD the newest Test4Sure CKS PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1QN1yCibnVghoGpUhcTGFjgj2XDiCKQhX>