


CKS日本語認定対策 & CKS問題数



BONUS!!! TopexamCKSダンプの一部を無料でダウンロード: https://drive.google.com/open?id=11Ryg_zDYsUj2lBshg9DKLyU6BfR590t

すべての人々がCKS試験に合格し、関連する認定を短時間で取得できるように、3つの異なるバージョンのCKS学習教材を設計しました。製品は、すべての人が同時に学習とテストを行うための実際の試験をシミュレートすることを試みることができ、学習コースでの学習不足に適した環境を提供することができます。当社からCKS学習教材を購入して使用すると、実際の試験のようにCKS学習テストを練習し、CKS試験に簡単に合格できます。

CKS試験は、Kubernetesセキュリティに関連する実世界の問題を解決する候補者の能力をテストする実践的でパフォーマンスベースの試験です。候補者は、所定の時間内に Kubernetes クラスターとアプリケーションのセキュリティに関連するタスクを実行する必要があります。試験はオンラインで実施され、候補者は世界中から受験できます。

CKS認定の対象となるには、候補者は、現在の認定されたKubernetes管理者（CKA）認定またはKubernetes Fundamentals（LFS258）コースの合格スコアを持つ必要があります。CKS認定試験は、15～20のパフォーマンスベースのタスクで構成される、提案されたオンライン試験です。候補者には試験を完了するのに2時間があり、合格するには少なくとも66%を獲得する必要があります。この試験は複数の言語で利用でき、世界のどこからでも撮影できます。

>> CKS日本語認定対策 <<

CKS問題数 & CKS日本語関連対策

Topexam世界は急速に変化しており、Linux Foundation従業員に対する要件はこれまでにない高くなっています。理想的な仕事を見つけて高収入を得たい場合は、優れた労働能力と深い知識を高めなければなりません。CKSのCertified Kubernetes Security Specialist (CKS)認定に合格すると、夢を実現できます。製品を購入すると、最高のCertified Kubernetes Security Specialist (CKS)学習教材が提供され、Certified Kubernetes Security Specialist (CKS)認定の取得に役立ちます。当社の製品は高品質であり、当社のサービスは完璧です。

Linux Foundation Certified Kubernetes Security Specialist (CKS) 認定 CKS 試験問題 (Q49-Q54):

質問 # 49

SIMULATION

A container image scanner is set up on the cluster.

Given an incomplete configuration in the directory

/etc/Kubernetes/conf/control and a functional container image scanner with HTTPS endpoint `https://acme.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 the latest.

- A. Send us the Feedback on it.

正解: A

質問 # 50

You are running a multi-tenant Kubernetes cluster where different teams manage their own applications. You want to ensure that each team's applications are isolated from each other to prevent potential security risks.

How would you use Network Policies to achieve this isolation?

正解:

解説:

Solution (Step by Step) :

1. Create separate namespaces for each team: This is the foundation for network isolation.

2. Define Network Policies for each namespace:

Restrict inbound traffic: Only allow specific protocols and ports from trusted sources to access the namespace.

Control outbound traffic: Limit outbound connections from pods within the namespace to specific destinations.

Example Network Policy (ingress):

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: team-a-ingress
  namespace: team-a
spec:
  podSelector: {}
  ingress:
    - from:
      - podSelector:
          matchLabels:
            app: allowed-service
      namespaceSelector:
          matchLabels:
            team: trusted-team
    egress: []
```

3. Apply the Network Policies: `bash kubectl apply -f team-a-ingress.yaml kubectl apply -f team-b-ingress.yaml` # Apply policies for other namespaces Example Scenario: Team A: Runs a web application accessible only from within its namespace. Team B: Runs a database service that can be accessed by Team A's application but not by other teams. Network Policies: Team A Network Policy: Ingress: Only allow traffic from Team B's database service- Egress Allow outbound traffic to the internet for updates and dependencies. Team B Network Policy: Ingress: Allow traffic from Team A's web application Egress Limit outbound traffic to specific servers for backups and maintenance. Note: Network Policies are a powerful tool for achieving network isolation in Kubernetes. They allow you to fine-tune the communication patterns between pods and namespaces, enhancing security and mitigating potential risks.

質問 # 51

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 statefulsets. 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.

正解:

解説:

```
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:~$

```

質問 #52

You are managing a Kubernetes cluster with a deployment named 'web-app' that runs multiple pods. These pods access a database hosted on a separate Kubernetes service named 'database-service'. You need to ensure that the pods can only connect to the database service and are restricted from accessing other services in the cluster.

正解:

解説:

Solution (Step by Step) :

1. Create a Network Policy:

- Create a Network Policy YAML file named 'web-app-policy.yaml' with the following content

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: web-app-policy
  namespace: default # Replace "default" with your namespace
spec:
  podSelector:
    matchLabels:
      app: web-app
  ingress:
    - from:
      - podSelector:
          matchLabels:
            app: database-service # Allow communication from pods with label "app: database-service"
  egress:
    - to:
      - podSelector:
          matchLabels:
            app: database-service # Allow communication to pods with label "app: database-service"

```

2. Apply the Network Policy: - Apply the policy using Skubectl apply -f web-app-policy.yaml 3. Verify the Network Policy: -

Check the status of the Network Policy using 'kubectl get networkpolicies' to confirm that it has been applied. 4. Test the

Connectivity: - Try to access the database service from a pod in the 'web-apps deployment. - Attempt to access other services from the same pod. - You should only be able to access the 'database-service' due to the Network Policy restrictions.

質問 # 53

You are running a Kubernetes cluster with a deployment named "my-app" that uses a container image from a private registry- You suspect that a recent deployment update may have introduced a vulnerability in one of the containers. Explain how you would use pod security policies and runtime security tools like Falco to investigate and mitigate this potential security risk.

正解:

解説:

Solution (Step by Step) :

1. Enable Pod Security Policies (PSPs):

- Create a PSP that restricts the container's capabilities, resource usage, and network access. For example, limit the container's resource requests and limits, restrict access to sensitive host resources, and disallow the creation of privileged containers.

```
apiVersion: policy/v1beta1
kind: PodSecurityPolicy
metadata:
  name: restricted-pod-security-policy
spec:
  # Define the allowed capabilities, resource limits, and network access
  # For example, disallow privileged containers and limit resource requests
  # ...
```

- Apply the PSP to your deployment using the 'podSecurityPolicy' field in the deployment's spec:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-app
spec:
  template:
    # ...
    spec:
      # ...
      securityContext:
        podSecurityPolicy: restricted-pod-security-policy
```

2. Implement Runtime Security with Falco: - Deploy Falco on your Kubernetes cluster using a Daemon set Falco uses rules to detect suspicious activity within containers. - Configure Falco rules to alert on events like: - Container privilege escalation attempts: Monitor for attempts to run containers with elevated privileges. - Suspicious network connections: Detect unexpected network connections originating from within containers- - File system modifications: Identify unauthorized modifications to files within container images. - Define a Falco rule that detects attempts to access sensitive files:

```
- rule: Container accesses sensitive file
desc: Detect attempts to access sensitive files in the container's filesystem
condition:
  - event: container.file.read
  - k8s.ns: my-app # Namespace of the deployment
  - k8s.pod: my-app # Pod name
  - container.file.path: /etc/passwd
output: Container in Pod %k8s.pod in Namespace %k8s.ns attempted to read sensitive file: %container.file.path
priority: high
```

3. Investigate the Vulnerability: - Review the Falco alerts to identify specific suspicious activities related to the deployment update. - Analyze the logs and metrics of the affected container to gather further evidence of the vulnerability. - Inspect the container image for any unauthorized changes or vulnerabilities. 4. Mitigate the Vulnerability: - If a vulnerability is identified, update the container image with a fix- - Redeploy the deployment with the updated container image. - Strengthen the PSP to prevent similar vulnerabilities from arising in the future. - Review and refine the Falco rules based on the findings of the investigation to improve the overall security posture.

質問 # 54

.....

製品がどれほど優れていても、ユーザーは使用過程でいくつかの難しい問題に遭遇します。CKSの実際の試験資料も例外ではありません。最高の製品体験を楽しむために、ユーザーが使用中のプロセスで問題が見つかった場合は、CKSを初めてチェックして、試験問題のパフォーマンス、ユーザーが問題を解決するのに役立つ専門のメンテナンススタッフ。CKSラーニングリファレンスファイルには、効率の良い製品メンテナンスチームがあり、数分でCKS試験の質問を送信できます。

CKS問題数: https://www.topexam.jp/CKS_shiken.html

- [illegible]

無料でクラウドストレージから最新のTopexamCKS PDFダンプをダウンロードする: <https://drive.google.com/open?id=11RygzDYsUji2lBshe9DKLyU6Bfr590t>