

# Latest CKS Exam Registration | 100% Free Valid Certified Kubernetes Security Specialist (CKS) Reliable Test Guide



P.S. Free & New CKS dumps are available on Google Drive shared by Getcertkey: <https://drive.google.com/open?id=1c8ZZ1KQskCxdZ4Lt75cjBRLGT9YEFFvH>

When you click into Getcertkey's site, you will see so many people daily enter the website. You can not help but be surprised. In fact, this is normal. Getcertkey is provide different training materials for alot of candidates. They are using our training materials tto pass the exam. This shows that our Linux Foundation CKS Exam Training materials can really play a role. If you want to buy, then do not miss Getcertkey website, you will be very satisfied.

The Certified Kubernetes Security Specialist (CKS) certification exam is a new credential offered by the Linux Foundation. It is designed to test the knowledge and skills of professionals who are responsible for securing Kubernetes-based systems. Certified Kubernetes Security Specialist (CKS) certification is essential for individuals who seek to demonstrate their mastery of best practices in security and compliance within Kubernetes environments.

>> **Latest CKS Exam Registration** <<

## CKS Reliable Test Guide, CKS Latest Braindumps Sheet

There may be a lot of people feel that the preparation process for exams is hard and boring, and hard work does not necessarily mean good results, which is an important reason why many people are afraid of examinations. Today, our CKS study materials will radically change this. High question hit rate makes you no longer aimless when preparing for the exam, so you just should review according to the content of our CKS Study Materials prepared for you. Instant answer feedback allows you to identify your vulnerabilities in a timely manner, so as to make up for your weaknesses.

## Linux Foundation Certified Kubernetes Security Specialist (CKS) Sample Questions (Q177-Q182):

### NEW QUESTION # 177

Use the kubesecc docker images to scan the given YAML manifest, edit and apply the advised changes, and passed with a score of 4 points.

```
kubesecc-test.yaml
```

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
name: kubesecc-demo
```

```
spec:
```

```
containers:
```

```
- name: kubesecc-demo
```

```
image: gcr.io/google-samples/node-hello:1.0
```

```
securityContext:
```

readOnlyRootFilesystem: true

- A. Hint: `docker run -i kubesecc/kubesecc:512c5e0 scan /dev/stdin < kubesecc-test.yaml`

**Answer: A**

### NEW QUESTION # 178

You have a Kubernetes cluster with multiple namespaces, each representing a different department. You need to ensure that resources in one namespace cannot access resources in another namespace, even if they are running as the same user. How would you implement this isolation policy and what are the potential risks if this isolation is not implemented effectively?

**Answer:**

Explanation:

Solution (Step by Step) :

1. Use Network Policies: Define network policies at the namespace level to control communication between pods. Each namespace will have its own set of policies.

- Example Network Policy (Namespace A):

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-internal-traffic
  namespace: namespace-a
spec:
  podSelector: {}
  ingress:
  - from:
    - podSelector: {}
    - namespaceSelector:
        matchLabels:
          name: namespace-a
```

2. Enable Pod Security Policies (PSPs) PSPs allow you to define security constraints for pods running in your cluster. You can restrict the use of specific resources, capabilities, and network access. - Example PSP:

```
apiVersion: policy/v1beta1
kind: PodSecurityPolicy
metadata:
  name: restricted-psp
spec:
  runAsUser:
  - "1000"
  - "1001"
  fsGroup:
  - "1000"
  seLinux:
  - level: S0
  hostNetwork: false
  hostPID: false
  hostIPC: false
  privileged: false
  allowPrivilegeEscalation: false
  readOnlyRootFilesystem: true
  volumes:
  - 'configMap'
  - 'secret'
  - 'emptyDir'
  - 'projected'
  - 'downwardAPI'
  - 'persistentVolumeClaim'
  - 'hostPath'
  - 'projected'
  - 'secret'
  - 'persistentVolumeClaim'
```

3. Isolate Resources: Ensure resources are not shared between namespaces, such as storage (persistent volumes) and configuration (config maps, secrets). - Example: Create separate persistent volumes and claims for each namespace. 4. Monitoring and Auditing: Implement monitoring and auditing tools to detect any unauthorized access attempts or violations of your isolation policy. 5. Potential Risks of Insufficient Isolation: - Data Breaches: Data in one namespace could be compromised by applications in another









```

    cidr: 192.168.0.0/14
  except:
  - to:
    - ipBlock:
      cidr: 192.168.0.0/16 # Exclude a specific range for security reasons (optional)
- ipBlock:
  cidr: 192.168.99.0/24
- ipBlock:
  cidr: 172.17.0.0/24
- ipBlock:
  cidr: 10.0.0.0/24
- ipBlock:
  cidr: 10.128.0.0/14 # Allow communication to logging service
  except:
  - to:
    - ipBlock:
      cidr: 10.128.0.0/16 # Exclude a specific range for security reasons (optional)
- ipBlock:
  cidr: 192.168.0.0/14
  except:
  - to:
    - ipBlock:
      cidr: 192.168.0.0/16 # Exclude a specific range for security reasons (optional)
- ipBlock:
  cidr: 192.168.99.0/24
- ipBlock:
  cidr: 172.17.0.0/24
- ipBlock:
  cidr: 10.0.0.0/24
- ipBlock:
  cidr: 100.64.0.0/10 # Allow communication to monitoring service
  except:
  - to:
    - ipBlock:
      cidr: 100.64.0.0/12 # Exclude a specific range for security reasons (optional)
- ipBlock:
  cidr: 172.16.0.0/12
  except:
  - to:
    - ipBlock:
      cidr: 172.16.0.0/14 # Exclude a specific range for security reasons (optional)
- ipBlock:
  cidr: 192.168.0.0

```

### NEW QUESTION # 180

You're tasked with securing a Kubernetes cluster running on Google Kubernetes Engine (GKE). One of the key security objectives is to ensure that only authorized users can access the cluster's API server and that communication between components within the cluster is encrypted. You need to configure the cluster's network policy and authentication mechanism to enforce these security controls. Explain step-by-step how you would configure GKE's network policies and authentication mechanisms to achieve these objectives.

#### Answer:

Explanation:

Solution (Step by Step) :

1. Configure Network Policies:

- Create Network Policies: use the 'kubectl' command to create network policies that define the communication rules between pods and services. For

example:

```
bash
```

```
kubectl apply -f network-policy.yaml
```

- Define Rules: Specify the rules in the network policy. For example:

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-http
spec:
  podSelector:
    matchLabels:
      app: my-app
  ingress:
  - from:
    - podSelector:
        matchLabels:
          app: my-other-app
  ports:
  - protocol: TCP
    port: 80

```

- Apply Policy: Apply the policy using `kubectl apply -f network-policy.yaml`- 2. Configure Authentication: - Enable Service Account Authentication: Configure service accounts to access the API server. In GKE, you can enable service account authentication by creating a service account key. - Create Service Account Key: Create a service account key with the following command: `bash gcloud iam service-accounts keys create service-account-key.json --iam-account service-account@project_iam_gserviceaccount.com`- Restrict Access: Configure the service account's permissions to minimize the risk of unauthorized access. Use the IAM policy to grant the service account access only to the required resources. 3. Configure Encryption: - Enable HTTPS for the API Server: In GKE, the API server runs over HTTPS by default. Verify that this is enabled in your cluster configuration. - Configure TLS Certificates: Ensure that the API server uses a valid TLS certificate for secure communication. In GKE, this is typically managed automatically. - Use Mutual TLS: For more robust authentication, configure mutual TLS between the API server and other components. You can use a certificate authority (CA) to issue certificates for each component and configure them for mutual authentication.

**NEW QUESTION # 181**  
SIMULATION

Create a new NetworkPolicy named deny-all in the namespace testing which denies all traffic of type ingress and egress traffic

**Answer:**

Explanation:

You can create a "default" isolation policy for a namespace by creating a NetworkPolicy that selects all pods but does not allow any ingress traffic to those pods.

```

---
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: default-deny-ingress
spec:
  podSelector: {}
  policyTypes:
  - Ingress

```

You can create a "default" egress isolation policy for a namespace by creating a NetworkPolicy that selects all pods but does not allow any egress traffic from those pods.

```

---
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-all-egress
spec:
  podSelector: {}
  egress:
  - {}

```



P.S. Free 2026 Linux Foundation CKS dumps are available on Google Drive shared by Getcertkey: <https://drive.google.com/open?id=1c8ZZ1KQskCxdZ4Lt75cjBRLGT9YEFFvH>