



# Linux-Foundation

## Exam Questions KCSA

Kubernetes and Cloud Native Security Associate (KCSA)

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

What information is stored in etcd?

- A. Etcd manages the configuration data, state data, and metadata for Kubernetes.
- B. Application logs and monitoring data for auditing and troubleshooting purposes.
- C. Sensitive user data such as usernames and passwords.
- D. Pod data contained in Persistent Volume Claims (e. hostPath).
- E. hostPath).

**Answer:** A

#### Explanation:

- etcd is Kubernetes key-value store for cluster state.
- Stores: ConfigMaps, Secrets, Pod definitions, Deployments, RBAC policies, and metadata.
- Exact extract (Kubernetes Docs – etcd):
- etcd is a consistent and highly-available key-value store used as Kubernetes backing store for all cluster data.
- Clarifications:
- B: Logs/metrics are handled by logging/monitoring solutions, not etcd.
- C: Secrets may be stored here but encoded in base64, not specifically "usernames/passwords" as primary use.
- D: Persistent Volumes are external storage, not stored in etcd.

References:

Kubernetes Docs — etcd: <https://kubernetes.io/docs/concepts/overview/components/#etcd>

### NEW QUESTION 2

Which of the following represents a baseline security measure for containers?

- A. Implementing access control to restrict container access.
- B. Configuring a static IP for each container.
- C. Configuring persistent storage for containers.
- D. Run containers as the root user.

**Answer:** A

#### Explanation:

- Access control (RBAC, least privilege, user restrictions) is a baseline container security best practice.
- Exact extract (Kubernetes Pod Security Standards – Baseline):
- The baseline profile is designed to prevent known privilege escalations. It prohibits running privileged containers or containers as root.
- Other options clarified:
- B: Static IPs not a security measure.
- C: Persistent storage is functionality, not security.
- D: Running as root is explicitly insecure.

References:

Kubernetes Docs — Pod Security Standards (Baseline): <https://kubernetes.io/docs/concepts/security/pod-security-standards/>

### NEW QUESTION 3

Which other controllers are part of the kube-controller-manager inside the Kubernetes cluster?

- A. Job controller, CronJob controller, and DaemonSet controller
- B. Pod, Service, and Ingress controller
- C. Namespace controller, ConfigMap controller, and Secret controller
- D. Replication controller, Endpoints controller, Namespace controller, and ServiceAccounts controller

**Answer:** D

#### Explanation:

- kube-controller-manager runs a set of controllers that regulate the cluster's state.
- Exact extract (Kubernetes Docs): "The kube-controller-manager runs controllers that are core to Kubernetes. Examples of controllers are: Node controller, Replication controller, Endpoints controller, Namespace controller, and ServiceAccounts controller."
- Why D is correct: All listed are actual controllers within kube-controller-manager.
- Why others are wrong:

- > A: Job and CronJob controllers are managed by kube-controller-manager, but DaemonSet controller is managed by the kube-scheduler/deployment logic.
  - > B: Pod, Service, Ingress controllers are not part of kube-controller-manager.
  - > C: ConfigMap and Secret do not have dedicated controllers.
- [References:, Kubernetes Docs — kube-controller-manager: <https://kubernetes.io/docs/reference/command-line-tools-reference/kube-controller-manager/>, ]

#### NEW QUESTION 4

In a Kubernetes cluster, what are the security risks associated with using ConfigMaps for storing secrets?

- A. Storing secrets in ConfigMaps does not allow for fine-grained access control via RBAC.
- B. Storing secrets in ConfigMaps can expose sensitive information as they are stored in plaintext and can be accessed by unauthorized users.
- C. Using ConfigMaps for storing secrets might make applications incompatible with the Kubernetes cluster.
- D. ConfigMaps store sensitive information in etcd encoded in base64 format automatically, which does not ensure confidentiality of data.

**Answer: B**

#### Explanation:

- > ConfigMaps are explicitly not for confidential data.
  - > Exact extract (ConfigMap concept): "A ConfigMap is an API object used to store non-confidential data in key-value pairs."
  - > Exact extract (ConfigMap concept): "ConfigMaps are not intended to hold confidential data. Use a Secret for confidential data."
  - > Why this is risky: data placed into a ConfigMap is stored as regular (plaintext) string values in the API and etcd (unless you deliberately use binaryData for base64 content you supply). That means if someone has read access to the namespace or to etcd/API Server storage, they can view the values.
  - > Secrets vs ConfigMaps (to clarify distractor D):
  - > Exact extract (Secret concept): "By default, secret data is stored as unencrypted base64-encoded strings. You can enable encryption at rest to protect Secrets stored in etcd."
  - > This base64 behavior applies to Secrets, not to ConfigMap data. Thus option D is incorrect for ConfigMaps.
  - > About RBAC (to clarify distractor A): Kubernetes does support fine-grained RBAC for both ConfigMaps and Secrets; the issue isn't lack of RBAC but that ConfigMaps are not designed for confidential material.
  - > About compatibility (to clarify distractor C): Using ConfigMaps for secrets doesn't make apps "incompatible"; it's simply insecure and against guidance.
- [References:, Kubernetes Docs — ConfigMaps: <https://kubernetes.io/docs/concepts/configuration/configmap/>, Kubernetes Docs — Secrets: <https://kubernetes.io/docs/concepts/configuration/secret/>, Kubernetes Docs — Encrypting Secret Data at Rest: <https://kubernetes.io/docs/tasks/administer-cluster/encrypt-data/>, Note: The citations above are from the official Kubernetes documentation and reflect the stated guidance that ConfigMaps are for non-confidential data, while Secrets (with encryption at rest enabled) are for confidential data, and that the 4C's map to defense in depth., ]

#### NEW QUESTION 5

What was the name of the precursor to Pod Security Standards?

- A. Container Runtime Security
- B. Kubernetes Security Context
- C. Container Security Standards
- D. Pod Security Policy

**Answer: D**

#### Explanation:

Kubernetes originally had a feature called PodSecurityPolicy (PSP), which provided controls to restrict pod behavior.

Official docs:

"PodSecurityPolicy was deprecated in Kubernetes v1.21 and removed in v1.25."

"Pod Security Standards (PSS) replace PodSecurityPolicy (PSP) with a simpler, policy-driven approach."

PSP was often complex and hard to manage, so it was replaced by Pod Security Admission (PSA) which enforces Pod Security Standards.

[References:, Kubernetes Docs — PodSecurityPolicy (deprecated): <https://kubernetes.io/docs/concepts/security/pod-security-policy/>, Kubernetes Blog — PodSecurityPolicy Deprecation: <https://kubernetes.io/blog/2021/04/06/podsecuritypolicy-deprecation-past-present-and-future/>, ]

#### NEW QUESTION 6

When using a cloud provider's managed Kubernetes service, who is responsible for maintaining the etcd cluster?

- A. Kubernetes administrator
- B. Namespace administrator
- C. Cloud provider
- D. Application developer

**Answer: C**

#### Explanation:

In managed Kubernetes services (EKS, GKE, AKS), the control plane is operated by the cloud provider.

This includes etcd, API server, controller manager, scheduler.

Users manage worker nodes (in some models) and workloads, but not the control plane.

Exact extract (GKE Docs):

"The control plane, including the API server and etcd database, is managed and maintained by Google."

Similarly for EKS and AKS, etcd is fully managed by the provider.

[References: GKE Architecture: <https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-architecture>, EKS Architecture: <https://docs.aws.amazon.com/eks/latest/userguide/eks-architecture.html>, AKS Docs: <https://learn.microsoft.com/en-us/azure/aks/concepts-clusters-workloads>, ]

#### NEW QUESTION 7

A Kubernetes cluster tenant can launch privileged Pods in contravention of the restricted Pod Security Standard mandated for cluster tenants and enforced by the built-in PodSecurity admission controller.

The tenant has full CRUD permissions on the namespace object and the namespaced resources. How did the tenant achieve this?

- A. The scope of the tenant role means privilege escalation is impossible.
- B. By tampering with the namespace labels.
- C. By deleting the PodSecurity admission controller deployment running in their namespace.
- D. By using higher-level access credentials obtained reading secrets from another namespace.

**Answer: B**

#### Explanation:

The PodSecurity admission controller enforces Pod Security Standards (Baseline, Restricted, Privileged) based on namespace labels.

If a tenant has full CRUD on the namespace object, they can modify the namespace label to remove or weaken the restriction (e.g., setting `pod-security.kubernetes.io/enforce=privileged`).

This allows privileged Pods to be admitted despite the security policy.



Incorrect options:

(A) is false — namespace-level access allows tampering.

(C) is invalid — PodSecurity admission is not namespace-deployed, it's a cluster-wide admission controller.

(D) is unrelated — Secrets from other namespaces wouldn't directly bypass PodSecurity enforcement.

References:

Kubernetes Documentation – Pod Security Admission

CNCF Security Whitepaper – Admission control and namespace-level policy enforcement weaknesses.

#### NEW QUESTION 8

As a Kubernetes and Cloud Native Security Associate, a user can set up audit logging in a cluster. What is the risk of logging every event at the full RequestResponse level?

- A. No risk, as it provides the most comprehensive audit trail.
- B. Increased storage requirements and potential impact on performance.
- C. Improved security and easier incident investigation.
- D. Reduced storage requirements and faster performance.

**Answer: B**

#### Explanation:

Audit logging records API server requests and responses for security monitoring.

The RequestResponse level logs the full request and response bodies, which can:

Significantly increase storage and performance overhead.

Potentially log sensitive data (including Secrets).

Therefore, while comprehensive, it introduces risks of performance degradation and excessive log volume.

References:

Kubernetes Documentation – Auditing

CNCF Security Whitepaper – Logging and monitoring: trade-offs between verbosity, storage, and security.

#### NEW QUESTION 9

When should soft multitenancy be used over hard multitenancy?

- A. When the priority is enabling resource sharing and efficiency between tenants.
- B. When the priority is enabling complete isolation between tenants.
- C. When the priority is enabling fine-grained control over tenant resources.
- D. When the priority is enabling strict security boundaries between tenants.

**Answer: A**

#### Explanation:

Soft multitenancy (Namespaces, RBAC, Network Policies) # assumes some level of trust between tenants, focuses on resource sharing and efficiency.

Hard multitenancy (separate clusters or strong virtualization) # strict isolation, used when tenants are untrusted.

Exact extract (CNCF TAG Security Multi-Tenancy Whitepaper):

??Soft multi-tenancy refers to multiple workloads running in the same cluster with some trust assumptions. It provides resource sharing and operational efficiency.

Hard multi-tenancy requires stronger isolation guarantees, typically separate clusters.??

References:

CNCF Security TAG — Multi-Tenancy Whitepaper: <https://github.com/cncf/tag-security/tree/main/multi-tenancy>

#### NEW QUESTION 10

In which order are the validating and mutating admission controllers run while the Kubernetes API server processes a request?

- A. The order of execution varies and is determined by the cluster configuration.
- B. Validating admission controllers run before mutating admission controllers.
- C. Validating and mutating admission controllers run simultaneously.
- D. Mutating admission controllers run before validating admission controllers.

**Answer: D**

**Explanation:**

The admission control flow in Kubernetes:

Mutating admission controllers run first and can modify incoming requests.

Validating admission controllers run after mutations to ensure the final object complies with policies.

This ensures policies validate the final, mutated object.

References:

Kubernetes Documentation – Admission Controllers CNCF Security Whitepaper – Admission control workflow.

**NEW QUESTION 10**

Which label should be added to the Namespace to block any privileged Pods from being created in that Namespace?

- A. privileged: false
- B. privileged: true
- C. pod-security.kubernetes.io/enforce: baseline
- D. pod.security.kubernetes.io/privileged: false

**Answer: C**

**Explanation:**

Kubernetes Pod Security Admission (PSA) enforces Pod Security Standards by applying labels on Namespaces.

Exact extract (Kubernetes Docs – Pod Security Admission):

??You can label a namespace with pod-security.kubernetes.io/enforce: baseline to enforce the Baseline policy.??

The baseline profile explicitly disallows privileged pods and other unsafe features.

Why others are wrong:

A & D: These labels do not exist in Kubernetes.

B: Setting privileged: true would allow privileged pods, not block them.

References:

Kubernetes Docs — Pod Security Admission: <https://kubernetes.io/docs/concepts/security/pod-security-admission/>

Kubernetes Docs — Pod Security Standards: <https://kubernetes.io/docs/concepts/security/pod-security-standards/>

**NEW QUESTION 15**

What is the reasoning behind considering the Cloud as the trusted computing base of a Kubernetes cluster?

- A. The Cloud enforces security controls at the Kubernetes cluster level, so application developers can focus on applications only.
- B. A Kubernetes cluster can only be trusted if the underlying Cloud provider is certified against international standards.
- C. A vulnerability in the Cloud layer has a negligible impact on containers due to Linux isolation mechanisms.
- D. A Kubernetes cluster can only be as secure as the security posture of its Cloud hosting.

**Answer: D**

**Explanation:**

The 4C??s of Cloud Native Security (Cloud, Cluster, Container, Code) model starts with Cloud as the base layer.

If the Cloud (infrastructure layer) is compromised, every higher layer (Cluster, Container, Code) inherits that compromise.

Exact extract (Kubernetes Security Overview):

??The 4C??s of Cloud Native security are Cloud, Clusters, Containers, and Code. You can think of the 4C??s as a layered approach. A Kubernetes cluster can only be as secure as the cloud infrastructure it is deployed on.??

This means the cloud is part of the trusted computing base of a Kubernetes cluster.

References:

Kubernetes Docs — Security Overview (4C??s): <https://kubernetes.io/docs/concepts/security/overview/#the-4cs-of-cloud-native-security>

**NEW QUESTION 18**

A container image is trojanized by an attacker by compromising the build server. Based on the STRIDE threat modeling framework, which threat category best defines this threat?

- A. Repudiation
- B. Spoofing
- C. Denial of Service
- D. Tampering

**Answer: D**

**Explanation:**

In STRIDE, Tampering is the threat category for unauthorized modification of data or code/artifacts. A trojanized container image is, by definition, an attacker??s modification of the build output (the image) after compromising the CI/build system—i.e., tampering with the artifact in the software supply chain.

Why not the others?

Spoofing is about identity/authentication (e.g., pretending to be someone/something).

Repudiation is about denying having performed an action without sufficient audit evidence.

Denial of Service targets availability (exhausting resources or making a service unavailable). The scenario explicitly focuses on an altered image resulting from a compromised build server—this squarely maps to Tampering.

Authoritative references (for verification and deeper reading):

Kubernetes (official docs) – Supply Chain Security (discusses risks such as compromised CI/CD pipelines leading to modified/poisoned images and emphasizes verifying image integrity/signatures).

Kubernetes Docs #Security #Supply chain security and Securing a cluster (sections on image provenance, signing, and verifying artifacts).

CNCF TAG Security – Cloud Native Security Whitepaper (v2) – Threat modeling in cloud-native and software supply chain risks; describes attackers modifying build outputs (images/artifacts) via CI

/CD compromise as a form of tampering and prescribes controls (signing, provenance, policy).

CNCF TAG Security – Software Supply Chain Security Best Practices – Explicitly covers CI/CD compromise leading to maliciously modified images and recommends SLSA, provenance attestation, and signature verification (policy enforcement via admission controls).

Microsoft STRIDE (canonical reference) – Defines Tampering as modifying data or code, which directly fits a trojanized image produced by a compromised build system.

### NEW QUESTION 21

How do Kubernetes namespaces impact the application of policies when using Pod Security Admission?

- A. Namespaces are ignored; Pod Security Admission policies apply cluster-wide only.
- B. Different policies can be applied to specific namespaces.
- C. Each namespace can have only one active policy.
- D. The default namespace enforces the strictest security policies by default.

**Answer: B**

#### Explanation:

Pod Security Admission (PSA) enforces policies by applying labels on namespaces, not globally across the cluster.

Exact extract (Kubernetes Docs – Pod Security Admission):

??You can apply Pod Security Standards to namespaces by adding labels such as pod-security.kubernetes.io/enforce. Different namespaces can enforce different policies.??

Clarifications:

A: Incorrect, namespaces are the unit of enforcement.

C: Misleading — a namespace can have multiple enforcement modes (enforce, audit, warn).

D: Default namespace does not enforce strict policies unless labeled.

References:

Kubernetes Docs — Pod Security Admission: <https://kubernetes.io/docs/concepts/security/pod-security-admission/>

### NEW QUESTION 23

What kind of organization would need to be compliant with PCI DSS?

- A. Retail stores that only accept cash payments.
- B. Government agencies that collect personally identifiable information.
- C. Non-profit organizations that handle sensitive customer data.
- D. Merchants that process credit card payments.

**Answer: D**

#### Explanation:

PCI DSS (Payment Card Industry Data Security Standard) applies to any entity that stores, processes, or transmits cardholder data.

Exact extract (PCI DSS official summary):

"PCI DSS applies to all entities that store, process or transmit cardholder data (CHD) and /or sensitive authentication data (SAD)."

Therefore, merchants who process credit card payments must comply.

Why others are wrong:

A: No card payments, so no PCI scope.

B: This falls under FISMA / NIST 800-53, not PCI DSS.

C: Non-profits may handle sensitive data, but PCI only applies if they process credit cards.

References:

PCI Security Standards Council — PCI DSS Summary: [https://www.pcisecuritystandards.org/pci\\_security/](https://www.pcisecuritystandards.org/pci_security/)

### NEW QUESTION 27

Which of the following snippets from a RoleBinding correctly associates user bob with Role pod-reader ?

- A. subjects:- kind: Username: pod-readerapiGroup: rbac.authorization.k8s.io roleRef: kind: Role name: bobapiGroup: rbac.authorization.k8s.io
- B. subjects:- kind: Username: bobapiGroup: rbac.authorization.k8s.io roleRef: kind: Role name: pod-readerapiGroup: rbac.authorization.k8s.io
- C. subjects:- kind: Username: bobapiGroup: rbac.authorization.k8s.io roleRef: kind: ClusterRole name: pod-readerapiGroup: rbac.authorization.k8s.io
- D. subjects:- kind: Group name: bobapiGroup: rbac.authorization.k8s.io roleRef: kind: Role name: pod-readerapiGroup: rbac.authorization.k8s.io

**Answer: B**

#### Explanation:

Kubernetes RBAC uses RoleBinding to grant permissions defined in a Role to a subject (user, group, or service account) within a namespace. The official example shows binding user jane to Role pod-reader:

"A RoleBinding grants the permissions defined in a Role to a user or set of users??."

Example:

subjects:

- kind: User

name: jane

apiGroup: rbac.authorization.k8s.io

roleRef:

kind: Role

name: pod-reader

apiGroup: rbac.authorization.k8s.io

— Kubernetes docs, RBAC: RoleBinding and ClusterRoleBinding

Option B matches this pattern exactly, with name: bob as the User subject and roleRef pointing to the Role named pod-reader.

A swaps the names (subject is pod-reader, role is bob) ?? incorrect.

C references a ClusterRole, not a Role (the question asks for Role).

D uses kind: Group even though we need the User bob.

[References:, Kubernetes Docs — Using RBAC Authorization ?? RoleBinding and ClusterRoleBinding: <https://kubernetes.io/docs/reference/access-authn-authz/rbac/#rolebinding-and-clusterrolebinding>, ]

### NEW QUESTION 30

How can a user enforce the Pod Security Standard without third-party tools?

- A. Through implementing Kyverno or OPA Policies.

- B. Use the PodSecurity admission controller.
- C. It is only possible to enforce the Pod Security Standard with additional tools within the cloud native ecosystem.
- D. No additional measures have to be taken to enforce the Pod Security Standard.

**Answer: B**

**Explanation:**

The PodSecurity admission controller (built-in as of Kubernetes v1.23+) enforces the Pod Security Standards (Privileged, Baseline, Restricted). Enforcement is namespace-scoped and configured through namespace labels.

Incorrect options:

- (A) Kyverno/OPA are external policy tools (useful but not required).
- (C) Not true, PodSecurity admission provides native enforcement.
- (D) Enforcement requires explicit configuration, not automatic.

[References: , Kubernetes Documentation – Pod Security Admission, CNCF Security Whitepaper – Policy enforcement and admission control., ]

**NEW QUESTION 32**

What is Grafana?

- A. A cloud-native distributed tracing system for monitoring microservices architectures.
- B. A container orchestration platform for managing and scaling applications.
- C. A platform for monitoring and visualizing time-series data.
- D. A cloud-native security tool for scanning and detecting vulnerabilities in Kubernetes clusters.

**Answer: C**

**Explanation:**

Grafana: An open-source analytics and visualization platform widely used with Prometheus, Loki, etc.

Exact extract (Grafana Docs): ??Grafana is the open-source analytics and monitoring solution for every database. It allows you to query, visualize, alert on, and understand your metrics no matter where they are stored.??

A is wrong: That describes Jaeger (distributed tracing).

B is wrong: That??s Kubernetes itself.

D is wrong: That??s Trivy/Aqua/Prisma type tools.

References:

Grafana Docs: <https://grafana.com/docs/grafana/latest/>

**NEW QUESTION 37**

What is the purpose of the Supplier Assessments and Reviews control in the NIST 800-53 Rev. 5 set of controls for Supply Chain Risk Management?

- A. To evaluate and monitor existing suppliers for adherence to security requirements.
- B. To conduct regular audits of suppliers' financial performance.
- C. To establish contractual agreements with suppliers.
- D. To identify potential suppliers for the organization.

**Answer: A**

**Explanation:**

In NIST SP 800-53 Rev. 5, SR-6: Supplier Assessments and Reviews requires evaluating and monitoring suppliers' security and risk practices.

Exact extract (NIST SP 800-53 Rev. 5, SR-6):

"The organization assesses and monitors suppliers to ensure they are meeting the security requirements specified in contracts and agreements."

This is about ongoing monitoring of supplier adherence, not financial audits, not contract creation, and not supplier discovery.

References:

NIST SP 800-53 Rev. 5, Control SR-6 (Supplier Assessments and Reviews): <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>

**NEW QUESTION 41**

Which of the following is a control for Supply Chain Risk Management according to NIST 800-53 Rev. 5?

- A. Access Control
- B. System and Communications Protection
- C. Supply Chain Risk Management Plan
- D. Incident Response

**Answer: C**

**NEW QUESTION 46**

.....

## Relate Links

**100% Pass Your KCSA Exam with Exambible Prep Materials**

<https://www.exambible.com/KCSA-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>