

Exam Questions KCSA

Kubernetes and Cloud Native Security Associate (KCSA)

<https://www.2passeasy.com/dumps/KCSA/>



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 information does a user need to verify a signed container image?

- A. The image's SHA-256 hash and the private key of the signing authority.
- B. The image's digital signature and the private key of the signing authority.
- C. The image's SHA-256 hash and the public key of the signing authority.
- D. The image's digital signature and the public key of the signing authority.

Answer: D

Explanation:

- > Container image signing (e.g., with cosign, Notary v2) uses asymmetric cryptography.
- > Verification process:
- > Retrieve the image's digital signature.
- > Validate the signature with the public key of the signer.
- > Exact extract (Sigstore Cosign Docs):
- > Verification of an image requires the signature and the signer's public key. The signature proves authenticity and integrity.
- > Why others are wrong:
- > A & B: The private key is only used by the signer, never shared.
- > C: The hash alone cannot prove authenticity without the digital signature.

References:

Sigstore Cosign Docs: <https://docs.sigstore.dev/cosign/overview>

NEW QUESTION 3

In Kubernetes, what is Public Key Infrastructure (PKI) used for?

- A. To manage certificates and ensure secure communication in a Kubernetes cluster.
- B. To automate the scaling of containers in a Kubernetes cluster.
- C. To manage networking in a Kubernetes cluster.
- D. To monitor and analyze performance metrics of a Kubernetes cluster.

Answer: A

Explanation:

Kubernetes uses PKI certificates extensively to secure communication between control plane components (API server, etcd, kube-scheduler, kube-controller-manager) and with kubelets.

Certificates enable mutual TLS authentication and encryption across components.

PKI does not handle scaling, networking, or monitoring.

References:

Kubernetes Documentation – Certificates

CNCF Security Whitepaper – Cluster communication security and the role of PKI.

NEW QUESTION 4

A cluster administrator wants to enforce the use of a different container runtime depending on the application a workload belongs to.

- A. By manually modifying the container runtime for each workload after it has been created.
- B. By modifying the kube-apiserver configuration file to specify the desired container runtime for each application.
- C. By configuring a validating admission controller webhook that verifies the container runtime based on the application label and rejects requests that do not comply.
- D. By configuring a mutating admission controller webhook that intercepts new workload creation requests and modifies the container runtime based on the application label.

Answer: D

Explanation:

- Kubernetes supports workload-specific runtimes via `RuntimeClass`.
- A mutating admission controller can enforce this automatically by:
 - Intercepting workload creation requests.
 - Modifying the Pod spec to set `runtimeClassName` based on labels or policies.
- Incorrect options:

- (A) Manual modification is not scalable or secure.
- (B) kube-apiserver cannot enforce per-application runtime policies.
- (C) A validating webhook can only reject, not modify, the runtime.

[References: , Kubernetes Documentation – `RuntimeClass`, CNCF Security Whitepaper – Admission controllers for enforcing runtime policies.,]

NEW QUESTION 5

You want to minimize security issues in running Kubernetes Pods. Which of the following actions can help achieve this goal?

- A. Sharing sensitive data among Pods in the same cluster to improve collaboration.
- B. Running Pods with elevated privileges to maximize their capabilities.
- C. Implement Pod Security standards in the Pod's YAML configuration.
- D. Deploying Pods with randomly generated names to obfuscate their identities.

Answer: C

Explanation:

- Pod Security Standards (PSS):
Kubernetes provides Pod Security Admission (PSA) to enforce security controls based on policies.
Official extract: Pod Security Standards define different isolation levels for Pods. The standards focus on restricting what Pods can do and what they can access.
The three standard profiles are:
Privileged: unrestricted (not recommended).
Baseline: minimal restrictions.
Restricted: highly restricted, enforcing least privilege.

- Why option C is correct:
Applying Pod Security Standards in YAML ensures Pods adhere to best practices like:
No root user.
Restricted host access.
No privilege escalation.
Seccomp/AppArmor profiles.
This directly minimizes security risks.

- Why others are wrong:
A: Sharing sensitive data increases risk of exposure.
B: Running with elevated privileges contradicts least privilege principle.
D: Random Pod names do not contribute to security.
[References: , Kubernetes Docs — Pod Security Standards: <https://kubernetes.io/docs/concepts/security/pod-security-standards/>, Kubernetes Docs — Pod Security Admission: <https://kubernetes.io/docs/concepts/security/pod-security-admission/>,]

NEW QUESTION 6

In the event that kube-proxy is in a `CrashLoopBackOff` state, what impact does it have on the Pods running on the same worker node?

- A. The Pods cannot communicate with other Pods in the cluster.
- B. The Pod cannot mount persistent volumes through CSI drivers.
- C. The Pod's security context restrictions cannot be enforced.
- D. The Pod's resource utilization increases significantly.

Answer: A

Explanation:

kube-proxy: manages cluster network routing rules (via iptables or IPVS). It enables Pods to communicate with Services and Pods across nodes. If kube-proxy fails (`CrashLoopBackOff`), service IP routing and cluster-wide pod-to-pod networking breaks. Local Pod-to-Pod communication within the same node

may still work, but cross-node communication fails.

Exact extract (Kubernetes Docs – kube-proxy):

"kube-proxy maintains network rules on nodes. These rules allow network communication to Pods from network sessions inside or outside of the cluster."

[References: Kubernetes Docs — kube-proxy: <https://kubernetes.io/docs/reference/command-line-tools-reference/kube-proxy/>,]

NEW QUESTION 7

Which way of defining security policy brings consistency, minimizes toil, and reduces the probability of misconfiguration?

- A. Using a declarative approach to define security policies as code.
- B. Relying on manual audits and inspections for security policy enforcement.
- C. Manually configuring security controls for each individual resource, regularly.
- D. Implementing security policies through manual scripting on an ad-hoc basis.

Answer: A

Explanation:

Defining policies as code (declarative) is a best practice in Kubernetes and cloud-native security.

This is aligned with GitOps and Policy-as-Code principles (OPA Gatekeeper, Kyverno, etc.).

Exact extract (CNCF Security Whitepaper):

Policy-as-Code enables declarative definition and enforcement of security policies, bringing consistency, automation, and reducing misconfiguration risk.

Manual audits, ad-hoc scripting, or individual configurations are error-prone and inconsistent.

References:

CNCF Security Whitepaper: <https://github.com/cncf/tag-security>

Kubernetes Docs — Policy as Code (OPA, Kyverno): <https://kubernetes.io/docs/concepts/security/>

NEW QUESTION 8

Which of the following statements on static Pods is true?

- A. The kubelet can run static Pods that span multiple nodes, provided that it has the necessary privileges from the API server.
- B. The kubelet can run a maximum of 5 static Pods on each node.
- C. The kubelet schedules static Pods local to its node without going through the kube-scheduler, making tracking and managing them difficult.
- D. The kubelet only deploys static Pods when the kube-scheduler is unresponsive.

Answer: C

Explanation:

Static Pods are managed directly by the kubelet on each node.

They are not scheduled by the kube-scheduler and always remain bound to the node where they are defined.

Exact extract (Kubernetes Docs – Static Pods):

Static Pods are managed directly by the kubelet daemon on a specific node, without the API server. They do not go through the Kubernetes scheduler.

➤ Clarifications:

A: Static Pods do not span multiple nodes.

B: No hard limit of 5 Pods per node.

D: They are not a fallback mechanism; kubelet always manages them regardless of scheduler state.

References:

Kubernetes Docs — Static Pods: <https://kubernetes.io/docs/tasks/configure-pod-container/static-pod/>

NEW QUESTION 9

What mechanism can I use to block unsigned images from running in my cluster?

- A. Enabling Admission Controllers to validate image signatures.
- B. Using PodSecurityPolicy (PSP) to enforce image signing and validation.
- C. Using Pod Security Standards (PSS) to enforce validation of signatures.
- D. Configuring Container Runtime Interface (CRI) to enforce image signing and validation.

Answer: A

Explanation:

Kubernetes Admission Controllers (particularly Validating Admission Webhooks) can be used to enforce policies that validate image signatures.

This is commonly implemented with tools like Sigstore/cosign, Kyverno, or OPA Gatekeeper.

PodSecurityPolicy (PSP): deprecated and never supported image signature validation.

Pod Security Standards (PSS): only apply to pod security fields (privilege, users, host access), not image signatures.

CRI: while runtimes (containerd, CRI-O) may integrate with signature verification tools, enforcement in Kubernetes is generally done via Admission Controllers at the API layer.

Exact extract (Admission Controllers docs):

Admission webhooks can be used to enforce custom policies on the objects being admitted. (e.g., validating signatures).

References:

Kubernetes Docs — Admission Controllers: <https://kubernetes.io/docs/reference/access-authn-authz/admission-controllers/>

Sigstore Project (cosign): <https://sigstore.dev/>

Kyverno ImageVerify Policy: <https://kyverno.io/policies/pod-security/require-image-verification/>

NEW QUESTION 10

Which of the following statements best describe container image signing and verification in the cloud environment?

- A. Container image signatures and their verification ensure their authenticity and integrity against tampering.
- B. Container image signatures are concerned with defining developer ownership of applications within multi-tenant environments.
- C. Container image signatures are mandatory in cloud environments, as cloud providers would deny the execution of unsigned container images.

D. Container image signatures affect the performance of containerized applications, as they increase the size of images with additional metadata.

Answer: A

Explanation:

Image signing (with Notary, cosign, or similar tools) ensures that images are from a trusted source and have not been modified. Exact extract (Sigstore cosign docs):
Cosign allows you to sign and verify container images to ensure authenticity and integrity.

Why others are wrong:

B: Ownership can be inferred but it's about authenticity & integrity not tenancy.

C: Not mandatory; enforcement requires admission controllers.

D: Metadata size is negligible and has no runtime performance impact.

References:

Sigstore Project: <https://docs.sigstore.dev/cosign/overview>

CNCF Security Whitepaper

NEW QUESTION 10

In a cluster that contains Nodes with multiple container runtimes installed, how can a Pod be configured to be created on a specific runtime?

A. By using a command-line flag when creating the Pod.

B. By modifying the Docker daemon configuration.

C. By setting the container runtime as an environment variable in the Pod.

D. By specifying the container runtime in the Pod's YAML file.

Answer: D

Explanation:

Kubernetes supports multiple container runtimes on a node via the RuntimeClass resource.

To select a runtime, you specify the runtimeClassName field in the Pod's YAML manifest. Example:

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
name: example
```

```
spec:
```

```
runtimeClassName: gvisor
```

```
containers:
```

```
- name: app
```

```
image: nginx
```

Incorrect options:

(A) You cannot specify container runtime through a kubectl command-line flag.

(B) Modifying the Docker daemon config does not direct Kubernetes Pods to a runtime.

(C) Environment variables inside a Pod spec do not control container runtimes.

References:

Kubernetes Documentation – RuntimeClass

CNCF Security Whitepaper – Workload isolation via different runtimes (e.g., gVisor, Kata) for enhanced security.

NEW QUESTION 12

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 15

Which security knowledge-base focuses specifically on offensive tools, techniques, and procedures?

A. MITRE ATT&CK

B. OWASP Top 10

C. CIS Controls

D. NIST Cybersecurity Framework

Answer: A

Explanation:

MITRE ATT&CK is a globally recognized knowledge base of adversary tactics, techniques, and procedures (TTPs). It is focused on describing offensive behaviors attackers use.

Incorrect options:

(B) OWASP Top 10 highlights common application vulnerabilities, not attacker techniques.

(C) CIS Controls are defensive best practices, not offensive tools.
(D) NIST Cybersecurity Framework provides a risk-based defensive framework, not adversary TTPs.

References:

MITRE ATT&CK Framework

CNCF Security Whitepaper – Threat intelligence section: references MITRE ATT&CK for describing attacker behavior.

NEW QUESTION 17

Which of the following is a valid security risk caused by having no egress controls in a Kubernetes cluster?

- A. Denial of Service
- B. Data exfiltration
- C. Increased attack surface
- D. Unauthorized access to external resources

Answer: B

Explanation:

Egress NetworkPolicies restrict outbound traffic from Pods.

Without egress restrictions, a compromised Pod could exfiltrate sensitive data (secrets, logs, customer data) to an attacker-controlled server.

Exact extract (Kubernetes Docs – Network Policies):

"Egress rules control outbound connections from Pods. Without such restrictions, compromised workloads can connect freely to external endpoints."

Other options clarified:

A: DoS is more about flooding, not egress absence.

C: ??Increased attack surface?? is vague but not the main risk.

D: True in a sense, but the precise and most common risk is data exfiltration.

[References: , Kubernetes Docs — Network Policies: <https://kubernetes.io/docs/concepts/services-networking/network-policies/>,]

NEW QUESTION 19

What is the purpose of an egress NetworkPolicy?

- A. To control the incoming network traffic to a Kubernetes cluster.
- B. To control the outbound network traffic from a Kubernetes cluster.
- C. To secure the Kubernetes cluster against unauthorized access.
- D. To control the outgoing network traffic from one or more Kubernetes Pods.

Answer: D

Explanation:

NetworkPolicy controls network traffic at the Pod level.

Ingress rules: control incoming connections to Pods.

Egress rules: control outgoing connections from Pods.

Exact extract (Kubernetes Docs – Network Policies):

"An egress rule controls outgoing connections from Pods that match the policy."

Clarifying wrong answers:

A/B: Too broad (cluster-level); policies apply per Pod/namespace.

C: Security against unauthorized access is broader than egress policies.

[References: , Kubernetes Docs — Network Policies: <https://kubernetes.io/docs/concepts/services-networking/network-policies/>,]

NEW QUESTION 23

Which of the following statements is true concerning the use of microVMs over user-space kernel implementations for advanced container sandboxing?

- A. MicroVMs allow for easier container management and orchestration than user-space kernel implementation.
- B. MicroVMs offer higher isolation than user-space kernel implementations at the cost of a higher per-instance memory footprint.
- C. MicroVMs provide reduced application compatibility and higher per-system call overhead than user-space kernel implementations.
- D. MicroVMs offer lower isolation and security compared to user-space kernel implementations.

Answer: B

Explanation:

MicroVM-based runtimes (e.g., Firecracker, Kata Containers) use lightweight VMs to provide strong isolation between workloads.

Compared to user-space kernel implementations (e.g., gVisor), microVMs generally:

Offer higher isolation and security (due to VM-level separation).

Come with a higher memory and resource overhead per instance than user-space approaches.

Incorrect options:

(A) Orchestration is handled by Kubernetes, not inherently easier with microVMs.

(C) Compatibility is typically better with microVMs, not worse.

(D) Isolation is stronger, not weaker.

[References: , CNCF Security Whitepaper – Workload isolation: microVMs vs. user-space kernel sandboxes., Kata Containers Project – isolation trade-offs.,]

NEW QUESTION 25

You are responsible for securing the kubelet component in a Kubernetes cluster.

Which of the following statements about kubelet security is correct?

- A. Kubelet runs as a privileged container by default.
- B. Kubelet does not have any built-in security features.
- C. Kubelet supports TLS authentication and encryption for secure communication with the API server.
- D. Kubelet requires root access to interact with the host system.

Answer: C

Explanation:

The kubelet is the primary agent that runs on each node in a Kubernetes cluster and communicates with the control plane. Kubelet supports TLS (Transport Layer Security) for both authentication and encryption when interacting with the API server. This is a core security feature that ensures secure node-to-control-plane communication.

Incorrect options:

- (A) Kubelet does not run as a privileged container by default; it runs as a system process (typically systemd-managed) on the host.
- (B) Kubelet does include built-in security features such as TLS authentication, authorization modes, and read-only vs secured ports.
- (D) While kubelet interacts with the host system (e.g., cgroups, container runtimes), it does not inherently require root access for communication security; RBAC and TLS handle authentication.

[References: Kubernetes Documentation – Kubelet authentication/authorization, CNCF Security Whitepaper – Cluster Component Security (discusses TLS and mutual authentication between kubelet and API server),]

NEW QUESTION 28

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 33

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual KCSA Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the KCSA Product From:

<https://www.2passeasy.com/dumps/KCSA/>

Money Back Guarantee

KCSA Practice Exam Features:

- * KCSA Questions and Answers Updated Frequently
- * KCSA Practice Questions Verified by Expert Senior Certified Staff
- * KCSA Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * KCSA Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year