Exam Questions KCSA

Kubernetes and Cloud Native Security Associate (KCSA)

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

NEW QUESTION 1
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., withcosign, Notary v2) uses asymmetric cryptography.
- Verification process:
- Retrieve theimage??s digital signature.
- Validate the signature with thepublic keyof 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 2
Is it possible to restrict permissions so that a controller can only change the image of a deployment (without changing anything else about it, e.g., environment variables, commands, replicas, secrets)?

A. Yes, by granting permission to the /image subresource.
B. Not with RBAC, but it is possible with an admission webhook.
C. No, because granting access to the spec.containers.image field always grants access to the rest of the spec object.
D. Yes, with a 'managed fields' annotation.

Answer: B

Explanation:
RBAC in Kubernetesis coarse-grained: it controlsverbs(get, update, patch, delete) onresources(e.g., deployments), butnot individual fieldswithin a resource.
There isno /image subresource for deployments(there is one for pods but only for ephemeral containers).
Therefore,RBAC cannot restrict changes only to the image field.
Admission Webhooks(mutating/validating)canenforce fine-grained policies (e.g., deny updates that change anything other than spec.containers[*].image).
Exact extract (Kubernetes Docs – Admission Webhooks):
"Admission webhooks can be used to enforce custom policies on objects being admitted."
[References:, Kubernetes Docs — RBAC: https://kubernetes.io/docs/reference/access-authn-authz/rbac/, Kubernetes Docs — Admission Webhooks: https://kubernetes.io/docs/reference/access-authn-authz/extensible-admission-controllers/,]

NEW QUESTION 3
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 providesPod 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 tobest practiceslike:
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 donotcontribute 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 4
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, butcross-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 5
Which of the following statements correctly describes a container breakout?

A. A container breakout is the process of escaping the container and gaining access to the Pod's network traffic.
B. A container breakout is the process of escaping a container when it reaches its resource limits.
C. A container breakout is the process of escaping the container and gaining access to the cloud provider's infrastructure.
D. A container breakout is the process of escaping the container and gaining access to the host operating system.

Answer: D

Explanation:
Container breakoutrefers to an attacker escaping container isolation and reaching thehost OS.
Once the host is compromised, the attacker can accessother containers, Kubernetes nodes, or escalate further.
Exact extract (Kubernetes Security Docs):
??If an attacker gains access to a container, they may attempt a container breakout to gain access to the host system.??

Other options clarified:
A: Network access inside a Pod # breakout.
B: Resource exhaustion is aDoS, not a breakout.
C: Cloud infrastructure compromise is possibleafterhost compromise, but not the definition of breakout.
References:
Kubernetes Security Concepts: https://kubernetes.io/docs/concepts/security/
CNCF Security Whitepaper (Threats section):https://github.com/cncf/tag-security

NEW QUESTION 6
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 policiesas code (declarative)is a best practice in Kubernetes and cloud-native security.
This is aligned withGitOpsandPolicy-as-Codeprinciples (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 7
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 Podsare managed directly by thekubeleton each node.
They arenot scheduled by the kube-schedulerand 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 8
In which order are thevalidating and mutating admission controllersrun 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:
Theadmission control flowin Kubernetes:
Mutating admission controllersrun first and can modify incoming requests.
Validating admission controllersrun after mutations to ensure the final object complies with policies.
This ensures policies validate thefinal, mutated object.
References:
Kubernetes Documentation – Admission Controllers CNCF Security Whitepaper – Admission control workflow.

NEW QUESTION 9
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 applyinglabels 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 doesnotenforce strict policies unless labeled.
References:
Kubernetes Docs — Pod Security Admission: https://kubernetes.io/docs/concepts/security/pod-security- admission/

NEW QUESTION 10
An attacker compromises a Pod and attempts to use its service account token to escalate privileges within the cluster. Which Kubernetes security feature is designed tolimit what this service account can do?

A. PodSecurity admission
B. NetworkPolicy
C. Role-Based Access Control (RBAC)
D. RuntimeClass

Answer: C

Explanation:
When a Pod is created, Kubernetes automatically mounts aservice account tokenthat can authenticate to the API server.
TheRole-Based Access Control (RBAC)system defines what actions a service account can perform.
By carefully restricting Roles and RoleBindings, administrators limit the blast radius of a compromised
Pod.
Incorrect options:
(A)PodSecurity admissionenforces workload-level security settings but does not control API access.
(B)NetworkPolicycontrols network communication, not API privileges.
(D)RuntimeClassselects container runtimes, unrelated to privilege escalation through API tokens.
References:
Kubernetes Documentation – Using RBAC Authorization
CNCF Security Whitepaper – Identity & Access Management: limiting lateral movement by constraining service account permissions.

NEW QUESTION 10
Which security knowledge-base focuses specifically onoffensive tools, techniques, and procedures?

A. MITRE ATT&CK
B. OWASP Top 10
C. CIS Controls
D. NIST Cybersecurity Framework

Answer: A

Explanation:
MITRE ATT&CKis a globally recognizedknowledge base of adversary tactics, techniques, and procedures (TTPs). It is focused on describingoffensive behaviorsattackers use.
Incorrect options:
(B)OWASP Top 10highlights common application vulnerabilities, not attacker techniques.
(C)CIS Controlsare defensive best practices, not offensive tools.
(D)NIST Cybersecurity Frameworkprovides 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 11
Which of the following statements is true concerning the use ofmicroVMsover 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 touser-space kernel implementations(e.g., gVisor), microVMs generally:
Offerhigher isolation and security(due to VM-level separation).
Come with ahigher memory and resource overhead per instancethan 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 15
What is the main reason an organization would use a Cloud Workload Protection Platform (CWPP) solution?

A. To protect containerized workloads from known vulnerabilities and malware threats.
B. To automate the deployment and management of containerized workloads.
C. To manage networking between containerized workloads in the Kubernetes cluster.
D. To optimize resource utilization and scalability of containerized workloads.

Answer: A

Explanation:
CWPP (Cloud Workload Protection Platform):As defined by Gartner and adopted across cloud
security practices, CWPPs are designed tosecure workloads(VMs, containers, serverless functions) in
hybrid and cloud environments.
They providevulnerability scanning, runtime protection, compliance checks, and malware detection.
Exact extract (Gartner CWPP definition):??Cloud workload protection platforms protect workloads regardless of location, including physical machines, VMs, containers, and serverless workloads. They provide vulnerability management, system integrity protection, intrusion detection and prevention, and malware protection.??
References:
Gartner: Cloud Workload Protection Platforms Market Guide (summary): https://www.gartner.com/reviews
/market/cloud-workload-protection-platforms
CNCF Security Whitepaper:https://github.com/cncf/tag-security

NEW QUESTION 16
.....

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