Exam Questions KCSA

Kubernetes and Cloud Native Security Associate (KCSA)

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

NEW QUESTION 1
On a client machine, what directory (by default) contains sensitive credential information?

A. /etc/kubernetes/
B. $HOME/.kube
C. /opt/kubernetes/secrets/
D. $HOME/.config/kubernetes/

Answer: B

Explanation:

- Thekubectlclient uses configuration from$HOME/.kube/configby default.
- This file contains: cluster API server endpoint, user certificates, tokens, or kubeconfigs #sensitive credentials.
- Exact extract (Kubernetes Docs – Configure Access to Clusters):
- ??By default, kubectl looks for a file named config in the $HOME/.kube directory. This file contains configuration information including user credentials.??
- Other options clarified:
- A: /etc/kubernetes/ exists on nodes (control plane) not client machines.
- C: /opt/kubernetes/secrets/ is not a standard path.
- D: $HOME/.config/kubernetes/ is not where kubeconfig is stored by default.

References:
Kubernetes Docs — Configure Access to Clusters: https://kubernetes.io/docs/concepts/configuration/organize- cluster-access-kubeconfig/

NEW QUESTION 2
Which of the following statements best describes the role of the Scheduler in Kubernetes?

A. The Scheduler is responsible for monitoring and managing the health of the Kubernetes cluster.
B. The Scheduler is responsible for ensuring the security of the Kubernetes cluster and its components.
C. The Scheduler is responsible for managing the deployment and scaling of applications in the Kubernetes cluster.
D. The Scheduler is responsible for assigning Pods to nodes based on resource availability and other constraints.

Answer: D

Explanation:

- TheKubernetes Schedulerassigns Pods to nodes based on:
- Resource requests & availability (CPU, memory, GPU, etc.)
- Constraints (affinity, taints, tolerations, topology, policies)
- Exact extract (Kubernetes Docs – Scheduler):
- ??The scheduler is a control plane process that assigns Pods to Nodes. Scheduling decisions take into account resource requirements, affinity/anti-affinity, constraints, and policies.??
- Other options clarified:
- A: Monitoring cluster health is theController Manager??s/kubelet??s job.
- B: Security is enforced throughRBAC, admission controllers, PSP/PSA, not the scheduler.
- C: Deployment scaling is handled by theController Manager(Deployment/ReplicaSet controller).

References:
Kubernetes Docs — Scheduler: https://kubernetes.io/docs/concepts/scheduling-eviction/kube-scheduler/

NEW QUESTION 3
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 4
A cluster is failing to pull more recent versions of images from k8s.gcr.io. Why may this be?

A. There is a network connectivity issue between the cluster and k8s.gcr.io.
B. There is a bug in the container runtime or the image pull process.
C. The authentication credentials for accessing k8s.gcr.io are incorrectly scoped.
D. The container image registry k8s.gcr.io has been deprecated.

Answer: D

Explanation:

- k8s.gcr.iowas the historic Kubernetes image registry.
- It has beendeprecatedand replaced withregistry.k8s.io.
- Exact extract (Kubernetes Blog):
- ??The k8s.gcr.io image registry will be frozen from April 3, 2023 and fully deprecated. All Kubernetes project images are now served from registry.k8s.io.??
- Pulling newer versions from k8s.gcr.io fails because the registry no longer receives updates.

References:
Kubernetes Blog — Image Registry Update: https://kubernetes.io/blog/2023/02/06/k8s-gcr-io-freeze- announcement/

NEW QUESTION 5
In a Kubernetes environment, what kind of Admission Controller can modify resource manifests when applied to the Kubernetes API to fix misconfigurations automatically?

A. ValidatingAdmissionController
B. PodSecurityPolicy
C. MutatingAdmissionController
D. ResourceQuota

Answer: C

Explanation:

- Kubernetes Admission Controllers can eithervalidateormutateincoming requests.
- MutatingAdmissionWebhook (Mutating Admission Controller):
- Canmodify or mutate resource manifestsbefore they are persisted in etcd.
- Used for automatic injection of sidecars (e.g., Istio Envoy proxy), setting default values, or fixing misconfigurations.
- ValidatingAdmissionWebhook (Validating Admission Controller):only allows/denies but doesnot change requests.
- PodSecurityPolicy:deprecated; cannot mutate requests.
- ResourceQuota:enforces resource usage, but does not mutate manifests.

Exact Extract:

- ??Mutating admission webhooks are invoked first, and can modify objects to enforce defaults.

Validating admission webhooks are invoked second, and can reject requests to enforce invariants.??
References:
Kubernetes Docs — Admission Controllers: https://kubernetes.io/docs/reference/access-authn-authz
/admission-controllers/
Kubernetes Docs — Admission Webhooks: https://kubernetes.io/docs/reference/access-authn-authz
/extensible-admission-controllers/

NEW QUESTION 6
Which technology can be used to apply security policy for internal cluster traffic at the application layer of the network?

A. Network Policy
B. Ingress Controller
C. Container Runtime
D. Service Mesh

Answer: D

Explanation:

- Service Mesh (e.g., Istio, Linkerd, Consul):operates atLayer 7 (application layer), enforcing policies like mTLS, authorization, and routing between services.
- NetworkPolicy:works atLayer 3/4 (IP/port), not Layer 7.
- Ingress Controller:handles external traffic ingress, not internal service-to-service traffic.
- Container Runtime:responsible for running containers, not enforcing application-layer security.
- Exact extract (Istio docs):

"Istio provides security by enforcing authentication, authorization, and encryption of service-to-service communication."
[References:, Kubernetes Docs — Network Policies: https://kubernetes.io/docs/concepts/services-networking/network-policies/, Istio Security Docs: https://istio.io/latest/docs/concepts/security/,]

NEW QUESTION 7
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 8
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 calledPodSecurityPolicy (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 enforcesPod 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 9
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 10
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 onresource 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
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 13
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 thatstores, 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 paymentsmust comply.
Why others are wrong:
A: No card payments, so no PCI scope.
B: This falls underFISMA / NIST 800-53, not PCI DSS.
C: Non-profits may handle sensitive data, but PCI only applies if they processcredit cards.
References:
PCI Security Standards Council — PCI DSS Summary: https://www.pcisecuritystandards.org/pci_security/

NEW QUESTION 17
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 21
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 NetworkPoliciesrestrict 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 isdata exfiltration.
[References:, Kubernetes Docs — Network Policies: https://kubernetes.io/docs/concepts/services-networking/network-policies/,]

NEW QUESTION 23
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:
NetworkPolicycontrols network trafficat the Pod level.
Ingress rules:controlincomingconnections to Pods.
Egress rules:controloutgoingconnectionsfrom 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 28
In order to reduce the attack surface of the Scheduler, which default parameter should be set to false?

A. --scheduler-name
B. --profiling
C. --secure-kubeconfig
D. --bind-address

Answer: B

Explanation:
Thekube-schedulerexposes aprofiling/debugging endpointwhen --profiling=true (default).
This can unnecessarily increase the attack surface.
Best practice: set --profiling=false in production.
Exact extract (Kubernetes Docs – kube-scheduler flags):
"--profiling (default true): Enable profiling via web interface host:port/debug/pprof/."
Why others are wrong:
--scheduler-name: just identifies the scheduler, not a security risk.
--secure-kubeconfig: not a valid flag.
--bind-address: changing it limits exposure but is not the default risk parameter for profiling.
[References:, Kubernetes Docs — kube-scheduler options: https://kubernetes.io/docs/reference/command-line-tools-reference/kube-scheduler/, ,]

NEW QUESTION 31
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 Reviewsrequires 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 aboutongoing monitoringof 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 34
.....

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