Boost up Your Certification Score

Huawei

H13-531_V3.0

HCIE-Cloud Computing (Written) V3.0



For More Information – Visit link below:

https://www.examsboost.com/

Product Version

✓ Up to Date products, reliable and verified.
✓ Questions and Answers in PDF Format.

Latest Version: 6.0

Question: 1

Kata Containers is an open-source container runtime that provides enhanced container security. Which of the following describes how Kata Containers runs in the cloud-native ecosystem?

A. Kata Containers is essentially a VM. Since VMs are not an element of cloud native, Kata Containers cannot be cloud native.
B. When using Kata Containers, you run VMs in containers. Cloud native container orchestration and management tools can be called through CRI.
C. When using Kata Containers, you run containers on VMs. Cloud native tools interact with the hypervisor where VMs are created to orchestrate and manage containers.
D. Kata Containers is a runtime that complies with Open Container Initiative (OCI). Cloud native container orchestration and management tools can be called through Container Runtime Interface (CRI).

Answer: D

Explanation:
Kata Containers combines the container-management model with isolation using independent lightweight virtual machines. Each Pod can execute within a micro-VM with its own guest kernel, improving separation from other tenants. OCI runtime compatibility and CRI integration allow Kubernetes and related tools to manage those workloads through established container interfaces. It is neither a VM running inside a regular container nor a requirement that orchestrators directly manipulate the hypervisor for each operation. In the V3.0 container curriculum, the precise command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. For examination accuracy, absolute phrases such as ' always ' , ' must ' , ' only ' , and ' cannot ' deserve special attention. The source sometimes demonstrates a supported example without establishing that every other deployment is prohibited. The distinction also matters operationally: selecting a setting outside the supported architecture can create a failure that resembles a network, storage, or application fault, even though the root cause is configuration incompatibility.
Exact source extract: “container (pod) runs on an independent micro-VM, has an independent OS”
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 858.

Question: 2

HUAWEI CLOUD Stack offers a range of security measures to help users safeguard privacy. However, technology cannot completely prevent social engineering attacks. Security awareness among all users must be improved to reduce the risk of information leakage.

A. TRUE
B. FALSE

Answer: A

Explanation:
Cloud technical controls such as authentication, encryption, network security, audit, and access policies reduce data disclosure risks but cannot guarantee protection against human deception. Social engineering can exploit credential handling, approvals, or trust in impersonated staff. Security awareness and organizational procedures must complement technical controls. Huawei ' s shared responsibility model and data lifecycle protection framework support this layered approach, although the inspected source does not specifically use the term social engineering. For security questions, distinguish the scope of the enforcement policy from the point where the service is managed. Tenant, subnet, interface, VPC, regional, and perimeter protections have different responsibilities. The distinction also matters operationally: selecting a setting outside the supported architecture can create a failure that resembles a network, storage, or application fault, even though the root cause is configuration incompatibility. The correct interpretation should be limited to the product version and topology in the cited training material. Later product releases or alternative deployment modes should not silently be substituted for the V3.0 curriculum.
Exact source extract: “compliance certification, data protection, and privacy security”
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 1115, 1108,1115.
; social engineering not explicit

Question: 3

Which of these statements about Docker are accurate?

A. Files cannot be shared between containers.
B. After a container is migrated, its IP address and port change.
C. Containers can be dynamically created and deleted.
D. Containers can be migrated between host machines.

Answer: C

Explanation:
Docker supports dynamic creation and deletion of containers, including run, create, stop, and rm. Containers can share files through volumes and bind mounts, so file-sharing prohibition is false. IPs and host ports do not universally change on workload relocation. Workloads can be recreated on a different host, but ordinary Docker commands do not guarantee live process-state migration. The question should specify whether migration means image-based redeployment or live migration. In the V3.0 container curriculum, the precise command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. A practical verification step is to identify which component owns the operation, which network or storage plane carries it, and which resource object records the state. This prevents confusing an architectural dependency with an optional capability. For examination accuracy, absolute phrases such as ' always ' , ' must ' , ' only ' , and ' cannot ' deserve special attention. The source sometimes demonstrates a supported example without establishing that every other deployment is prohibited.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.

Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 12,15,1047.
; D if workload recreation is called migration

Question: 4

In Kubernetes, you can use the kubectl exec command to work with a container. When executing this command against a multi-container Pod, you are advised to include the -c flag.
Which of the following explains why this flag is added?

A. It transfers the commands to the container.
B. It gives containers the permissions to execute the commands.
C. It specifies the one to operate.
D. It allows specific commands to be executed.

Answer: C

Explanation:
kubectl exec executes a command in an existing Pod container. When a Pod has multiple containers, -c identifies the intended container by name. It does not grant RBAC permissions, change executable privileges, or choose the command syntax. A complete invocation includes the Pod name, selected container, and the command following --. The lab demonstrates exec against a running container; the selection flag becomes important when several containers share a Pod. In the V3.0 container curriculum, the precise command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. The correct interpretation should be limited to the product version and topology in the cited training material. Later product releases or alternative deployment modes should not silently be substituted for the V3.0 curriculum. A practical verification step is to identify which component owns the operation, which network or storage plane carries it, and which resource object records the state. This prevents confusing an architectural dependency with an optional capability.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 35,971.

Question: 5

In HUAWEI CLOUD Stack Layer 3 networking, can all leaf switches be deployed only in M-LAG mode?

A. TRUE
B. FALSE

Answer: B

Explanation:
The Layer 3 networking design uses spine and leaf switches, but not every leaf type has identical connection and redundancy requirements. Border and service leaf roles can vary, and supported deployment modes depend on the topology. M-LAG provides link and device redundancy where

appropriate; it is not a blanket requirement for every leaf device. The statement ' s word only makes it too restrictive. The correct architecture should be selected using the specific Huawei Layer 3 topology and the interfaces attached to each leaf. In a Stack delivery design, supported combinations depend on the selected region architecture, physical network topology, resource pool, and product release. A broadly possible technology choice is not automatically a valid deployment configuration. A practical verification step is to identify which component owns the operation, which network or storage plane carries it, and which resource object records the state. This prevents confusing an architectural dependency with an optional capability. For examination accuracy, absolute phrases such as ' always ' , ' must ' , ' only ' , and ' cannot ' deserve special attention. The source sometimes demonstrates a supported example without establishing that every other deployment is prohibited.
Exact source extract: “1. In Layer 3 networking, 4,000 PMs can be deployed, including 2,000 PMs running”
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 132, 128,132,135.

Question: 6

Which statement below is false about a source Windows server migrated to HUAWEI CLOUD Stack by using Rainbow?

A. Only NTFS file systems are supported.
B. Externally mounted shared partitions (such as CIFS and NFS) cannot be migrated.
C. Migration service ports 445 and 8899 need to be enabled.
D. Shared disks can be migrated.

Answer: D

Explanation:
Rainbow Windows source migration supports NTFS local volumes and does not migrate shared disks. Huawei lists port 445 for SMB and port 8899 for agent management, with another data transmission port as needed. External CIFS/NFS shared mounts are not automatically copied as local NTFS disks. The assertion that shared disks can be migrated directly is explicitly contradicted by the Rainbow restriction slide. Rainbow migration decisions depend on the source operating system, disk and filesystem characteristics, network reachability, and the stage of synchronization. An operation valid for Linux migration is not necessarily valid for Windows block-level migration. For examination accuracy, absolute phrases such as ' always ' , ' must ' , ' only ' , and ' cannot ' deserve special attention. The source sometimes demonstrates a supported example without establishing that every other deployment is prohibited. The distinction also matters operationally: selecting a setting outside the supported architecture can create a failure that resembles a network, storage, or application fault, even though the root cause is configuration incompatibility.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 668,671.

Question: 7

Which of the following cloud-native technologies is a recommended vehicle for carrying microservices?

A. Declarative API
B. Container
C. Service mesh
D. Immutable infrastructure

Answer: B

Explanation:
Huawei explicitly recommends containers as an appropriate carrier for microservices and continuous integration. Containers package a service ' s code and dependencies for deployment while sharing a host kernel. Declarative APIs specify desired state; a service mesh manages communication, security, and observability; immutable infrastructure prescribes replacement-based deployment. These complementary techniques are not the primary runtime vehicle on which a microservice executes. In the V3.0 container curriculum, the precise command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. The distinction also matters operationally: selecting a setting outside the supported architecture can create a failure that resembles a network, storage, or application fault, even though the root cause is configuration incompatibility. The correct interpretation should be limited to the product version and topology in the cited training material. Later product releases or alternative deployment modes should not silently be substituted for the V3.0 curriculum.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 829.

Question: 8

When you run the kubectl get command to query an object and no information is returned, what could be the causes?
Select All that Apply

A. The object type specified in the query command is inconsistent with the actual type.
B. The command must be executed on the master node instead of the worker node.
C. The queried object is not in the Active state, and related flags are not added in the query.
D. You are not querying a global resource object and do not specify its namespace.

Answer: A, D

Explanation:
kubectl get must target the correct resource kind and appropriate namespace; an incorrect kind or namespace can produce no matching objects. The command need not run only on a master node when the client and credentials are configured correctly. Objects in Failed, Pending, or other states can usually still be listed, so a universal requirement for an Active state is incorrect. The source thus supports A/D, not an unqualified three-choice key; C requires a special query or filter condition absent from the question. That issue is material to the answer key and must be retained rather than silently treating an intended selection as an unconditional technical rule. In the V3.0 container curriculum, the precise

command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. The correct interpretation should be limited to the product version and topology in the cited training material. Later product releases or alternative deployment modes should not silently be substituted for the V3.0 curriculum.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 33,35,970.

Question: 9

When building a container image, are you advised to use a base image that includes multiple system tools or system libraries to ensure stable application operation and easy O & M?

A. TRUE
B. FALSE

Answer: B

Explanation:
Huawei explicitly recommends lightweight container base images. Bundling many unnecessary system utilities increases image size, download time, maintenance workload, and potential security exposure. Choose only the runtime dependencies needed by the application, and use multistage builds or separate diagnostic containers when appropriate. A large image does not automatically become more stable or easier to operate merely because it includes more system tools. In the V3.0 container curriculum, the precise command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. A practical verification step is to identify which component owns the operation, which network or storage plane carries it, and which resource object records the state. This prevents confusing an architectural dependency with an optional capability. For examination accuracy, absolute phrases such as ' always ' , ' must ' , ' only ' , and ' cannot ' deserve special attention. The source sometimes demonstrates a supported example without establishing that every other deployment is prohibited.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 931,932,933.
Evidence status: source-inferred

Question: 10

The docker network create command creates a network and defines network attributes. The docker network delete command removes a network.

A. TRUE
B. FALSE

Answer: B

Explanation:
docker network create can create a named network using a selected driver, as shown in Huawei ' s bridge-network lab. docker network delete is not the standard removal subcommand; docker network rm removes an eligible network. Active container attachments can prevent removal. The test is intentionally about exact CLI syntax rather than whether Docker supports creating and deleting networks conceptually. Because one command is invalid, the combined statement is false. In the V3.0 container curriculum, the precise command, Kubernetes API field, runtime behavior, and selected node or namespace matter. Similar-looking terms often refer to separate lifecycle stages and are not interchangeable. For examination accuracy, absolute phrases such as ' always ' , ' must ' , ' only ' , and ' cannot ' deserve special attention. The source sometimes demonstrates a supported example without establishing that every other deployment is prohibited. The distinction also matters operationally: selecting a setting outside the supported architecture can create a failure that resembles a network, storage, or application fault, even though the root cause is configuration incompatibility.
Source-extract limitation: No directly matching text could be verified in the selected PDF text layer; consult the cited diagram/slide.
Official reference: HCIE-Cloud Computing V3.0 Training Material, PDF pages 15,17,890.

Thank You for Trying Our Product

For More Information – Visit link below:

https://www.examsboost.com/

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ 90 Days Free Updates
- ✓ Money Back Pass Guarantee
- ✓ Instant Download or Email Attachment
- ✓ 24/7 Live Chat Support
- ✓ PDF file could be used at any Platform
- ✓ 50,000 Happy Customer

