

ACA Cloud Native Certification Exam

Alibaba Cloud ACA-CloudNative

Version Demo

Total Demo Questions: 16

Total Premium Questions: 169

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Native Fundamentals	3
Topic 2, Container Service for Kubernetes	96
Topic 3, Serverless Kubernetes	12
Topic 4, Microservices	20
Topic 5, DevOps	5
Topic 6, Observability	31
Topic 7, Mix Questions	2
Total	169

QUESTION NO: 1

When we use beats to collect Alibaba Cloud Kubernetes data, to what level should we deploy it?

- A. Deployment
- B. DaemonSet
- C. Sidecar
- D. StatefulSet

ANSWER: B

Explanation:

DaemonSet is the appropriate deployment level for Beats when collecting node-level Kubernetes data. A DaemonSet ensures that Kubernetes schedules one collector Pod on each eligible node in the cluster. This placement lets the Beat access node-local sources such as container logs, kubelet endpoints, cgroup statistics, filesystem information, and runtime metadata. As cluster nodes are added, Kubernetes automatically creates a corresponding collector Pod; when nodes are removed, the associated Pod is removed as well. This provides consistent collection coverage across an ACK cluster without manually scaling the collector workload.

This pattern is particularly important for Beats used for infrastructure observability, such as Metricbeat and Filebeat, because the information to be gathered is distributed across worker nodes rather than being limited to one application workload. Kubernetes defines DaemonSets specifically for running a copy of a Pod on all, or a selected subset of, nodes: [Kubernetes DaemonSet documentation](#). Elastic's Kubernetes guidance likewise describes deploying Metricbeat as a DaemonSet to collect host and container metrics from every Kubernetes node: [Run Metricbeat on Kubernetes](#).

QUESTION NO: 2

Which of the following are true, regarding Helm Charts? (Number of correct answers: 3)

Score 2

- A. Helm Charts are declarative but not procedural
- B. Helm Charts can achieve "immutable infrastructure"
- C. Helm Charts make it easy to install Kubernetes objects, such as ConfigMap, Services, Pods, etc.
- D. Helm Charts make it easy to achieve CD (continuous deployment) for applications

ANSWER: A C D

Explanation:

Helm Charts are declarative because a chart packages Kubernetes resource templates together with values that describe the intended application configuration. Rendering a chart produces Kubernetes manifests representing the desired state, rather than a step-by-step procedure for manually creating every resource. Helm tracks releases and applies chart revisions, allowing teams to manage an application's Kubernetes resources consistently through versioned chart packages. Helm's chart model is documented in the [Helm Charts guide](#).

Helm Charts also simplify installation of Kubernetes objects because a single chart can include templated manifests for resources such as ConfigMaps, Services, Deployments, Pods, and other API objects. Values files make it practical to reuse the same chart across environments while varying settings such as image tags, replica counts, endpoints, and resource requests. Helm's [Using Helm documentation](#) describes installing, upgrading, and rolling back releases. These repeatable, version-controlled installation and upgrade workflows integrate naturally with continuous-deployment pipelines: automation can validate a chart, package it, publish it to a registry, and deploy or upgrade a specified chart version into a target Kubernetes cluster.

QUESTION NO: 3

Which types of networks can ACK provide^ (Number of correct answers: 2)

- A. Flannel
- B. Terway
- C. VPN
- D. GA

ANSWER: A B

Explanation:

ACK supports Flannel and Terway (shown in the question as “Tcrway”) as Kubernetes container-network plug-ins. These plug-ins implement the cluster network plane: they assign Pod IP addresses, establish Pod-to-Pod connectivity, and enable workloads to communicate with Kubernetes Services. Flannel is an overlay-network solution. It is commonly used where a straightforward cluster network configuration is preferred, and it encapsulates Pod traffic to provide connectivity across nodes. Terway is Alibaba Cloud’s VPC-native networking plug-in. It can assign elastic network interfaces or secondary IP addresses from the VPC to Pods, allowing Pod addresses to be routable within the VPC and enabling tighter integration with Alibaba Cloud VPC network capabilities. ACK lets users select the appropriate network plug-in when creating a cluster, based on factors such as required network performance, IP-address planning, and VPC integration requirements. Therefore, the two listed ACK-provided network types are Flannel and Terway. Alibaba Cloud documents the available ACK network plug-ins and their deployment models in its [ACK network plug-ins documentation](#) and provides configuration guidance for [Terway](#).

QUESTION NO: 4

The success of deploying and scaling microservices is heavily dependent on the adoption of_____.

- A. Virtual machine
- B. Big data analytics
- C. Containers
- D. Cloud provider
- E. RDBMS

ANSWER: C

Explanation:

Containers are fundamental to successfully deploying and scaling microservices because they package each service together with its required runtime, libraries, and configuration in a portable, isolated unit. This consistency lets the same service artifact run reliably across development, testing, and production environments, reducing deployment differences and operational friction. Containers are lightweight compared with full virtual machines, so multiple independently deployable microservices can use infrastructure resources efficiently.

Container orchestration platforms, such as Kubernetes and Alibaba Cloud Container Service for Kubernetes, build on this model to automate scheduling, service discovery, health checks, rolling updates, replica management, and horizontal scaling. These capabilities align directly with microservice architectures, where individual services must be released, replaced, and scaled independently according to demand. Alibaba Cloud describes ACK as a managed Kubernetes service for deploying and managing containerized applications, while Kubernetes documents that Deployments manage replicated application workloads and support controlled updates. See [Alibaba Cloud Container Service for Kubernetes](#) and [Kubernetes Deployments documentation](#).

QUESTION NO: 5

Which of the following statements are correct regarding image distribution in ACR EE

(Container Registry Enterprise Edition)? (Number of correct answers: 3)

Score 2

- A. It enable large-scale distribution and policy-based global distribution
- B. P2P based distribution supports docker image snapshots without pulling
- C. P2P based method supports image distribution over more than 10,000 nodes
- D. The snapshot method speeds up image pulling, with pull times as low as 1 second

ANSWER: A C D

Explanation:

ACR Enterprise Edition image distribution is designed for organizations that must deliver container images efficiently across large Kubernetes and compute environments. "It enable large-scale distribution and policy-based global distribution" is correct because ACR EE provides enterprise image-distribution capabilities, including distribution policies that control how images are replicated and delivered across regions or instances. This helps organizations place images closer to workloads and apply distribution rules consistently at scale.

"P2P based method supports image distribution over more than 10,000 nodes" is also correct. The peer-to-peer distribution capability reduces the pressure on a central registry or network egress point by allowing nodes to obtain image data through a distributed delivery mechanism. This architecture is intended for very large cluster rollouts, including scenarios involving more than 10,000 nodes.

"The snapshot method speeds up image pulling, with pull times as low as 1 second" is correct because image snapshots support rapid image startup and delivery by preparing image data for efficient use on compute nodes. This is especially valuable for elastic scaling, batch workloads, and large-scale deployments where reducing container initialization time is important. See the [Alibaba Cloud Container Registry overview](#) and the [ACR image distribution documentation](#).

QUESTION NO: 6

Which of the following statements are correct about Simple Log Service alerting? (Number of correct answers: 3)

- A. Alert rules can be created from query and analysis results
- B. Alert notifications can be sent to notification channels
- C. Alert rules can use a monitoring interval and trigger conditions
- D. Alert rules automatically modify application source code

ANSWER: A B C

Explanation:

Alert rules can be created from query and analysis results is correct. Log Service can periodically execute configured query and analysis statements against logs in a Logstore and evaluate the results against alert conditions.

Alert notifications can be sent to notification channels is also correct. Alerting can notify relevant personnel or systems after an alert is triggered, allowing operational teams to respond to application and infrastructure events. **Alert rules can use a monitoring interval and trigger conditions** is correct because administrators define the evaluation frequency, query time range, and conditions that determine when an alert is raised.

Log Service alerting is intended for log-based monitoring and analysis. It does not automatically modify application source code when an alert occurs. For more information, see the [Log Service alerting documentation](#).

QUESTION NO: 7

Which of the following are correct statements about Alibaba Cloud Container Service for Kubernetes (ACK)'s Sandboxed Pods? (Number of correct answers: 3)

Score 2

- A. Sandboxed pods are a Docker runtime container type, and can run any docker image
- B. Sandboxed Pods provide an independent kernel to enable enhanced security isolation
- C. Sandboxed Pods can be run on both ECS and ECS Bare Metal (EBM) instances
- D. Sandboxed Pods can also support multiple running containers per Pod

ANSWER: A B D

Explanation:

Sandboxed Pods are designed to provide stronger workload isolation than standard containers while retaining Kubernetes Pod behavior. They use a sandboxed-container runtime and remain compatible with Docker/OCI container images, allowing applications packaged as Docker images to be deployed without requiring a different application packaging model. This makes the option stating that Sandboxed Pods are a Docker runtime container type and can run Docker images correct in the context of ACK's supported sandboxed runtime.

The enhanced isolation is achieved by assigning an independent kernel to the sandbox environment. Separating the kernel boundary substantially reduces the risk that a workload can affect the host or neighboring workloads, which is valuable for untrusted, security-sensitive, or multi-tenant applications. A Pod sandbox can also host multiple containers, preserving the standard Kubernetes pattern in which application containers and helper containers, such as sidecars, work together within one Pod.

These capabilities align with ACK's documented sandboxed-container support and Kubernetes Pod model. See the Alibaba Cloud documentation on [using sandboxed containers in ACK](#) and the Kubernetes [Pod concepts documentation](#).

QUESTION NO: 8

About the flow to design microservices architecture for Agile Development, which one is the correct answer?

- A. Requirement Gathering and analysis, System Design, Implementation, Integration and Testing, Deployment of system, Maintenance
- B. Code Repository, Build + Unit Test+Code Quality, Test Script.Test Suite, Provisioning to Containers, deploy as Microservices
- C. Design Thinking, Backlog, User Story Mapping, Epic, List of Microservices
- D. Requirement of Gathering and analysis as well as modular System Design per module, and then deploy applications to Cloud Provider

ANSWER: C

Explanation:

Design Thinking, Backlog, User Story Mapping, Epic, List of Microservices is the appropriate flow because it begins with understanding user and business needs before deciding service boundaries. Design thinking helps the team explore the problem space and identify valuable outcomes. Those needs are then captured and prioritized in a product backlog. User-story mapping organizes the backlog around the user journey, making it easier to identify cohesive business capabilities and slices of functionality. Related stories can be grouped into epics, which provide a useful level of abstraction for discussing domain responsibilities. From those business-oriented epics, the team can derive an initial list of microservices with bounded, independently deliverable responsibilities.

This approach supports Agile development because architectural decisions evolve from prioritized user value and domain behavior rather than from implementation tasks alone. The resulting service candidates can then be refined through iterative delivery, feedback, and domain analysis. The Agile Manifesto emphasizes early and continuous delivery of valuable software and close collaboration around changing requirements; user-story mapping is a practical technique for turning that principle into a release-oriented plan. See the [Agile Manifesto principles](#) and Alibaba Cloud's [Microservices Engine overview](#) for microservices context.

QUESTION NO: 9

Terraform and Helm Chart are the same tool, and are used for infrastructure orchestration.

Score 2

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because Terraform and Helm serve different purposes and are not the same tool. Terraform is an infrastructure-as-code tool that uses declarative configuration files to provision and manage infrastructure resources through providers. In Alibaba Cloud environments, it can create and manage resources such as virtual networks, ECS instances, databases, and ACK clusters, while tracking the intended infrastructure state.

Helm is a Kubernetes package manager. A Helm chart packages Kubernetes application manifests and templates, allowing an application and its Kubernetes resources to be installed, upgraded, configured, and versioned within a cluster. Helm therefore operates primarily at the application deployment layer of Kubernetes, rather than being a general infrastructure orchestration tool.

The tools can be used together: Terraform can provision an ACK cluster and its supporting cloud infrastructure, and Helm can then deploy workloads into that cluster. Their complementary use does not make them equivalent. Alibaba Cloud describes Terraform as a tool for defining and provisioning cloud infrastructure in its [Terraform documentation](#). Helm's role as the Kubernetes package manager is defined in the [official Helm documentation](#).

QUESTION NO: 10

Which of the following features options has released on Alibaba Cloud Elasticsearch? (The number of correct answer: 3)

- A. Elasticsearch
- B. Kibana
- C. Grafana
- D. Logstash

ANSWER: A B D

Explanation:

Alibaba Cloud Elasticsearch provides a managed Elastic Stack environment that includes the core Elasticsearch search and analytics engine, Kibana, and Logstash. **Elasticsearch** is the distributed engine used to store, index, search, and analyze data. In an Alibaba Cloud Elasticsearch cluster, Alibaba Cloud manages underlying operational tasks such as deployment, monitoring, maintenance, and version-related lifecycle functions, while users work with Elasticsearch APIs and cluster capabilities.

Kibana is available as the associated visualization and management interface. It enables users to explore indexed data, build visualizations and dashboards, and access Elastic Stack management functions through a browser. Alibaba Cloud documentation describes accessing Kibana from an Elasticsearch cluster as part of using the service.

Logstash is also supported as a managed data collection and processing component. It can ingest data from varied sources, transform or enrich events through pipeline configurations, and send the processed data to destinations such as Elasticsearch. This makes it a suitable companion for building log and data ingestion pipelines around an Alibaba Cloud Elasticsearch deployment.

See the [Alibaba Cloud Elasticsearch product overview](#) and the [Kibana user guide](#) for service and component details.

QUESTION NO: 11

Does GPU can be monitor by ARMS-Prometheus of alibaba cloud?

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. Alibaba Cloud Managed Service for Prometheus, a component of Application Real-Time Monitoring Service (ARMS), can monitor GPU resources by collecting NVIDIA GPU metrics through DCGM Exporter. This exporter exposes Prometheus-compatible metrics for GPU utilization, memory usage, temperature, power consumption, clock frequency, and error-related status. ARMS then stores and visualizes the collected metrics, allowing teams to create dashboards and configure alert rules for GPU workloads.

For Kubernetes environments such as Container Service for Kubernetes, GPU monitoring is typically enabled by deploying or configuring the NVIDIA DCGM Exporter and configuring the Prometheus instance to discover and scrape its metrics endpoint. This makes GPU observability available alongside cluster, node, container, and application metrics in the same ARMS Prometheus workflow. GPU metric collection is especially useful for identifying underutilized accelerators, capacity bottlenecks, and unhealthy devices in AI, machine-learning, and high-performance-computing workloads.

See Alibaba Cloud's [GPU monitoring documentation for Managed Service for Prometheus](#) and the [Managed Service for Prometheus documentation](#) for configuration details and supported monitoring capabilities.

QUESTION NO: 12

Which of the following statements are correct about Kubernetes liveness and startup probes? (Number of correct answers: 3)

- A. A liveness probe can cause Kubernetes to restart an unhealthy container
- B. A startup probe prevents liveness and readiness probes from running until startup succeeds
- C. HTTP, TCP socket, and exec checks can be used as probe handlers
- D. A readiness probe deletes a Pod when the probe fails
- E. Probes can be configured only for Pods that use host networking

ANSWER: A B C

Explanation:

A liveness probe can cause Kubernetes to restart an unhealthy container is correct. When a container repeatedly fails its liveness probe, the kubelet restarts the container according to the Pod restart policy. This is useful when an application is running but has become deadlocked or otherwise unable to recover itself.

A startup probe prevents liveness and readiness probes from running until startup succeeds is correct. It is intended for applications that require a long initialization time. **HTTP, TCP socket, and exec checks can be used as probe handlers** is also correct. Kubernetes supports these probe mechanisms to check an application according to its protocol and runtime behavior. See the [Kubernetes probes documentation](#).

QUESTION NO: 13

What are the scenarios in managing multiple K8s Cluster in Hybrid & Multi Cloud

Environment?

- A. Kubernetes Cluster Federation
- B. Auto Scaling to Serverless Kubernetes
- C. Register External Kubernetes Cluster to be managed from Alibaba Cloud Container Service Console
- D. Manage every clusters manually but to connect to the same image repository

ANSWER: A

Explanation:

Kubernetes Cluster Federation is the applicable scenario because federation is designed to coordinate workloads and resources across multiple independent Kubernetes clusters. In a hybrid and multi-cloud environment, clusters can reside in different cloud providers, regions, or on-premises locations while remaining separate Kubernetes control planes. A federation layer provides a consistent way to distribute applications, apply common policies, and coordinate multi-cluster service behavior rather than requiring each workload to be deployed and maintained independently in every cluster.

This model is particularly useful when an organization needs workload placement based on locality, availability, capacity, or compliance requirements. It supports the central management objective implied by a multi-cluster environment while preserving the autonomy and failure isolation of individual member clusters. Alibaba Cloud ACK One similarly provides fleet-oriented multi-cluster capabilities for centrally managing and distributing applications across Kubernetes clusters. See the Alibaba Cloud overview of [ACK One](#) and the Kubernetes project's archived reference for [Kubernetes Federation](#).

QUESTION NO: 14

Which of the following options are correct about scale up in Hybrid-cloud solution? (The number of correct answer: 3)

- A. The pods can be scale up to ASK (serverless cluster)
- B. The pods can be scale up to Virtual Node
- C. The pods can be scale up to normal worker nodes, such as ECS
- D. The pods can be scale up to Zero-load Node

ANSWER: A C D

Explanation:

In an ACK hybrid-cloud scale-up design, workloads can be expanded into ACK Serverless (ASK), onto conventional worker nodes such as ECS instances, or through zero-load nodes. These approaches provide different execution targets while allowing Kubernetes scheduling and workload definitions to remain consistent as capacity is extended beyond the original environment. ASK supplies a serverless Kubernetes execution environment, so workloads can be deployed without the cluster operator provisioning or managing underlying worker nodes. This makes it suitable as an elastic target when additional application capacity is needed quickly. ECS worker nodes provide the standard node-based expansion path: a node pool can add ECS instances and the scheduler can place pending pods on the newly available node capacity. Zero-load nodes support hybrid-cloud burst capacity by representing an elastic scheduling target without maintaining continuously active compute nodes, allowing capacity to be obtained when demand requires it rather than being permanently reserved. Together, these scale-up choices support workload growth while balancing operational control, elasticity, and resource efficiency. Alibaba Cloud documents ACK Serverless and its serverless workload model at [What is ACK Serverless](#), and documents node-pool based worker-node management at [Overview of node pools](#).

QUESTION NO: 15

When the asynchronous messages produced by one producer should go to more than one consumer, what tools we should use?

- A. Apache Tomcat
- B. Apache Maven
- C. Apache Kafka
- D. MySQL Server
- E. CDN

ANSWER: C

Explanation:

Apache Kafka is designed for asynchronous, event-driven messaging between producers and consumers. A producer publishes records to a Kafka topic, which durably stores the records in partitions for a configured retention period. Multiple consumers can then read the published messages independently when they belong to separate consumer groups. This provides the required one-to-many delivery pattern: for example, an order-created event can be consumed by one group for inventory processing, another for notifications, and another for analytics. Within each individual consumer group, Kafka distributes partition records among group members, allowing that group's processing workload to scale while maintaining coordinated consumption. Kafka also supports offset tracking, replay of retained messages, high-throughput partitioned topics, and fault-tolerant clustered deployment, making it suitable for decoupling services that communicate asynchronously. Alibaba Cloud's ApsaraMQ for Kafka is a managed Kafka-compatible service and supports the same topic, producer, consumer, and consumer-group model for building these event-streaming architectures. See the [Apache Kafka consumer documentation](#) and [Alibaba Cloud ApsaraMQ for Kafka documentation](#).

QUESTION NO: 16

For database migration, if the source database uses Oracle database features, which

Alibaba Cloud

database should you choose for migration?

Score 2

- A. PolarDB
- B. RDS
- C. AnalyticDB

ANSWER: A

Explanation:

PolarDB is the appropriate choice when the source workload relies on Oracle database features because Alibaba Cloud provides PolarDB for Oracle, an Oracle-compatible edition designed to help migrate Oracle applications with minimal changes. Its Oracle compatibility includes support for familiar Oracle-oriented SQL syntax, procedural language capabilities, data types, and application interfaces, allowing organizations to move compatible workloads to a cloud-native relational database while reducing the effort involved in rewriting schemas and application code.

PolarDB also provides a distributed shared-storage architecture that separates compute and storage, delivering high availability, elastic scaling, and strong performance for enterprise database workloads. These cloud-native characteristics make it suitable not only for compatibility-focused migration but also for modernization after migration. Before a production migration, teams should assess the specific Oracle features, packages, SQL statements, and third-party application dependencies in use, then validate them through testing in a PolarDB for Oracle environment. Alibaba Cloud documents the

Oracle-compatible PolarDB edition and its migration guidance in the [PolarDB for Oracle documentation](#) and the [Data Transmission Service documentation](#).