

# DUMPSBOSS.

**HCIE-Cloud Service Solutions Architect  
(Written) V2.0**

**Huawei H13-831 V2.0**

**Version Demo**

**Total Demo Questions: 34**

**Total Premium Questions: 346**

**Buy Premium PDF**

**<https://dumpsboss.co>**

**[support@dumpsboss.co](mailto:support@dumpsboss.co)**

**support@dumpsboss.co  
dumpsboss.co**

## Topic Break Down

| Topic                                                    | No. of Questions |
|----------------------------------------------------------|------------------|
| Topic 1, Cloud Service Trends and Technologies           | 85               |
| Topic 2, Cloud Service Solution Planning and Design      | 143              |
| Topic 3, Cloud Service Solution Integration and Delivery | 53               |
| Topic 4, Cloud Service Solution O&M and Optimization     | 65               |
| Total                                                    | 346              |

QUESTION NO: 1

Apply performance data for APM and which of the following services can provide performance data input? Multiple selection

- A. Cloud performance testing CPTS
- B. Cloud measurement CloudTest
- C. Application operation and maintenance WSAOM
- D. machine learning service MLS

**ANSWER: A B C**

**Explanation:**

APM relies on performance-related telemetry from both application operations and test activities to build a complete view of application behavior. **Cloud performance testing CPTS** is appropriate because performance tests generate measurable workload and transaction results, such as response time, throughput, concurrency, and error information. These results provide valuable performance-test input when assessing an application's behavior under load. **Cloud measurement CloudTest** is also appropriate because its testing capabilities produce test execution and performance results that can be used alongside APM observations to evaluate application quality and performance. **Application operation and maintenance WSAOM** provides operational monitoring data from running resources and applications. This runtime information complements APM by supplying the operational metrics and monitoring context needed for ongoing performance analysis, troubleshooting, and service health assessment. Together, test-generated performance data from CPTS and CloudTest and production operational data from AOM support APM across pre-release validation and live-service monitoring. Huawei Cloud documentation describes the APM service and its observability role at [APM Documentation](#), and documents AOM's monitoring capabilities at [AOM Documentation](#).

QUESTION NO: 2

A department's service data query system is already running on HUAWEI CLOUD but the delay rate of user feedback query response is high, which affects the user experience. As a cloud service architect, what of the following actions should you take to address this issue? Multiple selection

- A. The best solution for a particular system will vary depending on the workload type, using data to optimize the choice of resource type and configuration options
- B. Use caching middleware to put frequently accessed data in the distributed cache service Redis
- C. Eliminate single points of failure by introducing redundancy, and when one of them fails, the rest can simply absorb a larger share of the amount of work

**ANSWER: A B C**

**Explanation:**

High query-response latency should be addressed through measurement-driven resource planning, caching, and resilient scale-out capacity. "The best solution for a particular system will vary depending on the workload type, using data to optimize the choice of resource type and configuration options" is appropriate because architects should first use monitoring data—such as CPU utilization, memory pressure, disk I/O, network throughput, connection counts, and database-query metrics—to identify the actual bottleneck and select suitable compute, storage, and network specifications. "Use caching middleware to put frequently accessed data in the distributed cache service Redis" reduces repeated database reads and enables low-latency retrieval of hot data. Huawei Cloud Distributed Cache Service provides managed Redis-compatible caching, which is designed for high-throughput, low-latency access; cache keys, expiration policies, and invalidation must be designed to keep returned data suitably current. See [Huawei Cloud Distributed Cache Service](#). "Eliminate single points of failure by introducing redundancy, and when one of them fails, the rest can simply absorb a larger share of the amount of work" supports a load-balanced, horizontally scalable service tier. Multiple healthy instances distribute normal query traffic and retain adequate capacity during a failure, preventing a single overloaded or failed node from causing response delays. Huawei Cloud Elastic Load Balance supports distributing requests across backend servers; see [Huawei Cloud Elastic Load Balance](#).

### QUESTION NO: 3

A live streaming company deploys its live video streaming services by using cloud service resources such as EC, VPC, ELB, EIP, and other cloud services on HUAWEI CLOUD to ensure convenient access for viewers. At the same time, it provides low-latency, high-concurrency, high-definition and smooth video live streaming services. Which of the following cloud services can the live streaming company use to achieve the above goals? Multiple selection

- A.** Use CloudMonitor CES and Message Notification Service to create alarm rules that monitor the number of abnormal hosts and concurrent connections in the ELB when an ELB is abnormal. Those The message notification service will notify the resource abnormality in real time at the first time, and remind the operation and maintenance personnel of the live broadcast company to deal with it in time to avoid business losses as a result
- B.** Use CloudMonitor CES and Auto Scaling to flexibly respond to spikes or drops in business volume by configuring agents, creating alarm rules, and configuring scaling policies for ECS installation
- C.** By using CES and Message Notification Service, CES creates reasonable alarm rules and alarm notification parameters for the EIP and bandwidth, and when bandwidth abnormalities occur, CES will notify the resource source of the exception in real time through the message notification service at the first time. Remind the operation and maintenance personnel of the live broadcast company to deal with it in time to avoid business losses
- D.** By using CloudMonitor CES, by configuring corresponding scaling policies for ELBs and EIPs, the number of ELBs and EIPs can be reached when the network access flow is surge. The purpose of EIP's bandwidth is to automatically expand the capacity

**ANSWER: A B C**

#### Explanation:

Using CloudMonitor CES with Message Notification Service to monitor abnormal ELB hosts and concurrent connections is appropriate because CES can collect ELB monitoring metrics and trigger alarms when configured thresholds are met. Alarm actions can publish notifications through SMN, allowing operations personnel to receive timely alerts and respond before an ELB-related condition affects viewer access. Huawei documents CES alarm-rule creation and notification mechanisms at [Cloud Eye alarm rules](#).

Using CES together with Auto Scaling is also appropriate for a live-streaming workload whose traffic can rise and fall rapidly. CES metrics can serve as scaling triggers, while Auto Scaling policies add or remove ECS instances to maintain sufficient compute capacity and control costs as demand changes. This supports high concurrency and stable service delivery. Huawei describes the service's policy-based, metric-driven capacity adjustment in its [Auto Scaling documentation](#).

Monitoring EIP and bandwidth conditions with CES and sending notifications through SMN is likewise suitable. Bandwidth saturation or abnormal network behavior can degrade stream delivery, so prompt alerts enable the operations team to investigate and mitigate network-capacity issues quickly. Together, these monitoring, notification, and compute-scaling capabilities strengthen the availability and elasticity required for smooth live video services.

### QUESTION NO: 4

The cloud search service provides users with multi-condition retrieval, statistics, and reports of structured and unstructured text. Which of the following scenarios is not suitable for using the cloud search service? Single selection

- A.** Real-time data processing for IOT devices
- B.** Analyze and operate the business
- C.** Search and recommend products on e-commerce websites
- D.** Conduct operation and maintenance analysis and fault location of IT equipment

**ANSWER: A**

#### Explanation:

**Real-time data processing for IOT devices** is not a suitable primary scenario for Cloud Search Service (CSS). CSS is a managed search and analytics service based on Elasticsearch/OpenSearch capabilities. Its core value is indexing data so that users can perform full-text and multi-condition searches, aggregations, visual analysis, and report-oriented queries. It is therefore well suited when IoT data has already been collected and needs to be searched or analyzed, but it is not designed to be the stream-processing engine that performs real-time device-message processing, event handling, rule execution, or time-sensitive control actions.

An IoT real-time processing architecture normally uses dedicated ingestion, messaging, stream-processing, and rule-processing services to receive high-volume telemetry and react to events with low latency. Search storage can be added downstream for historical retrieval, troubleshooting, and analytics, but that is distinct from processing the incoming device data itself. Huawei Cloud describes CSS as an Elasticsearch-based service for search and analysis workloads, while its IoT offerings are intended to connect and manage devices and process their communications. See the [Huawei Cloud Search Service overview](#) and the [Huawei Cloud IoT Device Access service](#).

#### QUESTION NO: 5

In the HUAWEI CLOUD game containerized deployment solution, we recommend that the container image be stored in the DockerHub service to efficiently call the image.

- A. TRUE
- B. FALSE

**ANSWER: B**

**Explanation:**

**FALSE** is correct because HUAWEI CLOUD containerized deployments are designed to use Software Repository for Container (SWR) as the managed image repository. SWR integrates directly with Huawei Cloud services such as Cloud Container Engine (CCE), allowing workloads to pull private images through cloud-internal access paths and using IAM-based authorization. This provides efficient image distribution for game services, whose elastic scaling and frequent version releases can require many concurrent image pulls. SWR also supports image repositories, organization and permission management, image scanning, lifecycle controls, and image synchronization capabilities that fit enterprise container delivery practices. Keeping images in the same Huawei Cloud environment as the CCE cluster reduces external-network dependency and gives administrators centralized control over repository access and image governance. Huawei Cloud documentation describes SWR as its container image repository service and documents its integration with CCE for pulling images during workload deployment. See the [Huawei Cloud SWR documentation](#) and the [CCE image-pull documentation](#).

#### QUESTION NO: 6

Huawei provides a high-performance, open source microservice solution, which of the following microservices supports the application of PHP language? Single selection

- A. Spring microservices
- B. c、 DockerW Service
- C. ServiceMesh microservices

**ANSWER: C**

**Explanation:**

ServiceMesh microservices is the correct choice because a service mesh provides application-language-independent microservice governance. Rather than requiring the PHP application itself to use a Java-specific framework or library, the service mesh deploys a sidecar proxy alongside the application workload. The proxy handles service-to-service traffic and applies capabilities such as traffic routing, load balancing, service discovery integration, observability, mutual TLS, and policy enforcement. As a result, a PHP service can participate in the same governed microservice environment as services developed in Java, Go, Python, or other languages.

Huawei Cloud Application Service Mesh is based on the open-source Istio service-mesh architecture and uses sidecar proxies to manage communications between workloads. This architecture is especially suitable when an organization has polyglot services, including PHP, because governance is delivered at the network layer instead of being tied to a particular application framework. Huawei documents ASM as a platform for managing service communication and traffic in cloud-native applications; Istio also describes the data-plane sidecar model that enables language-neutral integration. See [Huawei Cloud Application Service Mesh documentation](#) and [Istio architecture documentation](#).

#### QUESTION NO: 7

In order to support the sharing and exchange of government affairs data of various commissions and offices in a province, and aggregate them into a unified platform for storage and high-line analysis, which of the following cloud services is recommended? Multiple selection

- A. Yarn
- B. Hive
- C. Spark
- D. DLI-FLink
- E. MapReduce

**ANSWER: A B C D E**

#### Explanation:

For a provincial government-data platform, the recommended capabilities span distributed storage, resource scheduling, SQL-based warehousing, batch processing, interactive analysis, and stream processing. **Yarn** provides cluster resource management and scheduling so that multiple data-processing workloads can share a common big-data platform in a controlled manner. **Hive** supplies a data-warehouse layer and SQL-oriented analysis model, which is useful for organizing and querying aggregated data from many commissions and offices. **MapReduce** provides reliable large-scale batch processing, while **Spark** supports faster general-purpose distributed computation and analytical workloads. Together, these components are commonly available through Huawei Cloud MapReduce Service (MRS), which is intended for deploying and managing big-data clusters and processing large data volumes. **DLI-FLink** adds stream-processing capability, enabling timely processing and analysis of continuously exchanged government-affairs data; it complements offline analysis with real-time analytical pipelines. A unified platform that must accommodate both accumulated historical data and continuously arriving shared data benefits from this full set of processing and analysis technologies. Huawei Cloud describes MRS big-data capabilities at [Huawei Cloud MRS](#) and serverless SQL and Flink processing at [Huawei Cloud DLI](#).

#### QUESTION NO: 8

A company needs to use HUAWEI CLOUD DATABASE (DS) DRS to back up offline data to ApsaraDB for RDS DRS can be as small as the transport granularity down to which of the following levels? Single selection

- A. Record
- B. Example
- C. Table
- D. Library

**ANSWER: C**

#### Explanation:

**Table** is correct. Huawei Cloud Data Replication Service (DRS) lets users define the migration or synchronization scope by selecting database objects, including individual tables. This makes table-level granularity the smallest applicable unit in this question: a task can include only the required tables rather than requiring an entire database or instance to be transferred.

Selecting the specific tables also supports more controlled data movement, reduces unnecessary transferred data, and enables a migration or replication plan to focus on the application data that is actually needed at the destination. During task configuration, DRS provides object selection functionality for choosing the databases and tables that participate in the task; the selected objects determine the scope of full and incremental processing. This capability is particularly useful when an organization is moving a subset of an existing production database, consolidating data, or preparing selected business tables for use by another database service. Huawei documents object selection as a core DRS task-configuration step and describes selecting tables within source databases for migration and synchronization. See the [Huawei Cloud DRS User Guide](#) and the [Huawei Cloud DRS product page](#).

#### QUESTION NO: 9

Cloud Data Migration CDM is an efficient, easy-to-use bulk data migration service. Which of the following security capabilities does the service provide to protect user data?

- A. Instance isolation: Users of the CDM service can only use the instance created by themselves, and the instance is isolated from each other and cannot be accessed by each other
- B. No intermediate storage: During the data migration process, CDM only handles data mapping and transformation, and does not store any user data or fragments
- C. Operating system security: CDM instances run a secure, hardened operating system
- D. Key encryption: CDM strongly encrypts connection credentials for data sources entered by users and stores them in the CDM database

**ANSWER: A B C D**

#### Explanation:

CDM protects migration workloads through several complementary security controls. “Instance isolation: Users of the CDM service can only use the instance created by themselves, and the instance is isolated from each other and cannot be accessed by each other” is a core tenancy protection mechanism: each customer operates its own CDM instance, preventing one tenant from directly accessing another tenant’s migration environment. “No intermediate storage: During the data migration process, CDM only handles data mapping and transformation, and does not store any user data or fragments” minimizes the exposure of business data because CDM processes and transfers the data rather than retaining migrated source records as an intermediate repository.

“Operating system security: CDM instances run a secure, hardened operating system” is also part of platform protection. A hardened and maintained underlying operating environment reduces the risk that the migration engine can be compromised. Finally, “Key encryption: CDM strongly encrypts connection credentials for data sources entered by users and stores them in the CDM database” protects sensitive connection information, such as source and destination access credentials, at rest. Together, isolation, non-persistence of business data, operating-system hardening, and encryption of connection secrets provide defense in depth for data migration. Huawei Cloud’s CDM documentation is available at [Huawei Cloud CDM Documentation](#), and Huawei Cloud’s broader security practices are described at [Huawei Cloud Security](#).

#### QUESTION NO: 10

Which of the following options is a factor to consider in designing for cost on the cloud?

- A. Scaling policy settings
- B. Instance specifications
- C. Select
- D. Instance type as calculation method

**ANSWER: A B D**

#### Explanation:

Cost-conscious cloud architecture requires selecting resources and operational controls that match actual workload demand. **Scaling policy settings** are a key consideration because autoscaling can add capacity during demand peaks and remove unneeded capacity when demand falls. Designing sensible minimum and maximum instance counts, triggers, and scaling thresholds helps avoid both overprovisioning and capacity-related service disruption. Huawei Cloud Auto Scaling supports policies based on scheduled actions and monitoring metrics, making these settings directly relevant to ongoing cost management.

**Instance specifications** are also central to cost design. CPU, memory, storage, network, and accelerator requirements should be aligned with measured application demand; appropriately sized instances prevent paying for capacity that the workload does not use. Finally, the **instance type as calculation method** concerns the resource model and charging approach selected for a workload. Matching predictable, long-running usage to suitable pricing choices and using flexible consumption methods for variable workloads can materially affect total cost. Huawei Cloud provides pricing information and calculators to evaluate these choices before deployment. See [Huawei Cloud Auto Scaling documentation](#) and the [Huawei Cloud pricing page](#).

#### QUESTION NO: 11

In the cloud migration solution in the financial field, the following description of the advantages of the financial zone solution, which is wrong? (Single selection)

- A. Multiple disaster recovery solutions: meet the different choices of various customers.
- B. Safe disaster recovery environment: level 4, T3+ computer room.
- C. One-stop disaster recovery: professional disaster recovery implementation and consulting team.
- D. High data security: The tool SDRS service synchronizes data across AZ dataA, and the RPO > 0. The

#### ANSWER: D

##### Explanation:

**High data security: The tool SDRS service synchronizes data across AZ dataA, and the RPO > 0. The** is the incorrect statement because cross-AZ protection delivered through Huawei Cloud Storage Disaster Recovery Service (SDRS) is designed around synchronous replication for protected EVS disks. With synchronous replication, a write is considered complete only after the replication target has also received it. This architecture provides a recovery point objective of zero, meaning that, after a planned or unplanned failover, there should be no committed-data loss between the protected availability zones. Therefore, describing the service's RPO as greater than zero misstates the key data-protection characteristic that makes this type of disaster recovery suitable for financial workloads with stringent continuity and data-consistency requirements. SDRS also supports disaster recovery across availability zones, helping isolate services from an individual AZ failure while retaining recoverability of the protected workload. Huawei describes SDRS capabilities and cross-AZ disaster-recovery positioning at [Huawei Cloud SDRS](#). The service documentation further describes SDRS concepts and replication behavior at [Huawei Cloud SDRS Product Description](#).

#### QUESTION NO: 12

Operating expenses (Operating Expense, OPE X) refers to the ongoing, expendable expenditure to run the business, as opposed to the capital expenditure (C). Capital Expenditure, CAPEX). Which of the following is not OPEX?(Single).

- A. Engineering design cost
- B. Maintenance contract
- C. Construction cost
- D. Migration cost

#### ANSWER: C

##### Explanation:

Construction cost is not an operating expense because it is incurred to create, acquire, or materially improve a long-lived asset, such as a building, data center facility, or other infrastructure. Rather than being recognized wholly as an expense in the period in which it is paid, directly attributable construction expenditure is generally capitalized as part of the asset's cost. The capitalized amount is then allocated over the asset's useful life through depreciation, matching the cost with the periods that benefit from use of the completed asset. This treatment reflects the central distinction between CAPEX and OPEX: CAPEX produces or significantly enhances a resource that delivers value over multiple accounting periods, while OPEX supports ordinary, ongoing business operations during the current period. The IFRS standard on property, plant, and equipment states that directly attributable costs necessary to bring an asset to the location and condition required for operation are included in its cost; this commonly includes construction-related expenditure. See [IFRS IAS 16: Property, Plant and Equipment](#) and [AccountingCoach: Capital Expenditures and Revenue Expenditures](#).

#### QUESTION NO: 13

Which of the following items is not included in the storage scalability design on HUAWEI CLOUD? Single selection

- A. OBS unlimited capacity
- B. SFS elastic expansion
- C. EVS online expansion
- D. DDM horizontal scaling

**ANSWER: D**

**Explanation:**

**DDM horizontal scaling** is the item that is not part of a storage scalability design. Distributed Database Middleware is a database-service component that provides distributed database capabilities, including database and table sharding, read/write splitting, and scaling of database workloads. Its horizontal scaling concept concerns the database layer: data and requests can be distributed across multiple database instances or shards to improve database processing capacity and concurrency. That is fundamentally an application data-management and database-architecture capability rather than expansion of an underlying cloud storage service.

In a storage scalability design, the focus is on increasing or elastically obtaining storage capacity for objects, files, or block volumes while retaining the required availability, performance, and data-access characteristics. Huawei Cloud documents DDM as a middleware service for building distributed relational databases and managing sharded database resources, which establishes the appropriate service boundary for this choice. See the [Huawei Cloud DDM product description](#). For the distinction between database-layer functions and block-storage capacity management, Huawei Cloud's [EVS disk expansion documentation](#) describes storage expansion as an operation on an elastic volume.

#### QUESTION NO: 14 - (SIMULATION)

The e-commerce platform of an e-commerce company is built based on microservices, but its logistics information is provided by a third-party company, so the e-commerce platform needs to integrate with the third-party logistics system to obtain logistics information, for the communication methods of these two systems, there are two kinds of synchronous communication and asynchronous communication, and the recommended communication method for this scenario is ().

(fill in the blank).

**ANSWER: See the explanation for the answer**

**Explanation:**

**Answer: Asynchronous communication.**

The e-commerce platform integrates with an external third-party logistics system. Because the third-party service may have network latency, temporary unavailability, or variable processing speed, asynchronous communication is recommended.

It decouples the e-commerce platform from the logistics provider.

It avoids blocking e-commerce service requests while waiting for a logistics response.

Messages can be buffered and retried through a message queue if the third-party system is temporarily unavailable. It improves system reliability and resilience for cross-system integration.

#### QUESTION NO: 15

An enterprise needs private connectivity between its on-premises data center and an application VPC on HUAWEI CLOUD. The enterprise cannot use the public Internet to transmit application traffic. After the connection is established, on-premises users must access only the required application port on ECS instances. Which measures should be taken? Select three.

- A. Configure a VPN connection between the on-premises gateway and the VPC VPN gateway
- B. Configure routes for the remote CIDR ranges on both the VPC and on-premises sides
- C. Configure ECS security group rules for the required application port and on-premises CIDR range
- D. Bind EIPs to all ECS instances to establish private connectivity
- E. Allow all inbound ports from the Internet in the ECS security group

#### ANSWER: A B C

##### Explanation:

A VPN connection provides encrypted private connectivity between the on-premises gateway and the VPC VPN gateway. Routing must be configured on both sides so that packets for the remote CIDR ranges are forwarded through the VPN connection. In addition, the ECS security group must permit only the required application port from the on-premises CIDR range.

An EIP provides public network access and does not establish the required private hybrid-cloud connection. Opening all inbound ports from the Internet violates the stated least-access requirement and unnecessarily expands the attack surface.

#### QUESTION NO: 16

To ensure business continuity, an enterprise will deploy and implement intra-city disaster recovery for W-E-A P-D B three-tier architecture on HUAWEI CLOUD, with the following goals: 1. Cross-AZ disaster recovery RPO=0 RTO< 30 min; 2. In order to reduce the cost of disaster recovery, the computing resources of the disaster recovery center are not started. 3. For convenient management, one-click disaster recovery switching is realized; Which of the following solutions is recommended for the above goals? Multiple selection

- A. Use ELB for service access load control, set the weight of the backend server, and under normal circumstances, 100% weight matches the production center. In the event of a disaster, 100% weight is matched to the disaster recovery center
- B. Web services and application servers are synchronized through SDRS Storage Synchronization Service
- C. Use RDS Database Service on the cloud to deploy primary/standby across AZs and synchronize data across AZs
- D. The intelligent DNS service is used to control the traffic of business access traffic and 100% of the user traffic is diverted to the production center under normal conditions, and 100% of the user traffic is switched to the disaster recovery center when the production center fails

#### ANSWER: A B C

##### Explanation:

The recommended design combines ELB traffic steering, SDRS protection for the web and application tiers, and cross-AZ RDS high availability. "Use ELB for service access load control, set the weight of the backend server, and under normal circumstances, 100% weight matches the production center. In the event of a disaster, 100% weight is matched to the disaster recovery center" enables the application entry point to direct normal traffic to the active site and move traffic to recovered backend instances during a DR event. ELB supports distributing requests to backend servers and health-based service availability across availability zones.

“Web services and application servers are synchronized through SDRS Storage Synchronization Service” provides server-level cross-AZ replication and the switchover capability needed to bring up protected workloads at the DR site, allowing the secondary compute instances to remain stopped during normal operation. “Use RDS Database Service on the cloud to deploy primary/standby across AZs and synchronize data across AZs” protects the stateful database tier through an RDS HA deployment. Synchronous replication between the primary and standby database instances supports the required zero data-loss objective, while automatic failover improves recovery time. Together, these services cover application recovery, database continuity, and controlled restoration of client traffic. See Huawei’s [SDRS User Guide](#) and [RDS High Availability documentation](#).

#### QUESTION NO: 17

An enterprise is transforming a monolithic application into multiple microservices. The services need centralized registration and discovery, unified configuration management, and traffic-governance capabilities such as rate limiting and circuit breaking. Which Cloud Service Engine (CSE) capabilities are applicable? Select three.

- A. Service registration and discovery
- B. Centralized configuration management
- C. Traffic governance, rate limiting, and circuit breaking
- D. Provisioning container nodes as a container runtime
- E. Migrating physical servers to ECS instances

**ANSWER: A B C**

#### Explanation:

CSE provides microservice governance capabilities that address these requirements. Service registration and discovery allow service instances to locate each other without hard-coded addresses. Configuration management centralizes application configuration and supports controlled configuration updates. Service governance provides controls such as traffic management, rate limiting, and fault isolation mechanisms including circuit breaking.

CSE is not a replacement for a container runtime or virtual-machine service. It does not itself provide server migration capabilities; SMS is the service designed to migrate servers to HUAWEI CLOUD.

#### QUESTION NO: 18

After you purchase basic cloud services such as HUAWEI CLOUD EPS EVRS and deploy an open source MySQL database on ECS, you will find that the SQL read and write of the database is slow. Which of the following descriptions of user actions are correct? ( Many)

- A. The customer database reads and writes slowly, and the preliminary analysis of the hard disk reads and writes slowly, and the failure to meet the performance indicators promised by the hard disk is reported to Huawei engineers to cooperate with the solution
- B. The customer can check which table is slow to read and write, and if a suitable index has not been established, create a new index.
- C. The customer database is slow to read and write, the preliminary analysis network delay is large, and there is a bottleneck in the bandwidth within the bandwidth purchase range, then the fault is reported to Huawei engineers to cooperate and solve
- D. In order to solve the fault as soon as possible, the customer does not need to analyze the preliminary analysis and directly report the fault

**ANSWER: A B C**

#### Explanation:

When a self-managed open-source MySQL database is deployed on ECS, the customer is responsible for initial fault isolation and database-level optimization, while Huawei Cloud can investigate suspected faults in the underlying cloud resources. Therefore, identifying that disk I/O is slow and opening a service request after verifying that the observed performance does not meet the purchased disk-performance commitment is an appropriate escalation path. The customer should provide the preliminary observations, such as I/O measurements, affected instances, and time range, so support can efficiently investigate the EVS and ECS layers.

Checking the tables and SQL statements that exhibit poor read/write performance, then creating an appropriate index where one is absent, is also a valid customer action. Proper indexes allow MySQL to reduce the rows examined and avoid unnecessary scans, provided the index matches the query predicates and access pattern. MySQL documents this as a core query-optimization technique. Likewise, where preliminary analysis identifies excessive network latency or an apparent bandwidth bottleneck despite traffic remaining within the purchased bandwidth specification, reporting the issue to Huawei Cloud with diagnostic evidence is appropriate. Huawei Cloud support provides the channel for submitting and tracking such resource-related issues. See [MySQL Index Optimization documentation](#) and [Huawei Cloud ECS documentation and service information](#).

#### QUESTION NO: 19

Biogenetics, drug research and development and other fields require high-performance, intensive computing, and are sensitive to cost. In this scenario, you can recommend that you select the Cloud Container Engine CCE service. Single selection

- A. TRUE
- B. FALSE

**ANSWER: B**

#### Explanation:

**FALSE** is correct because the workload characteristics described—large-scale, compute-intensive processing for biogenetics and pharmaceutical research—identify a high-performance computing scenario rather than, by themselves, a container-orchestration scenario. These workloads commonly require substantial CPU or GPU capacity, high-throughput and low-latency networking, scalable parallel execution, and storage performance appropriate for large research datasets. Huawei Cloud's HPC offerings are intended to provide elastic high-performance compute resources for compute-intensive scientific and engineering workloads, while allowing customers to control cost through resource selection and elastic scaling. Cloud Container Engine is a managed Kubernetes service. It is useful when an organization needs to deploy, manage, and scale containerized applications, but choosing it is not the primary recommendation solely from the stated need for intensive computing and cost sensitivity. A suitable architecture may use containers as part of an HPC workflow, but the scenario first calls for an HPC-oriented compute solution selected for the performance requirements of the research jobs. See Huawei Cloud's [High Performance Computing solution](#) and the [Cloud Container Engine product page](#).

#### QUESTION NO: 20

Intelligent edge cloud IEC Service is an extension of HUAWEI CLOUD infrastructure and is generally deployed in remote areas to meet low traffic requirements Low-latency business needs IEC Sites are farther away from the cloud architecture that It is also built by HUAWEI CLOUD establish Operations and operations.

- A. TRUE
- B. FALSE

**ANSWER: A**

#### Explanation:

TRUE is correct. Intelligent EdgeCloud (IEC) extends HUAWEI CLOUD infrastructure from centralized Regions to distributed edge sites. This model places compute, storage, and networking resources geographically nearer to end users, devices, and data sources, including locations that are distant from the central cloud Region. As a result, workloads that are sensitive to network delay can process data locally at an IEC site rather than sending every request back to a centralized cloud location.

This helps deliver the low-latency response needed by edge scenarios such as industrial applications, video services, content delivery, and connected devices.

IEC retains the cloud operating model: Huawei Cloud provides the underlying edge infrastructure and manages the platform's deployment, operations, and maintenance. Customers can therefore use cloud-like resource management and service capabilities while placing workloads at edge locations. The key idea expressed in the statement is that IEC is not an unrelated private-edge product; it is a managed extension of HUAWEI CLOUD designed to bring cloud resources closer to distributed business environments. Huawei Cloud describes IEC as an edge-cloud service that deploys cloud services close to users and data sources. See the [Huawei Cloud Intelligent EdgeCloud product page](#) and the [IEC product description](#).

#### QUESTION NO: 21

A customer purchases both HUAWEI CLOUD and a third-party cloud ECS、EIP、VPC and other cloud service resources, Deploy its e-commerce network As an expert in HUAWEI CLOUD O&M, you need to comprehensively monitor your business operation resources We recommend that you activate HUAWEI CLOUD Application O&M Management serve AOM, Customer use AOM Which of the following features can be implemented? Multiple selection

- A. You can use AOM to monitor ECS instances on HUAWEI CLOUD and other cloud service providers
- B. Customers can use AOM to monitor cloud hosts and docker containers, and also achieve depth of storage and networking
- C. Customers use AOM services to collect metrics, logs and application performance data through ICAGENT Agent must be manual Install to the host
- D. Customers use AOM, which can be collected by installing ICA GENT to the host, or can be provided through AOM of OpenAPI interface or Exporter interface, business indicators as custom indicators, connected to AOM

#### ANSWER: A B C D

##### Explanation:

Huawei Cloud Application Operations Management (AOM) provides unified observability for applications and their underlying operating environments. It can centrally monitor cloud hosts and application workloads deployed across Huawei Cloud, private-cloud, and third-party-cloud environments, provided the relevant agents and connectivity are configured. This supports the operational need to view a distributed e-commerce system through a common monitoring service rather than treating each cloud environment separately.

AOM collects and presents host and container observability data, including resource-consumption indicators associated with compute, disks, and network interfaces. Its ICAgent is deployed on the monitored host to collect metrics, logs, and application-performance-related data. Installation is performed on the host using the installation procedure or command supplied by AOM, after which the agent sends collected data to the service.

In addition to built-in collection, AOM supports custom metric ingestion. Business systems can expose metrics through an Exporter, and applications or integration tools can report custom metrics through AOM APIs. This enables business-specific indicators, such as order throughput, payment success rate, or cart activity, to be monitored together with infrastructure and application data. See the [Huawei Cloud AOM User Guide](#) and [AOM API Reference](#).

#### QUESTION NO: 22

A city has set up a government website to support government affairs sharing, but the website needs to support full-text search and fuzzy query scenarios such as year/month, and day by condition Single selection

- A. Cloud search service CSS
- B. MapReduce service MRS
- C. Table Store service CloudTable
- D. Recommendation System Service RES

**ANSWER: A**

**Explanation:**

Cloud search service CSS is the appropriate service because the stated requirement is centered on searching website content rather than merely storing or processing it. CSS is Huawei Cloud's managed search and analytics offering based on Elasticsearch, which is designed to create indexes over large bodies of text and return relevant documents rapidly. A government-affairs portal can index announcements, policies, service guides, and metadata such as publication year, month, and day. It can then combine full-text matching against document content with structured filter conditions on date fields, allowing users to narrow results to a specified time period while retaining search relevance.

CSS also supports Elasticsearch query capabilities that are important for this type of public website, including full-text queries, match-oriented search behavior, and fuzzy matching for input variations or misspellings. Its distributed architecture enables the search workload to scale as the volume of government information and the number of portal visitors increase. As a fully managed cloud service, CSS also reduces the operational burden of deploying, maintaining, and monitoring a search cluster. Huawei Cloud describes CSS and its Elasticsearch-based search capabilities in the [Cloud Search Service Product Description](#) and documents available query mechanisms in the [Elasticsearch full-text query documentation](#).

**QUESTION NO: 23 - (SIMULATION)**

In scenarios of massive data analysis in the environmental protection industry, you can activate the HUAWEI CLOUD MapReduce service MRS to utilize object storage

OBS implements low-cost storage, uses () to achieve terabytes/petabytes of data analysis, and then exports data to the data warehouse service through the visual import and export tool Loader DWS, complete the BL analysis. (Please fill in the corresponding English abbreviation) (fill in the blank)

**ANSWER: See the explanation for the answer**

**Explanation:**

**Answer: Hive**

In this MRS big-data analysis architecture, OBS provides low-cost object storage, and Hive is used for TB/PB-scale data warehousing and analysis. The processed data can then be exported to DWS through the Loader visual import/export tool for BI analysis.

**QUESTION NO: 24**

Which of the following components can be used to provide storage for massive amounts of IoT data while providing the ability to query in real time?

- A. H i v e
- B. i K af k a
- C. Hbase
- D. HDFS

**ANSWER: C**

**Explanation:**

Hbase is designed for large-scale, distributed storage where applications require low-latency, random read and write access to very large datasets. It is a column-family NoSQL database built on Hadoop Distributed File System storage, allowing data to be distributed across a cluster and expanded by adding nodes. This makes it well suited to Internet of Things workloads, in which devices continuously generate high volumes of time-oriented records such as measurements, states, locations, and events.

For real-time use cases, Hbase supports direct access to rows through its APIs and provides efficient lookups and range scans based on row-key design. An application can therefore ingest incoming device data and query the relevant records without first waiting for a batch-processing job to complete. Appropriate row-key and table design are important for distributing write load and ensuring that queries target the required data efficiently. Huawei Cloud MRS documents HBase as a distributed, column-oriented database intended to provide real-time random read/write access to massive structured data. See the [Huawei Cloud MRS component overview](#) and the [Apache HBase documentation](#).

#### QUESTION NO: 25

The following describes how to select HUAWEI CLOUD DATA WAREHOUSE SERVICE GAU S D B (DWS) or HUAWEI CLOUD Description of RDS, which are the correct items? Multiple selection

- A. DWS is more suitable for online transaction processing (OLTP), and insufficient read support for large amounts of data
- B. DWS uses the scale and resources of multiple nodes and uses various optimization methods (column storage, vector engine, distributed framework, etc.), focusing on online analytical processing (OLAP).
- C. RDS is more suitable for online transaction processing (OLTP), and insufficient read support for large amounts of data
- D. If you need to use custom code to process and analyze very large data sets through big data processing frameworks such as Apache Spark, Hadoop or HBase, you should make it with DWS

**ANSWER: B C**

#### Explanation:

DWS uses the computing and storage resources of multiple nodes and applies data-warehouse optimizations such as column-oriented storage, vectorized execution, and distributed processing. These architectural characteristics are intended for online analytical processing (OLAP): queries that scan, aggregate, join, and analyze large volumes of historical or business data. Huawei describes DWS as an enterprise-class cloud data warehouse designed for massive data analytics, making this selection appropriate when analytical performance and scalable warehouse processing are the core requirements. See the [Huawei Cloud DWS product description](#).

RDS is more suitable for online transaction processing (OLTP) workloads. Its managed relational database model is intended for application transactions that require reliable, concurrent reads and writes, such as order processing, account records, and operational application data. Although RDS can serve application queries and use database read-replica capabilities where applicable, it is not the primary service choice for warehouse-style analysis involving very large data volumes and extensive aggregations. Huawei positions RDS as a managed relational database service for databases such as MySQL, PostgreSQL, and SQL Server; see the [Huawei Cloud RDS product description](#). Therefore, selecting DWS for distributed OLAP analytics and RDS for OLTP applications correctly distinguishes the two services by workload type.

#### QUESTION NO: 26

An Internet ticketing application is deployed on multiple ECS instances. The application must provide one public access address, distribute requests only to healthy instances, and add ECS capacity automatically when the number of concurrent requests increases. Which solution is the most appropriate? Single selection.

- A. Bind an EIP to every ECS instance and allow clients to select an address
- B. Use ELB to distribute traffic to an Auto Scaling group of ECS instances
- C. Use DNS round-robin records for the ECS private IP addresses
- D. Deploy one active ECS instance and one standby ECS instance with a shared EIP

**ANSWER: B**

#### Explanation:

**Use ELB to distribute traffic to an Auto Scaling group of ECS instances** is correct. ELB provides a stable public access endpoint, distributes requests across backend ECS instances, and uses health checks to prevent failed instances from receiving traffic. Auto Scaling can add or remove ECS instances according to configured monitoring metrics and scaling policies. The ECS instances can be registered with the load balancer so that newly added capacity participates in request processing.

Binding an EIP to each ECS instance does not provide a single entry point, health-based traffic distribution, or automatic capacity adjustment. DNS round-robin alone cannot perform the same backend health management and request scheduling as ELB. A standby ECS instance improves fault recovery but does not address concurrent-request distribution or elastic scale-out.

#### QUESTION NO: 27

In the construction of the government cloud platform, with the help of the big data platform to provide public data for third parties to use, and develop big data applications that are conducive to people ' s livelihood services and economic development, which of the following options belong to the application scenarios of government affairs big data? Multiple selection

- A. Data retrieval
- B. Offline processing
- C. Unified analysis
- D. Intelligent decision-making
- E. Real-time risk control

**ANSWER: A B C D E**

#### Explanation:

Government-affairs big data platforms support the complete lifecycle of public-data use: collecting and governing cross-departmental data, processing it at scale, exposing usable data services, analyzing operational conditions, and turning findings into timely public-service and governance actions. **Data retrieval** is a core public-data-service capability, enabling authorized agencies, enterprises, and citizens to locate approved datasets and information resources. **Offline processing** supports large-volume historical data preparation, cleansing, aggregation, and periodic reporting. **Unified analysis** enables data from otherwise separate government domains to be analyzed using common standards and shared analytical capabilities. **Intelligent decision-making** applies models, indicators, and analytical results to support policy evaluation, resource allocation, and public-service improvement. **Real-time risk control** extends these capabilities to streaming or frequently updated data so that abnormal conditions and risks can be detected and handled promptly. Together, these scenarios reflect the practical value chain of a government big-data platform rather than isolated technical functions. Huawei Cloud describes government-cloud solutions as supporting digital government and public services, while its Data Lake Insight service provides batch and real-time analysis capabilities applicable to these data workloads. See [Huawei Cloud Government Solution](#) and [Huawei Cloud Data Lake Insight](#).

#### QUESTION NO: 28

An enterprise uses HUAWEI CLOUD ' s service department to sign its own e-commerce system and stores daily transaction records in Huawei

In cloud OBS, where storage space is used over time, and transactions older than one year are hardly used further, which of the following ways can tenants change the storage type to minimize costs?

- A. Use the lifecycle management configuration of OBS Service to automatically convert objects older than one year to archive storage
- B. Call the OBS SDK to convert the object storage type from one year ago to archive storage
- C. Use the lifecycle management configuration of OBS Service to automatically convert objects one year ago to IA storage
- D. By calling the SDK of OBS, the object storage type from one year ago is converted to the infrequent access storage

**ANSWER: A B**

**Explanation:**

Use the lifecycle management configuration of OBS Service to automatically convert objects older than one year to archive storage is correct because an OBS lifecycle rule can evaluate an object's age and automatically transition eligible data to a less expensive storage class. A rule with a 365-day transition to Archive is well suited to daily transaction records that are retained for compliance or historical purposes but are rarely read after one year. This approach is scalable and eliminates the operational effort and risk of manually identifying aged objects.

Call the OBS SDK to convert the object storage type from one year ago to archive storage is also correct. Applications can use the OBS SDK to programmatically perform the required object storage-class change, typically as part of a scheduled archival workflow. This is useful where the enterprise needs custom selection logic, integration with transaction-processing systems, or explicit control over when individual objects are transitioned. OBS Archive storage is designed for long-term, rarely accessed data, making it the appropriate cost-optimization destination for the stated access pattern. Huawei documents OBS storage classes and their intended access characteristics in the [OBS Product Description](#), and documents automated object management through [OBS lifecycle rules](#).

**QUESTION NO: 29**

The event-driven architecture uses event-triggered services and communication between applications, and there are three key components in its architecture: event producers, event message queues and event consumers, and producer services and consumer services are separated allowing independent scaling, updates, and independent deployment. Judgment

A. True

B. False

**ANSWER: A**

**Explanation:**

**True** is correct. Event-driven architecture organizes application interaction around events: an event producer detects or generates a state change, an event messaging component transports and buffers the event, and an event consumer receives the event and performs the required processing. This asynchronous interaction means that the producer does not have to directly invoke, wait for, or be deployed together with the consumer. As a result, each service can evolve independently and can be scaled according to its own workload characteristics. For example, consumers can be increased when event-processing demand rises, while producers can continue publishing events without being tightly coupled to consumer availability or release schedules. Huawei Cloud EventGrid describes this model as connecting event sources with event targets through event channels, enabling event-driven integration between cloud services and applications. Message-oriented services such as Distributed Message Service also provide the decoupled communication layer commonly used to deliver events reliably and asynchronously. These characteristics support independent deployment, maintenance, upgrades, and elastic scaling of producer and consumer services. See [Huawei Cloud EventGrid Product Introduction](#) and [Huawei Cloud Distributed Message Service for Kafka Product Introduction](#).

**QUESTION NO: 30**

Certain Company offline IDC Small systems have been deployed in the data center, which recently need to be migrated to the HUAWEI CLOUD environment. Deploy the business. When resources are limited and Judeng personnel are familiar with the deployment architecture, O & M personnel use it 5 A physical host completed the business department g, In order to improve the utilization of the host, multiple modules are deployed in each physical host

block web side. The middleware, application, and database are deployed on the same host. When migrating to the cloud. Which of the following methods is the choice of O & M personnel Radio

A. the ports opened by the offline service system, according to the guidance on the cloud set the appropriate ports required by the service module when the security module is opened

B. After the service is migrated to the cloud O & M personnel open each host for maintenance 22 and 3389 port, so that it works Operate O & M anywhere, without geographical restrictions, higher work efficiency, use diverse cloud services, and

flexible resources Features, choose the right specifications, put web. The end and application are deployed using the same virtual host, middleware and database Deploy with cloud services provided by a cl

**C.** From the perspective of architecture and maintenance convenience, directly copy the offline model to purchase 5 bare metal servers on the cloud to deploy services

**ANSWER: A**

**Explanation:**

The appropriate approach is to identify the ports that the migrated services actually require and configure only those ports in HUAWEI CLOUD security groups. A security group acts as a stateful virtual firewall for cloud resources: inbound and outbound traffic is permitted according to explicit rules. During migration, the operations team should map each service dependency—for example, web access, application-to-middleware communication, and application-to-database communication—and allow the required protocol, port, and source range only for the relevant resource or subnet. This retains the existing service communication requirements while applying cloud network isolation controls.

Using least-privilege security-group rules is particularly important when a former on-premises host contains multiple business components, because migration can expose communications that were previously confined to an internal IDC network. The rules should be reviewed as the workload is later split into separate compute instances or managed services, so that access remains limited to known service flows rather than broad network reachability. HUAWEI CLOUD documents security groups as the mechanism for controlling traffic to and from Elastic Cloud Servers and recommends defining rules based on actual access requirements. See the [HUAWEI CLOUD Virtual Private Cloud User Guide](#) and the [security group rules documentation](#).

**QUESTION NO: 31**

An enterprise deploys a set of enterprise websites on HUAWEI CLOUD and to ensure the high availability of the website web service, the enterprise uses the elastic load balancing service ELB and the backend passes 3 ECS ECS. Now that the company needs to implement monitoring log management for the website, which of the following is the right way?

- A.** To configure ELB access logs, you need to associate both a log group and a log stream
- B.** Configure a tracker in Cloud Trace Service (CTS) for operation audit logging
- C.** To implement log management, you need to create a log group and then create a log stream
- D.** Resource monitoring needs to set resource groups first, and then set alarm rules, and the cloud server must set alarm rules separately
- E.** For log collection of the cloud server, you need to install ICAgent

**ANSWER: A B C E**

**Explanation:**

Website log management on HUAWEI CLOUD combines service access logging, centralized log storage, host log collection, and—where operational audit records are required—cloud-operation tracing. ELB access logs should be delivered to Log Tank Service (LTS) through an LTS log group and log stream. This creates a defined destination for requests received by the load balancer, allowing the enterprise to search, analyze, retain, and create alarms from access records centrally. In LTS, the log group is the top-level logical container and the log stream is the collection unit within it, so creating the log group before its log stream is the required resource hierarchy. For application, web-server, operating-system, or other custom logs produced on the ECS instances, ICAgent must be installed and configured; it collects the selected files and sends their contents to the designated LTS log stream. A Cloud Trace Service tracker complements these logs by recording auditable management operations on cloud resources and can be configured for centralized tracking. Together, these capabilities provide visibility into incoming website requests, backend ECS log events, and cloud-resource operations. Huawei documentation for [ELB access logs](#) and [ICAgent-based LTS collection](#) describes these integration patterns.

**QUESTION NO: 32**

The following description of the usage cost and billing model of the cloud container engine CCE, which one is wrong? Single selection

- A. After the CCE cluster is hibernated, the worker nodes (ECS) in the cluster will automatically shut down
- B. CCE supports pay-as-you-go billing and annual subscription/ Monthly billing model
- C. CCE itself does not charge any fees, but related resources (such as nodes, bandwidth, etc.) created during use need to be paid
- D. CCE cluster is closed, the control node resource fee will stop charging

**ANSWER: A**

**Explanation:**

**After the CCE cluster is hibernated, the worker nodes (ECS) in the cluster will automatically shut down** is the incorrect description. CCE cluster hibernation is designed to suspend the cluster control plane so that control-plane-related charges can be reduced or stopped according to the applicable CCE billing model. Hibernation does not automatically power off or release the worker-node ECS instances. Those nodes are independent cloud resources and remain provisioned unless an administrator explicitly stops, deletes, or otherwise manages them. Consequently, their ECS-related charges and any associated resource charges can continue while the cluster is hibernated. This behavior protects node data and configuration and allows the cluster to be resumed without requiring the worker-node infrastructure to be recreated. Operationally, an administrator who wants to reduce the costs of worker nodes must perform the relevant ECS lifecycle action separately, while considering workload availability, local storage, and dependencies such as attached disks and network resources. Huawei Cloud describes CCE as a managed Kubernetes service and provides its current product and billing information at [Huawei Cloud CCE](#) and [CCE Pricing Details](#).

### QUESTION NO: 33

In order to meet the requirements of business agility, an enterprise plans to cloud some application systems, and requires the overall solution design off application to the cloud Regarding the design principles of cloud architecture, which of the following aspects should you consider as an I T architect? Multiple selection

- A. The business application layer ' s own capabilities such as retry, isolation, and elegant failure
- B. In order to ensure the migration effect, try to use high-specification configurations
- C. the system has no single point,HAS clusterDR Disaster recovery ( across THAT and cross Region Disaster recovery sounds), data backup, standardized O & M
- D. Latency and throughput of the system

**ANSWER: A C D**

**Explanation:**

Cloud architecture must address application resilience, availability and disaster recovery, and measurable performance from the beginning of the design. The business application layer ' s own capabilities such as retry, isolation, and elegant failure are essential because cloud environments use distributed services and can experience transient network or service faults. Applications should therefore handle recoverable failures safely, isolate faults to prevent cascading impact, and degrade gracefully when a noncritical dependency is unavailable. This increases service continuity and supports agile delivery without making a single component failure a business outage.

The system has no single point,HAS clusterDR Disaster recovery ( across THAT and cross Region Disaster recovery sounds), data backup, standardized O & M expresses core reliability and operational requirements. High availability across failure domains or regions, recoverable backups, and standardized monitoring and operations procedures enable the system to withstand infrastructure failures and be restored predictably. Latency and throughput of the system must also be considered because they are key nonfunctional requirements that determine user experience, capacity planning, resource sizing, and whether the migrated design can meet business service-level objectives under normal and peak load. Huawei Cloud's architecture guidance emphasizes resilient, secure, high-performing, and operationally manageable designs; see the [Huawei Cloud Architecture Center](#) and [Huawei Cloud white papers](#).

## QUESTION NO: 34

The following description of DCKER, WHICH ONE IS WRONG? The performance loss of radio

- A. Docker containers is relatively low compared to virtual machines
- B. The cker container is essentially an F process
- C. Docker containers do not need to depend on the Linux kernel to run
- D. Docker container is the f running instance of its image

**ANSWER: C**

### Explanation:

“Docker containers do not need to depend on the Linux kernel to run” is the incorrect description. A Docker container is an isolated process (or group of processes) that runs on a host operating system kernel rather than carrying its own complete guest operating system and kernel. On Linux, Docker uses capabilities provided by the Linux kernel, including namespaces for isolation and control groups (cgroups) for resource management. This shared-kernel design is a fundamental reason containers can start quickly and have substantially less overhead than traditional virtual machines, which emulate or virtualize hardware and run separate guest operating systems. A container image packages application code, libraries, configuration, and filesystem layers, but it does not package an independently booted Linux kernel for the container. Therefore, Linux containers require a compatible Linux kernel supplied by the host; Docker Desktop on non-Linux platforms uses a lightweight Linux virtual machine to provide that required kernel environment. Docker’s documentation describes containers as isolated processes and explains that the Docker Engine uses Linux kernel isolation features. See [Docker: What is a container?](#) and [Docker Engine security](#).