

DUMPSBOSS.

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

Huawei H13-831 V.20

Version Demo

Total Demo Questions: 34

Total Premium Questions: 356

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud computing technology trends	46
Topic 2, Cloud service solution architecture design	164
Topic 3, Cloud service solution integration	56
Topic 4, Cloud service solution O&M	74
Topic 5, Cloud service solution operation	16
Total	356

QUESTION NO: 1

A company is an Internet product development company, and the company ' s architect decides to use the grayscale release strategy to iterate the product version, which of the following options cannot be used as the reason why the company adopts grayscale release? Single selection

- A. Teams that adopt agile development, with high release frequency, basically maintain a weekly release frequency
- B. The new version can directly allow end users to test product features
- C. C . System upgrade may be a risk that both old and new versions can be tolerated
- D. D . The answer to the risk of user loss caused by sudden changes in user habits

ANSWER: C

Explanation:

“C . System upgrade may be a risk that both old and new versions can be tolerated” cannot justify adopting a grayscale release strategy. A grayscale release is principally a risk-control mechanism: traffic is progressively directed to a new version, allowing the organization to observe actual production behavior and limit blast radius before wider rollout. It is most valuable where a change introduces uncertainty or impact that warrants controlled validation, rapid rollback, and gradual expansion. If the upgrade risk is already tolerable for both the old and new versions, there is no risk-driven need to add the operational complexity of maintaining parallel versions, configuring traffic splitting, monitoring version-specific metrics, and managing rollback boundaries. In that situation, a normal release process may be sufficient, subject to the organization's change-management requirements. Huawei Cloud's deployment documentation describes grayscale approaches as a way to release versions in batches and manage traffic during application upgrades, supporting controlled production rollout. See the Huawei Cloud [Cloud Container Engine documentation](#) and the Huawei Cloud [Cloud Container Engine product overview](#) for deployment and traffic-management context.

QUESTION NO: 2

A provincial data platform needs to process historical government data stored in a distributed environment. The workload includes SQL-based warehouse analysis, large-scale data transformation, and scheduled offline batch computation. Select three suitable big-data technologies.

- A. Hive
- B. Spark
- C. MapReduce
- D. Distributed Cache Service for Redis as the primary data warehouse

ANSWER: A B C

Explanation:

Hive provides a data-warehouse and SQL-query layer for large datasets. Spark supports distributed transformation and analytical processing, particularly where multi-stage processing is required. MapReduce is suitable for scheduled, large-scale offline batch computation. DCS for Redis is an in-memory caching service and should not be used as the primary platform for persistent government-data warehousing and large offline analytics.

QUESTION NO: 3

Which of the following is not an optimization practice for low-latency services? Single selection

- A. Use heterogeneous computing resources and use differentiated computing architectures for different types of services, such as GPU and FPGA computing resource pools

B. B. Use intelligent DNS distribution deployment to shorten network latency, reduce data transmission, and improve efficiency

C. Statement Discipline Optimization

D. D. Containerized applications are deployed in CCE clusters, and application performance management APM is used to analyze the call status of applications and optimize the architecture after understanding application bottlenecks

E. Use the database intermediate DDM to split the database and table, reduce the whole to zero, and reduce the query delay of the huge amount of data in a single database

ANSWER: C

Explanation:

Statement Discipline Optimization is the correct selection because it is not a recognized or actionable Huawei Cloud low-latency optimization practice. The phrase does not identify a defined technology, architecture pattern, configuration method, or operational activity that an engineer can apply to reduce end-to-end response time. Effective latency optimization requires concrete measures that address identifiable contributors to delay, such as network path selection, compute performance, application call chains, database access, and resource locality. A valid practice should therefore specify what is being optimized and how it affects service response time, enabling the team to measure outcomes through latency metrics, traces, and performance monitoring. Huawei Cloud Application Performance Management provides observability of distributed application requests and call relationships, supporting identification of performance bottlenecks and targeted optimization. Likewise, Huawei Cloud's cloud-native guidance emphasizes using measurable application and infrastructure performance information to tune workloads. A vague expression such as **Statement Discipline Optimization** provides neither an established optimization mechanism nor sufficient implementation detail, so it cannot be treated as a low-latency service optimization practice. See Huawei Cloud's [APM User Guide](#) and [Cloud Container Engine best practices](#).

QUESTION NO: 4

A company needs to use hadoop for data analysis and processing, with an initial data volume of about 30PB and an average increase of 10PB per year in the next two years, using M RS / D The I+ O BS memory and computing separation architecture realizes flexible computing and storage separation and expansion capabilities, which can take into account scalability and single selection

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. For a Hadoop-oriented analytics platform beginning at approximately 30 PB and increasing by roughly 10 PB each year, separating persistent data storage from compute resources is an appropriate architecture. Huawei Cloud OBS provides highly scalable object storage for the data lake layer, allowing capacity to grow independently as the data volume expands. Compute services can then be provisioned and scaled according to the processing workload rather than being permanently coupled to storage capacity. This avoids having to add compute nodes merely to obtain additional disk space and supports more flexible resource planning for batch analysis, SQL analytics, and other big-data jobs.

Huawei Cloud MRS supports big-data ecosystem components and can use OBS as persistent storage in a storage-compute decoupled design. DLI is a serverless data analytics service that can analyze data stored in OBS, enabling elastic compute without cluster-capacity planning for every workload. Therefore, an MRS/DLI plus OBS architecture provides the required independent scaling of compute and storage while accommodating continued multi-petabyte data growth. See Huawei's [MRS product overview](#) and [DLI product overview](#).

QUESTION NO: 5

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 topic
- B. Diameter" for configuration
- 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 ICA gent

ANSWER: A C E

Explanation:

To configure ELB access logs, you need to associate both a log group and a log topic because ELB delivers access-log records to Cloud Log Tank Service (LTS), where the destination is organized by a log group and a log stream/topic. This provides centralized visibility into client requests, listener and backend access behavior, and traffic-related events for the website service. To implement log management, you need to create a log group and then create a log stream because these are the fundamental LTS containers: the log group provides the management boundary, while the stream stores records from a specific source or collection configuration. For log collection of the cloud server, you need to install ICA gent because the ICAgent installed on an ECS instance collects configured host log files and sends them to the designated LTS log stream. Together, ELB access-log delivery and ECS agent-based collection give operations staff a unified log-management design covering both the load-balancing entry point and the web servers. Huawei Cloud product guidance for [Log Tank Service](#) describes centralized collection and management of cloud logs, while the [Elastic Load Balance](#) service documentation covers load-balancing capabilities and operational integration.

QUESTION NO: 6

DDM supports both compute layer (DDM) scaling (increasing the number of nodes or increasing node specifications) and storage layer (MYS). QL) is expanding online, and storage layer expansion is divided into translation expansion (the total number of shards remains unchanged) and doubling expansion (doubling the number of shards). Computing layer expansion affects services in seconds, and storage layer expansion is completely transparent to services. Single selection

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. Huawei Cloud Distributed Database Middleware (DDM) separates the computing layer, provided by DDM nodes, from the storage layer, provided by underlying MySQL shard instances. This separation enables elastic scaling at both layers. At the computing layer, an instance can be expanded by adding nodes or changing node specifications. The related adjustment can cause a brief, seconds-level impact while DDM resources and connections are switched or rebalanced. At the storage layer, DDM supports online expansion strategies that either retain the existing shard count while increasing capacity/specifications or increase the shard count according to the supported shard-doubling approach. DDM coordinates routing and metadata so that applications continue accessing the logical database rather than needing to change connection endpoints, database schemas, or SQL statements. Consequently, storage-layer expansion is designed to be transparent from the application-service perspective. This layered elastic architecture is a key DDM capability: compute resources can be adjusted for routing and processing demand, while storage resources can be expanded to accommodate data volume and database throughput. Huawei's DDM product overview describes its distributed architecture and elastic scaling capabilities; see [Huawei Cloud DDM](#) and the [DDM Documentation](#).

QUESTION NO: 7

What is wrong with the following description of Docker?

- A. The Docker image contains the packaged application and the environment it depends on
- B. B. You can package the image through the docker save command
- C. C. You can build container images through the docker build command
- D. A writable layer is added to the top layer of the image

ANSWER: D

Explanation:

“A writable layer is added to the top layer of the image” is the incorrect description because a Docker image is not given a writable layer. An image consists of a stack of immutable, read-only filesystem layers, which makes it portable and suitable for reuse as a template. The writable layer belongs to a *container* created from that image, not to the image itself. When Docker starts or creates a container, it places a thin writable container layer above the image's read-only layers. Files created or modified during the container's runtime are stored in this container layer, while the original image layers remain unchanged. This layered design enables multiple containers to share the same underlying image layers efficiently, while each container has isolated runtime filesystem changes. Docker's storage documentation describes the writable container layer and its relationship to immutable image layers; Docker's image-layer documentation likewise explains that image layers are read-only. See [Docker storage drivers](#) and [Understanding image layers](#).

QUESTION NO: 8

When you use HUAWEI CLOUD OBJECT STORAGE MIGRATION SERVICE OMS to migrate data, the migration is listed in lexicographical order by default with the selected prefix. Whether the new data can be migrated successfully during the migration process It depends on whether the migration task has completed the list of relevant river turns when it is added.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. For an Object Storage Migration Service (OMS) object migration, OMS enumerates source objects using their names in lexicographical order, subject to the prefix selected for the task. This listing process determines the set of objects that OMS will subsequently process for migration. Therefore, an object added while the task is running can be migrated if it is added before OMS has passed the relevant point in its source-object listing process. In practical terms, the outcome for newly added data depends on the task's listing progress at the time the object is created, including the object name's position in the lexicographical sequence under the configured prefix. This behavior matters when source buckets continue receiving writes during a migration: administrators should plan an incremental migration or a final synchronization run when a consistent final data set is required. Huawei's OMS documentation describes OMS as a service for migrating data between object-storage systems and provides task configuration and operational guidance at [Huawei Cloud OMS Documentation](#). Additional service concepts and API documentation are available at [OMS API Reference](#).

QUESTION NO: 9

In the HPC scenario, the data of enterprise users can be uploaded to O B S by direct upload or data express , scenario architecture, the following description of the scenario, which are the correct items ? (Many)

- A. The REST API provided by OBS has been well developed and the console, SDK, and various tools are provided for users to create the OBS bucket and upload objects to the bucket
- B. The file semantics and HDFS semantics provided by OBS support directly mounting OBS to HPC flavors nodes and large data Under the application of AI-I analysis, it provides convenient and efficient data read. write and storage capabilities for all aspects of high-performance computing

- C. C . Large file storage can also replace OBS with SFS
- D. D . Data express support online mailing to deliver data to OBS bucket

ANSWER: A B D

Explanation:

“The REST API provided by OBS has been well developed and the console, SDK, and various tools are provided for users to create the OBS bucket and upload objects to the bucket” is correct because OBS offers multiple standard methods for creating buckets and transferring objects. These include the web console, RESTful APIs, language SDKs, and client-side tools, allowing enterprise data-ingestion processes to be automated or operated interactively according to the workload’s requirements. In an HPC architecture, this enables input datasets and generated results to be placed in durable, scalable object storage.

“The file semantics and HDFS semantics provided by OBS support directly mounting OBS to HPC flavor nodes” is correct in the intended architecture context. Huawei Cloud provides integration mechanisms for accessing OBS from compute and big-data workloads, enabling HPC and AI analysis applications to use centrally stored data while maintaining efficient access patterns for large-scale processing.

“Data Express supports online mailing to deliver data to OBS bucket” is correct because Data Express is intended for secure bulk data migration where network upload is impractical or too slow. Data is copied to a physical storage device, shipped through the service workflow, and imported into the target OBS bucket. See [Huawei Cloud OBS](#) and [Huawei Cloud Data Express](#).

QUESTION NO: 10

On HUAWEI CLOUD which of the following aggregation methods is not supported by CloudMonitor CES? Single selection

- A. maximum value
- B. difference value
- C. mean
- D. variance

ANSWER: B

Explanation:

difference value is the aggregation method not supported by Huawei Cloud Cloud Eye (CES). CES collects raw metric samples and rolls them up over a selected period using statistical aggregation functions. The supported functions include common statistical results such as maximum, minimum, average, sum, sample count, and variance. These functions summarize the set of samples observed during an aggregation interval and enable users to view a metric consistently at different periods and resolutions.

A difference value, in contrast, describes a calculated change between two readings rather than a statistical aggregation of the samples in one period. Where a change between successive values is needed, it must be derived from the metric data or handled through an appropriate metric/reporting design; it is not one of the CES aggregation values available for metric queries or chart aggregation. Huawei’s Cloud Eye API documentation defines the metric-data query aggregation field and its supported statistical values, including `variance`. The Cloud Eye user documentation also describes how monitored metrics are displayed and aggregated over periods. See [Huawei Cloud Eye API Reference: Querying Metric Data](#) and [Huawei Cloud Eye User Guide](#).

QUESTION NO: 11 - (SIMULATION)

HUAWEI CLOUD game O & M and log analysis scheme—Which of the following cloud services is generally involved? Multiple selection

- A , Cloud supervisor ridicule Gengwu THOSE
- B. Application Performance Management Service APM
- C. Log Analysis Service LTS
- D . Data encryption service DEW

ANSWER: See the explanation for the answer

Explanation:

Correct answers: A, B, C

A. Cloud Eye Service (CES) — used for infrastructure/resource monitoring, alarms, and operational visibility.

B. Application Performance Management (APM) — used to trace application requests and identify application performance bottlenecks.

C. Log Tank Service (LTS) — used to collect, search, analyze, and visualize game/application logs.

D. Data Encryption Workshop (DEW) is mainly for key management and encryption-related protection; it is not a general core service for an O&M monitoring and log-analysis solution.

QUESTION NO: 12

Pay only for the required computing resources and optimize the billing model according to business needs, accurately identify the usage and cost of the system , and then clearly allocate the cost to each business system is the main focus of cloud cost design. Single selection

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. Cloud cost design aims to make spending directly traceable to the business value and resource consumption that produce it. This starts with selecting resources that match actual demand and using appropriate charging modes, such as pay-per-use for variable or short-term workloads and longer-term commitment models where stable demand justifies them. Cost design also requires continuous visibility into consumption: teams need to identify which cloud services, accounts, projects, and workloads generate charges, rather than treating the cloud bill as an undifferentiated total.

Clear allocation of costs to individual business systems is equally important because it enables accountability, budgeting, showback or chargeback, and informed architectural decisions. Consistent resource organization and cost-allocation metadata, such as enterprise projects, tags, and cost-center identifiers, allow bills and usage reports to be associated with the responsible application or department. This supports ongoing optimization as workload demand changes. Huawei Cloud billing guidance describes how customers can review expenditure and usage details, while its enterprise-management capabilities support organizing resources for management and accounting purposes. See the [Huawei Cloud Billing documentation](#) and [Huawei Cloud Enterprise Management documentation](#).

QUESTION NO: 13

During the research stage of business migration to the cloud enterprises need to collect source-end system information. The following description of the source information that needs to be collected. which is wrong?

- A. You need to collect the effective data amount and disk size of the source host teaching disk
- B. You need to collect the kernel version of the source host operating system
- C. The network bandwidth from the source host to the target platform needs to be collected

D. You need to collect the average usage during the CPU-day period of the source host

ANSWER: D

Explanation:

You need to collect the average usage during the CPU-day period of the source host is the incorrect statement because migration assessment and target resource sizing must account for workload peaks, not merely an average CPU-use figure. An average over a day can conceal short but business-critical surges in CPU demand. If the target environment is sized from that average alone, the migrated application may have insufficient vCPU capacity during busy periods, causing performance degradation after cutover. A sound assessment therefore gathers CPU utilization trends across an appropriate observation window and identifies peak utilization, peak-time patterns, and sustained load together with other resource metrics. This gives architects enough information to select ECS specifications and scaling measures that satisfy application performance requirements under real production demand. Huawei Cloud Server Migration Service guidance treats assessment and preparation of source-server resources as key prerequisites for a successful migration, while Huawei Cloud ECS documentation emphasizes selecting compute specifications that fit the workload. See the [Huawei Cloud Server Migration Service best practices](#) and the [Elastic Cloud Server product documentation](#).

QUESTION NO: 14

In the HUAWEI CLOUD microservice engine CSE, ServiceCenter is a microservice component with microservice instance registration and discovery capabilities, the following description of ServiceCenter, which are the correct items?(Many)

- A. Supports persistent connections to monitor the state changes of microservice instances
- B. Provide CLI to view microservices, and do not support Web Portal to display microservice management interfaces
- C. ServiceComb The ability to register and dynamically discover services also relies on ServiceCenter Reality
- D. Manage microservice metadata through RESTful API

ANSWER: A C D

Explanation:

ServiceCenter is the service-registry component used by CSE and ServiceComb-based applications to register microservices and their instances, locate available instances, and maintain service information dynamically. "Supports persistent connections to monitor the state changes of microservice instances" is correct because the registry supports mechanisms for clients to obtain and track instance-status and endpoint changes, allowing consumers to react when providers are added, removed, or become unavailable. This dynamic registry behavior is fundamental to service discovery in a distributed microservice environment.

"ServiceComb The ability to register and dynamically discover services also relies on ServiceCenter Reality" is also correct: ServiceComb clients use ServiceCenter as the central registry for publishing instance information and querying discoverable service instances. "Manage microservice metadata through RESTful API" is correct because ServiceCenter exposes REST APIs for registry operations, including managing services, instances, schemas, and related metadata. These APIs allow platform tools and microservice clients to perform registration, discovery, heartbeat, and metadata-management operations programmatically. See the [Apache ServiceComb Service Center overview](#) and the [Service Center REST API documentation](#).

QUESTION NO: 15

A production database must be protected against both an AZ fault and a regional fault. The business accepts an RPO of several minutes for a regional disaster and plans to start application compute resources in the remote region only after a disaster occurs. Select three appropriate design measures.

- A. Deploy the production RDS instance in primary/standby mode across two AZs in the production region.
- B. Use DRS to continuously synchronize production database changes to a database instance in the remote region.

- C. Prepare and regularly test a recovery runbook that starts remote resources and switches service traffic after recovery validation.
- D. Rely only on EVS snapshots in the production region for protection against a regional disaster.
- E. Use asynchronous cross-region synchronization and treat it as a guaranteed zero-RPO mechanism.

ANSWER: A B C

Explanation:

An RDS primary/standby deployment across AZs provides local high availability for database-instance or AZ failures in the production region. DRS continuous synchronization to an RDS instance in a remote region provides the remote database copy needed for regional disaster recovery; because cross-region synchronization is asynchronous, its achievable RPO must be monitored and managed against the business target. A documented traffic-switching and recovery procedure is required to start remote application resources and direct users to the recovered service. A local snapshot alone does not protect against a regional event, and asynchronous replication cannot be described as a zero-RPO solution.

QUESTION NO: 16

With the rapid growth of user access, when the system database has a performance bottleneck, you can do read and write separation, according to the AKF expansion cube method. this practice is based on which of the following expansion ideas?

- A. X axis
- B. Y axis
- C. Z axis
- D. D . Split answers based on different businesses

ANSWER: A

Explanation:

X axis is correct. In the AKF Scaling Cube, X-axis scaling means cloning or replicating the same service so that incoming workload can be distributed across multiple equivalent instances. Read/write separation applies this replication principle to database access: a primary database instance processes writes, while replicated read instances absorb read traffic. This increases aggregate read capacity without requiring every request to compete for resources on the write-processing instance. A routing layer, database proxy, or application data-access logic directs queries to the appropriate database role, while replication propagates committed changes from the write source to read replicas. The approach is especially appropriate where read traffic is substantially larger than write traffic, which is common for user-facing systems. It is therefore an example of expanding capacity through duplicated instances that share the same functional responsibility, rather than dividing the application into distinct functions or partitioning users or data sets. The AKF model describes this replication-based growth approach as X-axis scaling; see [AKF Partners: AKF Scaling Cube](#). Huawei Cloud's database architecture guidance likewise describes read replicas as a way to offload read workloads and improve read performance; see [Huawei Cloud RDS Best Practices](#).

QUESTION NO: 17

The following description of the basic objects in Kubernetes, which are correct? Multiple selection

- A. Job: Used to control batch-task objects.
- B. Pod: The smallest deployable unit that can be created, scheduled, and managed.
- C. Service: Defines a logical collection of Pods and policies for accessing that collection.
- D. Deployment: An object used to manage stateless applications.

ANSWER: A B C D

Explanation:

All four statements describe standard Kubernetes workload or service-discovery objects when the fragmented wording is normalized. A Job is a workload controller intended for finite, batch-oriented work: it creates one or more Pods and tracks their successful completion according to the requested completion criteria. A Pod is Kubernetes' smallest deployable and manageable unit; it can contain one or more tightly coupled containers that share networking and, where configured, storage resources. A Service provides a stable logical abstraction over a selected set of Pods, defining how that group is reached through mechanisms such as a virtual IP address and DNS name. Its selector-based model allows the backing Pods to change without requiring clients to track individual Pod addresses. A Deployment is a workload controller for stateless applications. It manages ReplicaSets and declaratively maintains the desired number of Pod replicas while supporting controlled rolling updates and rollbacks. These objects are fundamental building blocks for running batch work, deploying applications, exposing application endpoints, and maintaining desired application state in Kubernetes. See the Kubernetes documentation for [Jobs](#) and for [Services](#).

QUESTION NO: 18

When planning a VPC, if there is no network isolation requirement between services and it is only deployed in a single region, which of the following architectures is most appropriate?

- A. Multi-VPC architecture
- B. Multi-Region Single VPC Architecture
- C. Multi-Region Multi-VPC Architecture
- D. Single-VPC architecture

ANSWER: D

Explanation:

Single-VPC architecture is the appropriate design when all services are deployed within one region and do not require network-level isolation from each other. A VPC is a logically isolated virtual network within a single Huawei Cloud region, and it can contain multiple subnets, Elastic Cloud Servers, load balancers, NAT gateways, and other resources. In this scenario, placing the workloads in one VPC simplifies IP address planning, route management, security-group administration, and connectivity between services. Subnets can still be used to organize resources by tier or function—for example, web, application, and database tiers—without introducing separate VPC boundaries. Huawei Cloud guidance describes VPCs as regional resources and supports subnet-based network segmentation within a VPC, making this architecture sufficient for a single-region deployment with no inter-service isolation requirement. Appropriate security groups and network ACLs can still enforce the required access controls while retaining the operational simplicity of one VPC. See the [Huawei Cloud Virtual Private Cloud User Guide](#) and [Huawei Cloud VPC product overview](#) for VPC scope and networking concepts.

QUESTION NO: 19

In the research stage of the implementation of the cloud migration plan, you need to collect customer source information, which of the following items belongs to the customer source information that needs to be collected at this stage ? Multiple selection

- A. The amount of data stored and the file type
- B. Network architecture, high reliability, disaster recovery, deployment requirements
- C. Source host specifications, IP, disks, and other information
- D. The user name, authentication email
- E. operating system version, and database engine version of all source system guests

ANSWER: A B C E

Explanation:

A migration research and assessment phase must establish an accurate inventory of the source environment so that the target cloud architecture, migration tools, migration sequence, and validation plan can be designed appropriately. **The amount of data stored and the file type** are essential for estimating migration duration, bandwidth demand, transfer method, storage service selection, and compatibility requirements. **Network architecture, high reliability, disaster recovery, deployment requirements** capture the application's operational design and availability objectives, which are needed to plan connectivity, target topology, resilience, and cutover arrangements.

Source host specifications, IP, disks, and other information provide the server-level inventory required for sizing cloud compute and storage resources, mapping network addresses, and determining migration dependencies. **operating system version, and database engine version of all source system guests** identifies platform compatibility constraints and helps select suitable migration approaches and target service versions. Together, these discovery items form the technical baseline for assessing workloads before migration and for creating a controlled, low-risk migration plan. Huawei Cloud Migration Center is intended to support migration assessment and planning based on source-resource discovery; see [Huawei Cloud Migration Center](#) and [Huawei Cloud Migration Center User Guide](#).

QUESTION NO: 20

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, from finding and preparing data through deriving insight and acting on it. **Data retrieval** is an essential scenario because government departments, authorized organizations, and public-service developers need to discover available datasets and obtain relevant information efficiently. **Offline processing** supports scheduled batch aggregation, cleansing, transformation, and historical-data computation, which are fundamental when integrating large volumes of data from multiple government systems. **Unified analysis** is applicable because a government data platform consolidates data across domains and enables consistent cross-department analysis, improving the value of shared public data. **Intelligent decision-making** follows from analytical results: models, dashboards, and data-driven insights can assist public-service planning, resource allocation, and governance decisions. **Real-time risk control** is also a government big-data scenario, as streaming data and near-real-time analytics can identify operational, public-safety, service, or compliance risks quickly enough to support timely intervention. These capabilities align with cloud big-data services that provide data integration, governance, development, analytics, and batch or real-time processing. Huawei Cloud describes these platform capabilities in its [DataArts Studio overview](#) and its [Data Lake Insight overview](#).

QUESTION NO: 21

When building CDN edge nodes, which of the following is not a mandatory server role?

- A. CACHE server
- B. PMGR server

- C. ELB server
- D. network management server

ANSWER: D

Explanation:

network management server is the correct answer because it is not a mandatory role deployed within every CDN edge-node build. An edge node must provide the functions that directly receive, schedule, and serve cached content: content caching, local request processing and balancing, and the node-management functions needed to operate those content-serving components. These functions correspond to the CACHE, PMGR, and ELB roles in the described CDN edge-node architecture. A network-management capability, by contrast, belongs to the broader operations and management domain. It can be centrally deployed and used to supervise network infrastructure and connectivity across multiple nodes rather than being installed as an intrinsic server role on each CDN edge node. This separation lets CDN edge nodes remain focused on efficient content delivery while management systems provide centralized monitoring and administration. Huawei Cloud CDN likewise describes CDN as an acceleration service based on distributed edge resources that cache and deliver content close to users, with service management handled through the cloud service's management interfaces. See Huawei's [CDN product description](#) and [CDN user guide](#).

QUESTION NO: 22

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. Set performance metrics and monitoring metrics to capture key performance indicators, and use visualization technology to identify performance problems , hot spots, and low usage rates in waiting statusCc . Use messaging middleware to put frequently accessed data in the distributed cache service Red iS
- C. D . 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

Explanation:

"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 query latency must be assessed against the application's actual workload profile. Capacity and resource selection should be based on observed CPU, memory, network, storage, database, and connection-use characteristics. This enables the architect to right-size compute resources and select configurations that remove the real performance bottleneck rather than applying a generic scaling change.

"Set performance metrics and monitoring metrics to capture key performance indicators, and use visualization technology to identify performance problems, hot spots, and low usage rates in waiting status" together with using Redis for frequently accessed data is also appropriate. Monitoring establishes a measurable latency baseline and helps locate slow query paths, saturated resources, and hotspots. Caching repeatedly requested data in a distributed Redis cache can substantially reduce database reads and query-processing time, improving response latency for cacheable high-frequency requests. Huawei Cloud Cloud Eye supports metric monitoring, alarms, and visualization for cloud resources, while Distributed Cache Service for Redis is designed for low-latency, high-throughput data access. See [Huawei Cloud Eye User Guide](#) and [Huawei Cloud Distributed Cache Service for Redis](#).

QUESTION NO: 23

As an important part of the daily application management of enterprises, alarm management can be carried out through the AOM alarm management function of application operation and maintenance management after business applications are migrated to the cloud The following describes the alarm threshold rule settings in AOM, which are correct?

- A. B. The default batch static threshold rule monitors all hosts or all components
- B. Static threshold rules can set threshold conditions on indicators of applications (such as application resource occupancy, throughput, and application request errors).
- C. D. Once the set threshold condition is triggered, a threshold alarm will be generated synchronously, and the user can receive the corresponding notification through email or SMS

ANSWER: A B C

Explanation:

AOM static threshold rules are used to continuously evaluate monitored metric data against administrator-defined fixed thresholds. The default batch static threshold rule can be applied broadly to all hosts or all components, which provides a practical baseline for monitoring resource and service health without creating a separate rule for every monitored object. Static threshold rules also support application-related metrics, enabling alert conditions for indicators such as resource usage, throughput, and application request errors. This allows operations teams to set measurable service-health objectives and detect abnormal behavior promptly. When a configured threshold condition is met, AOM generates a threshold alarm. Alarm notification can be configured through the alarm rule's notification settings and delivered by channels including email and SMS, allowing responsible personnel to receive timely operational alerts. In practice, the notification recipient and channel must be configured appropriately in the associated notification settings, while the threshold rule provides the condition that triggers the alarm. Huawei documents AOM as a monitoring and alarm-management service for cloud applications and resources, including configurable metric alarm rules and notifications. See the [Huawei Cloud Application Operations Management product page](#) and the [Huawei Cloud AOM documentation portal](#).

QUESTION NO: 24

After the host connection configuration is completed, a user does not collect logs through Log Service. The following description of the causes of the effect, which one is wrong?

- A. The collection path configured for the same host is unique.
- B. After configuration is completed, wait for logs to start being reported.
- C. The same log collection path on the host has already been configured in AOM.
- D. The host access configuration may not be standardized.

ANSWER: A

Explanation:

The collection path configured for the same host being unique is the incorrect cause of a failure to collect logs. A single, unambiguous collection path is the normal and desirable configuration state for host log collection. It tells Log Tank Service exactly which file or directory content is intended for collection and does not, by itself, prevent the installed collection agent from reading and reporting logs. In practice, each collection configuration should clearly associate a host group, log stream, and path pattern so that the agent can apply the configuration predictably. After host access and collection configuration are completed, the agent uses this path information to detect appended log content and forward it to the configured log stream. Therefore, uniqueness of the configured path is not an abnormal condition and cannot be treated as a reason that log collection has no effect. Huawei Cloud describes Log Tank Service host log collection and its host-access workflow in its [Log Tank Service documentation](#) and provides the service overview at [Huawei Cloud Log Tank Service](#).

QUESTION NO: 25

In the HUAWEI CLOUD microservice engine CSE, ServiceCenter is a microservice component with microservice instance registration and discovery capabilities, the following description of ServiceCenter, which are the correct items?(Many)

- A. Supports persistent connections to monitor the state changes of microservice instances

- B. Provide CLI to view microservices, and do not support Web Portal to display microservice management interfaces
- C. ServiceComb The ability to register and dynamically discover services also relies on ServiceCenter Registry
- D. Manage microservice metadata through RESTful API

ANSWER: A C D

Explanation:

ServiceCenter is the service-registry component used by ServiceComb/CSE to maintain microservice and instance information and make that information available for service discovery. It supports long-lived watch or persistent connections so consumers can receive instance-state and instance-list change notifications promptly. This enables clients to react to registrations, removals, and availability changes without relying solely on repeated polling. ServiceComb's registration and dynamic-discovery functions use ServiceCenter as the registry: providers register their instance data, while consumers query or watch the registry to locate suitable provider instances. ServiceCenter also exposes RESTful interfaces for managing and retrieving registry resources, including microservices, versions, instances, schemas, rules, and related metadata. These APIs allow the registry functions to be integrated with platform tools and operational automation. The ServiceComb ServiceCenter project documentation describes it as a service registry and identifies its REST API and registry/discovery responsibilities; its source project also documents the ServiceCenter implementation and interfaces. See the [Apache ServiceComb ServiceCenter documentation](#) and the [Apache ServiceComb ServiceCenter project](#).

QUESTION NO: 26

HUAWEI CLOUD AUDIT SERVICE CTS is a professional log audit service in the cloud security solution that collects, stores, and queries various cloud resource operation records

- A. Security analysis
- B. Compliance audit
- C. Problem location
- D. Intelligence collection

ANSWER: D

Explanation:

Intelligence collection is the correct selection because it is not a defined Cloud Trace Service (CTS) application scenario or core capability. CTS is Huawei Cloud's operation-audit service: it records management activity performed on cloud resources, retains those traces according to configured settings, and lets authorized users query and analyze them. This audit trail supplies reliable historical evidence about operations, such as who performed an action, when it was performed, and the relevant resource or API activity. Huawei documents CTS use in operational auditing contexts, including security analysis, compliance auditing, resource tracking, and problem locating. These uses depend on the trace records collected by CTS and are directly aligned with its purpose as a cloud-operation logging and audit service. "Intelligence collection" instead suggests gathering threat intelligence or external security intelligence feeds, which is not the function described for CTS. Therefore, in the context of CTS capabilities and application scenarios, intelligence collection is the applicable answer. See Huawei Cloud's [Cloud Trace Service product description](#) and its [CTS application scenarios documentation](#).

QUESTION NO: 27

Create an application in a CCE cluster that supports mounting multiple types of cloud storage volumes, including cloud storage types that support data sharing and multiple servers can mount the same file system: EVS disk storage volume, extreme file storage volume, object storage volume, and file volume storage. Judgment

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because shared, simultaneous mounting by multiple servers depends on the access mode and storage service, rather than being a common capability of every listed CCE volume type. In CCE, EVS disks are block-storage volumes and are normally used with the `ReadWriteOnce` access mode: the disk can be mounted read-write by one node at a time. Consequently, an EVS-backed volume cannot be treated as a shared file system that multiple servers mount concurrently for data sharing. Conversely, file services such as SFS and SFS Turbo are designed to provide shared file-system semantics and can be used where multiple workloads need concurrent access. CCE also supports several storage integrations, but their mounting behavior, Kubernetes access modes, and applicable workload patterns differ. A design for shared application data must therefore select a shared file-storage service and its appropriate access mode, rather than grouping EVS disk volumes together with shared file-system storage. Huawei's CCE storage documentation describes the supported storage choices and persistent-volume usage, while the EVS documentation describes EVS as elastic block storage and its attachment characteristics. See [Huawei Cloud CCE User Guide](#) and [Huawei Cloud EVS Product Description](#).

QUESTION NO: 28

What are the demands of insurance companies in the marketing process of HUAWEI CLOUD? Multiple selection

- A. Assist bancassurance business development
- B. Precision marketing
- C. Cross-selling
- D. Customized product development

ANSWER: A B C D

Explanation:

Insurance marketing modernization requires capabilities that improve both distribution effectiveness and the relevance of offers throughout the customer lifecycle. "Assist bancassurance business development" is a valid demand because insurers need digital support for bank-insurance channels, including coordinated customer engagement, sales enablement, and efficient channel operations. "Precision marketing" is equally central: cloud data and analytics capabilities help identify customer segments, understand needs, and deliver timely, relevant marketing activities. "Cross-selling" addresses the need to recognize opportunities to offer complementary insurance products to existing customers, increasing customer value while making interactions more relevant. "Customized product development" also belongs in the marketing-oriented demand set because product design must respond to segmented customer needs and market feedback; cloud-based data processing and agile digital capabilities support faster development and refinement of targeted products. Together, these capabilities form a connected marketing model: channel development reaches customers, precision insights identify needs, cross-selling expands suitable coverage, and customized products ensure the insurer has appropriate offerings. Huawei Cloud presents financial-services solutions around digital transformation, data-driven innovation, and agile service delivery; see [Huawei Cloud Financial Services](#) and the [Huawei Cloud Documentation](#).

QUESTION NO: 29

A user on a cloud needs to interact with resources outside the cloud through the public network, and the overall bandwidth is limited but there is an occasional meteor. From the perspective of cost savings, the user needs to choose to charge by bandwidth. Judgment

- A. True
- B. False

ANSWER: B

Explanation:

False is correct. The description indicates that outbound public-network demand is generally limited but has occasional sudden traffic spikes (the apparent meaning of “occasional meteor”). For this type of variable and bursty Internet egress pattern, charging by traffic is normally the cost-conscious billing model. It aligns charges with the volume of data actually transferred and avoids reserving and paying for a fixed bandwidth level that is needed only during infrequent peaks. Huawei Cloud Elastic IP billing supports both bandwidth-based and traffic-based approaches, allowing the billing method to be selected according to the workload’s traffic characteristics. Traffic billing is therefore appropriate where normal utilization is low or unpredictable and high usage is temporary; monitoring can then be used to assess actual consumption over time. Huawei’s Elastic IP pricing information describes the available billing dimensions and charging concepts: [Huawei Cloud Elastic IP Pricing](#). The Elastic IP product documentation also provides the service context for Internet-facing access and bandwidth configuration: [Huawei Cloud Elastic IP Product Description](#).

QUESTION NO: 30

The FPGA acceleration cloud server provides tools and environments for FPGA development and use, allowing users to easily develop FPGA accelerators and deploy services based on FPGA acceleration . Which of the following scenarios is suitable for using FPGAs to accelerate cloud servers?

- A. 3D personal somatosensory games, scenes that require a large number of graphics rendering
- B. Scenarios requiring massive computing and extremely low latency requirements in machine learning
- C. Scenarios requiring massive data processing such as gene sequencing and analysis
- D. Answers to scenarios where the financial industry has high requirements for computing power, ultra-low latency, high throughput, and response

ANSWER: B C D

Explanation:

FPGA-accelerated cloud servers are well suited to workloads whose algorithms can be expressed as dedicated, deeply parallel hardware pipelines and where predictable, very low processing latency is important. “Scenarios requiring massive computing and extremely low latency requirements in machine learning” is appropriate because FPGA logic can be customized for inference-oriented operations and data paths, reducing processing delay for suitable models. “Scenarios requiring massive data processing such as gene sequencing and analysis” is also appropriate: bioinformatics commonly includes repetitive, highly parallel operations that can be offloaded to an FPGA accelerator to improve throughput. “Answers to scenarios where the financial industry has high requirements for computing power, ultra-low latency, high throughput, and response” is a core FPGA use case as well. Financial workloads such as market-data processing and latency-sensitive analytics benefit from deterministic hardware processing and high-throughput network-adjacent acceleration. Huawei Cloud describes FPGA acceleration as providing programmable hardware acceleration for demanding computing tasks, while its Elastic Cloud Server portfolio provides the infrastructure on which these specialized flavors are delivered. See the [Huawei Cloud FPGA product page](#) and [Huawei Cloud Elastic Cloud Server overview](#).

QUESTION NO: 31

Which of the following is not an optimization practice for low-latency services? Single selection

- A. Use heterogeneous computing resources and use differentiated computing architectures for different types of services, such as GPU and FPGA computing resource pools
- B. B . Use intelligent DNS distribution deployment to shorten network latency, reduce data transmission, and improve efficiency
- C. C . Use relational database service RDS and use database management DAS to analyze SQL statements and obtain the corresponding slow SQL Statement Discipline Optimization
- D. D . Containerized applications are deployed in CCE clusters, and application performance management APM is used to analyze the call status of applications and optimize the architecture after understanding application bottlenecks
- E. Use the database intermediate DDM to split the database and table, reduce the whole to zero, and reduce the query delay of the huge amount of data in a single database

ANSWER: B

Explanation:

Use intelligent DNS distribution deployment to shorten network latency, reduce data transmission, and improve efficiency is the item that is not, by itself, a low-latency service optimization practice in this context. DNS is fundamentally a naming and endpoint-resolution mechanism: a client queries DNS to obtain an address before it establishes a connection to the application. Intelligent DNS policies can select an endpoint according to geography, carrier, or health status, so they can support traffic scheduling and improve the user's route to an appropriate service endpoint. However, DNS deployment does not optimize the application request path itself, database query execution, data partitioning, compute execution, or interservice calls. It also does not inherently reduce the amount of application data transmitted after a connection has been established. Therefore, describing DNS distribution alone as a mechanism that reduces data transmission and provides end-to-end low-latency optimization overstates its role. For low-latency workloads, DNS is an auxiliary traffic-entry capability, while latency optimization requires measures in the compute, application, data, and request-processing paths. Huawei Cloud DNS is described as a DNS and intelligent resolution service, which confirms its role in domain resolution and traffic direction rather than application-performance processing: [Huawei Cloud DNS](#). The DNS protocol's core purpose is likewise name-to-resource resolution, as specified in [RFC 1035](#).

QUESTION NO: 32

An enterprise stores customer records in an RDS database and requires centralized control and audit of administrator operations. Select four measures that should be implemented.

- A. Grant IAM permissions according to the least-privilege principle.
- B. Use Cloud Bastion Host to centrally manage and audit administrator access.
- C. Enable Database Audit for database access and SQL-operation records.
- D. Enable CTS to trace management operations on related cloud resources.
- E. Provide all administrators with one shared root account to simplify emergency operations.

ANSWER: A B C D

Explanation:

IAM should grant only the permissions required by each administrator, reducing the impact of compromised or inappropriate accounts. Cloud Bastion Host centralizes O&M access and records administrator sessions and commands. Database Audit records database access and SQL operations, which is required for accountability around sensitive records. CTS records management-plane operations on cloud resources and helps identify who changed a resource or configuration. A shared root account removes identity accountability and is unsuitable for controlled O&M.

QUESTION NO: 33

Which of the following items is optimized for data access in low-latency services? Multiple selection

- A. Analyze how business functions and users use system data to reasonably plan database partition tables
- B. Select appropriate databases and sharding algorithms for database sharding
- C. Infrequently accessed historical data classification (generally classified by date) is stored
- D. Use middleware to establish a sharding cluster
- E. Horizontal split of large tables according to field update frequency or field range frequency

ANSWER: A B C D

Explanation:

Low-latency data access depends on designing data distribution, storage placement, and request routing around actual business access patterns. “Analyze how business functions and users use system data to reasonably plan database partition tables” is correct because partition boundaries should reflect the most common query dimensions and traffic paths, allowing requests to reach a smaller, relevant data set. “Select appropriate databases and sharding algorithms for database sharding” is also correct: selecting a suitable database service and shard key/algorithm helps spread load, avoid hotspots, and support efficient targeted queries. “Infrequently accessed historical data classification (generally classified by date) is stored” is a valid hot/cold data practice. Separating date-based historical data from actively queried data reduces the working data volume and index pressure for latency-sensitive services. “D . Use middleware to establish a sharding cluster” is correct because sharding middleware can provide routing, distributed connection management, and transparent access to multiple database shards, supporting scalable low-latency access when properly designed. These measures align with Huawei Cloud relational-database and distributed-database service capabilities for scalable, highly available data workloads. See [Huawei Cloud RDS](#) and [Huawei Cloud DDS](#).

QUESTION NO: 34

Which of the following is the basic unit of the knowledge graph KG? Multiple selection

- A. Entity
- B. direction (Direction)
- C. relation (Relationship)
- D. network (Network)

ANSWER: A C

Explanation:

Knowledge graph content is constructed from entities and relations. An entity represents a uniquely identifiable real-world or conceptual object, such as a person, organization, device, product, or location. A relation expresses the meaningful semantic connection between entities, for example, an organization *employs* a person or a device *is located in* a site. Together, entities and relations form knowledge statements, commonly represented as subject–predicate–object triples. In graph implementations, entities are modeled as vertices and relations as edges; the relation may have a direction to preserve the meaning of the statement. Thus, direction is a characteristic used when representing a relationship rather than an independent knowledge unit, while a network is the larger graph structure composed from many entity-and-relation statements. Huawei Cloud Graph Engine Service describes graph data in terms of vertices and edges, which maps directly to entities and relationships in a knowledge graph. See the [Huawei Cloud Graph Engine Service product description](#). The RDF data model also defines graph statements through a subject, predicate, and object, reinforcing the entity-and-relation foundation of knowledge graphs: [W3C RDF 1.1 Concepts and Abstract Syntax](#).