AlibabaCloud.CAP-C01.v2026-11-23.q17

Exam Code:	CAP-C01
Exam Name:	Alibaba Cloud Certified Professional: Cloud Architect
Certification Provider:	Alibaba Cloud
Free Question Number:	17
Version:	v2026-11-23
# of views:	122
# of Questions views:	170
https://www.freeqas.com/qa/Alibaba-Cloud/CAP-C01/AlibabaCloud.CAP-C01.v2026-11-23.q17.html	

NEW QUESTION: 1
Clash for Victory is a popular online game that is hosted on Alibaba Cloud. As latency can have a huge impact on gameplay, the developers for the game want to implement a mechanism to route all requests to the closest access point, monitor the health of the game, and redirect traffic to healthy backend servers.
The game is deployed in multiple regions on Elastic Compute Service (ECS) instances that are part of Auto Scaling groups configured behind Application Load Balancers (ALBs).
As a Cloud Architect, what course of action would you take?
A. Create an Alibaba Cloud CDN distribution and specify the ALB as the origin server. Configure the cache behavior to use origin cache headers. Use Function Compute to optimize the traffic.
B. Configure an ApsaraDB RDS database to serve as the data store for the application. Create a Tair (Redis OSS-Compatible) cluster to act as the in-memory cache for hosting the application data.
C. Configure a Global Accelerator instance. In each region, add a listener for the port that the game listens on, and attach it to the endpoint group of that region. Add the ALB instances as the endpoint.
D. Create Alibaba Cloud CDN distribution and specify Object Storage Service (OSS) as the origin server.Configure the cache behavior to use origin cache headers.
Answer: C (LEAVE A REPLY)
Global Accelerator is designed for precisely this class of global interactive application. Clients connect through Alibaba Cloud ' s nearby acceleration access points, after which traffic traverses Alibaba Cloud ' s optimized global network toward the most appropriate backend region. This reduces Internet routing variability and improves end-to-end latency for latency-sensitive workloads such as gaming.

GA listeners define the application protocol and ports, while endpoint groups represent backend deployments in different regions. With intelligent routing, GA can select a nearby healthy endpoint group and forward requests accordingly. Alibaba Cloud also provides endpoint health checks; when an endpoint becomes unhealthy, new requests can automatically be directed to healthy endpoints.
CDN is highly effective for static cacheable content but is not the correct primary mechanism for arbitrary real-time game traffic and regional backend health-based routing. RDS and Tair improve application data performance but do not solve global network acceleration.
Therefore, GA in front of regional ALBs provides the correct architecture: nearest-entry acceleration, Alibaba Cloud backbone transport, regional endpoint groups, health monitoring, and failover.
Study Guide reference: Delivering Services and Content on Alibaba Cloud - Global Accelerator, ALB, multi-region traffic management, and high availability.

NEW QUESTION: 2
Jenny developed a healthcare analytics platform on Alibaba Cloud. To meet performance, disaster recovery, and compliance requirements, the platform requires continuous data replication across zones for fault tolerance and cost-efficient storage of historical patient data for regulatory retention. Jenny has recruited you to help build a solution to address these challenges.
Which of the following strategies would you propose? (Correct answers: 2)
A. Deploy DTS for live data migration and replication across zones for fault tolerance.
B. Deploy the databases on Elastic Compute Service (ECS) instances and use Cloud Backup to configure automatic backups.
C. Configure lifecycle rules in OSS for economically storing long-term backups.
D. Use DMS to manage all databases from a central location to cut down on management and operational costs.
Answer: A,C (LEAVE A REPLY)
Data Transmission Service and OSS lifecycle management address the two distinct requirements:
continuously replicated operational data and economical long-term retention.
DTS provides real-time data synchronization by continuously capturing changes from a source database and applying them to a destination. Alibaba Cloud identifies active geo-redundancy, disaster recovery, cross- region replication, and real-time data synchronization as standard DTS scenarios. Maintaining a secondary synchronized dataset in another failure domain therefore provides substantially better recovery characteristics than relying only on periodic backups.
Historical data has a different access pattern. Regulatory medical records can remain necessary for years while being rarely retrieved. OSS lifecycle policies can automatically transition aging objects from frequently accessed storage into lower-cost storage classes

such as Archive. Alibaba Cloud specifically describes lifecycle policies for medical records and other compliance datasets that become cold over time.
Option B provides backup protection but does not directly provide the requested continuous replication architecture and unnecessarily introduces self-managed databases on ECS. DMS simplifies database administration but is not a data-replication or archival-storage mechanism.
Consequently, DTS supplies the live replication/fault-tolerance layer, while OSS lifecycle policies supply automated and cost-efficient long-term retention.
Study Guide reference: Building Highly Available, Performant Cloud Architecture - DTS replication, disaster recovery, OSS lifecycle management, and compliance-oriented storage.

NEW QUESTION: 3
An educational technology company is migrating its learning management system (LMS) to Alibaba Cloud.
The company mainly serves enterprise-grade customers, hence their platform must support automated backups, high availability across multiple zones, and granular access control for sensitive student records.
You have been hired as the lead Cloud Architect to build their architecture on Alibaba Cloud.
Which of the following architectural components will fit their requirements? (Correct answers: 3)
A. Deploy PolarDB clusters across multiple availability zones for high availability.
B. Use Cloud Backup for automated backups of databases and application data.
C. Store static assets in unencrypted OSS buckets for cost efficiency.
D. Use Container Service for Kubernetes (ACK) with manual scaling.
E. Implement Resource Access Management (RAM) roles for granular access control.
Answer: A,B,E (LEAVE A REPLY)
Each selected option maps directly to one of the stated enterprise requirements.
A multi-zone PolarDB deployment protects the database against availability-zone-level infrastructure failure.
Alibaba Cloud ' s multi-AZ PolarDB architecture keeps replicated data in separate zones and can use standby resources for failover when the primary zone becomes unavailable, resulting in a higher database SLA than a single-zone deployment.
Cloud Backup provides centralized policy-driven backup and recovery for Alibaba Cloud and on-premises resources, including ECS data, databases, NAS, OSS, and other supported workloads. This addresses the requirement for automated application and database protection without building custom backup infrastructure.
RAM roles and associated policies provide fine-grained authorization. RAM policies define allowed actions, affected resources, and optional access conditions, enabling least-privilege access to sensitive student data and administrative services.

Leaving OSS data unencrypted conflicts with the handling of sensitive enterprise information, while manual ACK scaling does not contribute to the stated requirement for high availability or automated operations.
The resulting architecture therefore combines multi-AZ database resilience, managed backup protection, and fine-grained identity authorization.
Study Guide reference: Securing Workloads on Alibaba Cloud - RAM; Building Highly Available, Performant Cloud Architecture - PolarDB HA and Cloud Backup.

NEW QUESTION: 4
Clarence runs a technology services company that uses ApsaraDB RDS. The company wants to streamline their database management and set up a robust access control structure.
Which of the following solutions provides the required functionality with the LEAST operational overhead?
A. Use RAM policies to restrict access to databases and implement ActionTrail to audit actions performed on the database.
B. Implement access control via IP whitelist to control who has access to the database.
C. Use the permission management feature of Data Management Service (DMS) to control and audit user access to databases.
D. Set up a WAF instance in front of the database to block SQL injection attempts and ensure the security of the data.
Answer: C (LEAVE A REPLY)
Data Management Service is specifically designed to centralize database administration, authorization, and governance. Its permission model provides fine-grained control over instances, databases, tables, columns, rows, and sensitive data. Administrators can control logon, query, export, and change privileges while maintaining centralized governance rather than manually coordinating several independent services.
DMS is also appropriate where the organization needs auditable database access. Alibaba Cloud explicitly describes DMS authorization scenarios that include controlling RDS access and auditing employee data access.
RAM controls permissions to Alibaba Cloud resources and APIs, but combining RAM with ActionTrail would require multiple mechanisms and would not provide DMS ' s database-object-level permission model.
An RDS IP whitelist is useful as a network access restriction but cannot express granular user permissions or provide centralized database governance. WAF protects HTTP/HTTPS applications and is not the appropriate control plane for direct RDS database authorization.
DMS therefore provides the required access control and audit functionality with the lowest operational complexity by consolidating database security administration into a purpose-built managed service.

Study Guide reference: Securing Workloads on Alibaba Cloud - DMS access control, database governance, RDS security, and auditing.

NEW QUESTION: 5
Keanu uses Elastic Compute Service (ECS) instances to run batch jobs to process data at 00:00 AM and 04:00 AM every day. An Auto Scaling group manages the ECS instances, which run as pay-as-you-go instances. If a batch job fails on one instance, another instance will reprocess the job.
Which of the following solutions provides the MOST cost-effective way to run these jobs?
A. Create a new launch template for the Auto Scaling group. Configure the launch template to use higher specification instances. Set a policy to scale out based on CPU usage.
B. Purchase a 1-year Savings Plan for ECS that covers the instance family of the Auto Scaling group that the batch job uses.
C. Purchase a 1-year reserved instance for ECS that covers the instance family of the Auto Scaling group that the batch job uses.
D. Create a new launch template for the Auto Scaling group. Configure the launch template to use preemptible instances. Set a policy to scale out based on CPU usage.
Answer: D (LEAVE A REPLY)
The workload is an excellent candidate for preemptible, currently called spot, ECS capacity because it is short- lived, batch-oriented, horizontally scalable, and explicitly tolerant of interruption. The question states that if an instance fails, another instance can reprocess the job. That removes the main architectural objection to using interruptible compute.
Alibaba Cloud describes ECS spot instances as unused ECS capacity offered at substantially reduced prices compared with standard pay-as-you-go instances. The documented recommended scenarios include short- term tasks, data analytics, and applications that are highly fault tolerant and able to tolerate interruption.
Savings Plans and reserved-instance-style commitments can reduce costs for predictable long-term baseline consumption, but this workload runs only at two scheduled times each day. Purchasing committed capacity for a sporadic workload can leave paid capacity economically underutilized. Higher-specification pay-as-you- go instances also do not inherently minimize compute cost.
Using a launch template with preemptible/spot ECS instances allows Auto Scaling to add inexpensive compute when the batch workload requires it, while the application ' s retry mechanism handles interruptions.
Study Guide reference: Core Compute Infrastructure Deep Dive; Scalable Architecture on Alibaba Cloud; cost optimization for elastic compute.

NEW QUESTION: 6
Alisa ' s company recently migrated to Alibaba Cloud and wants to implement a solution to protect the traffic that flows in and out of the production VPC. Their previous on-premises

solution had an inspection server that performed specific operations such as traffic flow inspection and traffic filtering. The company wants to have the same functionalities on Alibaba Cloud.
Which of the following solutions will meet these requirements?
A. Set up Alibaba Cloud Log Service to collect traffic logs from the VPC to inspect traffic patterns.
B. Utilize Alibaba Cloud Traffic Mirroring to replicate traffic flow for analysis and filtering by other tools.
C. Deploy Alibaba Cloud Security Center for monitoring and inspecting VPC traffic.
D. Use Alibaba Cloud Cloud Firewall for direct traffic inspection and filtering within the production VPC.
Answer: D (LEAVE A REPLY)
Alibaba Cloud Cloud Firewall is the purpose-built managed network-security service for centralized traffic inspection, filtering, access control, and threat prevention. It protects multiple network boundaries, including traffic between cloud workloads and the Internet as well as east-west traffic involving VPC environments.
For Internet-facing assets, Cloud Firewall can inspect inbound and outbound traffic and apply access-control policies, DPI-based analysis, intrusion-prevention rules, and threat intelligence without requiring an organization to deploy and maintain a self-managed firewall appliance.
For private architectures, a VPC Firewall can inspect and control traffic traversing VPC boundaries through CEN or Express Connect. Access-control policies determine which traffic is permitted or blocked.
Simple Log Service can analyze network logs but does not itself enforce inline traffic filtering. Traffic Mirroring copies traffic to another inspection system and therefore requires additional tooling. Security Center focuses primarily on workload security posture, vulnerability management, malware protection, and host security rather than functioning as the VPC ' s inline network firewall.
Cloud Firewall therefore most closely reproduces the company ' s former inspection-and-filtering architecture as a fully managed cloud-native service.
Study Guide reference: Securing Workloads on Alibaba Cloud - Cloud Firewall, network inspection, access control, IPS, and VPC security.

NEW QUESTION: 7
Belinda is designing an API-driven cloud communications platform. The application is hosted on Elastic Compute Service (ECS) instances behind a Network Load Balancer (NLB). It leverages API Gateway to serve public-facing APIs to customers. For security reasons, Belinda wants to protect the platform against web exploits like SQL injection and large, sophisticated DDoS attacks.
Which combination of solutions provides the MOST protection? (Correct answers: 2)
A. Use Alibaba Cloud Web Application Firewall (WAF) to protect API Gateway.

B. Use Alibaba Cloud Web Application Firewall (WAF) to protect the NLB.
C. Use Alibaba Cloud Anti-DDoS Proxy to protect the NLB.
D. Use Alibaba Cloud Security Center with Anti-DDoS Basic.
Answer: A,C (LEAVE A REPLY)
The threats described exist at different layers, so the architecture should apply layered protection. WAF is the correct control for public HTTP/API traffic because it analyzes application-layer requests and blocks threats such as SQL injection, cross-site scripting, command injection, brute-force attacks, and other OWASP-style attacks. Alibaba Cloud explicitly supports integrating WAF with API Gateway and recommends disabling direct access paths afterward so clients cannot bypass WAF.
Large DDoS attacks require a dedicated volumetric mitigation service. Anti-DDoS Proxy scrubs malicious network and transport-layer traffic before clean traffic is forwarded toward the application infrastructure.
Alibaba Cloud ' s Anti-DDoS architecture supports protected services behind Server Load Balancer resources and provides port-forwarding mechanisms for non-HTTP workloads.
Alibaba Cloud specifically recommends combining Anti-DDoS and WAF when an application needs protection against both large volumetric attacks and sophisticated application-layer exploits.
WAF should protect the HTTP/API application boundary rather than being treated as a general Layer-4 NLB firewall. Security Center provides workload security and posture management, while basic DDoS protection alone is less suitable for the question ' s explicitly stated large and sophisticated attacks.
Study Guide reference: Securing Workloads on Alibaba Cloud - WAF, Anti-DDoS, API Gateway security, and defense-in-depth.

NEW QUESTION: 8
A financial services firm is deploying a real-time analytics platform on Alibaba Cloud to analyze sensitive customer transaction data. The platform must ensure data encryption at rest and in transit, centralized encryption key management, and compliance with internal security policies.
As a Cloud Architect, which Alibaba Cloud services would you recommend for the firm to utilize to meet these requirements? (Correct answers: 3)
A. Deploy Redis on ECS instances for session management.
B. Implement Alibaba Cloud Key Management Service (KMS) to manage encryption keys securely.
C. Configure Security Groups to restrict inbound traffic to ECS instances.
D. Use SSL/TLS certificates from SSL Certificates Service for encrypting data in transit.
E. Enable Transparent Data Encryption in ApsaraDB RDS for data at rest.
Answer: B,D,E (LEAVE A REPLY)
The three explicit requirements are centralized key management, encryption in transit, and encryption at rest.

KMS satisfies centralized cryptographic-key management. It provides secure key creation, storage, lifecycle management, access control, and integration with Alibaba Cloud services such as databases and storage.
Customer-managed software or hardware-protected keys can also be used where stricter compliance controls are required.
Certificate Management Service provides SSL/TLS certificates that establish encrypted HTTPS/TLS communication and verify server identity. TLS prevents sensitive transaction information from being transmitted as readable plaintext across network connections.
Transparent Data Encryption protects database data at rest. Alibaba Cloud RDS documentation identifies TDE as an encryption layer protecting database data files and backups without requiring applications to perform the underlying encryption themselves. RDS also separately supports SSL for database connection encryption.
Security Groups are an important network-access control, but the question requires exactly three controls specifically addressing encryption and centralized key management. Redis session management is unrelated.
The resulting security architecture therefore uses KMS for keys, TLS certificates for data in transit, and RDS TDE for persistent data at rest.
Study Guide reference: Securing Workloads on Alibaba Cloud - KMS, TLS/SSL, RDS encryption, TDE, and cryptographic key governance.

NEW QUESTION: 9
A startup is building a serverless backend for their mobile application on Alibaba Cloud. The backend includes APIs for user authentication, data storage, and real-time notifications. The startup expects variable traffic with unpredictable spikes and wants to minimize costs by paying only for actual usage.
Which combination of Alibaba Cloud services should the startup use to achieve this architecture? (Correct answers: 2)
A. Elastic Compute Service (ECS)
B. Function Compute
C. MaxCompute
D. Simple Message Queue (SMQ)
Answer: B,D (LEAVE A REPLY)
Function Compute and Simple Message Queue form an appropriate serverless, event-driven backend for bursty mobile workloads.
Function Compute is fully managed and automatically allocates elastic compute resources when functions are invoked. Alibaba Cloud bills Function Compute primarily according to actual configured resources and execution duration, allowing the startup to avoid maintaining continuously running application servers. This is particularly suitable for APIs and backend handlers experiencing unpredictable request spikes.
SMQ provides managed asynchronous messaging for decoupling backend components. It supports queue- based processing and publish-subscribe topics, making it suitable for

notification events, background jobs, workflow messages, and traffic buffering. Alibaba Cloud explicitly identifies load leveling during sudden traffic spikes as an SMQ use case. SMQ can also directly trigger Function Compute through queue or topic integrations, creating an event-driven design without dedicated message-consumer servers.
ECS requires persistent server resources and management, which conflicts with the requested serverless operating model. MaxCompute is a large-scale data warehouse and batch analytics platform rather than the application ' s transactional API or messaging layer.
Therefore, Function Compute supplies elastic execution while SMQ provides reliable asynchronous communication and notification-event handling.
Study Guide reference: Delivering Services and Content on Alibaba Cloud - Function Compute, serverless architectures, SMQ, event-driven integration, and elastic application services.

NEW QUESTION: 10
A company is deploying a mission-critical application on Alibaba Cloud ECS instances that require high availability and failover capabilities. The application must remain accessible even if an entire zone fails.
Which deployment strategy should the company use to ensure high availability?
A. Deploy all ECS instances in a single zone with a backup ECS instance.
B. Distribute ECS instances across multiple zones within a single region.
C. Use a single ECS instance with a high-performance EBS volume.
D. Deploy ECS instances in multiple regions for redundancy.
Answer: B (LEAVE A REPLY)
A multi-zone deployment within the same region is the standard architecture for surviving the complete failure of one availability zone. Alibaba Cloud zones within a region are physically isolated for fault tolerance while maintaining high-speed internal network connectivity. Alibaba Cloud explicitly recommends deploying resources across different zones when higher disaster-recovery capability is required.
For an application tier, multiple ECS instances should be placed in separate zones and normally fronted by a multi-zone ALB or another suitable load-balancing service. Health checks ensure that traffic is no longer directed to unavailable instances, allowing surviving instances in another zone to continue serving requests.
Alibaba Cloud ' s current high-availability guidance specifically describes multi-zone ECS and load-balancer deployments as the architecture for resilience against data-center or zone failures.
Option A remains vulnerable because all compute resources share the same zone failure domain. Option C retains a single compute point of failure regardless of storage performance. Option D can protect against an entire regional disaster, but it adds substantially greater networking, data synchronization, routing, and operational complexity than is necessary for the stated requirement of surviving one zone failure.

Therefore, distributing ECS instances across multiple zones within one region is the correct high-availability strategy.
Study Guide reference: Highly Available Architecture on Alibaba Cloud; Core Compute Infrastructure; multi- zone disaster recovery.

NEW QUESTION: 11
Kimiko works for a multinational retail chain that has just extended its internal business management services to Alibaba Cloud. They want to ensure the security and high-speed connectivity between its on-premises data center and multiple VPCs across different regions on Alibaba Cloud. The company wants to minimize network latency and ensure high availability without managing individual VPN connections for each VPC.
Which Alibaba Cloud service should the company use to achieve this goal?
A. Alibaba Cloud Express Connect with Transit Router
B. Apsara Stack (private cloud deployment)
C. Alibaba Cloud NAT Gateway
D. Alibaba Cloud VPN Gateway
Answer: A (LEAVE A REPLY)
Express Connect combined with Cloud Enterprise Network Transit Routers provides the appropriate enterprise hybrid-network architecture. Express Connect establishes a dedicated private circuit from the on- premises data center to Alibaba Cloud, providing lower latency, predictable performance, high bandwidth, and reduced exposure compared with Internet-based VPN connectivity.
The Express Connect circuit terminates through a Virtual Border Router. That VBR can be attached to CEN, where Transit Routers provide hub-and-spoke connectivity to multiple VPCs instead of requiring separate point-to-point connections for every VPC. Transit Routers can connect VPCs, VBRs, VPN attachments, and inter-region Transit Routers, enabling centralized routing across a geographically distributed enterprise network.
Alibaba Cloud also supports connecting an on-premises VBR and multiple VPCs to a Transit Router so that all attached networks communicate over private connectivity. Cross-region CEN connections then extend the architecture to VPCs in additional regions.
VPN Gateway relies on Internet-based encrypted tunnels and would require more tunnel management at scale.
NAT Gateway performs address translation and does not provide private hybrid interconnection. Apsara Stack is an on-premises cloud platform rather than the required connectivity service.
Study Guide reference: Building Enterprise-grade Networks on Alibaba Cloud - Express Connect, CEN, Transit Router, hybrid networking, and cross-region connectivity.

NEW QUESTION: 12
An online education company plans to migrate their services to Alibaba Cloud. They want to improve their video streaming quality during peak hours using Alibaba Cloud services.

As a Cloud Architect, you are tasked with building a solution to manage traffic spikes without compromising on video quality.
What combination of services would BEST fit their requirements? (Correct answers: 3)
A. Create Alibaba Cloud CDN distributions to cache video content at PoPs close to the audience.
B. Implement Tair (Redis OSS-Compatible) to reduce latency during user authentication.
C. Employ Auto Scaling to adjust server capacity dynamically based on user traffic.
D. Use Alibaba Cloud Global Accelerator for optimizing cross-region traffic paths.
Answer: A,C,D (LEAVE A REPLY)
CDN should form the content-delivery layer for cacheable video assets. Alibaba Cloud CDN distributes content to geographically dispersed edge nodes so viewers retrieve cached objects from nearby POPs rather than repeatedly connecting to a distant origin. Alibaba Cloud explicitly supports video-oriented CDN acceleration and edge delivery, reducing both end-user latency and origin bandwidth demand.
Auto Scaling addresses peak-time compute demand. It monitors workload conditions and automatically increases ECS capacity during traffic spikes, then scales resources back when traffic decreases. Alibaba Cloud specifically identifies streaming platforms with unpredictable demand as a typical Auto Scaling scenario.
Global Accelerator complements CDN and Auto Scaling where viewers or application endpoints are separated by long-distance or cross-border networks. GA routes traffic through Alibaba Cloud ' s optimized global BGP network, reducing latency, jitter, and packet loss. Alibaba documentation explicitly identifies video streaming as a latency-sensitive scenario that benefits from this architecture.
Tair can improve authentication or session performance, but it does not directly address video delivery quality or regional network performance. Therefore, CDN + Auto Scaling + Global Accelerator most directly satisfies the streaming requirement.
Study Guide reference: Delivering Services and Content on Alibaba Cloud - CDN, Global Accelerator, Auto Scaling, video delivery, and global traffic optimization.

NEW QUESTION: 13
Paul runs a successful online game. The game is deployed as a three-tier architecture in a VPC with ApsaraDB RDS for MySQL at the database layer. The game allows players to compete in real-time, and Paul wants to display a top-10 scoreboard in near-real time.
As a Cloud Architect, which of the following solutions would you propose for Paul ' s use case?
A. Place an Alibaba Cloud CDN distribution in front of the web application to cache the scoreboard in a section of the application.
B. Set up an Alibaba Cloud Tair (Redis OSS-Compatible) cluster to compute and cache the scores for the web application to display.
C. Set up an Alibaba Cloud ApsaraDB for PostgreSQL cluster to cache the scores for the web application to display.

D. Create a read replica in ApsaraDB for RDS MySQL to run queries to compute the scoreboard and serve the read traffic to the web application.

Answer: B (LEAVE A REPLY)

Tair is the appropriate architecture for a frequently changing game leaderboard because Redis-compatible sorted data structures are optimized for high-throughput score updates and ranking operations. Instead of repeatedly executing expensive aggregation and sorting queries against the relational database, the application can maintain player scores in Tair and retrieve the highest-ranked members directly.

Alibaba Cloud specifically documents leaderboard implementations based on Tair ' s sorted-set capabilities.

Tair ' s enhanced exZset functionality supports scalable ranking use cases and can distribute large leaderboards across multiple internal structures while preserving rank calculation.

This also protects RDS from high-frequency scoreboard queries. RDS remains the durable transactional database, while Tair operates as the low-latency serving and ranking layer for frequently accessed score data.

CDN caching is unsuitable because the leaderboard changes continually and requires near-real-time freshness.

A second relational PostgreSQL database would duplicate database infrastructure without providing the specialized in-memory ranking semantics required. An RDS read replica can reduce pressure on the primary database, but calculating and sorting rankings through SQL for every refresh remains less efficient than maintaining rankings directly in an in-memory sorted structure.

Study Guide reference: Building Highly Available, Performant Cloud Architecture - Tair, caching, low- latency architectures, and database performance optimization.

NEW QUESTION: 14

A retail chain is utilizing Alibaba Cloud ' s PolarDB to manage its inventory and sales databases. They need to ensure that their SQL queries have minimal impact on their production workload while enabling data analysts to run complex analytical queries.

How can the retail chain optimize their database operations for both transactional and analytical workloads?

(Correct answers: 2)

A. Utilize Elastic Compute Service (ECS) instances for running analytical queries externally.

B. Enable read/write splitting in PolarDB to redirect read queries to read-only nodes.

C. Implement ApsaraDB for MongoDB for handling complex queries separately.

D. Replicate data from PolarDB in real time to MaxCompute with Data Transmission Service (DTS), and perform analytics in MaxCompute.

Answer: (SHOW ANSWER)

Option B reduces interference between transaction processing and read-heavy workloads inside PolarDB.
PolarDB ' s read/write splitting automatically routes writes to the primary node while distributing eligible reads across read-only nodes. This allows reporting and ordinary SELECT traffic to consume separate compute capacity instead of competing with critical transactional writes on the primary node. Alibaba Cloud specifically identifies reporting, analytics, and read-heavy applications as appropriate read/write-splitting scenarios.
Option D provides stronger isolation for computationally expensive analytics. DTS can continuously synchronize PolarDB for MySQL data into MaxCompute using initial full synchronization followed by ongoing incremental change capture. Analysts can then execute large-scale analytical workloads against MaxCompute without placing those queries on the production transactional database.
Running queries on ECS does not inherently remove load from PolarDB because the SQL operations would still execute against the database. Introducing MongoDB merely to execute analytical queries creates a second database technology without an appropriate replication or analytical architecture.
Together, PolarDB read/write splitting optimizes operational read traffic while DTS-to-MaxCompute isolates heavyweight analytical processing, producing an effective HTAP-style separation of concerns.
Study Guide reference: Core Infrastructure Deep Dive - PolarDB, read/write splitting, DTS, MaxCompute, and analytical workload offloading.

NEW QUESTION: 15
A fintech startup is building a transaction monitoring application using Alibaba Cloud Container Service for Kubernetes (ACK). They are debating between adopting a dedicated cluster setup versus a serverless architecture to ensure scalability and cost efficiency. What factors should they consider to decide between a serverless or a dedicated ACK deployment? (Correct answers: 2)
A. A dedicated ACK cluster ensures lower latency as compared to a serverless setup due to exclusive resource allocation.
B. A dedicated ACK cluster is better suited if they need precise control over the infrastructure and networking configuration.
C. ACK Serverless offers better integration with Alibaba Cloud ' s network services.
D. ACK Serverless is ideal if they expect highly fluctuating traffic and prefer paying only for resources used.
Answer: (SHOW ANSWER)
The fundamental decision is the trade-off between infrastructure control and infrastructure abstraction.
A conventional managed/dedicated Kubernetes environment gives architects greater control over node pools, instance types, resource placement, operating-system configuration, network topology, specialized hardware, and related infrastructure choices.

This makes Option B appropriate when a regulated fintech application requires exact control over the compute and networking environment.
ACK Serverless takes the opposite approach. Alibaba Cloud manages the underlying node infrastructure, and workloads run without customers purchasing or maintaining worker nodes. Alibaba ' s documentation states that ACK Serverless uses pay-as-you-go resource consumption and can scale pods rapidly, making it particularly appropriate for variable or unpredictable workloads.
Alibaba Cloud also explicitly describes ACK Serverless as Kubernetes-compatible infrastructure where customers focus on application development while paying according to CPU and memory resources requested by workloads.
Option A is too absolute: dedicated Kubernetes infrastructure does not inherently guarantee lower application latency; latency depends on workload, topology, instance configuration, networking, and scaling behavior.
Option C is also incorrect because both managed and serverless ACK environments integrate deeply with Alibaba Cloud networking services.
Study Guide reference: Core Infrastructure Deep Dive - ACK deployment models, serverless containers, cluster control, elasticity, and cost optimization.

NEW QUESTION: 16
A global consultancy firm needs to ensure that their applications hosted on Alibaba Cloud have low latency for users distributed worldwide. They require a solution that can cache static content and dynamically route user requests to the nearest point-of-presence (PoP). Which Alibaba Cloud service should the firm implement to achieve this?
A. Alibaba Cloud CDN
B. Express Connect
C. Global Traffic Manager (GTM)
D. Server Load Balancer (SLB)
Answer: A (LEAVE A REPLY)
Alibaba Cloud CDN is specifically designed to accelerate content delivery by distributing and caching content at geographically distributed points of presence. When a user requests an accelerated resource, CDN DNS scheduling assigns an appropriate PoP. If the resource is cached there, the PoP returns it directly without requiring a round trip to the origin infrastructure. Alibaba Cloud documentation explicitly states that static resources are cached on the nearest PoP to reduce latency and improve access efficiency.
This directly satisfies both key requirements in the question: caching static content and directing users toward edge infrastructure close to their locations. If content is not cached or has expired, the CDN PoP performs an origin fetch and then serves the resource.
GTM can perform DNS-based geographical or latency routing among application endpoints and regions, but it is not itself a distributed static-content caching service. Express Connect provides private dedicated connectivity between enterprise/on-premises networks and

Alibaba Cloud. SLB distributes traffic among backend servers but does not provide a worldwide edge caching network.
Therefore, Alibaba Cloud CDN is the only listed service that simultaneously implements PoP-based edge delivery and static-content caching.
Study Guide reference: Delivering Content on Alibaba Cloud; content acceleration and global application- delivery architecture.

Valid CAP-C01 Dumps shared by PrepPdf.com for Helping Passing CAP-C01 Exam! PrepPdf.com now offer the **newest CAP-C01 exam dumps**, the PrepPdf.com CAP-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CAP-C01 dumps with Test Engine here: https://www.preppdf.com/Alibaba-Cloud/CAP-C01-prepaway-exam-dumps.html (**52** Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17
A company has a data ingestion workflow that consists of:
* An Alibaba Cloud Simple Message Queue topic for notifications about new data deliveries.
* An Alibaba Cloud Function Compute service to process the data and record metadata.
The company observes that the ingestion workflow fails occasionally because of network connectivity issues.
When such a failure occurs, the Function Compute instance does not ingest the corresponding data unless the company manually reruns the job.
Which combination of actions should a solutions architect take to ensure that the Function Compute instance ingests all data in the future? (Correct answers: 2)
A. Create a Simple Message Queue queue, and subscribe it to the existing Simple Message Queue topic.
B. Use Triggers to automatically invoke the Function Compute service.
C. Deploy the Function Compute instance in multiple Regions.
D. Increase the memory allocation for the Function Compute instance.
E. Modify the Function Compute service to read from the Simple Message Queue queue.
Answer: (SHOW ANSWER)
The core problem is durability during transient processing or network failures. A topic provides publish
/subscribe notification delivery, but placing an SMQ queue behind the topic creates a persistent buffer between the notification producer and the Function Compute consumer.
Alibaba Cloud supports subscribing a queue to an SMQ topic so that every message published to the topic is copied into the queue and retained for consumption.

Function Compute should then process records from the queue. Queue semantics provide the required failure recovery: after a consumer receives a message, it becomes temporarily invisible. The consumer deletes the message only after successful processing. If processing fails-for example because of a temporary connectivity problem-the message is not deleted and becomes visible again after the visibility timeout, allowing another processing attempt.
Function Compute also supports SMQ queue triggers with configurable retries, fault-tolerance behavior, and dead-letter handling, which can further strengthen this architecture.
Increasing memory does not fix network reliability, and multi-region deployment alone does not guarantee message persistence. A durable queue therefore prevents transient failures from silently losing ingestion work.
Study Guide reference: Core Infrastructure Deep Dive - Simple Message Queue, Function Compute, durable asynchronous processing, retries, and decoupled event-driven architecture.

Valid CAP-C01 Dumps shared by PrepPdf.com for Helping Passing CAP-C01 Exam! PrepPdf.com now offer the **newest CAP-C01 exam dumps**, the PrepPdf.com CAP-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CAP-C01 dumps with Test Engine here: https://www.preppdf.com/Alibaba-Cloud/CAP-C01-prepaway-exam-dumps.html (**52** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)