



Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.



CAP-C01 Practice Questions
CAP-C01 Practice Test
CAP-C01 Practice Exam
CAP-C01 Exam Questions
CAP-C01 Study Guide



killexams.com

Alibaba-Cloud

CAP-C01

Alibaba Cloud Certified Professional: Cloud Architect

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/CAP-C01>



Question: 1263

A content delivery network operator stores edge-cached metadata in OSS with ZRS in Singapore, CRR to Tokyo for failover. Server-side encryption with rotation keys fails sync due to key version desync. Which CRR key management setting, via console or PutBucketEncryption API, rotates and replicates KMS keys annually while ignoring /cache/temp/*?

- A. BucketEncryption: SSE-KMS with KeyRotation: true; CRR Rule Filter: Prefix /cache/temp/ exclude; ReplicateKeyVersions: All
- B. CRR Advanced: KMS Rotation Sync: Enabled; Annual Cycle; Exclude Prefix: /cache/temp/
- C. Policy: Replicate SSE-KMS; KeyUpdate: Auto; Ignore: temp prefix
- D. UpdateReplication: {"Encryption": {"Type": "SSE-KMS", "Rotation": "Yearly", "Filter": {"Prefix": "/cache/temp/ not"}}}

Answer: D

Explanation: Updating CRR with Encryption Type SSE-KMS, Rotation Yearly, and Prefix filter excluding /cache/temp/ ensures key versions sync across Singapore-Tokyo, supporting annual rotation for CDN metadata DR without temp data overhead. PutBucketEncryption complements this. Options A and D lack rotation specificity, B isn't API-based.

Question: 1264

Which encryption methods does Alibaba Cloud support to secure data stored in OSS?

- A. Client-side encryption before upload
- B. Server-side encryption with Alibaba Cloud KMS keys
- C. Use of third-party encryption services only

D. No encryption by default to save cost

Answer: A,B

Explanation: OSS supports server-side encryption managed by KMS and client-side encryption by users. Encryption is not disabled by default. Third-party encryption is an option but not mandatory.

Question: 1265

Which Alibaba Cloud ACK feature helps in optimizing the deployment of microservices by limiting the rate of requests handled by each microservice, thus preventing resource exhaustion?

- A.** Service mesh with Sentinel-based rate limiting
- B.** Kubernetes Horizontal Pod Autoscaler
- C.** Load balancer session persistence
- D.** Elastic Compute Service vertical scaling

Answer: A

Explanation: Alibaba Cloud service mesh includes Sentinel-based microservice rate limiting, which prevents overload by controlling the traffic to microservices. Horizontal Pod Autoscaler manages pod counts, but does not limit request rates per service. Load balancer session persistence and ECS vertical scaling do not provide rate limiting for microservices.

Question: 1266

Gaming studio's dev VPC uses PrivateZone for 'test.game.internal', caching external CDN queries. To mitigate dev DDoS tests (simulated 100 Gbps), use Anti-DDoS Origin with CreateDdosEvent for simulation, clear cache post-test,

and CLI for event creation. What simulation and clear?

A. No sim.

B. Simulate via CLI, set policy; call aliyun ddoso CreateDdosEvent --Resource 'dev-vpc-eip' --Protocol UDP --Pps 100000000 --Duration 300s; aliyun pvtz ClearCache --ZoneName 'test.game.internal' --Type A --RR 'cdn.*';

C. aliyun ddoso CreateDdosEvent --InstanceId 'origin-dev' --AttackType 'udp_flood' --Size 100G; aliyun pvtz FlushCache --ZoneId 'pz-test-001';

D. Basic event.

Answer: B

Explanation: CreateDdosEvent simulates 100 Gbps UDP for testing Origin policy. ClearCache post-test restores CDN cache. This validates dev resilience, with logs for tuning.

Question: 1267

During a cross-region failover test for ApsaraDB for Redis cluster (DRAM-based, 64 shards), a network partition isolates the primary proxy nodes, causing 30% slot unavailability. To achieve self-healing without manual intervention, what Tair-specific parameter in the instance config, combined with 'CLUSTER FAILOVER' command, restores quorum in under 60 seconds?

A. Set 'cluster-node-timeout=5000' and 'cluster-migration-barrier=2'; execute 'CLUSTER FAILOVER TAKEOVER' on replicas for majority vote

B. Adjust 'shard-proxy-count=8' and 'hotkey-cache.enable=true'; run 'CLUSTER SET-CONFIG-EPOCH' to increment epoch for failover

C. Enable 'tair.cluster.auto-healing=true' with 'proxy-timeout=10000'; use 'CLUSTER MEET ' to rejoin partitioned shards

D. Configure 'cluster-require-full-coverage=no' and 'replication-timeout=30000'; trigger 'CLUSTER FORGET ' post-partition

Answer: A

Explanation: In Tair's DRAM-based clusters, 'cluster-node-timeout=5000' detects partitions quickly, while 'cluster-migration-barrier=2' ensures replicas migrate slots only after majority quorum, preventing split-brain. The 'CLUSTER FAILOVER TAKEOVER' command on a replica promotes it if it holds >50% slots, restoring availability in <60s via gossip protocol. This self-healing leverages Tair's advanced proxy for load balancing, as documented in 2026 clustering updates, minimizing downtime in streaming scenarios without data loss, verified via 'CLUSTER INFO' metrics.

Question: 1268

OSS/ECS in hybrid via Cloud Enterprise Network encrypt with HSM-backed KMS. What kms HSM create, oss SSE-HSM, ecs disk, cen encrypt?

- A. OSS server-side only
- B. kms CreateHSMKey, oss SSE with HSM, ecs HSM key, cen encrypt
- C. No HSM
- D. Software keys only

Answer: B

Explanation: HSM: aliyun kms CreateKey --KeySpec "RSA_4096" --ProtectionMode HSM --KeyId "hsm-gov". OSS: aliyun oss PutBucketEncryption --SSEAlgorithm "aws:kms" --KMSKeyId "hsm-gov". ECS: aliyun ecs CreateDisk --KMSKeyId "hsm-gov" --Encrypted true. CEN: aliyun cen Update --EncryptTransit true. HSM FIPS 140-2 ensures quantum-resistant encryption.

Question: 1269

A multinational e-commerce company operates ECS instances across three VPCs

in the China (Beijing), China (Shanghai), and Singapore regions, each with CIDR blocks of 10.0.0.0/16, 172.16.0.0/16, and 192.168.0.0/16 respectively. To enable low-latency communication between these VPCs while integrating an on-premises data center in Tokyo with CIDR 10.1.0.0/16 via a VPN Gateway, the architect must configure CEN. During setup, the CEN instance reports an error due to route advertisement conflicts. What is the most effective resolution to ensure transitive routing without manual route table modifications?

- A. Attach all VPCs and the VPN Gateway to a single transit router in CEN, then enable automatic route propagation for all attachments while verifying non-overlapping CIDR blocks across regions.
- B. Deploy additional vSwitches in each VPC with secondary CIDR blocks and associate them with the VPN Gateway for isolated traffic flows.
- C. Manually add static routes in each VPC's route table pointing to the VPN Gateway's EIP, bypassing CEN's dynamic routing features.
- D. Create separate CEN instances per region and use Express Connect circuits to link them, propagating VPN routes via BGP dynamic routing.

Answer: A

Explanation: In Alibaba Cloud CEN, attaching multiple VPCs from different regions to a single transit router enables full-mesh connectivity with automatic route propagation, ensuring low-latency transitive routing across global networks. The error likely stems from overlapping CIDR blocks, which CEN rejects; verifying and adjusting non-overlapping ranges (e.g., ensuring Tokyo's 10.1.0.0/16 does not conflict with Beijing's 10.0.0.0/16) resolves this. Enabling automatic propagation advertises all attached network routes dynamically via BGP, eliminating manual configurations and supporting hybrid scenarios with VPN Gateways, as per CEN's latest multi-region hybrid cloud best practices for scalable enterprise networks.

Question: 1270

A company is integrating Alibaba Cloud with its existing identity provider (IdP) for single sign-on (SSO). What is the first step they should take to establish this integration?

- A. Enable MFA for all users
- B. Create a service-linked role in Alibaba Cloud
- C. Configure the IdP to trust Alibaba Cloud as a service provider
- D. Create a new RAM user in Alibaba Cloud

Answer: C

Explanation: The first step in integrating Alibaba Cloud with an existing identity provider for single sign-on is to configure the IdP to trust Alibaba Cloud as a service provider. This trust relationship is essential for enabling SSO functionality and ensuring secure access to resources.

Question: 1271

Which Alibaba Cloud feature allows the compliance team to respond quickly to suspicious activity detected in audit logs?

- A. Enable ECS security group logging with manual inspection
- B. Deploy static RAM policies with no alerting
- C. Use OSS bucket access logs with scheduled audits
- D. Integration of ActionTrail with Log Service analytics and CloudMonitor alerts

Answer: D

Explanation: ActionTrail with Log Service provides real-time analytics, and CloudMonitor can generate alerts on suspicious events, enabling rapid response to compliance or security incidents.

Question: 1272

You have created a custom metric in CloudMonitor to track the number of user logins to your application. What is the best practice for ensuring the accuracy of this metric?

- A. Push the login data to the metric in real-time.
- B. Disable the metric after initial setup to save resources.
- C. Manually update the metric every week.
- D. Use a third-party tool to track logins instead.

Answer: A

Explanation: The best practice for ensuring the accuracy of a custom metric tracking user logins in CloudMonitor is to push the login data to the metric in real-time. This approach ensures that the metric reflects the most current data, allowing for accurate monitoring and analysis.

Question: 1273

For secure code deployment, a DevSecOps team requires RAM roles for CodePipeline to assume "codeup:MergeRequest" with conditions on branch "main-2026" and MFA, plus tag inheritance for audit. What session tag condition propagates tags?

- A. Deny policy: "Condition": {"Null": {"mfa:Present": "true"}}
- B. Permission policy: "Effect": "Allow", "Condition": {"ExternalId": "BranchMain"}
- C. Trust policy: "Condition": {"ForAnyValue:StringEquals": {"codeup:Tag/Branch": "main-2026"}} without tags.
- D. In assume-role, use --tags Key=AuditTrail,Value=Deploy2026; session policy: "Condition": {"StringEquals": {"sts:TransitiveTagKeys": "AuditTrail"}}

Answer: D

Explanation: To propagate audit tags in RAM assumed roles for CodePipeline, pass --tags during "sts:AssumeRole" (e.g., Key=AuditTrail,Value=Deploy2026), and include in session policy "StringEquals": {"sts:TransitiveTagKeys": "AuditTrail"} to allow inheritance, combined with branch condition "ForAnyValue:StringEquals": {"codeup:Tag/Branch": "main-2026"} and MFA. This ensures traceable deployments. Option B misses tags; C uses external ID wrongly; D inverts MFA.

Question: 1274

A network engineer wants to monitor traffic performance between VPCs connected through CEN. What Alibaba Cloud service should be used?

- A. Security group inbound logs
- B. CloudMonitor with CEN flow logs enabled
- C. VPN Gateway log analysis
- D. NAT Gateway traffic monitoring

Answer: B

Explanation: CloudMonitor integrated with CEN flow logs can track traffic flows, performance, and detect anomalies between VPCs connected via CEN. VPN Gateway logs focus on VPN traffic. NAT Gateway and security group logs monitor different scopes.

Question: 1275

A development team is using ACK to manage their containerized applications. They want to ensure efficient resource utilization. What practices should they

adopt?

- A. Implement Horizontal Pod Autoscaler (HPA)
- B. Use a single node pool for all workloads
- C. Regularly monitor resource usage with CloudMonitor
- D. Define resource requests and limits for each container

Answer: A,C,D

Explanation: Defining resource requests and limits ensures that containers receive the resources they need without over-provisioning. HPA allows for dynamic scaling based on demand, and regular monitoring with CloudMonitor helps identify inefficiencies. Using a single node pool can lead to resource contention and should be avoided.

Question: 1276

How does Alibaba Cloud's auto-scaling improve cost efficiency for applications behind a CDN?

- A. By disabling scaling during peak hours.
- B. By launching all instances at once regardless of demand.
- C. By relying on manual intervention for scaling decisions.
- D. By scaling compute resources automatically based on traffic patterns, reducing over-provisioning.

Answer: D

Explanation: Auto-scaling adjusts resources dynamically according to traffic, optimizing cost by avoiding idle resources while maintaining performance.

Question: 1277

A retail company is using Alibaba Cloud for its e-commerce platform. They want to ensure that their data is protected against accidental deletions. What strategy should they implement?

- A. Rely on manual backups only
- B. Schedule backups once a month
- C. Enable versioning in OSS
- D. Use local storage for data

Answer: C

Explanation: Enabling versioning in OSS protects against accidental deletions by retaining previous versions of objects. This feature allows for easy recovery of data that may have been mistakenly deleted.

Question: 1278

Which features are essential when configuring Alibaba Cloud Function Compute for edge computing with CDN?

- A. Exclusive use of static IP addresses for function calls
- B. Integration with Anti-DDoS Pro for function security
- C. Ability to write serverless code executed at CDN edge nodes
- D. Use of RAM roles to limit access to backend resources

Answer: C,D

Explanation: Function Compute enables serverless code at CDN edges to reduce latency. RAM roles provide granular access control. Static IPs are not mandatory, and Anti-DDoS Pro is a network-level defense.

Question: 1279

An IoT platform balances MQTT over TCP at Layer 4 using CLB, with 10,000 devices connecting via UDP fallback. Backend ECS show 15% health check failures from port 1883 overload, and client IPs are lost in logs. Configuration lacks Proxy and uses round-robin. To support 50,000 connections with IP preservation and UDP/TCP hybrid, what SLB migration and tuning is advised?

- A. Upgrade CLB to high-spec, add Tengine for MQTT proxying, disable Proxy, and weighted round-robin.
- B. Retain CLB, enable QUIC fallback, uniform scheduling without IP preservation.
- C. Switch to ALB for Layer 7 MQTT, use X-Forwarded-For, fixed port checks every 10s.
- D. Migrate to NLB, enable Proxy protocol for TCP/UDP, configure least connections with hybrid listeners, and 2s health checks on port 1883.

Answer: D

Explanation: NLB's Layer 4 prowess handles 50,000 IoT connections with native TCP/UDP support via hybrid listeners, using Proxy protocol to preserve device IPs for auditing in logs, addressing the loss issue. Least connections scheduling evens load on port 1883, reducing 15% failures, with 2s checks for proactive removal. This leverages 2026 NLB auto-scaling for IoT surges, surpassing CLB's LVS limits (max ~10k connections) and ALB's Layer 7 overhead unsuitable for raw MQTT/UDP.

Question: 1280

You are tasked with aggregating logs from multiple Alibaba Cloud services for centralized analysis. Which of the following configurations would best utilize Simple Log Service (SLS) for this requirement?

- A. Set up SLS to collect logs from each service separately and analyze them individually.
- B. Use SLS to stream logs to a third-party analysis tool directly.
- C. Configure log collection from multiple services into a single SLS project for unified access.
- D. Disable log collection for all services to reduce costs.

Answer: C

Explanation: Simple Log Service (SLS) is designed to aggregate logs from various Alibaba Cloud services into a single project. This configuration allows for centralized access and analysis of logs, making it easier to identify trends, troubleshoot issues, and maintain compliance across services.

Question: 1281

A content platform with 2 PB OSS data in us-east-1 needs DR to ap-northeast-1 with RPO=0, integrating CMS alerts for sync failures. Which monitoring and DR setup?

- A. Set HBR for OSS snapshots with geo-copy to ap-northeast-1.
- B. Enable CMS metric alerts on OSS CRR lag >30s, dashboarding sync progress.
- C. Configure OSS CRR with same-region replication first, then cross with versioning.
- D. Use RDS active-active geo with multi-write for content metadata, alerted via CMS events.
- E. Integrate SLS with CMS for log-based alerts on replication errors.

Answer: B,E

Explanation: CMS alerts on CRR lag metrics provide real-time visibility, with dashboards tracking progress for zero RPO enforcement. SLS integration correlates replication logs with CMS for detailed error alerting. RDS active-active

suits writes but content is OSS; same-region CRR is intra; HBR snapshots are periodic, not continuous.

Question: 1282

When configuring Alibaba Cloud ActionTrail for compliance, what setting ensures that all API call records are available nearly in real-time for security monitoring?

- A. Disable detailed data events and monitor management events only
- B. Enable integration of ActionTrail logs with Log Service streaming alongside CloudMonitor alert rules
- C. Rely only on ECS native logging with batch upload
- D. Use delayed OSS bucket storage without real-time analysis

Answer: B

Explanation: Integration of ActionTrail with Log Service streaming enables near real-time processing of API call logs, and CloudMonitor alerts enable immediate attention to security or compliance events, essential for proactive monitoring.

Question: 1283

A gaming studio migrates Unity servers (100 VMs, Ubuntu, 50 TB) from AWS to ECS SMC quick, public multi-thread16, incremental, ESS with hook for Unity build verify, scale on player count >10k. Which quick thread and player rule?

- A. MigQuick --16thread AWS --50TB; Hook verify; PlayerRule>10k add3
- B. StartQuick --Thread16 Public Inc50TB; Lifecycle buildtest; Scale >10k
- C. Quick --AWS U16t Pub50TB Inc; Hook UnityVerify; Rule Players10k +4
- D. QuickMigration --AWS --Ubuntu --Threads16 --Public --50TB --Incremental;

```
CreateLifecycleHook --VerifyUnityBuild --ScalingGroupId asg-game;  
CreateScalingRule --PlayerCount>10000 add5
```

Answer: D

Explanation: Quick for AWS Ubuntu, 16 threads public incremental 50 TB. Hook verifies Unity build, rule adds 5 on >10k players.

Question: 1284

A company is concerned about unauthorized access to its resources and wants to implement a policy that automatically revokes access after a period of inactivity. What should they do?

- A. Manually review user activity monthly
- B. Implement a policy with session timeout settings
- C. Disable inactive accounts immediately
- D. Use a third-party tool to monitor user activity

Answer: B

Explanation: Implementing a policy with session timeout settings allows the company to automatically revoke access after a period of inactivity. This proactive approach enhances security by minimizing the risk of unauthorized access due to forgotten or unattended sessions.

Question: 1285

An insurance under Solvency II is using Alibaba Cloud's eu-central-1 for risk modeling on PAI. Audits model inputs. 2026 ESG factors. Which?

- A. Single.

- B. PAI-trails.
- C. Batch.
- D. Regional ActionTrail for PAI, SLS with ESG-tagged queries; KMS for inputs.

Answer: D

Explanation: Solvency II audits models. ActionTrail captures, SLS tags ESG, KMS secures.

Question: 1286

You are implementing a data migration strategy from RDS for MySQL to ApsaraDB for MongoDB. Which method would be most efficient for handling schema differences?

- A. Migrate data in JSON format directly
- B. Utilize a schema conversion tool before migration
- C. Manually adjust the schema in MongoDB after migration
- D. Use a direct data migration tool without adjustment

Answer: B

Explanation: Utilizing a schema conversion tool before migration is the most efficient method for handling schema differences between RDS for MySQL and ApsaraDB for MongoDB. This tool can help automate the conversion process, ensuring that the data is structured correctly in the target database and minimizing manual adjustments.

Question: 1287

In an ACK dedicated cluster (discontinued post-Aug 2024, but legacy), with 3 master nodes for HA, you integrate ASM for multi-cluster traffic (m1c1, m1c2 in

same VPC). To expose a unified ingress for bookinfo app across clusters using serverless gateway, while enforcing JWT auth on /productpage with custom ext-authz gRPC service, what IstioGateway CRD and AuthorizationPolicy ensures secure cross-cluster access?

A. apiVersion: istio.alibabacloud.com/v1beta1 kind: IstioGateway metadata: name: unified-ingress spec: gatewayType: ingress ports: - name: https port: 443; Apply policy: apiVersion: security.istio.io/v1beta1 kind: AuthorizationPolicy metadata: name: jwt-auth spec: rules: - from: - source: requestPrincipals: ["*"] to: - operation: paths: ["/productpage"] when: - key: request.auth.claims[iss] values: ["https://example.com"]

B. In ASM console: Create Gateway > Serverless > Add clusters m1c1,m1c2; VirtualService: hosts: productpage, http: - match: uri.prefix: /productpage route: destination: host: bookinfo, port:9080; Custom authz: deploy ext-authz grpc:9000, policy: selector: matchLabels: app: bookinfo, action: CUSTOM, provider: ext-authz

C. kubectl apply -f gateway.yaml: spec: selector: istio: ingressgateway; servers: - port: number:443 tls: https_redirect:true; Policy: kind: RequestAuthentication spec: selector: matchLabels: app: productpage jwtRules: - issuer: "test-issuer" jwksUri: "https://raw.githubusercontent.com/istio/istio/release-1.18/security/tools/jwt/samples/jwks.json"

D. CRD: kind: Gateway spec: - port:443 protocol:HTTPS; Attach to ASM instance mesh1; Auth: PeerAuthentication STRICT; AuthorizationPolicy: action: ALLOW, rules: - to: operation: METHODS: GET paths: /productpage when: request.headers[authorization] startsWith Bearer

Answer: B

Explanation: ASM's serverless ingress gateway provides a unified entry for multi-cluster (m1c1, m1c2) bookinfo exposure via console creation, routing /productpage to port 9080. Deploying ext-authz as gRPC on 9000 enables custom JWT validation, referenced in AuthorizationPolicy with CUSTOM action for granular path enforcement, ensuring secure cross-VPC traffic without redirect configs or peer auth that doesn't handle JWT claims. This leverages ASM v1.18+

for managed HA, avoiding deprecated dedicated cluster pitfalls.

Question: 1288

Consider a VPC with overlapping CIDR blocks between the VPC (192.168.1.0/24) and a connected on-premises network (192.168.1.0/24) via VPN Gateway. What is the recommended solution to enable successful hybrid communication?

- A. Use NAT over the VPN Gateway to translate IP addresses dynamically
- B. Change the VPC CIDR block to a non-overlapping range and update route tables
- C. Disable VPN Gateway and switch to Express Connect
- D. Allow overlapping CIDR blocks and rely on security groups

Answer: B

Explanation: Overlapping CIDR conflicts cause routing issues. Changing the VPC CIDR to a unique, non-overlapping address space is the proper solution to avoid routing conflicts. NAT over VPN (B) is complex and not standard practice in Alibaba Cloud hybrid architectures. Express Connect (C) does not solve CIDR conflicts. Security groups cannot resolve IP addressing conflicts (D).

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllabus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.