



*Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.*



SAE-C01 MCQs  
SAE-C01 TestPrep  
SAE-C01 Study Guide  
SAE-C01 Practice Test  
SAE-C01 Exam Questions



[killexams.com](http://killexams.com)

**Alibaba-Cloud**

# SAE-C01

*Alibaba Cloud Certified Expert: Solutions Architect*

ORDER FULL VERSION

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



**Question: 1592**

A global e-commerce platform experiences unpredictable spikes in write and read requests with massive amounts of semi-structured user data, requiring smooth handling of 10k+ operations per second without full table scans. Which feature of Table Store architecture optimally supports this workload?

- A. Multi-master replication
- B. Vertical scaling of storage nodes
- C. Single-table partitioning
- D. ACID-compliant transactions

Answer: C

Explanation: Table Store utilizes single-table partitioning to shard data across multiple nodes, allowing it to handle massive structured data and 10k+ operations per second efficiently. Vertical scaling or multi-master replication is less relevant here, while ACID transactions are limited in Table Store as a NoSQL database.

**Question: 1593**

A genomics research consortium is migrating a 5 PB exome sequencing dataset to Alibaba Cloud for real-time variant calling using a hybrid Spark-on-Kubernetes workload, where OSS-HPC serves as the intermediate scratch space for SAM/BAM file transformations. The architecture must handle 10,000 concurrent reads from E-HPC clusters with sub-10ms latency for random 4KB metadata accesses, while ensuring fault-tolerant striping of sequence alignments to underlying OSS buckets for long-term petabyte-scale data lake persistence. During peak variant discovery phases, the system experiences metadata contention leading to 20% throughput degradation. Which OSS-HPC capabilities and integrations would resolve this contention and enhance overall pipeline resilience?

- A. Implement OSS-HPC's dynamic scaling of metadata servers to auto-provision up to 1,000 MDS nodes based on IOPS demand, restoring throughput to 95% of baseline during contention spikes
- B. Utilize OSS-HPC's integration with Alibaba Cloud's Data Lake Analytics to offload

metadata indexing to serverless queries, reducing MDS load by 70% for BAM header lookups

C. Enable OSS-HPC's erasure coding with N+3 redundancy on OSS-backed stripes, ensuring zero data loss during 48-hour transformation jobs even with 3 concurrent node failures

D. Configure OSS-HPC with hybrid caching layers using local NVMe SSDs on ECS nodes, prefetching alignment metadata to achieve consistent sub-5ms access times across the Kubernetes pod fleet

Answer: A, C, D

Explanation: OSS-HPC addresses metadata contention in high-concurrency genomics workloads by dynamically scaling metadata servers (MDS) up to 1,000 instances based on real-time IOPS metrics, which intelligently redistributes load and recovers throughput to near-baseline levels without manual intervention. For resilience, its erasure coding integration with OSS applies N+3 parity across striped files, tolerating up to three node failures mid-job while maintaining data integrity for critical BAM transformations in petabyte pipelines. Hybrid caching in OSS-HPC leverages NVMe SSDs on ECS instances for metadata prefetching, delivering sub-5ms latencies for random accesses in distributed Spark jobs on Kubernetes. Offloading to Data Lake Analytics is feasible for analytics but does not directly mitigate OSS-HPC's internal MDS contention, as it operates outside the parallel file system's core metadata plane.

### Question: 1594

After enabling Prometheus monitoring with ARMS, why might some metrics not be visible in dashboards?

- A. CS container policies blocked metric endpoints
- B. Jenkins pipeline failed to deploy applications
- C. Incorrect Prometheus scraping configuration or missing metrics exposition
- D. Falco detection disabled metrics

Answer: C

Explanation: Prometheus requires correct scraping targets and that services expose metrics endpoints properly. Jenkins pipelines don't affect metrics visibility after deployment. CS and Falco focus on security, not metrics exposure.

### Question: 1595

A SaaS CRM provider scales customer data pipelines on Alibaba Cloud, blending pay-as-you-go ECS for ETL jobs with preemptible for data cleansing bursts, amid a TCO shift from 250 on-premises servers (4-year depreciation at 500,000 USD total, 0.16 USD/kWh power). The architect must layer discounts for 25% YoY growth. Which?

- A. TCO via Pricing Calculator projecting server refresh cycles against ECS with 3-year Savings Plans at 70% discount for ETL.
- B. Preemptible cleansing at 75% bid, with Data Lake Analytics for burst orchestration.
- C. 1-year Reserved for ETL, underprojecting growth TCO.
- D. Savings Plans 80% commitment + preemptible, TCO growth-scaled.

Answer: A, B, D

Explanation: Pricing Calculator TCO models on-premises (250 servers 125,000 USD/year + power) vs. ECS, 3-year Savings Plans 70% for ETL yielding 48% savings with growth. Preemptible 75% for bursts, Data Lake Analytics handles. Savings 80% + preemptible, TCO YoY scaled, over Reserved.

### Question: 1596

During a penetration test, an application behind Alibaba Cloud WAF was found vulnerable to a new zero-day OWASP attack vector not covered in existing managed rulesets. What is the best course of action?

- A. Disable WAF to avoid false positives
- B. Create a custom WAF rule to detect and block the zero-day attack pattern immediately
- C. Wait for Alibaba Cloud to update managed rulesets
- D. Handle traffic filtering manually

Answer: B

Explanation: Creating a custom WAF rule allows immediate mitigation of the new threat, protecting the application proactively. Disabling WAF exposes the app, waiting for vendor updates delays mitigation, and manual filtering is impractical at scale.

### Question: 1597

A multinational e-commerce platform experiences a surge in user-generated content during peak shopping seasons, resulting in 5 million daily writes of semi-structured product reviews, each with variable attributes like timestamps, ratings, and embedded

images metadata. The platform requires a solution to handle this at sub-10ms latency while integrating with real-time analytics pipelines for sentiment analysis. Existing infrastructure uses ECS instances with self-managed Cassandra, but hotspots cause 20% query failures. Which approaches should the solutions architect recommend to migrate to Alibaba Cloud while ensuring 99.999999999% data durability and seamless scaling to 10 PB without downtime?

- A. Configure Table Store with a composite partition key using user ID and timestamp hash to distribute hotspots, enabling single-table partitioning for massive structured data ingestion at 10K+ ops/sec per partition.
- B. Deploy PolarDB for MySQL in a multi-master cluster edition with shared storage, leveraging log-based replication for active geo-redundancy across three AZs to achieve zero RPO during regional failures.
- C. Utilize Table Store's streaming read interfaces combined with DataHub for incremental data tunnels to MaxCompute, allowing SQL-based real-time queries on 1 billion+ records without pre-defining schemas.
- D. Implement PolarDB-X for distributed scaling with Paxos protocol multi-replica consistency, partitioning review data across petabyte-scale storage while supporting HTAP for concurrent OLTP writes and OLAP analytics.

Answer: A, C

Explanation: For handling massive semi-structured data with variable attributes and high write throughput, Table Store's wide-column NoSQL model excels due to its schema-free design, sparse tables allowing differing column counts per row, and automated partitioning via composite keys to mitigate hotspots—achieving 10K+ ops/sec per partition and scaling to 10 PB per table with three-replica durability ensuring 99.999999999% reliability through Parallel-Raft consensus. Integrating with DataHub and MaxCompute via streaming reads enables efficient incremental synchronization for real-time analytics like sentiment processing on trillion-scale records, avoiding the need for schema pre-definition and reducing migration complexity from Cassandra. PolarDB options, while robust for relational workloads, introduce unnecessary overhead for semi-structured NoSQL ingestion, as their shared storage and log replication are optimized for ACID transactions rather than the platform's variable-attribute, high-velocity writes.

### **Question: 1598**

A security team wants to enforce runtime protection of containers on ACK (Alibaba Cloud Container Service). Which CS feature leverages Falco rules to detect suspicious system calls and behaviors?

- A. Vulnerability scanning of container images
- B. Runtime security monitoring
- C. Jenkins pipeline enforcement
- D. Prometheus metric security

Answer: B

Explanation: CS runtime security integrates Falco rules to monitor system calls and container behaviors in real time to detect breaches or anomalous activity. Image vulnerability scanning checks images before deployment, Jenkins pipelines manage CI/CD, and Prometheus metrics focus on monitoring, not runtime security.

**Question: 1599**

An enterprise experiences increased latency when encrypting data at the application layer using client-side encryption before uploading to OSS. What is a possible alternative to reduce encryption overhead while maintaining strong security?

- A. Use symmetric encryption with static keys
- B. Remove encryption entirely
- C. Use OSS SSE-KMS encryption
- D. Switch to TDE for object encryption

Answer: C

Explanation: OSS SSE-KMS offloads encryption to server side with managed keys, reducing client overhead while maintaining strong security guarantees. TDE applies only to databases, and removing encryption reduces security.

**Question: 1600**

A e-commerce peak hybrid, centers in Seoul and Dubai to Alibaba Cloud Korea/ME via CEN. ECMP 60 Gbps, flow anomalies, AI latency, CCN partners. Which CEN for peak e-com?

- A. ECMP 60 Gbps peaks.
- B. Flow anomalies, CCN.
- C. AI latency.
- D. No ECMP.

Answer: A, B, C

Explanation: ECMP for loads, flow/CCN detect/partner, AI fast. No ECMP bottlenecks.

**Question: 1601**

In a healthcare telemedicine app on Alibaba Cloud integrating with on-premises EMR systems, threats involve ransomware via web tampering on patient portals, SQLi in query endpoints, and SYN floods on ALB during virtual consults. Security Center's anti-ransomware must snapshot recoveries, WAF custom policies for healthcare schema compliance, and Anti-DDoS exception thresholds to avoid blackholing consult streams. Which scenario-specific configurations across services ensure HIPAA-aligned protection?

- A. Security Center Snapshot Recovery for tamper-proof patient data restoration
- B. WAF Healthcare Schema Policies for SQLi mitigation in EMR queries
- C. Anti-DDoS Exception Thresholds for SYN flood tolerance on ALB
- D. CTDR Cross-Service Alerting for integrated ransomware and DDoS tracing

Answer: A, B, C, D

Explanation: Security Center Snapshot Recovery provides automated backups and point-in-time restores for tampered portal files, integrated with web tamper proofing to alert and revert ransomware changes, supporting HIPAA data integrity. WAF Healthcare Schema Policies validate EMR API payloads against standards like FHIR, blocking SQLi via parameterized queries and input escaping tailored for sensitive health data exchanges. Anti-DDoS Exception Thresholds customize SYN flood limits per ALB IP, preventing blackhole disruptions to real-time consult video streams while scrubbing excess traffic. CTDR Cross-Service Alerting unifies Security Center, WAF, and Anti-DDoS logs for threat tracing, generating HIPAA-compliant incident reports on correlated ransomware-DDoS vectors.

**Question: 1602**

A legal doc review on Alibaba Cloud deploys EDAS HSF for analysis, auto-scaling threads>80%/latency, zero-downtime blue-green, GTS reviews, mTLS, SchedulerX indexing, Jaeger, RDS store. Which EDAS?

- A. HSF analysis, auto-scaling threads/latency, blue-green zero-downtime, GTS reviews.
- B. Professional mTLS, Jaeger traces, SchedulerX indexing, RDS.
- C. Standard no GTS, manual, no Jaeger.

D. Container no mTLS, basic.

Answer: A, B

Explanation: HSF for doc services, auto-scaling threads>80%/latency; blue-green zero-downtime, GTS consistency. Professional mTLS; Jaeger traces, SchedulerX indexes, RDS stores.

**Question: 1603**

Which statement correctly describes the route table behavior in Alibaba Cloud regarding VSwitches within a single VPC?

- A. Route tables must be manually configured for each VSwitch to enable communication between VSwitches
- B. Each VSwitch automatically inherits the main route table from the VPC for routing
- C. Route propagation in the route table must be disabled to allow VSwitch routing
- D. VSwitches in different AZs cannot communicate regardless of route tables

Answer: B

Explanation: By default, all VSwitches within a VPC use the main route table, allowing them to communicate without additional configuration. Route propagation settings control dynamic routes and are unrelated to default VSwitch routing.

**Question: 1604**

Gaming leaderboard resolves "score.game.leader" to SRV Redis in ACK VPC (Philippines Manila), TCP 6379 every 8 seconds. Forward to Redis Cloud 34.202.0.1, 200 TTL, CloudMonitor score updates. 2026, which ACK Redis?

- A. Acceleration zones SRV Redis, TCP 6379, ACK bind, forward to Redis Cloud, CloudMonitor updates.
- B. Regular zones leader, recursion Redis, subzones score.
- C. HTTP Redis, manual SRV, no forward.
- D. No SRV, disable ACK, ignore updates.

Answer: A

Explanation: Acceleration zones SRV for Redis services, TCP port checks at 8 seconds,

ACK binding auto; forwarding + CloudMonitor update scores. Regular leader phases, recursion risks; HTTP mismatches, manual SRV static, no forward isolates. No SRV limits, disable ACK manual, ignore unranked.

**Question: 1605**

Your organization must integrate ACK clusters with VPCs to ensure secure private network communication between pods and enterprise services. Which ACK networking feature supports this integration?

- A. Using public IPs for pod-to-service communication
- B. Using overlay networks without VPC integration
- C. Deploying nodes without private IPs
- D. VPC-native networking with ENI (Elastic Network Interface) mode

Answer: D

Explanation: ACK supports VPC-native networking using ENI mode, allowing pods to obtain IP addresses from the VPC subnet, facilitating secure, high-performance private connectivity within enterprise networks.

**Question: 1606**

A supply chain Oracle/OSS HBR Mumbai GRS Chennai, RAC consistency, OSS Cold Archive 40 days access, delete 3 years. Compression, DR ECS.

- A. HBR Oracle RAC, compression GRS.
- B. OSS Cold Archive 40, expire 1095.
- C. ECS annual DR.
- D. Dedup chain data.

Answer: A, B, D

Explanation: RAC consistency/compression for queries, Cold Archive for shipments. ECS DR, dedup savings.

**Question: 1607**

PolarDB's LogIndex technology offers two record replay modes: lazy replay and parallel

replay. For which scenario is parallel replay especially beneficial?

- A. Encrypting data-at-rest
- B. Increasing storage capacity dynamically
- C. Minimizing replication lag for large OLAP queries
- D. Scaling read-only nodes horizontally

Answer: C

Explanation: Parallel replay accelerates redo log application on replicas, reducing replication lag for complex analytical queries (OLAP) with heavy workloads. It does not directly relate to storage scaling, encryption, or horizontal scaling of nodes.

**Question: 1608**

Telecom lifts OSS/BSS from Amdocs to RDS, modernizing to 5G services on MSE. TM Forum compliance. Which?

- A. DTS Amdocs schemas.
- B. MSE TM APIs.
- C. Event streaming.
- D. Orchestration.

Answer: A, B, C

Explanation: DTS syncs OSS/BSS. MSE exposes TM. Streaming enables 5G.

**Question: 1609**

How does Alibaba Cloud ActionTrail assist in meeting HIPAA compliance requirements for covered entities using cloud resources?

- A. By encrypting all ECS data volumes by default
- B. By logging all API calls related to resource access and delivering encrypted logs for audit and forensic analysis
- C. By anonymizing user identities in cloud logs
- D. By automatically patching cloud services for vulnerabilities

Answer: B

Explanation: HIPAA compliance requires detailed, tamper-evident audit trails recording access to protected health information and cloud resources. ActionTrail helps by logging all relevant API calls and enabling secure log delivery. Encryption of data volumes and patching are important security controls but not direct functions of ActionTrail. Anonymizing logs would conflict with audit traceability.

### Question: 1610

For AR/VR content, creators upload 90 GB scenes multipart globally, accelerated via Transfer. Policies limit to verified creators, abort >4 hours. CDN with geo-routing. Which accel and route enhances?

- A. Enable Transfer Acceleration, SDK endpoint \*.accelerate.aliyuncs.com
- B. Policy "Allow" if "StringEquals": {"sts:RoleSessionName": "verified-creator-\*"}
- C. Lifecycle abort >4 hours incomplete
- D. CDN geo-routing to nearest POP for scene delivery

Answer: A, B, C, D

Explanation: Acceleration speeds global multipart. Session name verifies creators. Abort manages incompletes. Geo-CDN reduces latency. Powers 2026 AR/VR.

### Question: 1611

How does Alibaba Cloud Function Compute ensure high availability across zones?

- A. Re-deploy functions on ECS instances in one zone
- B. Deploy functions in distributed clusters across multiple availability zones
- C. Use static DNS to route requests
- D. Use manual failover scripts

Answer: B

Explanation: Function Compute is deployed on distributed clusters in multiple availability zones so that if one zone fails, another takes over automatically to ensure service continuity.

### Question: 1612

A financial analytics platform processes batch jobs on c6.xlarge ECS instances, facing memory exhaustion during quarterly reports that involve complex simulations. Vertical scaling to c7.2xlarge is planned during off-peak hours, but the workload uses GPU-accelerated libraries incompatible with standard c7 CPUs. The architect must balance compute upgrades with fallback options for burst scenarios. Which factors should be prioritized for a hybrid vertical scaling strategy?

- A. Schedule instance type changes via Auto Scaling's vertical scaling feature during maintenance windows, ensuring zero-downtime by using blue-green deployments with SLB routing.
- B. Select c7.2xlarge for its superior SPECint benchmark scores (up to 20% higher than c6 equivalents), optimizing CPU-bound simulations without altering NUMA configurations for multi-threaded tasks.
- C. Provision burstable t7 instances as a supplemental tier, configuring CloudMonitor alarms to trigger vertical scaling only when sustained CPU credits deplete below 20%.
- D. Validate application compatibility with c7's enhanced I/O engine, testing packet forwarding rates up to 100 Gbps to handle increased data ingress from ApsaraDB RDS during reports.

Answer: B, D

Explanation: c7.2xlarge instances deliver markedly higher SPECint performance due to Intel Xeon 6 P-cores and CIPU optimizations, making them ideal for CPU-intensive financial simulations in vertical scaling, with NUMA awareness already tuned for such workloads. Testing compatibility with the enhanced I/O engine ensures the upgraded packet forwarding capabilities (up to 100 Gbps) support escalated data flows from RDS without latency spikes. Auto Scaling vertical scaling supports scheduled changes but relies on SLB for zero-downtime, not inherently requiring blue-green for ECS upgrades. Burstable t7 instances suit variable loads but introduce credit management complexity unsuitable as primary for sustained batch vertical scaling.

### Question: 1613

Energy monitoring CLB L4 TCP sensors, GA anycast AU 35% latency. CLB IP hash, drain 55s, ACL data. Weights 60/40. Which?

- A. CLB TCP IP hash, 55s drain, ACL TCP
- B. GA TCP to CLB, Anycast, 60/40, CEN
- C. LCU, multicast sensors, geo AU
- D. GA checks 8s, 1.2 Gbps, failover 11s

Answer: A, B

Explanation: CLB TCP IP hash persist, drain, ACL. GA TCP Anycast, weights, CEN for 35% cut. LCU; multicast; geo; checks 5s; failover 10s.

**Question: 1614**

Your team wants to run serverless applications in Alibaba Cloud Function Compute using a custom runtime for a new language not natively supported. Which approach is the best practice to implement this?

- A. Use Alibaba Cloud provided Node.js runtime and transpile code
- B. Develop a custom runtime adhering to FC runtime API and package it with the function
- C. Run the code on ECS instances behind FC
- D. Use EDAS for custom language support

Answer: B

Explanation: Function Compute supports custom runtimes by allowing users to provide their own runtime implementations that comply with the FC runtime API. This enables running functions in any language by packaging the runtime alongside function code.

**Question: 1615**

If a NAT Gateway EIP is bound to multiple VPCs in different regions, what is the Alibaba Cloud recommendation?

- A. EIP bound to NAT Gateway must be unique per region; cross-region binding is unsupported
- B. NAT Gateway supports multi-region EIP binding by default
- C. VPC peering is automatically established to support multi-region EIP use
- D. Gateway Load Balancer distributes traffic to multiple NAT Gateways

Answer: A

Explanation: NAT Gateway EIPs are regional resources and cannot be bound across multiple regions. The architecture must design EIPs specific to each region's NAT Gateway.

**Question: 1616**

Which Alibaba Cloud feature allows customers to import their own cryptographic key material into KMS?

- A. BYOK (Bring Your Own Key)
- B. SSE-KMS default keys
- C. Client-side encryption only
- D. TDE key import

Answer: A

Explanation: BYOK enables customers to maintain control over key material by importing keys into Alibaba Cloud KMS, improving trust and compliance.

**Question: 1617**

A logistics firm's 250 on-premises hosts report network latency metrics to CloudMonitor, with dashboards for global views and alerts on >200ms spikes. Apsara Stack governance logs ACK pod evictions for audits. Which correct?

- A. CloudMonitor Hybrid latency metrics, global dashboards, spike alerts
- B. Apsara Stack logging for ACK pod evictions in governance audits
- C. CloudMonitor regional-only views
- D. Apsara Stack no eviction logging

Answer: A, B

Explanation: CloudMonitor Hybrid collects network latency from hybrid agents, renders global dashboards, and triggers alerts for spikes over 200ms to ensure logistics visibility. Apsara Stack captures detailed logs of ACK pod evictions for comprehensive governance and audit trails.

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, 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.



## Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

## Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

## Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

## Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

## Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.