AWS Certified Solutions Architect – Professional (SAP-C02)

100 Questions & Answers

*Welcome to your complete AWS Certified Solutions Architect – Professional (SAP-C02) practice question set. This collection is designed not just to quiz you, but to simulate the **real-world AWS architecture decision-making** tested on the exam.*

Learning Objectives and Expectations

You'll get:

- 100 realistic, scenario-based questions modeled after the actual SAP-C02 exam format
- Organized in batches of 10 questions, followed by 10 answers and explanations
- Clear, concise explanations to reinforce AWS architecture decisions, service behavior, and trade-off logic

SAP-C02 Exam Domains

Each domain represents a core architectural responsibility. Questions span networking, resilience, security, migration, and cost-optimized design:

- Domain 1: **Design Solutions for Organizational Complexity** – 26%
- Domain 2: **Design for New Solutions** – 29%
- Domain 3: **Continuous Improvement for Existing Solutions** – 25%
- Domain 4: **Accelerate Workload Migration and Modernization** – 20%

Quick Reminder: How the Exam Works

- **Number of Questions:** 75 total (65 scored + 10 unscored)
- **Format:** Multiple choice + multiple response, scenario-based
- **Time Limit:** 180 minutes
- **Passing Score:** 750/1000
- **Test Provider:** Pearson VUE (online proctored or in person)

Questions By Domain

Domain	Title	Questions Assigned	Question Numbers
Domain 1	Organizational Complexity (26%)	26 Questions	Q1, Q3, Q9–10, Q18, Q21–22, Q29, Q32, Q35, Q36, Q38, Q42, Q48, Q50, Q56–57, Q59, Q63, Q65, Q68, Q70, Q77, Q83, Q85, Q100
Domain 2	Design for New Solutions (29%)	29 Questions	Q2, Q5, Q7, Q12–13, Q15, Q17, Q19, Q23, Q25–26, Q31, Q33, Q39, Q41, Q43, Q45–47, Q49, Q51, Q53, Q58, Q64, Q66, Q69, Q75
Domain 3	Improve Existing Solutions (25%)	25 Questions	Q4, Q6, Q8, Q11, Q14, Q16, Q20, Q24, Q27–28, Q30, Q34, Q37, Q40, Q44, Q54–55, Q60–61, Q67, Q71, Q76, Q80, Q82, Q84
Domain 4	Migration & Modernization (20%)	20 Questions	Q72–74, Q78–79, Q81, Q86–87, Q88–89, Q90–99

Remember — You Don't Need to Be Perfect to Pass!

The SAP-C02 passing score is 750/1000, which means you can miss **up to 12–15 scored questions** and still succeed. The key is to **read carefully, understand the business need, and choose the best-fit solution based on trade-offs**. This exam doesn't test memorization — it tests your architectural judgment.

Questions 1–10

Q1.
A company needs to centralize CloudTrail logs from 50 AWS accounts spread across multiple Organizational Units (OUs). The solution must be scalable, secure, and enforceable across all accounts without requiring manual configuration in each. What is the best way to implement this?

A) Create a CloudTrail in each account and configure S3 bucket policies manually
B) Use AWS Control Tower to enable trails in member accounts
C) Use an AWS Organizations trail with centralized S3 bucket and apply service control policies
D) Use EventBridge to forward logs to a central logging account

---

Q2.
A financial institution is designing a payment processing system on AWS. The system must maintain low-latency, durability, and guarantee **exactly-once processing** of payment events. Which combination of services is most appropriate?

A) Amazon SQS + AWS Lambda + Amazon Aurora
B) Amazon Kinesis Data Streams + Lambda + DynamoDB
C) Amazon SNS + S3 + Step Functions
D) Amazon MQ + EC2 worker instances + RDS MySQL

Q3.
Your company has a Direct Connect connection between the corporate data center and an AWS region. They want to route traffic to **VPCs in multiple regions** using this connection. What should you recommend?

A) Use VPC Peering across regions
B) Use a Direct Connect Gateway attached to Transit Gateways in each region
C) Set up a VPN over the Direct Connect to each region's VPC
D) Configure a CloudHub VPN mesh with BGP

Q4.
A customer is migrating a monolithic web application that uses an on-prem relational database. They want to reduce ops overhead and gain fault tolerance in the cloud, but **cannot refactor the app**. Which migration strategy and service should they use?

A) Refactor using Lambda and DynamoDB
B) Rehost to EC2 and manage MySQL manually
C) Replatform to RDS for MySQL using AWS DMS
D) Repurchase with a third-party SaaS CRM

Q5.
A startup hosts a globally used API with unpredictable usage spikes. They want low operational overhead, high availability, and automatic scaling. Which architecture best meets the need?

A) EC2 instances in Auto Scaling Groups with Route 53 geo-routing
B) Lambda functions behind API Gateway with CloudFront distribution
C) ECS on EC2 with Application Load Balancer
D) EKS with node groups in each AWS region

Q6.
A healthcare provider uses Amazon RDS for PostgreSQL in one region. Due to compliance requirements, they must **replicate the database to another region** with minimal lag and support for read queries. What's the most efficient solution?

A) Enable cross-region read replicas on RDS
B) Use DMS with ongoing replication
C) Enable Multi-AZ for the RDS instance
D) Use S3 export and import for nightly sync

Q7.
A legacy document processing app was lifted to AWS and now runs on EC2. It frequently fails due to OS-level patching issues. The team wants to modernize while minimizing refactoring. What should they do?

A) Move to Lambda and S3 triggers
B) Containerize and run in ECS on Fargate
C) Migrate to Elastic Beanstalk
D) Replace with Amazon WorkDocs

Q8.
A workload running on Amazon EMR has increasing costs. Utilization of task nodes fluctuates heavily depending on time of day. What strategy optimizes cost without sacrificing performance?

A) Convert to EC2 Spot Instances for all nodes
B) Use EMR Instance Fleets with On-Demand master and Spot task nodes
C) Purchase RIs for the entire cluster
D) Use Amazon Redshift instead

Q9.
A customer wants to **enforce encryption** for all Amazon S3 uploads across all accounts in the org. What's the best method to meet this at scale?

A) Use bucket policies that deny unencrypted uploads
B) Apply default encryption on all buckets manually
C) Use IAM policies to require SSE-S3
D) Set up SCPs to enforce encrypted writes globally

Q10.
An internal app hosted behind a Network Load Balancer (NLB) must be protected against common web exploits (SQLi, XSS). What is the best AWS-native way to implement this?

A) Attach WAF to the NLB directly
B) Put a CloudFront distribution in front of the NLB and attach AWS WAF
C) Use Shield Advanced with the NLB
D) Use API Gateway in front of the NLB

Answers 1–10

A1.
Answer: C) Use an AWS Organizations trail with centralized S3 bucket and apply service control policies
Explanation: Organization-level trails allow you to centrally log all activity. Combined with SCPs, you can enforce logging and prevent modifications across all accounts.

A2.
Answer: B) Amazon Kinesis Data Streams + Lambda + DynamoDB
Explanation: Kinesis ensures ordered, durable ingestion. Paired with Lambda and idempotent DynamoDB writes, this combo supports near real-time, exactly-once processing.

A3.
Answer: B) Use a Direct Connect Gateway attached to Transit Gateways in each region
Explanation: Direct Connect Gateway lets you extend a single DX to multiple regions through regional Transit Gateways. It's scalable and avoids region-specific links.

A4.
Answer: C) Replatform to RDS for MySQL using AWS DMS
Explanation: Replatforming offloads database operations while maintaining compatibility. DMS supports minimal-downtime migrations and preserves schema and data.

A5.
Answer: B) Lambda functions behind API Gateway with CloudFront distribution
Explanation: This architecture auto-scales, has low ops overhead, and benefits from CloudFront's caching and global acceleration features.

A6.
Answer: A) Enable cross-region read replicas on RDS
Explanation: Cross-region replicas are built-in, require minimal setup, and provide low-latency read access plus DR support.

A7.
Answer: B) Containerize and run in ECS on Fargate
Explanation: Fargate eliminates server management and patching while keeping the app mostly intact — a good modernization step short of a full refactor.

A8.
Answer: B) Use EMR Instance Fleets with On-Demand master and Spot task nodes
Explanation: This combo balances stability (On-Demand master) with cost savings (Spot task nodes), adapting to workload variability.

A9.
Answer: A) Use bucket policies that deny unencrypted uploads

Explanation: Bucket policies can reject any PUT requests lacking server-side encryption headers, enforcing org-wide encryption.

A10.

Answer: B) Put a CloudFront distribution in front of the NLB and attach AWS WAF
Explanation: WAF integrates with CloudFront, not directly with NLB. This pattern enables edge-based filtering of malicious requests.

Questions 11–20

Q11.
An enterprise is deploying a global application with users in North America, Europe, and Asia. They need low-latency access and consistent routing. Which solution best meets these requirements?

A) Use Application Load Balancers in each region and Route 53 weighted routing
B) Deploy in one region and use CloudFront for global access
C) Deploy in multiple regions and use AWS Global Accelerator
D) Use Route 53 latency-based routing with a single-region backend

---

Q12.
A critical workload runs on EC2 in a single Availability Zone. Recently, an AZ failure caused a multi-hour outage. Management wants an HA design with **minimal code changes**. What should you do?

A) Move to AWS Lambda across AZs
B) Use Auto Scaling Groups across multiple AZs behind an ALB
C) Use EC2 Spot Instances in another region
D) Implement CloudFront caching to reduce origin traffic

Q13.
A company uses AWS Systems Manager Patch Manager across accounts. They want to ensure patch compliance and alert on non-compliant resources. What solution should be used?

A) Use CloudTrail and Athena to query patch logs
B) Aggregate patch compliance data in AWS Config with conformance packs
C) Manually export SSM reports and send them via email
D) Use CloudWatch metrics only to track instance health

Q14.
You are designing a document archival solution that must **store 500 TB** of rarely accessed data for **7 years** with the **lowest cost**. Access time can be several hours. What is the best solution?

A) S3 Standard with lifecycle policies
B) S3 Glacier Deep Archive

C) Amazon EBS Snapshots stored in S3
D) FSx for Lustre with data tiering

Q15.
A customer uses a third-party identity provider to authenticate users and wants to allow federated access to the AWS console without creating IAM users. What service should be used?

A) Amazon Cognito
B) IAM Identity Center (AWS SSO)
C) IAM Federation with long-lived access keys
D) AWS Directory Service

Q16.
You're migrating a legacy app that relies on **POSIX-compliant file storage** with shared access across multiple EC2 instances in different AZs. Which solution fits best?

A) Amazon S3 with Transfer Acceleration
B) Amazon EFS
C) Amazon FSx for Windows
D) Instance store volumes with replication

Q17.
A startup wants to minimize cost while building a data lake ingestion pipeline. The data arrives sporadically and needs to be processed within a few minutes. Which architecture is best?

A) EC2 workers polling S3 every 10 minutes
B) Lambda triggered by S3 with data processed and stored in DynamoDB
C) Kafka on EC2 with hourly batch jobs
D) EKS cluster with worker nodes scaling manually

Q18.
An enterprise requires that all application deployments be **blue/green** with automatic rollback on failure. Which combination of services should you use?

A) CodeDeploy + Lambda + CloudWatch Alarms
B) EC2 Auto Scaling + manual scripts
C) Elastic Beanstalk + NAT Gateway
D) CloudFormation with drift detection

Q19.
Your development team frequently pushes updates to an S3-hosted static site. They report that some users still see old versions after updates. What's the root cause and best fix?

A) IAM caching delays — update policies
B) S3 eventual consistency — use versioned buckets
C) CloudFront caching — invalidate distribution or use cache busting
D) Network ACLs blocking access — update rules

Q20.
A machine learning team needs to periodically run training jobs with GPU acceleration. The workload is sporadic and doesn't justify always-on resources. What's the best setup?

A) Use EC2 GPU instances with auto-scaling
B) Use ECS on Fargate with GPU tasks
C) Use AWS Batch with GPU-based instance types
D) Pre-purchase GPU EC2 Reserved Instances

Answers 11–20

A11.
Answer: C) Deploy in multiple regions and use AWS Global Accelerator
Explanation: Global Accelerator routes users through AWS edge locations to the nearest healthy regional endpoint, providing the lowest latency and automatic failover.

A12.
Answer: B) Use Auto Scaling Groups across multiple AZs behind an ALB
Explanation: This is the simplest way to achieve high availability for EC2 apps with minimal changes. It ensures distribution and failover across AZs.

A13.
Answer: B) Aggregate patch compliance data in AWS Config with conformance packs
Explanation: AWS Config can collect patch compliance from Systems Manager across

accounts, and conformance packs can generate centralized reports and remediation actions.

A14.
Answer: B) S3 Glacier Deep Archive
Explanation: Glacier Deep Archive is the lowest-cost storage class for long-term, rarely accessed data and is ideal for compliance archives.

A15.
Answer: B) IAM Identity Center (AWS SSO)
Explanation: IAM Identity Center integrates with external IdPs for SSO access to AWS accounts without creating IAM users. It supports SAML federation and fine-grained access control.

A16.
Answer: B) Amazon EFS
Explanation: EFS provides shared, scalable, POSIX-compliant file storage across AZs and is ideal for EC2 workloads requiring file system semantics.

A17.
Answer: B) Lambda triggered by S3 with data processed and stored in DynamoDB
Explanation: This is a low-cost, event-driven architecture that handles sporadic workloads without provisioning compute resources.

A18.
Answer: A) CodeDeploy + Lambda + CloudWatch Alarms
Explanation: CodeDeploy supports blue/green deployment with rollback triggered by CloudWatch alarm failures. It works with Lambda, ECS, and EC2.

A19.
Answer: C) CloudFront caching — invalidate distribution or use cache busting
Explanation: CloudFront caches S3 content at the edge. Without invalidation or versioned URLs, users may receive stale content.

A20.
Answer: C) Use AWS Batch with GPU-based instance types
Explanation: AWS Batch provisions GPU instances on-demand for batch jobs, eliminating the need to keep expensive GPU resources running 24/7.

Questions 21–30

Q21.
A company stores logs in Amazon S3. Analysts query this data regularly using Amazon Athena. They want to reduce query costs and improve performance. What is the best solution?

A) Compress files using Gzip before storing in S3
B) Store logs in JSON format
C) Convert data to Apache Parquet and partition it
D) Use S3 Intelligent-Tiering

Q22.
Your application runs on EC2 and requires access to an S3 bucket in another account. You want to **follow least privilege** and avoid hardcoding credentials. What should you do?

A) Create an IAM user in the target account and share access keys
B) Attach an inline policy granting access to the bucket
C) Create a resource-based bucket policy allowing access from a specific role
D) Copy the data from the target bucket to the local account

Q23.
An organization runs a stateful application on EC2 with EBS volumes. They want to improve **resilience and RTO** without moving to managed services. What should they implement?

A) Enable Multi-AZ for EBS volumes
B) Use EC2 Auto Recovery and regular snapshots
C) Add a second network interface for failover
D) Convert to instance store for faster performance

Q24.
A developer is building an API backend using API Gateway and Lambda. They need to authenticate users via a third-party IdP and pass identity to downstream services. What service should be integrated?

A) Amazon Cognito
B) IAM Identity Center
C) AWS Directory Service
D) GuardDuty

Q25.
You need to provide temporary, **audited, least-privilege access** for a third-party vendor to troubleshoot an issue in your AWS account. What is the best practice?

A) Create a new IAM user and share login details
B) Use STS with limited-duration role assumption
C) Grant full access and monitor CloudTrail
D) Create an SSO group for vendor access

Q26.
A real-time analytics system ingests and processes thousands of IoT device updates per second. Which solution provides scalable ingestion and stream processing?

A) Amazon MQ + Lambda
B) Amazon SQS FIFO + ECS
C) Amazon Kinesis Data Streams + Kinesis Data Analytics
D) Amazon SNS + DynamoDB

Q27.
An ecommerce platform uses RDS for MySQL. During major sales events, the application becomes unresponsive due to high read latency. What's the **simplest way** to improve performance?

A) Add a CloudFront distribution
B) Create a read replica and update app connections
C) Scale the database vertically
D) Use Aurora Global Database

Q28.
A government agency is designing a sensitive workload requiring **data residency** in a specific country. Which AWS feature ensures compliance with this requirement?

A) AWS Local Zones
B) Availability Zones
C) Data Transfer Acceleration
D) AWS Region selection

Q29.
A company wants to **enforce a tagging strategy** where EC2 instances must have a "Project" tag upon creation. What is the best way to enforce this policy?

A) Use AWS Config rules
B) Enable cost allocation tags
C) Use IAM policies with `aws:RequestTag` condition
D) Use CloudTrail to detect untagged instances

Q30.
A financial app requires high transaction throughput, ACID compliance, and fast failover in a single region. Which database service best meets this requirement?

A) Amazon RDS Single-AZ
B) Amazon Aurora Multi-AZ
C) DynamoDB On-Demand
D) Amazon DocumentDB

Answers 21–30

A21.
Answer: C) Convert data to Apache Parquet and partition it
Explanation: Parquet is a columnar format that reduces scan time and cost. Partitioning further limits query scope, improving Athena performance.

A22.
Answer: C) Create a resource-based bucket policy allowing access from a specific role
Explanation: This method enables cross-account access without credentials and follows the least privilege model using role assumption.

A23.
Answer: B) Use EC2 Auto Recovery and regular snapshots
Explanation: Auto Recovery restarts instances on hardware failure, and snapshots ensure recovery points. This improves RTO without rearchitecting.

A24.
Answer: A) Amazon Cognito
Explanation: Cognito integrates with third-party IdPs (via SAML/OIDC), issues tokens, and can pass identity to API Gateway and Lambda.

A25.
Answer: B) Use STS with limited-duration role assumption
Explanation: STS allows secure, temporary access with full CloudTrail visibility and avoids persistent credentials.

A26.
Answer: C) Amazon Kinesis Data Streams + Kinesis Data Analytics
Explanation: Kinesis is designed for scalable, real-time streaming ingestion and analytics, perfect for high-volume IoT data.

A27.
Answer: B) Create a read replica and update app connections
Explanation: Read replicas reduce read load on the primary DB. This is a simple, effective solution for performance under heavy read traffic.

A28.
Answer: D) AWS Region selection
Explanation: Only deploying resources in a compliant AWS region ensures data residency. Services do not cross region boundaries unless explicitly configured.

A29.
Answer: C) Use IAM policies with aws:RequestTag condition
Explanation: IAM condition keys like aws:RequestTag can enforce tagging requirements at resource creation time.

A30.
Answer: B) Amazon Aurora Multi-AZ
Explanation: Aurora offers high throughput, ACID compliance, and automatic failover within a region, suitable for financial workloads.

Questions 31–40

Q31.
A financial services company has strict regulatory requirements for **auditing all user activity** across dozens of AWS accounts. They want a centralized approach to **log management and retention**, and need to ensure logs are **immutable and retained for 7 years**. What should you recommend?

A) Enable AWS Config in each account and aggregate logs to Amazon S3
B) Create org-level CloudTrail with centralized S3, S3 Object Lock, and lifecycle rules
C) Store logs in CloudWatch Logs with retention set to 7 years
D) Use Athena to query individual account logs stored in local S3 buckets

---

Q32.
An application deployed on EC2 in a public subnet accesses an RDS database in a private subnet. You notice **database connections are intermittently failing** under load. After investigation, you suspect **source port exhaustion** on the NAT Gateway. What's the most cost-effective fix?

A) Add a second NAT Gateway in another AZ
B) Add a NAT instance for failover
C) Deploy the EC2 app in private subnets using interface endpoints
D) Use Application Load Balancer in front of the RDS

Q33.
You are migrating a legacy enterprise Java application to AWS. It uses WebLogic, runs 24/7, and stores data in an Oracle DB. The app cannot be refactored. The team wants to reduce operational overhead. What's the best migration target?

A) Amazon EC2 with WebLogic and Oracle installed
B) Amazon ECS with containers running WebLogic
C) AWS Marketplace AMI with WebLogic and Amazon RDS for Oracle
D) AWS Lambda with Amazon Aurora

Q34.
A video streaming platform stores millions of video files in Amazon S3. They need to enforce that **only CloudFront** can access the S3 bucket, **preventing direct access** via public URLs. What's the most secure solution?

A) Enable public access block and use signed URLs in the app
B) Use an origin access identity (OAI) in CloudFront and update the S3 bucket policy
C) Restrict S3 access with VPC endpoint policies
D) Enable S3 versioning and enable encryption

Q35.
A security team requires that **all IAM changes** (creation, deletion, policy updates) across AWS accounts are tracked and reviewed **daily**. What is the most efficient and scalable way to implement this?

A) Create CloudWatch Events for IAM API calls and push to an SNS topic
B) Enable AWS Config with IAM rules and email snapshots daily
C) Use AWS CloudTrail with an Athena query scheduled via EventBridge
D) Enable GuardDuty to detect IAM changes and trigger Lambda functions

---

Q36.
A global company runs its core e-commerce app in us-east-1. They want to enable **fast failover** to eu-west-1 in case of a regional disaster, while maintaining cost-efficiency. Which architecture supports RTO < 5 minutes and minimal idle cost?

A) Multi-site active-active with fully provisioned infrastructure in both regions
B) Warm standby with scaled-down resources in eu-west-1, Route 53 failover
C) Backup and restore using nightly S3 snapshots
D) Pilot light strategy with database replication only

Q37.
Your data engineering team ingests CSV files into S3, transforms them using EMR, and loads results into Redshift. They want to **reduce time and cost** of the ETL pipeline. What should you recommend?

A) Replace EMR with Glue ETL and convert files to Parquet
B) Use S3 event triggers to run transformation code on Lambda
C) Query S3 data directly from Redshift using Spectrum
D) Move the pipeline to Kinesis Data Analytics

Q38.
A SaaS provider manages isolated tenant environments across AWS accounts. They want to **share a central RDS database cluster** with read access to each tenant account. What should they do?

A) Share the RDS endpoint via AWS RAM
B) Set up VPC peering and allow access through security groups
C) Use RDS proxy and IAM roles for cross-account access
D) Attach the RDS cluster to a shared Transit Gateway and configure routing

Q39.
A developer team uses AWS CloudFormation to deploy infrastructure. Occasionally, stack updates accidentally delete critical resources. They want to **prevent accidental deletion** of production resources while still using automation. What should you implement?

A) Enable rollback triggers and snapshot protection
B) Use termination protection and stack policies
C) Add deletion lifecycle hooks in Auto Scaling groups
D) Use SSM Automation for all production deployments

Q40.
You manage a compliance-sensitive workload that encrypts all data at rest. The team wants to ensure **no one, including AWS employees**, can decrypt the data. What's the most secure and auditable option?

A) Use SSE-S3 with default AWS-managed keys
B) Use customer-managed CMKs and enable key rotation
C) Use customer-managed CMKs with key policies that deny AWS root access
D) Use client-side encryption with S3 and KMS envelope keys

Answers 31–40

A31.
Answer: B) Create org-level CloudTrail with centralized S3, S3 Object Lock, and lifecycle rules
Explanation: An org-wide CloudTrail with S3 Object Lock ensures log immutability and long-term retention. Lifecycle rules manage storage cost.

A32.
Answer: C) Deploy the EC2 app in private subnets using interface endpoints
Explanation: Moving EC2 to private subnets and using VPC endpoints avoids NAT Gateway entirely, eliminating port exhaustion and reducing cost.

A33.
Answer: C) AWS Marketplace AMI with WebLogic and Amazon RDS for Oracle
Explanation: This solution meets licensing and ops requirements without refactoring. RDS reduces DB overhead, and the Marketplace AMI simplifies app migration.

A34.
Answer: B) Use an origin access identity (OAI) in CloudFront and update the S3 bucket policy
Explanation: OAI restricts direct S3 access and allows only CloudFront to fetch content securely.

A35.
Answer: C) Use AWS CloudTrail with an Athena query scheduled via EventBridge
Explanation: CloudTrail logs IAM activity, and Athena can query logs daily. EventBridge can schedule queries, enabling automated daily reviews.

A36.
Answer: B) Warm standby with scaled-down resources in eu-west-1, Route 53 failover
Explanation: Warm standby provides lower idle cost than active-active and meets fast failover (RTO < 5 min) using Route 53 health checks.

A37.
Answer: A) Replace EMR with Glue ETL and convert files to Parquet
Explanation: Glue is serverless and cost-efficient. Using Parquet (columnar format) reduces I/O and processing time, speeding up the ETL pipeline.

A38.
Answer: D) Attach the RDS cluster to a shared Transit Gateway and configure routing
Explanation: Transit Gateway enables scalable cross-account, cross-VPC connectivity. RDS can be accessed securely with proper routing and SGs.

A39.
Answer: B) Use termination protection and stack policies

Explanation: Stack policies prevent updates to protected resources. Termination protection stops stacks from being accidentally deleted.

A40.
Answer: C) Use customer-managed CMKs with key policies that deny AWS root access
Explanation: CMKs can be explicitly restricted via key policies to prevent even AWS from accessing or managing keys, maximizing control.

Questions 41–50

Q41.
A company wants to migrate a legacy batch-processing workload to AWS. It runs once per day, processes 1 TB of data, and requires 8 vCPUs and 32 GB RAM. The goal is to minimize cost and operational overhead. Which solution should you choose?

A) EC2 Reserved Instance with a cron job
B) AWS Batch with Spot Instances and compute environment
C) EKS cluster with manually scaled node groups
D) Lambda with provisioned concurrency

Q42.
Your organization uses AWS Organizations and wants to enforce that **only specific instance types** are launched in production accounts. What is the best approach to implement this?

A) Use IAM policies with condition keys
B) Use SCPs with condition keys for instance types
C) Use a Lambda function to terminate unapproved instances
D) Use AWS Config to detect and alert on instance type usage

Q43.
A data science team builds models using Amazon SageMaker notebooks and wants to share read-only access with an external contractor. The contractor should **not be able to export or copy the data**. What is the best approach?

A) Share the notebook instance and restrict permissions with IAM
B) Use SageMaker Studio and disable copy/paste in IAM session policies
C) Share the notebook files via S3 pre-signed URLs
D) Create a temporary IAM user with MFA

Q44.
A customer has deployed an application in AWS that must support **hundreds of concurrent WebSocket connections** with low-latency and minimal infrastructure management. What AWS service should be used?

A) EC2 with Nginx WebSocket proxy
B) API Gateway with WebSocket API

C) Application Load Balancer with Lambda
D) AppSync with DynamoDB Streams

Q45.
A DevOps engineer wants to ensure that infrastructure deployed with CloudFormation is **always tagged** with the environment (e.g., `Environment=Prod`). How can this be enforced?

A) Use AWS Budgets with tag filtering
B) Use AWS Config managed rules for resource tagging
C) Use SSM Parameter Store and pull tags into templates
D) Enable tagging in CloudTrail

Q46.
A mobile application offloads image uploads to Amazon S3. To optimize performance and reduce load on the backend, how should uploads be handled securely and efficiently?

A) Use S3 Transfer Acceleration and pre-signed URLs
B) Upload to S3 via API Gateway
C) Store in Amazon EBS and sync to S3
D) Use an EC2 proxy that writes to S3

Q47.
A company hosts a mission-critical app on EC2 and uses S3 for storing large data files. Recently, S3 access latency increased significantly during peak hours. The files are read-only. What solution improves performance and reduces S3 traffic?

A) Enable S3 versioning
B) Add an ElastiCache Redis cluster as a cache
C) Use CloudFront with S3 as the origin
D) Use S3 Select for partial reads

Q48.
You are designing a secure multi-account architecture where each account must **log all console and API activity** centrally. The solution must prevent tampering. What should be implemented?

A) Enable AWS Config with aggregator
B) Create a centralized S3 bucket with bucket policy and object versioning

C) Use AWS CloudTrail with centralized logging and S3 Object Lock
D) Forward logs using Lambda to Elasticsearch

Q49.
A large retail company has frequent changes to IAM policies across accounts. They want to **automatically detect and prevent overly permissive policies** (e.g., wildcard actions on all resources). What is the best AWS-native solution?

A) Enable IAM Access Analyzer with archive rules
B) Use AWS Config with managed rules for IAM
C) Use GuardDuty to detect abnormal behavior
D) Create a CloudWatch alarm for policy size increases

Q50.
You're tasked with migrating an on-prem system to AWS. The application is composed of tightly coupled services and uses NFS shares. The goal is to **lift-and-shift** with minimal downtime. Which strategy best fits this case?

A) Use AWS Application Migration Service (MGN) for rehosting
B) Use DMS with ongoing replication and refactor services
C) Use Server Migration Service with Amazon EFS
D) Use Snowball to transfer data and recreate the environment manually

---

Answers 41–50

A41.
Answer: B) AWS Batch with Spot Instances and compute environment
Explanation: AWS Batch is ideal for scheduled batch workloads and integrates with Spot for cost savings. It requires minimal infrastructure management.

A42.
Answer: B) Use SCPs with condition keys for instance types
Explanation: Service Control Policies can prevent launching disallowed instance types organization-wide, enforcing compliance at the org level.

A43.
Answer: B) Use SageMaker Studio and disable copy/paste in IAM session policies
Explanation: Studio supports session policies where clipboard actions can be blocked, providing controlled, read-only access.

A44.
Answer: B) API Gateway with WebSocket API
Explanation: API Gateway supports fully managed WebSocket APIs, scaling automatically with minimal infrastructure for real-time connections.

A45.
Answer: B) Use AWS Config managed rules for resource tagging
Explanation: AWS Config includes rules to check for required tags on resources. It can alert or trigger remediation workflows.

A46.
Answer: A) Use S3 Transfer Acceleration and pre-signed URLs
Explanation: Pre-signed URLs let the app upload directly to S3 securely. Transfer Acceleration speeds up global uploads via edge locations.

A47.
Answer: C) Use CloudFront with S3 as the origin
Explanation: CloudFront caches S3 content at edge locations, reducing latency and lowering load on S3 during high-traffic periods.

A48.
Answer: C) Use AWS CloudTrail with centralized logging and S3 Object Lock
Explanation: Centralized CloudTrail with Object Lock ensures logs are immutable and meet audit compliance needs.

A49.
Answer: B) Use AWS Config with managed rules for IAM
Explanation: AWS Config includes managed rules like iam-policy-no-statements-with-admin-access to detect risky policies automatically.

A50.
Answer: A) Use AWS Application Migration Service (MGN) for rehosting
Explanation: MGN enables lift-and-shift migrations with minimal downtime and supports applications with complex dependencies.

Questions 51–60

Q51.
A startup is building a mobile backend that must handle thousands of concurrent requests during peak hours, with minimal operational overhead. The architecture should scale automatically and minimize idle costs. Which solution is most appropriate?

A) EC2 Auto Scaling with ALB and Elastic IPs
B) Lambda with API Gateway and DynamoDB
C) ECS on EC2 with fixed instance pool
D) EKS with Kubernetes Ingress Controller

Q52.
A security audit found that IAM users in multiple accounts have active access keys that haven't been rotated in over 90 days. What is the best solution to detect and prevent this issue in the future?

A) Create CloudWatch alarms based on key age
B) Use IAM credential report and enforce with SCP
C) Enable AWS Config rule `access-keys-rotated` and set compliance alerts
D) Use IAM Access Analyzer to detect unused keys

Q53.
An analytics team stores compressed CSV files in S3 and queries them with Athena. They notice query times and costs are increasing. What should you recommend to improve performance and reduce cost?

A) Store the data in JSON format instead of CSV
B) Use DynamoDB for all metadata
C) Convert the data to Apache Parquet and partition it by date
D) Enable S3 Standard-IA and Intelligent-Tiering

Q54.
A compliance requirement mandates that **critical production secrets must never traverse the public internet** and be stored securely with audit tracking. What AWS service satisfies this?

A) IAM Secure Token Service
B) AWS Secrets Manager with VPC endpoint and KMS encryption

C) S3 with signed URLs
D) Amazon Cognito with OIDC tokens

Q55.
A web app runs on EC2 and is fronted by a CloudFront distribution. The team deploys a new version but users still see the old assets. What is the most likely cause and best fix?

A) S3 inconsistency — enable versioning
B) CloudFront caching — invalidate objects or use cache-busting
C) EC2 instance caching — reboot the instances
D) ALB stickiness — disable session affinity

Q56.
A retail company uses multiple AWS accounts and needs to enforce encryption on all EBS volumes. What is the most scalable approach?

A) Use IAM policies to require encrypted EBS
B) Use a CloudTrail query to detect unencrypted volumes
C) Use AWS Config rules and remediate with Lambda
D) Encrypt EBS at runtime with EC2 user data

Q57.
You are designing a multi-account architecture and want to minimize inter-VPC peering complexity. VPCs across accounts need to connect securely and scale over time. What solution is most appropriate?

A) VPC Peering with route table isolation
B) Use AWS PrivateLink with shared endpoints
C) Create a centralized VPC and use NAT
D) Deploy AWS Transit Gateway and share via RAM

Q58.
An enterprise's security team wants to **block all root account activity** except for emergency cases. What is the best practice?

A) Delete the root user credentials
B) Enable MFA and create an SCP that denies all actions except for a break-glass role
C) Rotate the root access key every 30 days
D) Use IAM Access Analyzer to monitor root actions

Q59.
You manage a highly available web app deployed across three AZs. One AZ becomes unavailable, but the app becomes slow. Logs show that the Auto Scaling group isn't replacing failed instances. What is the likely root cause?

A) Elastic Load Balancer health checks are misconfigured
B) Launch template has AZ hard-coded
C) Auto Scaling cooldown period is too short
D) Cross-zone load balancing is disabled

Q60.
A company needs to **migrate 80 TB of data** from their on-prem NFS to S3 over a weekend. The available internet bandwidth is limited. What is the best solution to complete the transfer on time?

A) Use AWS Snowball Edge device and DataSync
B) Use S3 Transfer Acceleration
C) Mount S3 using s3fs and copy directly
D) Configure S3 replication with NFS

Answers 51–60

A51.
Answer: B) Lambda with API Gateway and DynamoDB
Explanation: Serverless architecture scales automatically, reduces idle cost, and is fully managed. Ideal for bursty mobile backend traffic.

A52.
Answer: C) Enable AWS Config rule access-keys-rotated and set compliance alerts
Explanation: AWS Config tracks key rotation across accounts and can trigger remediation or alerts when thresholds are breached.

A53.
Answer: C) Convert the data to Apache Parquet and partition it by date
Explanation: Parquet reduces I/O and improves performance for Athena. Partitioning further narrows scanned data.

A54.
Answer: B) AWS Secrets Manager with VPC endpoint and KMS encryption
Explanation: Secrets Manager stores encrypted secrets, supports audit logging, and can be accessed over private VPC endpoints.

A55.
Answer: B) CloudFront caching — invalidate objects or use cache-busting
Explanation: CloudFront caches S3 or EC2 responses. Invalidate the cache or version assets to force fresh fetch.

A56.
Answer: C) Use AWS Config rules and remediate with Lambda
Explanation: Config rules detect non-compliant volumes; Lambda can automate remediation by snapshotting and recreating encrypted volumes.

A57.
Answer: D) Deploy AWS Transit Gateway and share via RAM
Explanation: TGW simplifies large-scale, multi-account VPC connectivity. RAM allows cross-account attachment.

A58.
Answer: B) Enable MFA and create an SCP that denies all actions except for a break-glass role
Explanation: MFA plus an SCP can block all root activity unless explicitly permitted. SCPs apply to the root user too.

A59.
Answer: B) Launch template has AZ hard-coded
Explanation: If the launch template restricts AZs, the Auto Scaling group can’t launch in healthy ones. Templates should support multiple AZs.

A60.
Answer: A) Use AWS Snowball Edge device and DataSync

Explanation: Snowball enables offline bulk data transfer. DataSync can pre-stage metadata and manage the upload pipeline efficiently.

Questions 61–70

Q61.
A company hosts a customer-facing application on EC2 instances behind an ALB. The marketing team frequently updates static content. They want these updates to be reflected immediately worldwide, and reduce origin load. What's the best solution?

A) Store static assets in S3 with CloudFront and versioned URLs
B) Host all content on EC2 and use Route 53 latency routing
C) Use EBS volumes and attach them to new EC2s in each AZ
D) Use ElastiCache to cache static pages globally

Q62.
A team is deploying a Lambda function to handle user registrations. Occasionally, users experience timeouts. Investigation reveals that the function sometimes exceeds 15 seconds under load. What is the most scalable fix?

A) Increase Lambda timeout to the max allowed
B) Move the logic to a step-by-step workflow using AWS Step Functions
C) Attach a larger memory size to the function
D) Use Lambda provisioned concurrency

Q63.
You are designing an API that needs to authenticate users, support third-party identity providers, and allow user data synchronization across devices. What is the best AWS-native service for this use case?

A) Amazon Cognito
B) AWS IAM
C) SAML Federation with STS
D) AWS Directory Service

Q64.
A SaaS platform requires each tenant to have isolated environments. The team wants a scalable way to create and manage the same infrastructure stack for each tenant account. Which solution should you use?

A) CloudFormation StackSets with service-managed permissions
B) Use Terraform in each account with manual execution

C) Share a common VPC using Resource Access Manager
D) Use AWS CodeDeploy to bootstrap infrastructure

Q65.
A company wants to run Windows-based file services on AWS, integrated with on-prem Active Directory. Users need to map network drives as they did on-prem. Which service is best suited?

A) Amazon FSx for Windows File Server
B) Amazon EFS
C) AWS Storage Gateway (Tape Gateway)
D) S3 with Transfer Acceleration

Q66.
An EC2-based application stores logs locally. After a crash, the logs are lost. What is the most durable, scalable solution for long-term log retention?

A) Mount EBS to store logs
B) Write logs to instance store volumes
C) Stream logs to CloudWatch Logs or S3
D) Use a bastion host to collect logs manually

Q67.
A research institution uses AWS to process satellite images. Processing is CPU-intensive, must scale to 1,000s of concurrent tasks, and run only when needed. What is the best architecture?

A) Lambda with S3 triggers
B) AWS Batch with Spot Instances and job queues
C) EC2 with Auto Scaling
D) ECS on Fargate with scheduled tasks

Q68.
A customer has a Direct Connect connection to a VPC in us-west-2. They need to reach VPCs in eu-central-1 and ap-southeast-1 using the same connection. What solution supports this?

A) Use multiple VPN tunnels over Direct Connect
B) Attach a Direct Connect Gateway to Transit Gateways in each region

C) Create separate Direct Connect connections to each region
D) Use VPC peering and NAT gateway chaining

Q69.
A company receives financial transaction data over HTTPS. The app must validate the **integrity and origin** of each payload. What technique should be used?

A) Server-side encryption
B) Use digital signatures with a shared secret
C) Validate digital signatures using asymmetric cryptography
D) Store hashes of each payload in S3

Q70.
An application hosted on AWS needs to store sensitive documents. Compliance requires that encryption keys be rotated annually and access to key usage must be auditable. What service should be used?

A) S3 with SSE-S3
B) AWS KMS with customer-managed CMKs and CloudTrail enabled
C) EC2 with GPG key management
D) AWS Secrets Manager

Answers 61–70

A61.
Answer: A) Store static assets in S3 with CloudFront and versioned URLs
Explanation: This offloads content to CloudFront, ensuring global low-latency access. Versioning helps bypass stale cache issues.

A62.
Answer: B) Move the logic to a step-by-step workflow using AWS Step Functions
Explanation: Step Functions help break down long-running processes into manageable steps with retries and better scalability.

A63.
Answer: A) Amazon Cognito

Explanation: Cognito handles user pools, federated logins, synchronization, and token management for apps.

A64.
Answer: A) CloudFormation StackSets with service-managed permissions
Explanation: StackSets allow you to push consistent infrastructure stacks across accounts and regions efficiently.

A65.
Answer: A) Amazon FSx for Windows File Server
Explanation: FSx integrates with Active Directory and provides SMB access, ideal for Windows file sharing.

A66.
Answer: C) Stream logs to CloudWatch Logs or S3
Explanation: Both services offer durability and are ideal for centralized log retention beyond instance lifecycle.

A67.
Answer: B) AWS Batch with Spot Instances and job queues
Explanation: Batch is designed for large-scale, compute-intensive, scheduled or on-demand tasks with cost optimization via Spot.

A68.
Answer: B) Attach a Direct Connect Gateway to Transit Gateways in each region
Explanation: DX Gateway allows one physical DX link to connect to multiple regional VPCs via TGWs.

A69.
Answer: C) Validate digital signatures using asymmetric cryptography
Explanation: Digital signatures using public/private key pairs ensure data integrity and origin verification.

A70.
Answer: B) AWS KMS with customer-managed CMKs and CloudTrail enabled
Explanation: CMKs support rotation, granular policies, and audit logs for compliance and security visibility.

Questions 71–80

Q71.
A SaaS company is onboarding hundreds of new tenants. Each tenant needs an isolated environment for compute and storage. The company wants to reduce deployment time and enforce consistency. What is the best solution?

A) Use Lambda functions with environment variables per tenant
B) Deploy isolated stacks using AWS CloudFormation StackSets
C) Use EKS namespaces for tenant isolation
D) Use a single account with IAM user isolation per tenant

Q72.
A customer stores data in S3 that must be retained for legal reasons but is **rarely accessed**. They want the **lowest-cost** storage that still supports compliance audits and access within 12 hours. Which storage class fits?

A) S3 Standard-IA
B) S3 Glacier
C) S3 Glacier Deep Archive
D) S3 One Zone-IA

Q73.
Your team uses ECS with Fargate to run containerized apps. Recently, deployments have failed due to **container image pull issues**. You want to increase **resiliency and speed** of deployments. What should you do?

A) Store images in Amazon ECR and enable ECR pull-through cache
B) Use Docker Hub with retries enabled
C) Use S3 for image delivery and pre-load containers
D) Use Elastic File System (EFS) for container images

Q74.
A real estate company uses Lambda to process property data uploads. Each file is around 500 MB and triggers a Lambda function. They report timeouts and memory errors. What's the most scalable and cost-effective fix?

A) Increase Lambda memory to 10 GB and retry on error
B) Use S3 events to trigger an AWS Batch job instead of Lambda

C) Use multipart uploads and reprocess in chunks via Lambda
D) Convert the app to run in EC2 with user data scripts

Q75.
A global ecommerce company wants to protect against **cross-region latency** and **DNS-based DDoS attacks** while ensuring fast failover. Which solution should be implemented?

A) Use Route 53 with health checks and latency routing
B) Use AWS Global Accelerator with ALBs in each region
C) Use CloudFront with regional origins
D) Use NAT Gateway with Regional Failover

Q76.
A team stores API logs in S3 and queries them using Amazon Athena. They want **automated detection of anomalies**, like traffic spikes or unusual patterns. What should they use?

A) AWS Config with custom rules
B) Amazon Macie with Athena integration
C) CloudWatch Contributor Insights
D) Amazon QuickSight with anomaly detection

Q77.
A compliance team wants to prevent developers from **disabling CloudTrail logging** in production accounts. What's the most effective method to enforce this?

A) Enable a CloudWatch alarm on CloudTrail stop events
B) Use AWS Config rule to detect changes and send SNS alerts
C) Apply an SCP that denies cloudtrail:StopLogging action
D) Use CloudTrail insights to detect behavior patterns

Q78.
A startup wants to test code changes quickly by deploying isolated environments for each feature branch, and tear them down when done. What AWS service or strategy best supports this?

A) Use Lambda aliases with version control
B) Use ECS with blue/green deployments

C) Use AWS CDK to spin up ephemeral environments per branch
D) Create manual EC2 snapshots per deployment

Q79.
A SaaS company provides each customer with a dedicated subdomain (e.g., customer1.example.com). They want to route traffic based on the subdomain to different API endpoints. What's the best routing solution?

A) Use Route 53 with wildcard records and API Gateway custom domains
B) Use ALB host-based routing rules
C) Use S3 static hosting with subdomain matching
D) Use CloudFront with Lambda@Edge for domain parsing

Q80.
A research company is analyzing sensitive medical images using EC2. To comply with regulations, **data must be encrypted at rest, and keys rotated automatically**, with **access limited to specific roles**. What is the correct setup?

A) Use EC2 with EBS encrypted by default KMS key
B) Use customer-managed CMKs with automatic rotation and key policies
C) Use instance store volumes and encrypt manually
D) Use S3 with server-side encryption and IAM policies

Answers 71–80

A71.
Answer: B) Deploy isolated stacks using AWS CloudFormation StackSets
Explanation: StackSets allow scalable, consistent deployment of infrastructure across accounts or regions, perfect for tenant environments.

A72.
Answer: C) S3 Glacier Deep Archive
Explanation: This class offers the lowest cost for long-term storage with 12-hour retrieval, meeting compliance needs at minimal expense.

A73.
Answer: A) Store images in Amazon ECR and enable ECR pull-through cache
Explanation: ECR with pull-through caching improves reliability and speed of pulling images, especially compared to public registries.

A74.
Answer: B) Use S3 events to trigger an AWS Batch job instead of Lambda
Explanation: AWS Batch handles large, heavy compute jobs better than Lambda, which has size and timeout limits.

A75.
Answer: B) Use AWS Global Accelerator with ALBs in each region
Explanation: Global Accelerator improves performance, protects against DNS-based attacks, and offers automatic regional failover.

A76.
Answer: D) Amazon QuickSight with anomaly detection
Explanation: QuickSight supports ML-powered anomaly detection directly from Athena-queried datasets.

A77.
Answer: C) Apply an SCP that denies cloudtrail:StopLogging action
Explanation: SCPs enforce organization-wide controls and can block disabling CloudTrail regardless of IAM permissions.

A78.
Answer: C) Use AWS CDK to spin up ephemeral environments per branch
Explanation: CDK supports automation and versioned deployments per branch, ideal for isolated, short-lived test environments.

A79.
Answer: A) Use Route 53 with wildcard records and API Gateway custom domains
Explanation: This supports scalable, dynamic subdomain routing with API Gateway and wildcard DNS management.

A80.
Answer: B) Use customer-managed CMKs with automatic rotation and key policies
Explanation: CMKs support granular access control, key rotation, and compliance-friendly auditability.

Questions 81–90

Q81.
Your team is migrating an internal HR application to AWS. It relies on a legacy Oracle database and requires minimal downtime during cutover. The team also wants to move to a managed database platform. What migration approach is best?

A) Use DMS with continuous replication to Amazon RDS for Oracle
B) Rehost the Oracle DB on EC2 and then refactor later
C) Use Snowball Edge to copy the DB files and restore in AWS
D) Replace Oracle with DynamoDB and rewrite the application

Q82.
A multi-tenant SaaS app on ECS uses shared infrastructure. A large tenant is monopolizing compute resources, causing latency for others. What is the best way to resolve this?

A) Use EC2 dedicated hosts per tenant
B) Throttle API Gateway requests by source IP
C) Move the tenant to a dedicated ECS service with its own resource limits
D) Enable caching in CloudFront

Q83.
A financial firm is required to retain all logs for 10 years and ensure **no one can delete or modify them**. What solution ensures compliance?

A) Store logs in S3 with lifecycle policies
B) Store logs in S3 Glacier with object tagging
C) Use S3 with Object Lock in Compliance Mode
D) Stream logs to CloudWatch and enable retention

Q84.
An EC2-based analytics application in a private subnet must download datasets from the internet during scheduled jobs. What's the most cost-effective and scalable method?

A) Assign public IPs to EC2 instances
B) Use a NAT Gateway in one AZ for all traffic
C) Use a NAT instance with scheduled start/stop
D) Use VPC endpoints for S3 and host data in S3

Q85.
A global company wants a **consistent security baseline** applied automatically across all new AWS accounts created under its organization. What AWS service can achieve this?

A) AWS Config
B) AWS Systems Manager
C) AWS Control Tower
D) IAM Access Analyzer

Q86.
A web application stores temporary session data. It must be accessible with **microsecond latency** and scale under heavy read/write operations. What is the ideal storage solution?

A) Amazon DynamoDB
B) Amazon RDS with Multi-AZ
C) Amazon ElastiCache (Redis)
D) Amazon S3 with Lambda

Q87.
A data pipeline consumes messages from multiple producers. Order must be preserved **per producer**, and **exactly-once delivery** is required. Which messaging service best fits?

A) Amazon SNS
B) Amazon SQS Standard
C) Amazon SQS FIFO
D) Amazon Kinesis Data Firehose

Q88.
Your company wants to enforce MFA for all IAM users and prevent any API access if MFA is not enabled. What should you use?

A) Use a custom CloudTrail-based Lambda function to block calls
B) Use IAM policies with aws:MultiFactorAuthPresent condition
C) Disable IAM users and only use roles
D) Use a VPC endpoint policy to block non-MFA calls

Q89.
You manage a large VPC with hundreds of EC2 instances, and the security team needs to **monitor outbound network activity** to detect potential data exfiltration. What solution meets this need?

A) Enable VPC Flow Logs with destination to S3 and analyze via Athena
B) Use GuardDuty with Lambda integration
C) Enable CloudTrail Data Events
D) Create a WAF rule to log requests

Q90.
A containerized app in ECS Fargate needs to securely store and retrieve credentials to access an external API. What is the best AWS-native solution?

A) Hardcode credentials in the container image
B) Store credentials in S3 and encrypt with KMS
C) Use AWS Secrets Manager and inject via environment variables
D) Use instance metadata service v2

Answers 81–90

A81.
Answer: A) Use DMS with continuous replication to Amazon RDS for Oracle
Explanation: DMS supports near-zero downtime migration and moves to managed RDS. It's ideal when minimizing refactoring and outage.

A82.
Answer: C) Move the tenant to a dedicated ECS service with its own resource limits
Explanation: This isolates noisy tenants, enforces compute limits, and maintains fairness for other customers.

A83.
Answer: C) Use S3 with Object Lock in Compliance Mode
Explanation: Object Lock prevents deletion/modification for a specified duration, meeting WORM retention requirements.

A84.
Answer: C) Use a NAT instance with scheduled start/stop
Explanation: NAT instances are more cost-effective than NAT Gateways for infrequent use. Scheduling start/stop saves cost further.

A85.
Answer: C) AWS Control Tower
Explanation: Control Tower automates account creation and enforces guardrails, ensuring a consistent security and governance baseline.

A86.
Answer: C) Amazon ElastiCache (Redis)
Explanation: Redis supports microsecond latency, high throughput, and is ideal for ephemeral session data at scale.

A87.
Answer: C) Amazon SQS FIFO
Explanation: FIFO queues maintain message order per group (e.g., per producer) and support exactly-once processing.

A88.
Answer: B) Use IAM policies with aws:MultiFactorAuthPresent condition
Explanation: This condition enforces that users must authenticate with MFA for API calls, blocking access otherwise.

A89.
Answer: A) Enable VPC Flow Logs with destination to S3 and analyze via Athena
Explanation: VPC Flow Logs provide network traffic metadata, and Athena enables scalable analysis for security and audit teams.

A90.
Answer: C) Use AWS Secrets Manager and inject via environment variables
Explanation: Secrets Manager securely stores credentials and can inject them into containers at runtime using environment variables.

Questions 91–100

Q91.
A company is modernizing a monolithic application by migrating it to microservices on AWS. They want each service to scale independently and communicate asynchronously. What architecture pattern and AWS service combination is most appropriate?

A) Use EC2 Auto Scaling with Nginx for service discovery
B) Use API Gateway for synchronous HTTP calls between services
C) Use Amazon SQS and SNS for decoupled service communication
D) Use EFS for shared service state

Q92.
A compliance policy requires that **customer-uploaded data be scanned for malware** before being processed by backend systems. Which architecture should you implement?

A) Use AWS WAF with a custom rule to detect file types
B) Trigger a Lambda function on S3 PUT to scan with third-party AV software in ECS
C) Use Amazon Macie to detect malware signatures
D) Enable S3 versioning and manually inspect files

Q93.
A company runs sensitive workloads in a private VPC and requires **encrypted DNS resolution** with minimal latency. What solution meets this requirement?

A) Use public Route 53 DNS with DNSSEC
B) Use Route 53 Resolver with DNS over TLS
C) Use Route 53 Resolver endpoints with VPC DNS forwarding
D) Use NAT Gateway and external DNS servers

Q94.
Your company deploys infrastructure using AWS CloudFormation. Occasionally, stack updates cause unexpected downtime. You need a safer way to **test changes before applying them**. What feature should you use?

A) Stack policies
B) CloudFormation Change Sets

C) Drift Detection
D) StackSets with service control policies

Q95.
An enterprise is running hundreds of EC2 instances for multiple departments. They want to implement a **cost allocation strategy** by department. What is the simplest way to support this?

A) Create a separate AWS account for each department
B) Use AWS Budgets for each team
C) Use cost allocation tags and enforce tagging with IAM
D) Enable consolidated billing and assign IAM groups

Q96.
You are building a multi-region app that must store **low-latency, replicated key-value data**. Each region should accept writes and sync in near real time. What AWS service supports this?

A) Amazon RDS with read replicas
B) Amazon ElastiCache with global replication
C) Amazon DynamoDB Global Tables
D) Amazon S3 with CRR and EventBridge

Q97.
An IoT platform collects millions of sensor readings daily. The data must be **ingested in real time**, **processed**, and **stored long-term** in a cost-efficient format. Which architecture is best?

A) Kinesis Data Streams → Lambda → RDS
B) API Gateway → DynamoDB
C) Kinesis Data Firehose → S3 (Parquet) with lifecycle policies
D) MQTT broker on EC2 → EBS

Q98.
A team wants to run high-volume data science jobs on AWS without managing infrastructure. Jobs may run for hours and require scaling from 10 to 100 vCPUs. What service fits best?

A) Amazon EMR
B) AWS Batch with managed compute environments

C) Lambda with Step Functions
D) ECS with EC2 spot fleets

Q99.
You must design a **disaster recovery plan** for a legacy application that includes on-prem systems and EC2 instances. The RTO must be <1 hour, and cost should be minimized. What strategy is best?

A) Backup and restore with Amazon S3 and EC2 snapshots
B) Multi-site active-active architecture
C) Warm standby with smaller instance sizes in DR region
D) Pilot light with only critical services pre-provisioned

Q100.
A developer accidentally deleted a production S3 bucket. Management wants to ensure this never happens again, even from admin users. What should you implement?

A) Enable versioning and MFA delete on the S3 bucket
B) Use CloudTrail to detect and recover deleted buckets
C) Deny `s3:DeleteBucket` with bucket policy
D) Apply an SCP that blocks all `s3:DeleteBucket` actions organization-wide

Answers 91–100

A91.
Answer: C) Use Amazon SQS and SNS for decoupled service communication
Explanation: Asynchronous communication using SQS and SNS enables scalable, loosely-coupled microservices.

A92.
Answer: B) Trigger a Lambda function on S3 PUT to scan with third-party AV software in ECS
Explanation: S3 events can trigger scanning processes in a scalable, automated manner using ECS or Lambda as the engine.

A93.
Answer: C) Use Route 53 Resolver endpoints with VPC DNS forwarding
Explanation: Route 53 Resolver enables encrypted DNS within the VPC and can integrate with hybrid DNS setups.

A94.
Answer: B) CloudFormation Change Sets
Explanation: Change Sets preview what resources will change before applying a stack update, reducing the risk of downtime.

A95.
Answer: C) Use cost allocation tags and enforce tagging with IAM
Explanation: Tags allow tracking by department. IAM policies can require tags at resource creation time for cost allocation.

A96.
Answer: C) Amazon DynamoDB Global Tables
Explanation: Global Tables support multi-region writes and near real-time replication with low-latency access.

A97.
Answer: C) Kinesis Data Firehose → S3 (Parquet) with lifecycle policies
Explanation: Firehose delivers streaming data to S3 in efficient formats like Parquet. S3 lifecycle policies reduce long-term storage costs.

A98.
Answer: B) AWS Batch with managed compute environments
Explanation: AWS Batch handles high-volume, scalable compute jobs without provisioning infrastructure, ideal for data science workloads.

A99.
Answer: C) Warm standby with smaller instance sizes in DR region
Explanation: Warm standby provides rapid recovery (<1 hr RTO) at reduced cost by running scaled-down resources that can be scaled up during failover.

A100.
Answer: D) Apply an SCP that blocks all `s3:DeleteBucket` actions organization-wide
Explanation: SCPs prevent destructive actions at the org level—even admin users in child accounts cannot override them.