



"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

[www.atmicnetworks .com](http://www.atmicnetworks.com)

Warning: Keep connected with our support team
for latest updates

Question: 1

A company's software development team needs an Amazon RDS Multi-AZ cluster. The RDS cluster will serve as a backend for a desktop client that is deployed on premises. The desktop client requires direct connectivity to the RDS cluster.

The company must give the development team the ability to connect to the cluster by using the client when the team is in the office.

Which solution provides the required connectivity MOST securely?

- A. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Use AWS Site- to-Site VPN with a customer gateway in the company's office.
- B. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use AWS Site- to-Site VPN with a customer gateway in the company's office.
- C. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use RDS security groups to allow the company's office IP ranges to access the cluster.
- D. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Create a cluster user for each developer. Use RDS security groups to allow the users to access the cluster.

Answer: B

Explanation:

Requirement Analysis: Need secure, direct connectivity from an on-premises client to an RDS cluster, accessible only when in the office.

VPC with Private Subnets: Ensures the RDS cluster is not publicly accessible, enhancing security.

Site-to-Site VPN: Provides secure, encrypted connection between on-premises office and AWS VPC.

Implementation:

Create a VPC with two private subnets.

Launch the RDS cluster in the private subnets.

Set up a Site-to-Site VPN connection with a customer gateway in the office.

Conclusion: This setup ensures secure and direct connectivity with minimal exposure, meeting the requirement for secure access from the office.

Reference

AWS Site-to-Site VPN: [AWS Site-to-Site VPN Documentation](#)

Question: 2

A social media company wants to store its database of user profiles, relationships, and interactions in the AWS Cloud. The company needs an application to monitor any changes in the database. The application needs to analyze the relationships between the data entities and to provide recommendations to users.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Neptune to store the information. Use Amazon Kinesis Data Streams to process changes in the database.
- B. Use Amazon Neptune to store the information. Use Neptune Streams to process changes in the database.
- C. Use Amazon Quantum Ledger Database (Amazon QLDB) to store the information. Use Amazon Kinesis Data Streams to process changes in the database.
- D. Use Amazon Quantum Ledger Database (Amazon QLDB) to store the information. Use Neptune Streams to process changes in the database.

Answer: B

Explanation:

Amazon Neptune: Neptune is a fully managed graph database service that is optimized for storing and querying highly connected data. It supports both property graph and RDF graph models, making it suitable for applications that need to analyze relationships between data entities.

Neptune Streams: Neptune Streams captures changes to the graph and streams these changes to other AWS services. This is useful for applications that need to monitor and respond to changes in real-time, such as providing recommendations based on user interactions and relationships.

Least Operational Overhead: Using Neptune Streams directly with Amazon Neptune ensures that the solution is tightly integrated, reducing the need for additional components and minimizing operational overhead. This integration simplifies the architecture by eliminating the need for a separate service like Kinesis for change processing.

Reference:

[Amazon Neptune Documentation](#)

[Neptune Streams Documentation](#)

Question: 3

A company uses an Amazon S3 bucket as its data lake storage platform. The S3 bucket contains a massive amount of data that is accessed randomly by multiple teams and hundreds of applications. The company wants to reduce the S3 storage costs and provide immediate availability for frequently accessed objects.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an S3 Lifecycle rule to transition objects to the S3 Intelligent-Tiering storage class.
- B. Store objects in Amazon S3 Glacier. Use S3 Select to provide applications with access to the data.
- C. Use data from S3 storage class analysis to create S3 Lifecycle rules to automatically transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
- D. Transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Create an AWS Lambda function to transition objects to the S3 Standard storage class when they are accessed by an application.

Answer: A

Explanation:

Amazon S3 Intelligent-Tiering: This storage class is designed to optimize costs by automatically moving data between two access tiers (frequent and infrequent) when access patterns change. It provides cost savings without performance impact or operational overhead.

S3 Lifecycle Rules: By creating an S3 Lifecycle rule, the company can automatically transition objects to the Intelligent-Tiering storage class. This eliminates the need for manual intervention and ensures that objects are moved to the most cost-effective storage tier based on their access patterns.

Operational Efficiency: Intelligent-Tiering requires no additional management and delivers immediate availability for frequently accessed objects. This makes it the most operationally efficient solution for the given requirements.

Reference:

Amazon S3 Intelligent-Tiering

S3 Lifecycle Policies

Question: 4

A company needs to optimize its Amazon S3 storage costs for an application that generates many files that cannot be recreated. Each file is approximately 5 MB and is stored in Amazon S3 Standard storage.

The company must store the files for 4 years before the files can be deleted. The files must be immediately accessible. The files are frequently accessed in the first 30 days of object creation, but they are rarely accessed after the first 30 days.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an S3 Lifecycle policy to move the files to S3 Glacier Instant Retrieval 30 days after object creation. Delete the files 4 years after object creation.
- B. Create an S3 Lifecycle policy to move the files to S3 One Zone-Infrequent Access (S3 One Zone-IA) 30 days after object creation. Delete the files 4 years after object creation.
- C. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Delete the files 4 years after object creation.
- D. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Move the files to S3 Glacier Flexible Retrieval 4 years after object creation.

Answer: C

Explanation:

Amazon S3 Standard-IA: This storage class is designed for data that is accessed less frequently but requires rapid access when needed. It offers lower storage costs compared to S3 Standard while still providing high availability and durability.

Access Patterns: Since the files are frequently accessed in the first 30 days and rarely accessed afterward, transitioning them to S3 Standard-IA after 30 days aligns with their access patterns and reduces storage costs significantly.

Lifecycle Policy: Implementing a lifecycle policy to transition the files to S3 Standard-IA ensures automatic management of the data lifecycle, moving files to a lower-cost storage class without manual intervention. Deleting the files after 4 years further optimizes costs by removing data that is no longer needed.

Reference:

Amazon S3 Storage Classes

S3 Lifecycle Configuration

Question: 5

A company runs an AWS Lambda function in private subnets in a VPC. The subnets have a default route to the internet through an Amazon EC2 NAT instance. The Lambda function processes input data and saves its output as an object to Amazon S3.

Intermittently, the Lambda function times out while trying to upload the object because of saturated traffic on the NAT instance's network. The company wants to access Amazon S3 without traversing the internet.

Which solution will meet these requirements?

- A. Replace the EC2 NAT instance with an AWS managed NAT gateway.
- B. Increase the size of the EC2 NAT instance in the VPC to a network optimized instance type

- C. Provision a gateway endpoint for Amazon S3 in the VPC. Update the route tables of the subnets accordingly.
- D. Provision a transit gateway. Place transit gateway attachments in the private subnets where the Lambda function is running.

Answer: C

Explanation:

Gateway Endpoint for Amazon S3: A VPC endpoint for Amazon S3 allows you to privately connect your VPC to Amazon S3 without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.

Provisioning the Endpoint:

Navigate to the VPC Dashboard.

Select "Endpoints" and create a new endpoint.

Choose the service name for S3 (com.amazonaws.region.s3).

Select the appropriate VPC and subnets.

Adjust the route tables of the subnets to include the new endpoint.

Update Route Tables: Modify the route tables of the subnets to direct traffic destined for S3 to the newly created endpoint. This ensures that traffic to S3 does not go through the NAT instance, avoiding the saturated network and eliminating timeouts.

Operational Efficiency: This solution minimizes operational overhead by removing dependency on the NAT instance and avoiding internet traffic, leading to more stable and secure S3 interactions.

Reference:

VPC Endpoints for Amazon S3

Creating a Gateway Endpoint

Question: 6

A company needs to design a hybrid network architecture. The company's workloads are currently stored in the AWS Cloud and in on-premises data centers. The workloads require single-digit latencies to communicate. The company uses an AWS Transit Gateway transit gateway to connect multiple VPCs.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Establish an AWS Site-to-Site VPN connection to each VPC.
- B. Associate an AWS Direct Connect gateway with the transit gateway that is attached to the VPCs.
- C. Establish an AWS Site-to-Site VPN connection to an AWS Direct Connect gateway.

- D. Establish an AWS Direct Connect connection. Create a transit virtual interface (VIF) to a Direct Connect gateway.
- E. Associate AWS Site-to-Site VPN connections with the transit gateway that is attached to the VPCs

Answer: B,D

Explanation:

AWS Direct Connect: Provides a dedicated network connection from your on-premises data center to AWS, ensuring low latency and consistent network performance.

Direct Connect Gateway Association:

Direct Connect Gateway: Acts as a global network transit hub to connect VPCs across different AWS regions.

Association with Transit Gateway: Enables communication between on-premises data centers and multiple VPCs connected to the transit gateway.

Transit Virtual Interface (VIF):

Create Transit VIF: To connect Direct Connect with a transit gateway.

Setup Steps:

Establish a Direct Connect connection.

Create a transit VIF to the Direct Connect gateway.

Associate the Direct Connect gateway with the transit gateway attached to the VPCs.

Cost Efficiency: This combination avoids the recurring costs and potential performance variability of VPN connections, providing a robust, low-latency hybrid network solution.

Reference:

AWS Direct Connect

Transit Gateway and Direct Connect Gateway

Question: 7

A company is running a highly sensitive application on Amazon EC2 backed by an Amazon RDS database. Compliance regulations mandate that all personally identifiable information (PII) be encrypted at rest.

Which solution should a solutions architect recommend to meet this requirement with the LEAST amount of changes to the infrastructure?

- A. Deploy AWS Certificate Manager to generate certificates. Use the certificates to encrypt the database volume.

- B. Deploy AWS CloudHSM. generate encryption keys, and use the keys to encrypt database volumes.
- C. Configure SSL encryption using AWS Key Management Service (AWS KMS) keys to encrypt database volumes.
- D. Configure Amazon Elastic Block Store (Amazon EBS) encryption and Amazon RDS encryption with AWS Key Management Service (AWS KMS) keys to encrypt instance and database volumes.

Answer: D

Explanation:

EBS Encryption:

Default EBS Encryption: Can be enabled for new EBS volumes.

Use of AWS KMS: Specify AWS KMS keys to handle encryption and decryption of data transparently.

Amazon RDS Encryption:

RDS Encryption: Encrypts the underlying storage for RDS instances using AWS KMS.

Configuration: Enable encryption when creating the RDS instance or modify an existing instance to enable encryption.

Least Amount of Changes:

Both EBS and RDS support seamless encryption with AWS KMS, requiring minimal changes to the existing infrastructure.

Enables compliance with regulatory requirements without modifying the application.

Operational Efficiency: Using AWS KMS for both EBS and RDS ensures a consistent, managed approach to encryption, simplifying key management and enhancing security.

Reference:

Amazon EBS Encryption

Amazon RDS Encryption

AWS Key Management Service

Question: 8

A global ecommerce company runs its critical workloads on AWS. The workloads use an Amazon RDS for PostgreSQL DB instance that is configured for a Multi-AZ deployment.

Customers have reported application timeouts when the company undergoes database failovers. The company needs a resilient solution to reduce failover time

Which solution will meet these requirements?

- A. Create an Amazon RDS Proxy. Assign the proxy to the DB instance.
- B. Create a read replica for the DB instance Move the read traffic to the read replica.
- C. Enable Performance Insights. Monitor the CPU load to identify the timeouts.
- D. Take regular automatic snapshots Copy the automatic snapshots to multiple AWS Regions

Answer: A

Explanation:

Amazon RDS Proxy: RDS Proxy is a fully managed, highly available database proxy that makes applications more resilient to database failures by pooling and sharing connections, and it can automatically handle database failovers.

Reduced Failover Time: By using RDS Proxy, the connection management between the application and the database is improved, reducing failover times significantly. RDS Proxy maintains connections in a connection pool and reduces the time required to re-establish connections during a failover.

Configuration:

Create an RDS Proxy instance.

Configure the proxy to connect to the RDS for PostgreSQL DB instance.

Modify the application configuration to use the RDS Proxy endpoint instead of the direct database endpoint.

Operational Benefits: This solution provides high availability and reduces application timeouts during failovers with minimal changes to the application code.

Reference:

Amazon RDS Proxy

Setting Up RDS Proxy

Question: 9

A company is planning to deploy its application on an Amazon Aurora PostgreSQL Serverless v2 cluster. The application will receive large amounts of traffic. The company wants to optimize the storage performance of the cluster as the load on the application increases

Which solution will meet these requirements MOST cost-effectively?

- A. Configure the cluster to use the Aurora Standard storage configuration.
- B. Configure the cluster storage type as Provisioned IOPS.

- C. Configure the cluster storage type as General Purpose.
- D. Configure the cluster to use the Aurora I/O-Optimized storage configuration.

Answer: D

Explanation:

Aurora I/O-Optimized: This storage configuration is designed to provide consistent high performance for Aurora databases. It automatically scales IOPS as the workload increases, without needing to provision IOPS separately.

Cost-Effectiveness: With Aurora I/O-Optimized, you only pay for the storage and I/O you use, making it a cost-effective solution for applications with varying and unpredictable I/O demands.

Implementation:

During the creation of the Aurora PostgreSQL Serverless v2 cluster, select the I/O-Optimized storage configuration.

The storage system will automatically handle scaling and performance optimization based on the application load.

Operational Efficiency: This configuration reduces the need for manual tuning and ensures optimal performance without additional administrative overhead.

Reference:

Amazon Aurora I/O-Optimized

Question: 10

A company is preparing to store confidential data in Amazon S3. For compliance reasons, the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys must be rotated every year.

Which solution meets these requirements and is the MOST operationally efficient?

- A. Server-side encryption with customer-provided keys (SSE-C)
- B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- C. Server-side encryption with AWS KMS keys (SSE-KMS) with manual rotation
- D. Server-side encryption with AWS KMS keys (SSE-KMS) with automatic rotation

Answer: D

Explanation:

SSE-KMS: Server-side encryption with AWS Key Management Service (SSE-KMS) provides robust encryption of data at rest, integrated with AWS KMS for key management and auditing.

Automatic Key Rotation: By enabling automatic rotation for the KMS keys, the system ensures that keys are rotated annually without manual intervention, meeting compliance requirements.

Logging and Auditing: AWS KMS automatically logs all key usage and management actions in AWS CloudTrail, providing the necessary audit logs.

Implementation:

Create a KMS key with automatic rotation enabled.

Configure the S3 bucket to use SSE-KMS with the created KMS key.

Ensure CloudTrail is enabled for logging KMS key usage.

Operational Efficiency: This solution provides encryption, automatic key management, and auditing in a seamless, fully managed way, reducing operational overhead.

Reference:

AWS KMS Automatic Key Rotation

Amazon S3 Server-Side Encryption

Question: 11

A company has several on-premises Internet Small Computer Systems Interface (iSCSI) network storage servers. The company wants to reduce the number of these servers by moving to the AWS Cloud. A solutions architect must provide low-latency access to frequently used data and reduce the dependency on on-premises servers with a minimal number of infrastructure changes.

Which solution will meet these requirements?

- A. Deploy an Amazon S3 File Gateway
- B. Deploy Amazon Elastic Block Store (Amazon EBS) storage with backups to Amazon S3
- C. Deploy an AWS Storage Gateway volume gateway that is configured with stored volumes
- D. Deploy an AWS Storage Gateway volume gateway that is configured with cached volumes.

Answer: D

Explanation:

Storage Gateway Volume Gateway (Cached Volumes): This configuration allows you to store your primary data in Amazon S3 while retaining frequently accessed data locally in a cache for low-latency access.

Low-Latency Access: Frequently accessed data is cached locally on-premises, providing low-latency access while the less frequently accessed data is stored cost-effectively in Amazon S3.

Implementation:

Deploy a Storage Gateway appliance on-premises or in a virtual environment.

Configure it as a volume gateway with cached volumes.

Create volumes and configure your applications to use these volumes.

Minimal Infrastructure Changes: This solution integrates seamlessly with existing on-premises infrastructure, requiring minimal changes and reducing dependency on on-premises storage servers.

Reference:

AWS Storage Gateway Volume Gateway

Volume Gateway Cached Volumes

Question: 12

A marketing company receives a large amount of new clickstream data in Amazon S3 from a marketing campaign. The company needs to analyze the clickstream data in Amazon S3 quickly. Then the company needs to determine whether to process the data further in the data pipeline.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create external tables in a Spark catalog. Configure jobs in AWS Glue to query the data.
- B. Configure an AWS Glue crawler to crawl the data. Configure Amazon Athena to query the data.
- C. Create external tables in a Hive metastore. Configure Spark jobs in Amazon EMR to query the data.
- D. Configure an AWS Glue crawler to crawl the data. Configure Amazon Kinesis Data Analytics to use SQL to query the data.

Answer: B

Explanation:

AWS Glue Crawler: AWS Glue is a fully managed ETL (Extract, Transform, Load) service that makes it easy to prepare and load data for analytics. A Glue crawler can automatically discover new data and schema in Amazon S3, making it easy to keep the data catalog up-to-date.

Crawling the Data:

Set up an AWS Glue crawler to scan the S3 bucket containing the clickstream data.

The crawler will automatically detect the schema and create/update the tables in the AWS Glue Data Catalog.

Amazon Athena:

Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL.

Once the data catalog is updated by the Glue crawler, use Athena to query the clickstream data directly in S3.

Operational Efficiency: This solution leverages fully managed services, reducing operational overhead. Glue crawlers automate data cataloging, and Athena provides a serverless, pay-per-query model for quick data analysis without the need to set up or manage infrastructure.

Reference:

AWS Glue

Amazon Athena

Question: 13

A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects. According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.

Which solution will meet these requirements?

- A. Configure an S3 gateway endpoint.
- B. Create an S3 bucket in a private subnet.
- C. Create an S3 bucket in the same AWS Region as the EC2 instances.
- D. Configure a NAT gateway in the same subnet as the EC2 instances.

Answer: A

Explanation:

VPC Endpoint for S3: A gateway endpoint for Amazon S3 enables you to privately connect your VPC to S3 without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.

Configuration Steps:

In the VPC console, navigate to "Endpoints" and create a new endpoint.

Select the service name for S3 (com.amazonaws.region.s3).

Choose the VPC and the subnets where your EC2 instances are running.

Update the route tables for the selected subnets to include a route pointing to the endpoint.

Security Compliance: By configuring an S3 gateway endpoint, all traffic between the VPC and S3 stays within the AWS network, complying with the company's security regulations to avoid internet traversal.

Reference:

VPC Endpoints for Amazon S3

Question: 14

A company wants to isolate its workloads by creating an AWS account for each workload. The company needs a solution that centrally manages networking components for the workloads. The solution also must create accounts with automatic security controls (guardrails).

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Control Tower to deploy accounts. Create a networking account that has a VPC with private subnets and public subnets. Use AWS Resource Access Manager (AWS RAM) to share the subnets with the workload accounts.
- B. Use AWS Organizations to deploy accounts. Create a networking account that has a VPC with private subnets and public subnets. Use AWS Resource Access Manager (AWS RAM) to share the subnets with the workload accounts.
- C. Use AWS Control Tower to deploy accounts. Deploy a VPC in each workload account. Configure each VPC to route through an inspection VPC by using a transit gateway attachment.
- D. Use AWS Organizations to deploy accounts. Deploy a VPC in each workload account. Configure each VPC to route through an inspection VPC by using a transit gateway attachment.

Answer: A

Explanation:

AWS Control Tower: Provides a managed service to set up and govern a secure, multi-account AWS environment based on AWS best practices. It automates the setup of AWS Organizations and applies security controls (guardrails).

Networking Account:

Create a centralized networking account that includes a VPC with both private and public subnets.

This centralized VPC will manage and control the networking resources.

AWS Resource Access Manager (AWS RAM):

Use AWS RAM to share the subnets from the networking account with the other workload accounts.

This allows different workload accounts to utilize the shared networking resources without the need to manage their own VPCs.

Operational Efficiency: Using AWS Control Tower simplifies the setup and governance of multiple

AWS accounts, while AWS RAM facilitates centralized management of networking resources, reducing operational overhead and ensuring consistent security and compliance.

Reference:

AWS Control Tower

AWS Resource Access Manager

Question: 15

A company's SAP application has a backend SQL Server database in an on-premises environment. The company wants to migrate its on-premises application and database server to AWS. The company needs an instance type that meets the high demands of its SAP database. On-premises performance data shows that both the SAP application and the database have high memory utilization.

Which solution will meet these requirements?

- A. Use the compute optimized Instance family for the application Use the memory optimized instance family for the database.
- B. Use the storage optimized instance family for both the application and the database
- C. Use the memory optimized instance family for both the application and the database
- D. Use the high performance computing (HPC) optimized instance family for the application. Use the memory optimized instance family for the database.

Answer: C

Explanation:

Memory Optimized Instances: These instances are designed to deliver fast performance for workloads that process large data sets in memory. They are ideal for high-performance databases like SAP and applications with high memory utilization.

High Memory Utilization: Both the SAP application and the SQL Server database have high memory demands as per the on-premises performance data. Memory optimized instances provide the necessary memory capacity and performance.

Instance Types:

For the SAP application, using a memory optimized instance ensures the application has sufficient memory to handle the high workload efficiently.

For the SQL Server database, memory optimized instances ensure optimal database performance with high memory throughput.

Operational Efficiency: Using the same instance family for both the application and the database simplifies management and ensures both components meet performance requirements.

Reference:

Amazon EC2 Instance Types

SAP on AWS

Question: 16

A company plans to rehost an application to Amazon EC2 instances that use Amazon Elastic Block Store (Amazon EBS) as the attached storage

A solutions architect must design a solution to ensure that all newly created Amazon EBS volumes are encrypted by default. The solution must also prevent the creation of unencrypted EBS volumes

Which solution will meet these requirements?

- A. Configure the EC2 account attributes to always encrypt new EBS volumes.
- B. Use AWS Config. Configure the encrypted-volumes identifier Apply the default AWS Key Management Service (AWS KMS) key.
- C. Configure AWS Systems Manager to create encrypted copies of the EBS volumes. Reconfigure the EC2 instances to use the encrypted volumes
- D. Create a customer managed key in AWS Key Management Service (AWS KMS) Configure AWS Migration Hub to use the key when the company migrates workloads.

Answer: A

Explanation:

EC2 Account Attributes: Amazon EC2 allows you to set account attributes to automatically encrypt new EBS volumes. This ensures that all new volumes created in your account are encrypted by default.

Configuration Steps:

Go to the EC2 Dashboard.

Select "Account Attributes" and then "EBS encryption".

Enable default EBS encryption and select the default AWS KMS key or a customer-managed key.

Prevention of Unencrypted Volumes: By setting this account attribute, you ensure that it is not possible to create unencrypted EBS volumes, thereby enforcing compliance with security requirements.

Operational Efficiency: This solution requires minimal configuration changes and provides automatic enforcement of encryption policies,

reducing operational overhead.

Reference:

Amazon EC2 Default EBS Encryption

Question: 17

A weather forecasting company needs to process hundreds of gigabytes of data with sub-millisecond latency. The company has a high performance computing (HPC) environment in its data center and wants to expand its forecasting capabilities.

A solutions architect must identify a highly available cloud storage solution that can handle large amounts of sustained throughput. Files that are stored in the solution should be accessible to thousands of compute instances that will simultaneously access and process the entire dataset.

What should the solutions architect do to meet these requirements?

- A. Use Amazon FSx for Lustre scratch file systems
- B. Use Amazon FSx for Lustre persistent file systems.
- C. Use Amazon Elastic File System (Amazon EFS) with Bursting Throughput mode.
- D. Use Amazon Elastic File System (Amazon EFS) with Provisioned Throughput mode.

Answer: B

Explanation:

Amazon FSx for Lustre: Lustre is a high-performance file system designed for workloads that require fast storage with sustained high throughput and low latency. It integrates with Amazon S3, making it suitable for HPC environments.

Persistent File Systems:

Persistent Storage: Suitable for long-term storage and recurrent use, providing durability and availability.

High Throughput and Low Latency: Persistent Lustre file systems can handle large amounts of data with sub-millisecond latency, meeting the needs of high-performance computing workloads.

Simultaneous Access: FSx for Lustre allows thousands of compute instances to access and process large datasets concurrently, ensuring that the high volume of data is handled efficiently.

Highly Available: FSx for Lustre is designed to provide high availability and is managed by AWS, reducing the operational burden.

Reference:

Question: 18

A company plans to run a high performance computing (HPC) workload on Amazon EC2 Instances. The workload requires low-latency network performance and high network throughput with tightly coupled node-to-node communication.

Which solution will meet these requirements?

- A. Configure the EC2 instances to be part of a cluster placement group
- B. Launch the EC2 instances with Dedicated Instance tenancy.
- C. Launch the EC2 instances as Spot Instances.
- D. Configure an On-Demand Capacity Reservation when the EC2 instances are launched.

Answer: A

Explanation:

Cluster Placement Group: This type of placement group is designed to provide low-latency network performance and high throughput by grouping instances within a single Availability Zone. It is ideal for applications that require tightly coupled node-to-node communication.

Configuration:

When launching EC2 instances, specify the option to launch them in a cluster placement group.

This ensures that the instances are physically located close to each other, reducing latency and increasing network throughput.

Benefits:

Low-Latency Communication: Instances in a cluster placement group benefit from enhanced networking capabilities, enabling low-latency communication.

High Network Throughput: The network performance within a cluster placement group is optimized for high throughput, which is essential for HPC workloads.

Reference:

Placement Groups

High Performance Computing on AWS

Question: 19

A company needs a secure connection between its on-premises environment and AWS. This connection does not need high bandwidth and will handle a small amount of traffic. The connection should be set up quickly.

What is the MOST cost-effective method to establish this type of connection?

- A. Implement a client VPN
- B. Implement AWS Direct Connect.
- C. Implement a bastion host on Amazon EC2.
- D. Implement an AWS Site-to-Site VPN connection.

Answer: D

Explanation:

AWS Site-to-Site VPN: This provides a secure and encrypted connection between an on-premises environment and AWS. It is a cost-effective solution suitable for low bandwidth and small traffic needs.

Quick Setup:

Site-to-Site VPN can be quickly set up by configuring a virtual private gateway on the AWS side and a customer gateway on the on-premises side.

It uses standard IPsec protocol to establish the VPN tunnel.

Cost-Effectiveness: Compared to AWS Direct Connect, which requires dedicated physical connections and higher setup costs, a Site-to-Site VPN is less expensive and easier to implement for smaller traffic requirements.

Reference:

AWS Site-to-Site VPN

Question: 20

A company is designing a new multi-tier web application that consists of the following components:

- Web and application servers that run on Amazon EC2 instances as part of Auto Scaling groups
- An Amazon RDS DB instance for data storage

A solutions architect needs to limit access to the application servers so that only the web servers can access them. Which solution will meet these requirements?

- A. Deploy AWS PrivateLink in front of the application servers. Configure the network ACL to allow only the web servers to access the application servers.
- B. Deploy a VPC endpoint in front of the application servers. Configure the security group to allow only the web servers to access the application servers.
- C. Deploy a Network Load Balancer with a target group that contains the application servers' Auto Scaling group. Configure the network ACL to allow only the web servers to access the application servers.
- D. Deploy an Application Load Balancer with a target group that contains the application servers' Auto Scaling group. Configure the security group to allow only the web servers to access the application servers.

Answer: D

Explanation:

Application Load Balancer (ALB): ALB is suitable for routing HTTP/HTTPS traffic to the application servers. It provides advanced routing features and integrates well with Auto Scaling groups.

Target Group Configuration:

Create a target group for the application servers and register the Auto Scaling group with this target group.

Configure the ALB to forward requests from the web servers to the application servers.

Security Group Setup:

Configure the security group of the application servers to only allow traffic from the web servers' security group.

This ensures that only the web servers can access the application servers, meeting the requirement to limit access.

Benefits:

Security: Using security groups to restrict access ensures a secure environment where only intended traffic is allowed.

Scalability: ALB works seamlessly with Auto Scaling groups, ensuring the application can handle varying loads efficiently.

Reference:

Application Load Balancer

Security Groups for Your VPC

Question: 21

A company has separate AWS accounts for its finance, data analytics, and development departments. Because of costs and security

concerns, the company wants to control which services each AWS account can use

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Systems Manager templates to control which AWS services each department can use
- B. Create organization units (OUs) for each department in AWS Organizations. Attach service control policies (SCPs) to the OUs.
- C. Use AWS CloudFormation to automatically provision only the AWS services that each department can use.
- D. Set up a list of products in AWS Service Catalog in the AWS accounts to manage and control the usage of specific AWS services

Answer: B

Explanation:

AWS Organizations: AWS Organizations allows you to create multiple AWS accounts and manage them centrally. You can organize accounts into organizational units (OUs) and apply policies to these units.

Organizational Units (OUs):

Create separate OUs for each department: finance, data analytics, and development.

Place the respective AWS accounts for each department into their corresponding OUs.

Service Control Policies (SCPs):

SCPs are policies that can restrict which AWS services and actions are available to accounts in an OU.

Create SCPs to define which services each department can use and attach these policies to the appropriate OUs.

SCPs apply to all IAM users, groups, and roles within the accounts in the OU, providing centralized control over service usage.

Operational Efficiency: Using AWS Organizations and SCPs provides a scalable and centralized way to manage permissions across multiple accounts with minimal operational overhead.

Reference:

AWS Organizations

Service Control Policies

Question: 22

A company has a web application that has thousands of users. The application uses 8-10 user-uploaded images to generate AI images.

Users can download the generated AI Images once every 6 hours. The company also has a premium user option that gives users the ability to download the generated AI images anytime

The company uses the user-uploaded images to run AI model training twice a year. The company needs a storage solution to store

the images.

Which storage solution meets these requirements MOST cost-effectively?

- A. Move uploaded images to Amazon S3 Glacier Deep Archive. Move premium user-generated AI images to S3 Standard. Move non-premium user-generated AI images to S3 Standard-Infrequent Access (S3 Standard-IA).
- B. Move uploaded images to Amazon S3 Glacier Deep Archive. Move all generated AI images to S3 Glacier Flexible Retrieval.
- C. Move uploaded images to Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Move premium user-generated AI images to S3 Standard. Move non-premium user-generated AI images to S3 Standard-Infrequent Access (S3 Standard-IA).
- D. Move uploaded images to Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Move all generated AI images to S3 Glacier Flexible Retrieval.

Answer: C

Explanation:

S3 One Zone-IA:

Suitable for infrequently accessed data that doesn't require multiple Availability Zone resilience.

Cost-effective for storing user-uploaded images that are only used for AI model training twice a year.

S3 Standard:

Ideal for frequently accessed data with high durability and availability.

Store premium user-generated AI images here to ensure they are readily available for download at any time.

S3 Standard-IA:

Cost-effective storage for data that is accessed less frequently but still requires rapid retrieval.

Store non-premium user-generated AI images here, as these images are only downloaded once every 6 hours, making it a good balance between cost and accessibility.

Cost-Effectiveness: This solution optimizes storage costs by categorizing data based on access patterns and durability requirements, ensuring that each type of data is stored in the most cost-effective manner.

Reference:

Amazon S3 Storage Classes

S3 One Zone-IA

Question: 23

A global company runs its workloads on AWS. The company's application uses Amazon S3 buckets across AWS Regions for sensitive data storage and analysis. The company stores millions of objects in multiple S3 buckets daily. The company wants to identify all S3 buckets that are not versioning-enabled.

Which solution will meet these requirements?

- A. Set up an AWS CloudTrail event that has a rule to identify all S3 buckets that are not versioning-enabled across Regions
- B. Use Amazon S3 Storage Lens to identify all S3 buckets that are not versioning-enabled across Regions.
- C. Enable IAM Access Analyzer for S3 to identify all S3 buckets that are not versioning-enabled across Regions
- D. Create an S3 Multi-Region Access Point to identify all S3 buckets that are not versioning-enabled across Regions

Answer: B

Explanation:

Amazon S3 Storage Lens:

S3 Storage Lens provides organization-wide visibility into object storage usage and activity trends.

It can generate metrics and insights about your S3 buckets, including versioning status.

Configuration:

Enable S3 Storage Lens at the organization level.

Configure the dashboard to include the versioning status metric.

Identify Non-Versioned Buckets:

Use the S3 Storage Lens dashboard to filter and identify buckets that do not have versioning enabled.

Storage Lens provides detailed insights and reports which can be used to enforce compliance and manage storage effectively.

Operational Efficiency: Using S3 Storage Lens provides a centralized, easy-to-use interface for monitoring bucket configurations across multiple Regions and accounts, reducing the need for custom scripts or manual checks.

Reference:

Amazon S3 Storage Lens

S3 Storage Lens Metrics

Question: 24

A media company uses an Amazon CloudFront distribution to deliver content over the internet. The company wants only premium customers to have access to the media streams and file content. The company stores all content in an Amazon S3 bucket. The company also delivers content on demand to customers for a specific purpose, such as movie rentals or music downloads.

Which solution will meet these requirements?

- A. Generate and provide S3 signed cookies to premium customers
- B. Generate and provide CloudFront signed URLs to premium customers.
- C. Use origin access control (OAC) to limit the access of non-premium customers
- D. Generate and activate field-level encryption to block non-premium customers.

Answer: B

Explanation:

CloudFront Signed URLs: These URLs allow you to provide limited access to content that is being served through an Amazon CloudFront distribution. Signed URLs can be generated to grant **timelimited access** to premium customers.

Content Restriction:

By using CloudFront signed URLs, you can control access to your media streams and file content stored in S3.

These URLs can be customized with an expiration time, ensuring that access is only available for a specific period, which is useful for scenarios like movie rentals or music downloads.

Security and Flexibility:

Signed URLs ensure that only authenticated users (premium customers) can access the restricted content.

This approach integrates seamlessly with CloudFront and S3, providing an efficient way to manage access controls without additional overhead.

Operational Efficiency: Using CloudFront signed URLs leverages AWS managed services to handle the complexity of access control, reducing the need for custom implementation and maintenance.

Reference:

Serving Private Content with Signed URLs and Signed Cookies

Question: 25

A company uses Amazon API Gateway to manage its REST APIs that third-party service providers access. The company must protect the REST APIs from SQL injection and cross-site scripting attacks.

What is the MOST operationally efficient solution that meets these requirements?

- A. Configure AWS Shield.
- B. Configure AWS WAF.
- C. Set up API Gateway with an Amazon CloudFront distribution. Configure AWS Shield in CloudFront.
- D. Set up API Gateway with an Amazon CloudFront distribution. Configure AWS WAF in CloudFront.

Answer: D

Explanation:

Amazon API Gateway with CloudFront: API Gateway allows you to create, deploy, and manage APIs, while CloudFront provides a CDN to deliver content with low latency and high transfer speeds.

AWS WAF (Web Application Firewall):

AWS WAF can be configured in CloudFront to protect against common web exploits, including SQL injection and cross-site scripting (XSS).

WAF allows you to create custom rules to block specific attack patterns and can be managed centrally.

Configuration:

Deploy your APIs using Amazon API Gateway.

Set up an Amazon CloudFront distribution in front of the API Gateway.

Configure AWS WAF on the CloudFront distribution to apply security rules.

Operational Efficiency: This solution provides robust protection with minimal operational overhead by leveraging managed AWS services, ensuring that your APIs are secure without extensive custom implementation.

Reference:

Using AWS WAF to Protect Your APIs

How CloudFront Works with AWS WAF

Question: 26

A company runs its application on Oracle Database Enterprise Edition. The company needs to migrate the application and the database to AWS. The company can use the Bring Your Own License (BYOL) model while migrating to AWS. The application uses third-party database features that require privileged access.

A solutions architect must design a solution for the database migration.

Which solution will meet these requirements MOST cost-effectively?

- A. Migrate the database to Amazon RDS for Oracle by using native tools. Replace the third-party features with AWS Lambda.
- B. Migrate the database to Amazon RDS Custom for Oracle by using native tools. Customize the new database settings to support the third-party features.
- C. Migrate the database to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS). Customize the new database settings to support the third-party features.
- D. Migrate the database to Amazon RDS for PostgreSQL by using AWS Database Migration Service (AWS DMS). Rewrite the application code to remove the dependency on third-party features.

Answer: B

Explanation:

Amazon RDS Custom for Oracle: This service allows you to bring your own Oracle Database licenses and provides the flexibility to customize the database settings, making it suitable for applications that require privileged access and third-party database features.

BYOL (Bring Your Own License):

RDS Custom supports the BYOL model, allowing you to use your existing Oracle licenses and comply with licensing requirements.

This helps in leveraging existing investments and reducing migration costs.

Customization and Third-Party Features:

RDS Custom allows for deeper customization of the database environment compared to standard RDS instances.

This makes it possible to support the third-party features that your application relies on without significant changes.

Migration Process:

Use native Oracle tools like Data Pump or RMAN to migrate the database to RDS Custom.

Customize the database settings post-migration to ensure compatibility with third-party features. Reference:

Amazon RDS Custom for Oracle

Question: 27

A news company that has reporters all over the world is hosting its broadcast system on AWS. The reporters send live broadcasts to the broadcast system. The reporters use software on their phones to send live streams through the Real Time Messaging Protocol (RTMP).

A solutions architect must design a solution that gives the reporters the ability to send the highest quality streams. The solution must provide accelerated TCP connections back to the broadcast system.

What should the solutions architect use to meet these requirements?

- A. Amazon CloudFront
- B. AWS Global Accelerator
- C. AWS Client VPN
- D. Amazon EC2 instances and AWS Elastic IP addresses

Answer: B

Explanation:

AWS Global Accelerator: This service provides a global fixed entry point to your applications and optimizes the path to your application through the AWS global network, reducing latency and improving performance.

Accelerated TCP Connections:

Global Accelerator uses the AWS global network to route traffic to the nearest edge location, improving the performance and reliability of your live streams.

It provides static IP addresses that act as a fixed entry point to your application, simplifying DNS management.

High-Quality Streams:

By leveraging Global Accelerator, reporters can send live streams with the highest quality and low latency.

This service automatically reroutes traffic to the nearest available AWS Region, ensuring consistent performance even during traffic spikes or failures.

Operational Efficiency: Using Global Accelerator simplifies the network setup and provides an optimized path for live streams without the need for complex configurations, making it an efficient solution for real-time streaming applications.

Reference:

Question: 28

A company serves its website by using an Auto Scaling group of Amazon EC2 instances in a single AWS Region. The website does not require a database

The company is expanding, and the company's engineering team deploys the website to a second Region. The company wants to distribute traffic across both Regions to accommodate growth and for disaster recovery purposes The solution should not serve traffic from a Region in which the website is **unhealthy**.

Which policy or resource should the company use to meet these requirements?

- A. An Amazon Route 53 simple routing policy
- B. An Amazon Route 53 multivalue answer routing policy
- C. An Application Load Balancer in one Region with a target group that specifies the EC2 instance IDs from both Regions
- D. An Application Load Balancer in one Region with a target group that specifies the IP addresses of the EC2 instances from both Regions

Answer: B

Explanation:

Amazon Route 53 Multivalue Answer Routing: This routing policy allows Route 53 to return multiple values, such as IP addresses, in response to DNS queries. This can distribute traffic across multiple resources and includes health checks to ensure traffic is only routed to healthy instances.

Health Checks:

Configure health checks for each Region to monitor the health of the website instances.

Route 53 will only include healthy instances in the DNS responses, ensuring that traffic is not routed to an unhealthy Region.

Load Distribution and Disaster Recovery:

Multivalue answer routing helps balance the load between instances in different Regions.

If instances in one Region become unhealthy, Route 53 will route traffic to the healthy instances in the other Region.

Operational Simplicity: This solution does not require complex configurations or additional resources, making it a simple yet effective way to distribute traffic and ensure high availability.

Reference:

Amazon Route 53 Routing Policies

Multivalued Answer Routing

Question: 29

A company has an on-premises SFTP file transfer solution. The company is migrating to the AWS Cloud to scale the file transfer solution and to optimize costs by using Amazon S3. The company's employees will use their credentials for the on-premises Microsoft Active Directory (AD) to access the new solution. The company wants to keep the current authentication and file access mechanisms.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure an S3 File Gateway. Create SMB file shares on the file gateway that use the existing Active Directory to authenticate.
- B. Configure an Auto Scaling group with Amazon EC2 instances to run an SFTP solution. Configure the group to scale up at 60% CPU utilization.
- C. Create an AWS Transfer Family server with SFTP endpoints. Choose the AWS Directory Service option as the identity provider. Use AD Connector to connect the on-premises Active Directory.
- D. Create an AWS Transfer Family SFTP endpoint. Configure the endpoint to use the AWS Directory Service option as the identity provider to connect to the existing Active Directory.

Answer: C

Explanation:

AWS Transfer Family: This service provides fully managed support for file transfers directly into and out of Amazon S3 using the SFTP, FTPS, and FTP protocols.

SFTP Endpoints:

Set up an AWS Transfer Family server and configure SFTP endpoints to handle the file transfers.

This service is scalable and managed, reducing operational overhead compared to running an SFTP solution on EC2 instances.

Integration with Active Directory:

Choose the AWS Directory Service option as the identity provider for the Transfer Family server.

Use AD Connector to link the on-premises Active Directory with AWS, allowing employees to use their existing AD credentials to access the SFTP service.

Operational Efficiency: This solution leverages managed services for both file transfer and identity management, ensuring minimal changes to the current authentication mechanisms and reducing operational overhead.

Reference:

AWS Transfer Family

AWS Directory Service and AD Connector

Question: 30

A company has an application that serves clients that are deployed in more than 20,000 retail storefront locations around the world. The application consists of backend web services that are exposed over HTTPS on port 443. The application is hosted on Amazon EC2 Instances behind an Application Load Balancer (ALB). The retail locations communicate with the web application over the public internet. The company allows each retail location to register the IP address that the retail location has been allocated by its local ISP.

The company's security team recommends to increase the security of the application endpoint by restricting access to only the IP addresses registered by the retail locations.

What should a solutions architect do to meet these requirements?

- A. Associate an AWS WAF web ACL with the ALB. Use IP rule sets on the ALB to filter traffic. Update the IP addresses in the rule to include the registered IP addresses.
- B. Deploy AWS Firewall Manager to manage the ALB. Configure firewall rules to restrict traffic to the ALB. Modify the firewall rules to include the registered IP addresses.
- C. Store the IP addresses in an Amazon DynamoDB table. Configure an AWS Lambda authorization function on the ALB to validate that incoming requests are from the registered IP addresses.
- D. Configure the network ACL on the subnet that contains the public interface of the ALB. Update the ingress rules on the network ACL with entries for each of the registered IP addresses.

Answer: A

Explanation:

AWS WAF (Web Application Firewall): AWS WAF allows you to create custom rules to block or allow web requests based on conditions that you specify.

Web ACL (Access Control List):

Create a web ACL and associate it with the ALB.

Use IP rule sets to specify the IP addresses of the retail locations that are allowed to access the application.

Security and Flexibility:

AWS WAF provides a scalable way to manage access control, ensuring that only traffic from registered IP addresses is allowed.

You can dynamically update the IP rule sets to add or remove IP addresses as needed.

Operational Simplicity: Using AWS WAF with a web ACL is straightforward and integrates seamlessly with the ALB, providing an efficient solution for managing access control based on IP addresses.

Reference:

AWS WAF

How AWS WAF Works

Question: 31

A company hosts its application on several Amazon EC2 instances inside a VPC. The company creates a dedicated Amazon S3 bucket for each customer to store their relevant information in Amazon S3.

The company wants to ensure that the application running on EC2 instances can securely access only the S3 buckets that belong to the company's AWS account.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a gateway endpoint for Amazon S3 that is attached to the VPC Update the IAM instance profile policy to provide access to only the specific buckets that the application needs.
- B. Create a NAT gateway in a public subnet with a security group that allows access to only Amazon S3 Update the route tables to use the NAT Gateway.
- C. Create a gateway endpoint for Amazon S3 that is attached to the VPC Update the IAM instance profile policy with a Deny action and the following condition key:
- D. Create a NAT Gateway in a public subnet Update route tables to use the NAT Gateway Assign bucket policies for all buckets with a Deny action and the following condition key:

Answer: A

Explanation:

Question: 32

A company stores user data in AWS. The data is used continuously with peak usage during business hours. Access patterns vary, with some

data not being used for months at a time. A solutions architect must choose a cost-effective solution that maintains the highest level of durability while maintaining high availability.

Which storage solution meets these requirements?

- A. Amazon S3 Standard
- B. Amazon S3 Intelligent-Tiering
- C. Amazon S3 Glacier Deep Archive
- D. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B

Explanation:

Amazon S3 Intelligent-Tiering is the most cost-effective solution for this scenario, providing both high availability and durability while adjusting automatically to changing access patterns. It moves data across two access tiers: one optimized for frequent access and another for infrequent access, based on usage patterns. This tiering ensures that the company avoids paying for unused storage while also keeping frequently accessed data in a more accessible tier.

Key AWS references and benefits of S3 Intelligent-Tiering:

High Durability and Availability: Amazon S3 offers 99.999999999% durability and 99.9% availability for objects stored, ensuring data is always protected.

Automatic Tiering: Data is automatically moved between tiers based on access patterns, making it ideal for workloads with unpredictable or variable access patterns.

No Retrieval Fees: Unlike S3 One Zone-IA or Glacier, there are no retrieval fees, making this more cost-effective in scenarios where access patterns vary over time.

AWS Documentation: According to the AWS Well-Architected Framework under the Cost Optimization Pillar, S3 Intelligent-Tiering is recommended for storage when access patterns change over time, as it minimizes costs while maintaining availability.

Question: 33

A medical company wants to perform transformations on a large amount of clinical trial data that comes from several customers. The company must extract the data from a relational database that contains the customer data. Then the company will transform the data by using a series of complex rules. The company will load the data to Amazon S3 when the transformations are complete.

All data must be encrypted where it is processed before the company stores the data in Amazon S3. All data must be encrypted by using customer-specific keys.

Which solution will meet these requirements with the LEAST amount of operational effort?

- A. Create one AWS Glue job for each customer Attach a security configuration to each job that uses server-side encryption with Amazon S3 managed keys (SSE-S3) to encrypt the data.
- B. Create one Amazon EMR cluster for each customer Attach a security configuration to each cluster that uses client-side encryption with a custom client-side root key (CSE-Custom) to encrypt the data.
- C. Create one AWS Glue job for each customer Attach a security configuration to each job that uses client-side encryption with AWS KMS managed keys (CSE-KMS) to encrypt the data.
- D. Create one Amazon EMR cluster for each customer Attach a security configuration to each cluster that uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt the data.

Answer: C

Explanation:

AWS Glue jobs are designed for extract, transform, and load (ETL) operations, which are perfect for transforming clinical trial data. AWS Glue integrates with AWS Key Management Service (KMS), allowing for customer-specific encryption keys, fulfilling the encryption requirement with minimal operational effort. Client-side encryption with AWS KMS ensures that the data is encrypted before it is sent to S3, aligning with the security needs specified in the scenario.

Key aspects:

AWS Glue: This managed ETL service simplifies data transformation, reduces operational overhead, and integrates seamlessly with KMS.

CSE-KMS: Client-side encryption with KMS ensures that the data is encrypted with customer-specific keys before it is processed or stored in S3, offering robust security.

Minimal Operational Overhead: Compared to managing an EMR cluster, AWS Glue automates much of the process, making it a lower-effort solution.

AWS Documentation: According to the AWS Well-Architected Framework, encryption with AWS KMS offers strong security controls that meet the needs of industries requiring high levels of confidentiality.

Question: 34 A company uses an Amazon DynamoDB table to store data that the company receives from devices.

The DynamoDB table supports a customer-facing website to display recent activity on customer devices. The company configured the table with provisioned throughput for writes and reads.

The company wants to calculate performance metrics for customer device data on a daily basis. The solution must have minimal effect on the table's provisioned read and write capacity.

Which solution will meet these requirements?

- A. Use an Amazon Athena SQL query with the Amazon Athena DynamoDB connector to calculate performance metrics on a recurring schedule.
- B. Use an AWS Glue job with the AWS Glue DynamoDB export connector to calculate performance metrics on a recurring schedule.
- C. Use an Amazon Redshift COPY command to calculate performance metrics on a recurring schedule.
- D. Use an Amazon EMR job with an Apache Hive external table to calculate performance metrics on a recurring schedule.

Answer: A

Explanation:

Amazon Athena provides a cost-effective, serverless way to query data without affecting the performance of DynamoDB. By using the Athena DynamoDB connector, the company can perform the necessary SQL queries without consuming read capacity on the DynamoDB table, which is essential for minimizing impact on provisioned throughput.

Key benefits:

Minimal Impact on Provisioned Capacity: Athena queries do not directly impact DynamoDB's read capacity, making it ideal for running analytics without affecting the customer-facing workloads.

Cost-Effective: Athena is a serverless solution, meaning you pay only for the queries you run, making it highly cost-effective compared to running a dedicated cluster like Amazon EMR or Redshift.

AWS Documentation: The use of Athena to query DynamoDB through its connector aligns with AWS's best practices for performance efficiency and cost optimization.

Question: 35

An ecommerce company runs several internal applications in multiple AWS accounts. The company uses AWS Organizations to manage its AWS accounts.

A security appliance in the company's networking account must inspect interactions between applications across AWS accounts.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) in the networking account to send traffic to the security appliance. Configure the application accounts to send traffic to the NLB by using an interface VPC endpoint in the application accounts
- B. Deploy an Application Load Balancer (ALB) in the application accounts to send traffic directly to the security appliance.
- C. Deploy a Gateway Load Balancer (GWLB) in the networking account to send traffic to the security appliance. Configure the application accounts to send traffic to the GWLB by using an interface GWLB endpoint in the application accounts

D. Deploy an interface VPC endpoint in the application accounts to send traffic directly to the security appliance.

Answer: C

Explanation:

The Gateway Load Balancer (GWLB) is specifically designed to route traffic through a security appliance in a hub-and-spoke model, making it the ideal solution for inspecting traffic between multiple AWS accounts. GWLB enables you to simplify, scale, and deploy third-party virtual appliances transparently, and it can work across multiple VPCs or accounts using interface endpoints (Gateway Load Balancer Endpoints).

Key AWS features:

Traffic Inspection: The GWLB allows the centralized security appliance to inspect traffic between different VPCs, making it suitable for inspecting inter-account interactions.

Interface VPC Endpoints: By using interface endpoints in the application accounts, traffic can securely and efficiently be routed to the security appliance in the networking account.

AWS Documentation: The use of GWLB aligns with AWS's best practices for centralized network security, simplifying architecture and reducing operational complexity.

Question: 36

A company uses Amazon EC2 instances and stores data on Amazon Elastic Block Store (Amazon EBS) volumes. The company must ensure that all data is encrypted at rest by using AWS Key Management Service (AWS KMS). The company must be able to control rotation of the encryption keys.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a customer managed key. Use the key to encrypt the EBS volumes.
- B. Use an AWS managed key to encrypt the EBS volumes. Use the key to configure automatic key rotation.
- C. Create an external KMS key with imported key material. Use the key to encrypt the EBS volumes.
- D. Use an AWS owned key to encrypt the EBS volumes.

Answer: A

Explanation:

To meet the requirement of controlling key rotation with minimal operational overhead, creating a customer managed key (CMK) in AWS KMS is the optimal solution. With CMKs, you can define custom key rotation policies, ensuring that you retain control over the key lifecycle, including enabling automatic key rotation every year.

Key AWS features:

Custom Key Management: A customer managed key allows you to control the key policies, lifecycle, and enable key rotation for compliance.

Least Operational Overhead: Using a customer managed key simplifies encryption management while offering more flexibility than AWS managed or owned keys.

AWS Documentation: The AWS Well-Architected Framework recommends customer managed keys for environments where key control and flexibility are required.

Question: 37

A weather forecasting company collects temperature readings from various sensors on a continuous basis. An existing data ingestion process collects the readings and aggregates the readings into larger Apache Parquet files. Then the process encrypts the files by using client-side encryption with KMS managed keys (CSE-KMS). Finally, the process writes the files to an Amazon S3 bucket with separate prefixes for each calendar day.

The company wants to run occasional SQL queries on the data to take sample moving averages for a specific calendar day.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure Amazon Athena to read the encrypted files. Run SQL queries on the data directly in Amazon S3.
- B. Use Amazon S3 Select to run SQL queries on the data directly in Amazon S3.
- C. Configure Amazon Redshift to read the encrypted files. Use Redshift Spectrum and Redshift query editor v2 to run SQL queries on the data directly in Amazon S3.
- D. Configure Amazon EMR Serverless to read the encrypted files. Use Apache SparkSQL to run SQL queries on the data directly in Amazon S3.

Answer: A

Explanation:

Amazon Athena is a serverless query service that allows you to run SQL queries directly on data stored in Amazon S3 without the need for a data warehouse. It is cost-effective because you only pay for the queries you run, and it can handle Apache Parquet files efficiently. Additionally, Athena integrates with KMS, making it suitable for querying encrypted data.

Key AWS features:

Cost-Effective: Athena charges only for the data scanned by the queries, making it a more costeffective solution compared to Redshift or EMR for occasional queries.

Direct S3 Querying: Athena supports querying data directly in S3, including Parquet files, without needing to move the data.

AWS Documentation: Athena's compatibility with encrypted Parquet files in S3 makes it the ideal choice for this scenario, reducing both cost and complexity.

Question: 38

A company is planning to migrate a legacy application to AWS. The application currently uses NFS to communicate to an on-premises storage solution to store application data. The application cannot be modified to use any other communication protocols other than NFS for this purpose.

Which storage solution should a solutions architect recommend for use after the migration?

A. AWS DataSync

- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon Elastic File System (Amazon EFS)
- D. Amazon EMR File System (Amazon EMRFS)

Answer: C

Explanation:

Amazon Elastic File System (EFS) is the ideal solution for migrating legacy applications that require NFS (Network File System) communication. EFS provides fully managed, scalable NFS storage in the cloud, and it supports the standard NFS protocols, allowing the legacy application to continue using NFS without modification after migration to AWS.

Key AWS features:

NFS Support: EFS natively supports the NFSv4 protocol, which makes it the best solution for workloads that rely on NFS communication.

Scalability and Availability: EFS automatically scales as application demands grow, making it a highly available and reliable storage solution.

AWS Documentation: According to AWS's best practices for file storage, EFS is recommended for any workloads requiring NFS support in a cloud environment.

Question: 39

A company runs an environment where data is stored in an Amazon S3 bucket. The objects are accessed frequently throughout the day. The company has strict data encryption requirements for data that is stored in the S3 bucket. The company currently uses AWS Key Management Service (AWS KMS) for encryption.

The company wants to optimize costs associated with encrypting S3 objects without making additional calls to AWS KMS.

Which solution will meet these requirements?

- A. Use server-side encryption with Amazon S3 managed keys (SSE-S3).
- B. Use an S3 Bucket Key for server-side encryption with AWS KMS keys (SSE-KMS) on the new objects.
- C. Use client-side encryption with AWS KMS customer managed keys.
- D. Use server-side encryption with customer-provided keys (SSE-C) stored in AWS KMS.

Answer: B

Explanation:

An S3 Bucket Key reduces the cost of using AWS KMS for server-side encryption by decreasing the number of requests made to KMS. By enabling S3 Bucket Key, the company can meet its encryption requirements with KMS keys while optimizing costs by reducing the number of KMS API requests.

Key AWS features:

Cost Optimization: S3 Bucket Keys reduce the frequency of KMS calls, optimizing the cost associated with encryption while still using AWS KMS for key management.

Compliance with KMS Encryption: This solution continues to meet the strict encryption requirements of the company by using KMS managed keys.

AWS Documentation: Using an S3 Bucket Key is recommended for organizations looking to optimize encryption costs without compromising security.

Question: 40

A company is designing the architecture for a new mobile app that uses the AWS Cloud. The company uses organizational units (OUs) in AWS Organizations to manage its accounts. The company wants to tag Amazon EC2 instances with data sensitivity by using values of sensitive and nonsensitive IAM identities must not be able to delete a tag or create instances without a tag

Which combination of steps will meet these requirements? (Select TWO.)

- A. In Organizations, create a new tag policy that specifies the data sensitivity tag key and the required values. Enforce the tag values for the EC2 instances Attach the tag policy to the appropriate OU.
- B. In Organizations, create a new service control policy (SCP) that specifies the data sensitivity tag key and the required tag values Enforce the tag values for the EC2 instances. Attach the SCP to the appropriate OU.
- C. Create a tag policy to deny running instances when a tag key is not specified. Create another tag policy that prevents identities from deleting tags Attach the tag policies to the appropriate OU.
- D. Create a service control policy (SCP) to deny creating instances when a tag key is not specified. Create another SCP that prevents identities from deleting tags Attach the SCPs to the appropriate OU.
- E. Create an AWS Config rule to check if EC2 instances use the data sensitivity tag and the specified values. Configure an AWS Lambda function to delete the resource if a noncompliant resource is found.

Answer: A,D

Explanation:

To meet the requirements for tagging and preventing instance creation or deletion without proper tags, the company can use a combination of AWS Organizations tag policies and service control policies (SCPs).

Tag Policies: These enforce specific tag values across resources. Creating a tag policy with required values (e.g., sensitive, non-sensitive) and attaching it to the appropriate organizational unit (OU) ensures consistency in tagging.

SCPs: SCPs can be used to enforce compliance by preventing instance creation without a tag and preventing tag deletion. These policies control actions at the account level across the organization.

Key AWS features:

Tag Policies help standardize tags across accounts, and SCPs enforce governance by restricting actions that violate the policies.

AWS Documentation: AWS best practices recommend using tag policies and SCPs to enforce compliance across multiple accounts within AWS Organizations.

Question: 41

A company runs multiple workloads on virtual machines (VMs) in an on-premises data center. The company is expanding rapidly. The on-premises data center is not able to scale fast enough to meet business needs. The company wants to migrate the workloads to AWS.

The migration is time sensitive. The company wants to use a lift-and-shift strategy for non-critical workloads.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Use the AWS Schema Conversion Tool (AWS SCT) to collect data about the VMs.
- B. Use AWS Application Migration Service. Install the AWS Replication Agent on the VMs.
- C. Complete the initial replication of the VMs. Launch test instances to perform acceptance tests on the VMs.
- D. Stop all operations on the VMs. Launch a cutover instance.
- E. Use AWS App2Container (A2C) to collect data about the VMs.
- F. Use AWS Database Migration Service (AWS DMS) to migrate the VMs.

Answer: B,C,D

Explanation:

AWS Application Migration Service (AWS MGN) is the recommended tool for a lift-and-shift strategy, especially for time-sensitive migrations. It automates the replication of on-premises VMs to AWS, minimizing the effort required for migration and testing.

Key steps:

Replication with AWS MGN: The AWS Replication Agent is installed on the VMs to continuously replicate data to AWS, allowing you to manage migration easily.

Testing and Cutover: Initial replication allows for testing in AWS before performing the final cutover, ensuring that the migration process is smooth and data integrity is maintained.

AWS Documentation: AWS MGN is recommended for migrating virtual machines to the cloud with minimal downtime and disruption.

Question: 42

A company's application is deployed on Amazon EC2 instances and uses AWS Lambda functions for an event-driven architecture. The company uses nonproduction development environments in a different AWS account to test new features before the company deploys the features to production.

The production instances show constant usage because of customers in different time zones. The company uses nonproduction instances only during business hours on weekdays. The company does not use the nonproduction instances on the weekends. The company wants to optimize the costs to run its application on AWS.

Which solution will meet these requirements MOST cost-effectively?

- A. Use On-Demand Instances (or the production instances. Use Dedicated Hosts for the nonproduction instances on weekends only.
- B. Use Reserved instances for the production instances and the nonproduction instances Shut down the nonproduction instances when not in use.
- C. Use Compute Savings Plans for the production instances. Use On-Demand Instances for the nonproduction instances Shut down the nonproduction instances when not in use.
- D. Use Dedicated Hosts for the production instances. Use EC2 Instance Savings Plans for the nonproduction instances.

Answer: C

Explanation:

Compute Savings Plans offer the most flexible and cost-effective solution for the production instances, as they provide significant savings (up to 66%) for both EC2 and AWS Lambda usage, while allowing flexibility in the type of instance family, size, and even region. For nonproduction instances, using On-Demand Instances ensures you only pay for the instances when they are running, and shutting them

down during off-hours further optimizes cost.

Key AWS features:

Compute Savings Plans: Provide savings based on consistent usage, making it ideal for production environments with steady load.

On-Demand Instances: Suitable for nonproduction environments that are used intermittently. Shutting them down when not in use avoids unnecessary costs.

AWS Documentation: According to AWS's cost optimization best practices, using a combination of Savings Plans for production and On-Demand Instances for nonproduction environments that are used sparingly results in optimal cost savings.

Question: 43

A company runs database workloads on AWS that are the backend for the company's customer portals. The company runs a Multi-AZ database cluster on Amazon RDS for PostgreSQL.

The company needs to implement a 30-day backup retention policy. The company currently has both automated RDS backups and manual RDS backups. The company wants to maintain both types of existing RDS backups that are less than 30 days old.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure the RDS backup retention policy to 30 days for automated backups by using AWS Backup. Manually delete manual backups that are older than 30 days.
- B. Disable RDS automated backups. Delete automated backups and manual backups that are older than 30 days. Configure the RDS backup retention policy to 30 days for automated backups.
- C. Configure the RDS backup retention policy to 30 days for automated backups. Manually delete manual backups that are older than 30 days
- D. Disable RDS automated backups. Delete automated backups and manual backups that are older than 30 days automatically by using AWS CloudFormation. Configure the RDS backup retention policy to 30 days for automated backups.

Answer: A

Explanation:

Setting the RDS backup retention policy to 30 days for automated backups through AWS Backup allows the company to retain backups cost-effectively. Manual backups, however, are not automatically managed by RDS's retention policy, so they need to be manually deleted if they are older than 30 days to avoid unnecessary storage costs.

Key AWS features:

Automated Backups: Can be configured with a retention policy of up to 35 days, ensuring that older automated backups are deleted automatically.

Manual Backups: These are not subject to the automated retention policy and must be manually managed to avoid extra costs.

AWS Documentation: AWS recommends using backup retention policies for automated backups while manually managing manual backups.

Question: 44

A company hosts an application in a private subnet. The company has already integrated the application with Amazon Cognito. The company uses an Amazon Cognito user pool to authenticate users.

The company needs to modify the application so the application can securely store user documents in an Amazon S3 bucket.

Which combination of steps will securely integrate Amazon S3 with the application? (Select TWO.)

- A. Create an Amazon Cognito identity pool to generate secure Amazon S3 access tokens for users when they successfully log in.
- B. Use the existing Amazon Cognito user pool to generate Amazon S3 access tokens for users when they successfully log in.
- C. Create an Amazon S3 VPC endpoint in the same VPC where the company hosts the application.
- D. Create a NAT gateway in the VPC where the company hosts the application. Assign a policy to the S3 bucket to deny any request that is not initiated from Amazon Cognito.
- E. Attach a policy to the S3 bucket that allows access only from the users' IP addresses.

Answer: A,C

Explanation:

To securely integrate Amazon S3 with an application that uses Amazon Cognito for user authentication, the following two steps are essential:

Step 1: Create an Amazon Cognito Identity Pool (Option A)

Amazon Cognito Identity Pools allow users to obtain temporary AWS credentials to access AWS resources, such as Amazon S3, after successfully authenticating with the Cognito user pool. The identity pool bridges the gap between user authentication and AWS service access by generating temporary credentials using AWS Identity and Access Management (IAM).

Once a user logs in using the Cognito User Pool, the identity pool provides IAM roles with specific permissions that the application can use to access S3 securely. This ensures that each user has appropriate access controls while accessing the S3 bucket.

This is a secure way to ensure that users only have temporary and least-privilege access to the S3 bucket for their documents.

Step 2: Create an Amazon S3 VPC Endpoint (Option C)

By creating an Amazon S3 VPC endpoint, the company ensures that communication between the application (which is hosted in a private subnet) and the S3 bucket occurs over the AWS private network, without the need to traverse the internet. This enhances security and prevents exposure of data to public networks.

The VPC endpoint allows the application to access the S3 bucket privately and securely within the VPC. It also ensures that traffic stays within the AWS network, reducing attack surface and improving overall security.

Why the Other Options Are Incorrect:

Option B: This is incorrect because Amazon Cognito User Pools are used for user authentication, not for generating S3 access tokens. To provide S3 access, you need to use Amazon Cognito Identity Pools, which offer AWS credentials.

Option D: An NAT gateway is unnecessary in this scenario. Using a VPC endpoint for S3 access provides a more secure and cost-effective solution by keeping traffic within AWS.

Option E: Attaching a policy to restrict access based on IP addresses is not scalable or efficient. It would require managing users' dynamic IP addresses, which is not an effective security measure for this use case.

AWS Reference:

Amazon Cognito Identity Pools

Amazon VPC Endpoints for S3

Question: 45

A company is using AWS DataSync to migrate millions of files from an on-premises system to AWS. The files are 10 KB in size on average.

The company wants to use Amazon S3 for file storage. For the first year after the migration the files will be accessed once or twice and must be immediately available. After 1 year the files must be archived for at least 7 years.

Which solution will meet these requirements MOST cost-effectively?

A. Use an archive tool to group the files into large objects. Use DataSync to migrate the objects. Store the objects in S3 Glacier Instant Retrieval for the first year. Use a lifecycle configuration to transition the files to S3 Glacier Deep Archive after 1 year with a retention period of 7 years.

B. Use an archive tool to group the files into large objects. Use DataSync to copy the objects to S3 Standard-Infrequent Access (S3 Standard-IA). Use a lifecycle configuration to transition the files to S3 Glacier Instant Retrieval after 1 year with a retention period of 7 years.

C. Configure the destination storage class for the files as S3 Glacier Instant Retrieval. Use a lifecycle policy to transition the files to S3 Glacier Flexible Retrieval after 1 year with a retention period of 7 years.

D. Configure a DataSync task to transfer the files to S3 Standard-Infrequent Access (S3 Standard-IA) Use a lifecycle configuration to transition the files to S3. Deep Archive after 1 year with a retention period of 7 years.

Answer: A

Explanation:

Question: 46

A company sets up an organization in AWS Organizations that contains 10 AWS accounts. A solutions architect must design a solution to provide access to the accounts for several thousand employees. The company has an existing identity provider (IdP). The company wants to use the existing IdP for authentication to AWS.

Which solution will meet these requirements?

- A. Create IAM users for the employees in the required AWS accounts. Connect IAM users to the existing IdP. Configure federated authentication for the IAM users.
- B. Set up AWS account root users with user email addresses and passwords that are synchronized from the existing IdP.
- C. Configure AWS IAM Identity Center Connect IAM Identity Center to the existing IdP Provision users and groups from the existing IdP
- D. Use AWS Resource Access Manager (AWS RAM) to share access to the AWS accounts with the users in the existing IdP.

Answer: C

Explanation:

AWS IAM Identity Center:

IAM Identity Center provides centralized access management for multiple AWS accounts within an organization and integrates seamlessly with existing identity providers (IdPs) through SAML 2.0 federation.

It allows users to authenticate using their existing IdP credentials and gain access to AWS resources without the need to create and manage separate IAM users in each account.

IAM Identity Center also simplifies provisioning and de-provisioning users, as it can automatically synchronize users and groups from the external IdP to AWS, ensuring secure and managed access.

Integration with Existing IdP:

The solution involves configuring IAM Identity Center to connect to the company's IdP using SAML. This setup allows employees to log in with their existing credentials, reducing the complexity of managing separate AWS credentials.

Once connected, IAM Identity Center handles authentication and authorization, granting users access to the AWS accounts based on their assigned roles and permissions.

Why the Other Options Are Incorrect:

Option A: Creating separate IAM users for each employee is not scalable or efficient. Managing thousands of IAM users across multiple AWS accounts introduces unnecessary complexity and operational overhead.

Option B: Using AWS root users with synchronized passwords is a security risk and goes against AWS best practices. Root accounts should never be used for day-to-day operations.

Option D: AWS Resource Access Manager (RAM) is used for sharing AWS resources between accounts, not for federating access for users across accounts. It doesn't provide a solution for authentication via an external IdP.

AWS Reference:

AWS IAM Identity Center

SAML 2.0 Integration with AWS IAM Identity Center

By setting up IAM Identity Center and connecting it to the existing IdP, the company can efficiently manage access for thousands of employees across multiple AWS accounts with a high degree of operational efficiency and security. Therefore, Option C is the best solution.

Question: 47

A company currently stores 5 TB of data in on-premises block storage systems. The company's current storage solution provides limited space for additional data. The company runs applications on premises that must be able to retrieve frequently accessed data with low latency. The company requires a cloud-based storage solution.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Use Amazon S3 File Gateway. Integrate S3 File Gateway with the on-premises applications to store and directly retrieve files by using the SMB file system.
- B. Use an AWS Storage Gateway Volume Gateway with cached volumes as iSCSI targets.
- C. Use an AWS Storage Gateway Volume Gateway with stored volumes as iSCSI targets.
- D. Use an AWS Storage Gateway Tape Gateway. Integrate Tape Gateway with the on-premises applications to store virtual tapes in Amazon S3.

Answer: B

Explanation:

The company needs a cloud-based storage solution for frequently accessed data with low latency, while retaining their current on-premises infrastructure for some data storage. AWS Storage Gateway's Volume Gateway with cached volumes is the most appropriate solution for

this scenario.

AWS Storage Gateway - Volume Gateway (Cached Volumes):

Volume Gateway with cached volumes allows you to store frequently accessed data in the AWS Cloud while keeping the most recently accessed data cached locally on-premises. This ensures low-latency access to active data while providing scalability for the rest of the data in the cloud.

The cached volume option stores the primary data in Amazon S3 but caches frequently accessed data locally, ensuring fast access. This configuration is well-suited for applications that require fast access to frequently used data but can tolerate cloud-based storage for the rest.

Since the company is facing limited on-premises storage, cached volumes provide an ideal solution, as they reduce the need for additional on-premises storage infrastructure.

Why Not the Other Options?:

Option A (S3 File Gateway): S3 File Gateway provides a file-based interface (SMB/NFS) for storing data directly in S3. While it is great for file storage, the company's need for block-level storage with iSCSI targets makes Volume Gateway a better fit.

Option C (Volume Gateway - Stored Volumes): Stored volumes keep all the data on-premises and asynchronously back up to AWS. This would not address the company's storage limitations since they would still need substantial on-premises storage.

Option D (Tape Gateway): Tape Gateway is designed for archiving and backup, not for frequently accessed low-latency data.

AWS Reference:

AWS Storage Gateway - Volume Gateway

Question: 48

A company has a three-tier web application that processes orders from customers. The web tier consists of Amazon EC2 instances behind an Application Load Balancer. The processing tier consists of EC2 instances. The company decoupled the web tier and processing tier by using Amazon Simple Queue Service (Amazon SQS). The storage layer uses Amazon DynamoDB.

At peak times some users report order processing delays and halts. The company has noticed that during these delays, the EC2 instances are running at 100% CPU usage, and the SQS queue fills up. The peak times are variable and unpredictable.

The company needs to improve the performance of the application.

Which solution will meet these requirements?

- A. Use scheduled scaling for Amazon EC2 Auto Scaling to scale out the processing tier instances for the duration of peak usage times. Use the CPU Utilization metric to determine when to scale.
- B. Use Amazon ElastiCache for Redis in front of the DynamoDB backend tier. Use target utilization as a metric to determine when

to scale.

C. Add an Amazon CloudFront distribution to cache the responses for the web tier. Use HTTP latency as a metric to determine when to scale.

D. Use an Amazon EC2 Auto Scaling target tracking policy to scale out the processing tier instances. Use the ApproximateNumberOfMessages attribute to determine when to scale.

Answer: D

Explanation:

The issue in this case is related to the processing tier, where EC2 instances are overwhelmed at peak times, causing delays. Option D, using an Amazon EC2 Auto Scaling target tracking policy based on the ApproximateNumberOfMessages in the SQS queue, is the best solution.

Auto Scaling with Target Tracking:

Target tracking policies dynamically scale out or in based on a specific metric. For this use case, you can monitor the ApproximateNumberOfMessages in the SQS queue. When the number of messages (orders) in the queue increases, the Auto Scaling group will scale out more EC2 instances to handle the additional load, ensuring that the queue doesn't build up and cause delays.

This solution is ideal for handling variable and unpredictable peak times, as Auto Scaling can automatically adjust based on real-time load rather than scheduled times.

Why Not the Other Options?:

Option A (Scheduled Scaling): Scheduled scaling works well for predictable peak times, but this company experiences unpredictable peak usage, making scheduled scaling less effective.

Option B (ElastiCache for Redis): Adding a caching layer would help if DynamoDB were the bottleneck, but in this case, the issue is CPU overload on EC2 instances in the processing tier.

Option C (CloudFront): CloudFront would help cache static content from the web tier, but it wouldn't resolve the issue of the processing tier's overloaded EC2 instances.

AWS Reference:

Amazon EC2 Auto Scaling Target Tracking

Amazon SQS ApproximateNumberOfMessages

Question: 49

An online gaming company hosts its platform on Amazon EC2 instances behind Network Load Balancers (NLBs) across multiple AWS Regions. The NLBs can route requests to targets over the internet. The company wants to improve the customer playing experience by reducing end-to-end load time for its global customer base.

Which solution will meet these requirements?

- A. Create Application Load Balancers (ALBs) in each Region to replace the existing NLBs. Register the existing EC2 instances as targets for the ALBs in each Region.
- B. Configure Amazon Route 53 to route equally weighted traffic to the NLBs in each Region.
- C. Create additional NLBs and EC2 instances in other Regions where the company has large customer bases.
- D. Create a standard accelerator in AWS Global Accelerator. Configure the existing NLBs as target endpoints.

Answer: D

Explanation:

The company wants to reduce end-to-end load time for its global customer base. AWS Global Accelerator provides a network optimization service that reduces latency by routing traffic to the nearest AWS edge locations, improving the user experience for globally distributed customers.

AWS Global Accelerator:

Global Accelerator improves the performance of your applications by routing traffic through AWS's global network infrastructure. This reduces the number of hops and latency compared to using the public internet.

By creating a standard accelerator and configuring the existing NLBs as target endpoints, Global Accelerator ensures that traffic from users around the world is routed to the nearest AWS edge location and then through optimized paths to the NLBs in each region. This significantly improves end-to-end load time for global customers.

Why Not the Other Options?:

Option A (ALBs instead of NLBs): ALBs are designed for HTTP/HTTPS traffic and provide layer 7 features, but they wouldn't solve the latency issue for a global customer base. The key problem here is latency, and Global Accelerator is specifically designed to address that.

Option B (Route 53 weighted routing): Route 53 can route traffic to different regions, but it doesn't optimize network performance. It simply balances traffic between endpoints without improving latency.

Option C (Additional NLBs in more regions): This could potentially improve latency but would require setting up infrastructure in multiple regions. Global Accelerator is a simpler and more efficient solution that leverages AWS's existing global network.

AWS Reference:

By using AWS Global Accelerator with the existing NLBs, the company can optimize global traffic routing and improve the customer experience by minimizing latency. Therefore, Option D is the correct answer.

Question: 50

A company has 15 employees. The company stores employee start dates in an Amazon DynamoDB table. The company wants to send an email message to each employee on the day of the employee's work anniversary.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Create a script that scans the DynamoDB table and uses Amazon Simple Notification Service (Amazon SNS) to send email messages to employees when necessary. Use a cron job to run this script every day on an Amazon EC2 instance.
- B. Create a script that scans the DynamoDB table and uses Amazon Simple Queue Service (Amazon SQS) to send email messages to employees when necessary. Use a cron job to run this script every day on an Amazon EC2 instance.
- C. Create an AWS Lambda function that scans the DynamoDB table and uses Amazon Simple Notification Service (Amazon SNS) to send email messages to employees when necessary. Schedule this Lambda function to run every day.
- D. Create an AWS Lambda function that scans the DynamoDB table and uses Amazon Simple Queue Service (Amazon SQS) to send email messages to employees when necessary. Schedule this Lambda function to run every day.

Answer: C

Explanation:

AWS Lambda for Operational Efficiency:

AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers. It automatically scales based on the number of invocations and eliminates the need to maintain and monitor EC2 instances, making it far more operationally efficient compared to running a cron job on EC2.

By using Lambda, you pay only for the compute time that your function uses. This is especially beneficial when dealing with lightweight tasks, such as scanning a DynamoDB table and sending email messages once a day.

Amazon DynamoDB:

DynamoDB is a highly scalable, fully managed NoSQL database. The table stores employee start dates, and scanning the table to find the employees who have a work anniversary on the current day is a lightweight operation. Lambda can easily perform this operation using the DynamoDB Scan API or queries, depending on how the data is structured.

Amazon SNS for Email Notifications:

Amazon Simple Notification Service (SNS) is a fully managed messaging service that supports sending notifications to a variety of endpoints, including email. SNS is well-suited for sending out email messages to employees, as it can handle the fan-out messaging pattern (sending the same message to multiple recipients).

In this scenario, once Lambda identifies employees who have their work anniversaries, it can use SNS to send the email notifications efficiently. SNS integrates seamlessly with Lambda, and sending emails via SNS is a common pattern for this type of use case.

Event Scheduling:

To automate this daily task, you can schedule the Lambda function using Amazon EventBridge (formerly CloudWatch Events). EventBridge can trigger the Lambda function on a daily schedule (cron-like scheduling). This avoids the complexity and operational overhead of manually setting up cron jobs on EC2 instances.

Why Not EC2 or SQS?:

Option A & B suggest running a cron job on an Amazon EC2 instance. This approach requires you to manage, scale, and patch the EC2 instance, which increases operational overhead. Lambda is a better choice because it automatically scales and doesn't require server management.

Amazon Simple Queue Service (SQS) is ideal for decoupling distributed systems but isn't necessary in this context because the goal is to send notifications to employees on their work anniversaries. SQS adds unnecessary complexity for this straightforward use case, where SNS is the simpler and more efficient solution.

AWS Reference:

AWS Lambda

Amazon SNS

Amazon DynamoDB

Amazon EventBridge

Summary:

Using AWS Lambda combined with Amazon SNS to send notifications, and scheduling the function with Amazon EventBridge to run daily, is the most operationally efficient solution. It leverages AWS serverless technologies, which reduce the need for infrastructure management and provide automatic scaling. Therefore, Option C is the correct and optimal choice.

Question: 51

A company runs its production workload on an Amazon Aurora MySQL DB cluster that includes six Aurora Replicas. The company wants near-real-time reporting queries from one of its departments to be automatically distributed across three of the Aurora Replicas. Those three replicas have a different compute and memory specification from the rest of the DB cluster.

Which solution meets these requirements?

A. Create and use a custom endpoint for the workload.

- B. Create a three-node cluster clone and use the reader endpoint.
- C. Use any of the instance endpoints for the selected three nodes.
- D. Use the reader endpoint to automatically distribute the read-only workload.

Answer: A

Explanation:

In Amazon Aurora, a custom endpoint is a feature that allows you to create a load-balanced endpoint that directs traffic to a specific set of instances in your Aurora DB cluster. This is particularly useful when you want to route traffic to a subset of instances that have different configurations or when you want to isolate specific workloads (e.g., reporting queries) to certain instances.

Custom Endpoint: The correct solution is to create a custom endpoint that includes the three Aurora Replicas that the department wants to use for near-real-time reporting. This custom endpoint will distribute the reporting queries only across the three selected replicas with the specified compute and memory configurations, ensuring that these queries do not affect the rest of the DB cluster.

Other Options:

Option B (Create a three-node cluster clone): This would create a separate cluster with its own resources, but it is not necessary and could incur additional costs. Also, it doesn't leverage the existing replicas.

Option C (Use any of the instance endpoints): This would involve manually managing connections to individual instances, which is not scalable or automatic.

Option D (Use the reader endpoint): The reader endpoint would distribute the read queries across all replicas in the cluster, not just the selected three. This would not meet the requirement to limit the reporting queries to only three specific replicas.

AWS Reference:

Amazon Aurora Endpoints- Provides detailed information on the different types of endpoints available in Aurora, including custom endpoints.

Custom Endpoints in Amazon Aurora- Specific documentation on how to create and use custom endpoints to direct traffic to selected instances in an Aurora cluster.

Question: 52

A company is migrating its databases to Amazon RDS for PostgreSQL. The company is migrating its applications to Amazon EC2 instances. The company wants to optimize costs for long-running workloads.

Which solution will meet this requirement MOST cost-effectively?

- A. Use On-Demand Instances for the Amazon RDS for PostgreSQL workloads. Purchase a 1 year Compute Savings Plan with the No Upfront option for the EC2 instances.
- B. Purchase Reserved Instances for a 1 year term with the No Upfront option for the Amazon RDS for PostgreSQL workloads. Purchase a 1 year EC2 Instance Savings Plan with the No Upfront option for the EC2 instances.
- C. Purchase Reserved Instances for a 1 year term with the Partial Upfront option for the Amazon RDS for PostgreSQL workloads. Purchase a 1 year EC2 Instance Savings Plan with the Partial Upfront option for the EC2 instances.
- D. Purchase Reserved Instances for a 3 year term with the All Upfront option for the Amazon RDS for PostgreSQL workloads. Purchase a 3 year EC2 Instance Savings Plan with the All Upfront option for the EC2 instances.

Answer: D

Explanation:

Question: 53

A company is implementing a new application on AWS. The company will run the application on multiple Amazon EC2 instances across multiple Availability Zones within multiple AWS Regions. The application will be available through the internet. Users will access the application from around the world.

The company wants to ensure that each user who accesses the application is sent to the EC2 instances that are closest to the user's location.

Which solution will meet these requirements?

- A. Implement an Amazon Route 53 geolocation routing policy. Use an internet-facing Application Load Balancer to distribute the traffic across all Availability Zones within the same Region.
- B. Implement an Amazon Route 53 geoproximity routing policy. Use an internet-facing Network Load Balancer to distribute the traffic across all Availability Zones within the same Region.
- C. Implement an Amazon Route 53 multivalue answer routing policy Use an internet-facing Application Load Balancer to distribute the traffic across all Availability Zones within the same Region.
- D. Implement an Amazon Route 53 weighted routing policy. Use an internet-facing Network Load Balancer to distribute the traffic across all Availability Zones within the same Region.

Answer: A

Explanation:

The requirement is to route users to the nearest AWS Region where the application is deployed. The best solution is to use Amazon Route

53 with a geolocation routing policy, which routes traffic based on the geographic location of the user making the request.

Geolocation Routing: This routing policy ensures that users are directed to the resources (in this case, EC2 instances) that are geographically closest to them, thereby reducing latency and improving the user experience.

Application Load Balancer (ALB): Within each Region, an internet-facing Application Load Balancer (ALB) is used to distribute incoming traffic across multiple EC2 instances in different Availability Zones. ALBs are designed to handle HTTP/HTTPS traffic and provide advanced features like contentbased routing, SSL termination, and user authentication.

Why Not Other Options?:

Option B (Geoproximity + NLB): Geoproximity routing is similar but more complex as it requires finetuning the proximity settings. A Network Load Balancer (NLB) is better suited for TCP/UDP traffic rather than HTTP/HTTPS.

Option C (Multivalue Answer Routing + ALB): Multivalue answer routing does not direct traffic based on user location but rather returns multiple values and lets the client choose. This does not meet the requirement for geographically routing users.

Option D (Weighted Routing + NLB): Weighted routing splits traffic based on predefined weights and does not consider the user's geographic location. NLB is not ideal for this scenario due to its focus on lower-level protocols.

AWS Reference:

Amazon Route 53 Routing Policies- Detailed explanation of the various routing policies available in Route 53, including geolocation.

Elastic Load Balancing- Information on the different types of load balancers in AWS and when to use them.

Question: 54

A company recently migrated a monolithic application to an Amazon EC2 instance and Amazon RDS. The application has tightly coupled modules. The existing design of the application gives the application the ability to run on only a single EC2 instance.

The company has noticed high CPU utilization on the EC2 instance during peak usage times. The high CPU utilization corresponds to degraded performance on Amazon RDS for read requests. The company wants to reduce the high CPU utilization and improve read request performance.

Which solution will meet these requirements?

- A. Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Configure an RDS read replica for read requests.
- B. Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Add an RDS read replica and redirect all read/write traffic to the replica.
- C. Configure an Auto Scaling group with a minimum size of 1 and maximum size of 2. Resize the RDS DB instance to an instance type that has more CPU capacity.
- D. Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Resize the RDS DB instance to an instance type that has more CPU capacity.

Answer: A

Explanation:

To address the high CPU utilization on the EC2 instance and the degraded performance of Amazon RDS for read requests, the solution involves two key actions: resizing the EC2 instance and leveraging Amazon RDS read replicas.

Resizing the EC2 Instance: The first step is to resize the EC2 instance to a type with more CPU capacity to handle the higher computational demands during peak usage times. This helps to alleviate the immediate pressure on the CPU.

Auto Scaling Group with a Size of 1: Although the application can only run on a single EC2 instance due to its monolithic nature, creating an Auto Scaling group with a minimum and maximum size of 1 ensures that the instance is automatically restarted or replaced in case of failure, maintaining high availability.

RDS Read Replica: Configuring an RDS read replica allows the application to offload read requests to a separate instance, thus reducing the load on the primary RDS instance. This improves the performance of read operations, which were previously bottlenecked due to the high CPU usage on the EC2 instance.

Why Not Other Options?:

Option B: Redirecting all traffic to the RDS read replica is not recommended because replicas are meant for read traffic only, not for write operations. This could lead to data consistency issues.

Option C: Increasing the RDS instance type capacity helps, but it doesn't address the high CPU usage on the EC2 instance, nor does it provide a solution for scaling reads.

Option D: While resizing both the EC2 and RDS instances increases their capacities, it doesn't address the specific need to offload read traffic from the primary RDS instance.

AWS Reference:

Amazon RDS Read Replicas- Explains how to create and use read replicas to offload read traffic from the primary database instance.

Resizing Your EC2 Instance- Guidance on resizing EC2 instances to meet workload demands.

Question: 55

A company is building an application on AWS that connects to an Amazon RDS database. The company wants to manage the application configuration and to securely store and retrieve credentials for the database and other services.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use AWS AppConfig to store and manage the application configuration. Use AWS Secrets Manager to store and retrieve the credentials.
- B. Use AWS Lambda to store and manage the application configuration. Use AWS Systems Manager Parameter Store to store and

retrieve the credentials.

C. Use an encrypted application configuration file Store the file in Amazon S3 for the application configuration. Create another S3 file to store and retrieve the credentials.

D. Use AWS AppConfig to store and manage the application configuration. Use Amazon RDS to store and retrieve the credentials.

Answer: A

Explanation:

This solution meets the company's requirements with minimal administrative overhead and ensures security and ease of management.

AWS AppConfig: AWS AppConfig is a service designed to manage application configuration in a secure and validated way. It allows you to deploy configurations safely and quickly without affecting the application's performance or availability.

AWS Secrets Manager: AWS Secrets Manager is specifically designed to manage, retrieve, and rotate credentials for databases and other services. It integrates seamlessly with AWS services like Amazon RDS, making it an ideal solution for securely storing and retrieving database credentials. Secrets Manager also provides automatic rotation of credentials, reducing the operational burden.

Why Not Other Options?:

Option B (AWS Lambda + Parameter Store): While AWS Lambda can be used for managing configurations and AWS Systems Manager Parameter Store can store credentials, this approach involves more manual setup and does not offer the same level of integrated management and security as AppConfig and Secrets Manager.

Option C (Encrypted S3 Configuration File): Storing configuration and credentials in S3 files involves more manual management and security considerations, increasing the administrative overhead.

Option D (AppConfig + RDS for credentials): RDS is not designed for storing application credentials; it's better suited for managing database instances and their configurations.

AWS Reference:

AWS AppConfig- Describes how to use AWS AppConfig for managing application configurations.

AWS Secrets Manager- Provides details on securely storing and retrieving credentials using AWS Secrets Manager.

Question: 56

A company stores data in an on-premises Oracle relational database. The company needs to make the data available in Amazon Aurora PostgreSQL for analysis The company uses an AWS Site-to-Site VPN connection to connect its on-premises network to AWS.

The company must capture the changes that occur to the source database during the migration to Aurora PostgreSQL.

Which solution will meet these requirements?

A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the Oracle schema to Aurora

PostgreSQL schema. Use the AWS Database Migration Service (AWS DMS) full-load migration task to migrate the data.

B. Use AWS DataSync to migrate the data to an Amazon S3 bucket. Import the S3 data to Aurora PostgreSQL by using the Aurora PostgreSQL `aws_s3` extension.

C. Use the AWS Schema Conversion Tool (AWS SCT) to convert the Oracle schema to Aurora PostgreSQL schema. Use AWS Database Migration Service (AWS DMS) to migrate the existing data and replicate the ongoing changes.

D. Use an AWS Snowball device to migrate the data to an Amazon S3 bucket. Import the S3 data to Aurora PostgreSQL by using the Aurora PostgreSQL `aws_s3` extension.

Answer: C

Explanation:

For the migration of data from an on-premises Oracle database to Amazon Aurora PostgreSQL, this solution effectively handles schema conversion, data migration, and ongoing data replication.

AWS Schema Conversion Tool (SCT): SCT is used to convert the Oracle database schema to a format compatible with Aurora PostgreSQL. This tool automatically converts the database schema and code objects, like stored procedures, to the target database engine.

AWS Database Migration Service (DMS): DMS is employed to perform the data migration. It supports both full-load migrations (for initial data transfer) and continuous replication of ongoing changes (Change Data Capture, or CDC). This ensures that any updates to the Oracle database during the migration are captured and applied to the Aurora PostgreSQL database, minimizing downtime.

Why Not Other Options?:

Option A (SCT + DMS full-load only): This option does not capture ongoing changes, which is crucial for a live database migration to ensure data consistency.

Option B (DataSync + S3): AWS DataSync is more suited for file transfers rather than database migrations, and it doesn't support ongoing change replication.

Option D (Snowball + S3): Snowball is typically used for large-scale data transfers that don't require continuous synchronization, making it less suitable for this scenario where ongoing changes must be captured.

AWS Reference:

AWS Schema Conversion Tool- Guidance on using SCT for database schema conversions.

AWS Database Migration Service- Detailed documentation on using DMS for data migrations and ongoing replication.

Question: 57

A financial services company plans to launch a new application on AWS to handle sensitive financial transactions. The company will deploy the application on Amazon EC2 instances. The company will use Amazon RDS for MySQL as the database. The company's security policies mandate that data must be encrypted at rest and in transit.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure encryption at rest for Amazon RDS for MySQL by using AWS KMS managed keys. Configure AWS Certificate Manager (ACM) SSL/TLS certificates for encryption in transit.
- B. Configure encryption at rest for Amazon RDS for MySQL by using AWS KMS managed keys. Configure IPsec tunnels for encryption in transit
- C. Implement third-party application-level data encryption before storing data in Amazon RDS for MySQL. Configure AWS Certificate Manager (ACM) SSL/TLS certificates for encryption in transit.
- D. Configure encryption at rest for Amazon RDS for MySQL by using AWS KMS managed keys. Configure a VPN connection to enable private connectivity to encrypt data in transit.

Answer: A

Explanation:

This solution provides encryption at rest and in transit with the least operational overhead while adhering to the company's security policies.

Encryption at Rest: Amazon RDS for MySQL can be configured to encrypt data at rest by using AWS Key Management Service (KMS) managed keys. This encryption is applied automatically to all data stored on disk, including backups, read replicas, and snapshots. This solution requires minimal operational overhead because AWS manages the encryption and key management process.

Encryption in Transit: AWS Certificate Manager (ACM) allows you to provision, manage, and deploy SSL/TLS certificates seamlessly. These certificates can be used to encrypt data in transit by configuring the MySQL instance to use SSL/TLS for connections. This setup ensures that data is encrypted between the application and the database, protecting it from interception during transmission.

Why Not Other Options?:

Option B (IPsec tunnels): While IPsec tunnels encrypt data in transit, they are more complex to manage and require additional configuration and maintenance, leading to higher operational overhead.

Option C (Third-party application-level encryption): Implementing application-level encryption adds complexity, requires code changes, and increases operational overhead.

Option D (VPN for encryption): A VPN solution for encrypting data in transit is unnecessary and adds additional complexity without providing any benefit over SSL/TLS, which is simpler to implement and manage.

AWS Reference:

Amazon RDS Encryption- Information on how to configure and use encryption for Amazon RDS.

AWS Certificate Manager (ACM)- Details on using ACM to manage SSL/TLS certificates for securing data in transit.

Question: 58

A startup company is hosting a website for its customers on an Amazon EC2 instance. The website consists of a stateless Python application and a MySQL database. The website serves only a small amount of traffic. The company is concerned about the reliability of the instance and needs to migrate to a highly available architecture. The company cannot modify the application code.

Which combination of actions should a solutions architect take to achieve high availability for the website? (Select TWO.)

- A. Provision an internet gateway in each Availability Zone in use.
- B. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB instance.
- C. Migrate the database to Amazon DynamoDB, and enable DynamoDB auto scaling.
- D. Use AWS DataSync to synchronize the database data across multiple EC2 instances.
- E. Create an Application Load Balancer to distribute traffic to an Auto Scaling group of EC2 instances that are distributed across two Availability Zones.

Answer: B,E

Explanation:

To achieve high availability for the website, two key actions should be taken: Amazon RDS for MySQL Multi-AZ: By migrating the database to an RDS for MySQL Multi-AZ deployment, the database becomes highly available. Multi-AZ provides automatic failover from the primary database to a standby replica in another Availability Zone, ensuring database availability even in the case of an AZ failure.

Application Load Balancer and Auto Scaling: Deploying an Application Load Balancer (ALB) in front of the EC2 instances ensures that traffic is evenly distributed across the instances. Configuring an Auto Scaling group to run EC2 instances across multiple Availability Zones ensures that the application remains available even if one instance or one AZ becomes unavailable. This setup enhances fault tolerance and improves reliability.

Why Not Other Options?:

Option A (Internet Gateway per AZ): Internet Gateways are region-wide resources and do not need to be provisioned per Availability Zone. This option does not contribute to high availability.

Option C (DynamoDB + Auto Scaling): DynamoDB would require changes to the application code to switch from MySQL, which is not possible per the question's constraints.

Option D (DataSync): AWS DataSync is used for data transfer and synchronization, not for achieving high availability for a database.

AWS Reference:

Amazon RDS Multi-AZ Deployments- Explanation of how Multi-AZ deployments work in Amazon RDS.

Application Load Balancing- Details on how to configure and use ALB for distributing traffic across multiple instances.

Question: 59

A company has stored millions of objects across multiple prefixes in an Amazon S3 bucket by using the Amazon S3 Glacier Deep Archive storage class. The company needs to delete all data older than 3 years except for a subset of data that must be retained. The company has identified the data that must be retained and wants to implement a serverless solution.

Which solution will meet these requirements?

- A. Use S3 Inventory to list all objects. Use the AWS CLI to create a script that runs on an Amazon EC2 instance that deletes objects from the inventory list.
- B. Use AWS Batch to delete objects older than 3 years except for the data that must be retained
- C. Provision an AWS Glue crawler to query objects older than 3 years. Save the manifest file of old objects. Create a script to delete objects in the manifest.
- D. Enable S3 Inventory. Create an AWS Lambda function to filter and delete objects. Invoke the Lambda function with S3 Batch Operations to delete objects by using the inventory reports.

Answer: D

Explanation:

To meet the requirement of deleting objects older than 3 years while retaining certain data, this solution leverages serverless technologies to minimize operational overhead.

S3 Inventory: S3 Inventory provides a flat file that lists all the objects in an S3 bucket and their metadata, which can be configured to include data such as the last modified date. This inventory can be generated daily or weekly.

AWS Lambda Function: A Lambda function can be created to process the S3 Inventory report, filtering out the objects that need to be retained and identifying those that should be deleted.

S3 Batch Operations: S3 Batch Operations can execute tasks such as object deletion at scale. By invoking the Lambda function through S3 Batch Operations, you can automate the process of deleting the identified objects, ensuring that the solution is serverless and requires minimal operational management.

Why Not Other Options?:

Option A (AWS CLI script on EC2): Running a script on an EC2 instance adds unnecessary operational overhead and is not serverless.

Option B (AWS Batch): AWS Batch is designed for running large-scale batch computing workloads, which is overkill for this scenario.

Option C (AWS Glue + script): AWS Glue is more suited for ETL tasks, and this approach would add unnecessary complexity compared to the serverless Lambda solution.

AWS Reference:

Amazon S3 Inventory- Information on how to set up and use S3 Inventory.

S3 Batch Operations- Documentation on how to perform bulk operations on S3 objects using S3 Batch Operations.

Question: 60

A company runs several websites on AWS for its different brands. Each website generates tens of gigabytes of web traffic logs each day. A solutions architect needs to design a scalable solution to give the company's developers the ability to analyze traffic patterns across all the company's websites. This analysis by the developers will occur on demand once a week over the course of several months. The solution must support queries with standard SQL.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the logs in Amazon S3. Use Amazon Athena for analysis.
- B. Store the logs in Amazon RDS. Use a database client for analysis.
- C. Store the logs in Amazon OpenSearch Service. Use OpenSearch Service for analysis.
- D. Store the logs in an Amazon EMR cluster. Use a supported open-source framework for SQL-based analysis.

Answer: A

Explanation:

This solution is the most cost-effective and scalable for analyzing large amounts of web traffic logs.

Amazon S3: Storing the logs in Amazon S3 is highly scalable, durable, and cost-effective. S3 is designed to handle large-scale data storage, which is ideal for storing tens of gigabytes of log data generated daily by multiple websites.

Amazon Athena: Athena is a serverless, interactive query service that allows you to analyze data in S3 using standard SQL. It works directly with the data stored in S3, so there's no need to load the data into a database, which saves on costs and reduces complexity. Athena charges based on the amount of data scanned by queries, making it a cost-effective solution for on-demand analysis that only occurs once a week.

Why Not Other Options?:

Option B (Amazon RDS): Storing logs in a relational database like Amazon RDS would be more expensive due to the storage and I/O costs associated with RDS. Additionally, it would require more **management overhead**.

Option C (Amazon OpenSearch Service): OpenSearch is a good option for full-text search and analytics on log data, but it might be more costly and complex to manage compared to the simplicity and cost-effectiveness of Athena for periodic SQL-based queries.

Option D (Amazon EMR): While EMR can handle large-scale data processing, it involves more operational overhead and might be overkill for the type of ad-hoc, SQL-based analysis required here. Additionally, EMR costs can be higher due to the need to maintain a cluster.

AWS Reference:

Amazon S3- Information on how to store and manage data in Amazon S3.

Amazon Athena- Documentation on using Amazon Athena for querying data stored in S3 using SQL.

Question: 61

A company is migrating a legacy application from an on-premises data center to AWS. The application relies on hundreds of cron jobs that run between 1 and 20 minutes on different recurring schedules throughout the day.

The company wants a solution to schedule and run the cron jobs on AWS with minimal refactoring. The solution must support running the cron jobs in response to an event in the future.

Which solution will meet these requirements?

- A. Create a container image for the cron jobs. Use Amazon EventBridge Scheduler to create a recurring schedule. Run the cron job tasks as AWS Lambda functions.
- B. Create a container image for the cron jobs. Use AWS Batch on Amazon Elastic Container Service (Amazon ECS) with a scheduling policy to run the cron jobs.
- C. Create a container image for the cron jobs. Use Amazon EventBridge Scheduler to create a recurring schedule. Run the cron job tasks on AWS Fargate.
- D. Create a container image for the cron jobs. Create a workflow in AWS Step Functions that uses a Wait state to run the cron jobs at a specified time. Use the RunTask action to run the cron job tasks on AWS Fargate.

Answer: C

Explanation:

This solution is the most suitable for running cron jobs on AWS with minimal refactoring, while also supporting the possibility of running jobs in response to future events.

Container Image for Cron Jobs: By containerizing the cron jobs, you can package the environment and dependencies required to run the jobs, ensuring consistency and ease of deployment across different environments.

Amazon EventBridge Scheduler: EventBridge Scheduler allows you to create a recurring schedule that can trigger tasks (like running your cron jobs) at specific times or intervals. It provides finegrained control over scheduling and integrates seamlessly with AWS services.

AWS Fargate: Fargate is a serverless compute engine for containers that removes the need to manage EC2 instances. It allows you to run containers without worrying about the underlying infrastructure. Fargate is ideal for running jobs that can vary in duration, like cron jobs, as it scales automatically based on the task's requirements.

Why Not Other Options?:

Option A (Lambda): While AWS Lambda could handle short-running cron jobs, it has limitations in terms of execution duration (maximum of 15 minutes) and might not be suitable for jobs that run up to 20 minutes.

Option B (AWS Batch on ECS): AWS Batch is more suitable for batch processing and workloads that require complex job dependencies or orchestration, which might be more than what is needed for simple cron jobs.

Option D (Step Functions with Wait State): While Step Functions provide orchestration capabilities, this approach would introduce unnecessary complexity and overhead compared to the straightforward scheduling with EventBridge and running on Fargate.

AWS Reference:

Amazon EventBridge Scheduler- Details on how to schedule tasks using Amazon EventBridge Scheduler.

AWS Fargate- Information on how to run containers in a serverless manner using AWS Fargate.

Question: 62

A solutions architect needs to connect a company's corporate network to its VPC to allow onpremises access to its AWS resources. The solution must provide encryption of all trafficbetween the corporate network and the VPC at the network layer and the session layer. The solution also must provide security controls to prevent unrestricted access between AWS and the on-premises systems.

Which solution meets these requirements?

- A. Configure AWS Direct Connect to connect to the VPC. Configure the VPC route tables to allow and deny traffic between AWS and on premises as required.
- B. Create an IAM policy to allow access to the AWS Management Console only from a defined set of corporate IP addresses Restrict user access based on job responsibility by using an IAM policy and roles
- C. Configure AWS Site-to-Site VPN to connect to the VPC. Configure route table entries to direct traffic from on premises to the VPC. Configure instance security groups and network ACLs to allow only required traffic from on premises.
- D. Configure AWS Transit Gateway to connect to the VPC. Configure route table entries to direct traffic from on premises to the VPC. Configure instance security groups and network ACLs to allow only required traffic from on premises.

Answer: C

Explanation:

This solution meets the requirements of providing encryption at both the network and session layers while also allowing for controlled access between on-premises systems and AWS resources.

AWS Site-to-Site VPN: This service allows you to establish a secure and encrypted connection between your on-premises network and AWS VPC over the internet or via AWS Direct Connect. The VPN encrypts data at the network layer (IPsec) as it travels between the corporate network and AWS.

Routing and Security Controls: By configuring route table entries, you can ensure that only the traffic intended for AWS resources is directed to the VPC. Additionally, by setting up security groups and network ACLs, you can further restrict and control which traffic is allowed to communicate with the instances within your VPC. This approach provides the necessary security to prevent unrestricted access, aligning with the company's security policies.

Why Not Other Options?:

Option A (AWS Direct Connect): While Direct Connect provides a private connection, it does not inherently provide encryption. Additional steps would be required to encrypt traffic, and it doesn't address the session layer encryption.

Option B (IAM policies for Console access): This option does not meet the requirement for network level encryption and security between the corporate network and the VPC.

Option D (AWS Transit Gateway): Although Transit Gateway can help in managing multiple connections, it doesn't directly provide encryption at the network layer. You would still need to configure a VPN or use other methods for encryption.

AWS Reference:

AWS Site-to-Site VPN- Overview of AWS Site-to-Site VPN capabilities, including encryption.

Security Groups and Network ACLs- Information on configuring security groups and network ACLs to control traffic.

Question: 63

A company runs a Node.js function on a server in its on-premises data center. The data center stores data in a PostgreSQL database. The company stores the credentials in a connection string in an environment variable on the server. The company wants to migrate its application to AWS and to replace the Node.js application server with AWS Lambda. The company also wants to migrate to Amazon RDS for PostgreSQL and to ensure that the database credentials are securely managed.

Which solution will meet these requirements with the LEAST operational overhead?

A. Store the database credentials as a parameter in AWS Systems Manager Parameter Store. Configure Parameter Store to automatically rotate the secrets every 30 days. Update the Lambda function to retrieve the credentials from the parameter.

B. Store the database credentials as a secret in AWS Secrets Manager. Configure Secrets Manager to automatically rotate the credentials

every 30 days Update the Lambda function to retrieve the credentials from the secret.

C. Store the database credentials as an encrypted Lambda environment variable. Write a custom Lambda function to rotate the credentials. Schedule the Lambda function to run every 30 days.

D. Store the database credentials as a key in AWS Key Management Service (AWS KMS). Configure automatic rotation for the key. Update the Lambda function to retrieve the credentials from the KMS key.

Answer: B

Explanation:

AWS Secrets Manager is designed specifically to securely store and manage sensitive information such as database credentials. It integrates seamlessly with AWS services like Lambda and RDS, and it provides automatic credential rotation with minimal operational overhead.

AWS Secrets Manager: By storing the database credentials in Secrets Manager, you ensure that the credentials are securely stored, encrypted, and managed. Secrets Manager provides a built-in mechanism to automatically rotate credentials at regular intervals (e.g., every 30 days), which helps in maintaining security best practices without requiring additional manual intervention.

Lambda Integration: The Lambda function can be easily configured to retrieve the credentials from Secrets Manager using the AWS SDK, ensuring that the credentials are accessed securely at runtime.

Why Not Other Options?:

Option A (Parameter Store with Rotation): While Parameter Store can store parameters securely, Secrets Manager is more tailored for secrets management and automatic rotation, offering more features and less operational overhead.

Option C (Encrypted Lambda environment variable): Storing credentials directly in Lambda environment variables, even when encrypted, requires custom code to manage rotation, which increases operational complexity.

Option D (KMS with automatic rotation): KMS is for managing encryption keys, not for storing and rotating secrets like database credentials. This option would require more custom implementation to manage credentials securely.

AWS Reference:

AWS Secrets Manager- Detailed documentation on how to store, manage, and rotate secrets using AWS Secrets Manager.

Using Secrets Manager with AWS Lambda- Guidance on integrating Secrets Manager with Lambda for secure credential management.

Question: 64

A company wants to standardize its Amazon Elastic Block Store (Amazon EBS) volume encryption strategy. The company also wants to minimize the cost and configuration effort required to operate the volume encryption check.

Which solution will meet these requirements?

- A. Write API calls to describe the EBS volumes and to confirm the EBS volumes are encrypted. Use Amazon EventBridge to schedule an AWS Lambda function to run the API calls.
- B. Write API calls to describe the EBS volumes and to confirm the EBS volumes are encrypted. Run the API calls on an AWS Fargate task.
- C. Create an AWS Identity and Access Management (IAM) policy that requires the use of tags on EBS volumes. Use AWS Cost Explorer to display resources that are not properly tagged. Encrypt the untagged resources manually.
- D. Create an AWS Config rule for Amazon EBS to evaluate if a volume is encrypted and to flag the volume if it is not encrypted.

Answer: D

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. By creating a Config rule, you can automatically check whether your Amazon EBS volumes are encrypted and flag those that are not, with minimal cost and configuration effort.

AWS Config Rule: AWS Config provides managed rules that you can use to automatically check the compliance of your resources against predefined or custom criteria. In this case, you would create a rule to evaluate EBS volumes and determine if they are encrypted. If a volume is not encrypted, the rule will flag it, allowing you to take corrective action.

Operational Overhead: This approach significantly reduces operational overhead because once the rule is in place, it continuously monitors your EBS volumes for compliance, and there's no need for manual checks or custom scripting.

Why Not Other Options?:

Option A (Lambda with API calls and EventBridge): While this can work, it involves writing and maintaining custom code, which increases operational overhead compared to using a managed AWS Config rule.

Option B (API calls on Fargate): Running API calls on Fargate is more complex and costly compared to using AWS Config, which provides a simpler, managed solution.

Option C (IAM policy with Cost Explorer): This option does not directly enforce encryption compliance and involves manual intervention, making it less efficient and more prone to errors.

AWS Reference:

AWS Config Rules- Overview of AWS Config rules and how they can be used to evaluate resource configurations.

Amazon EBS Encryption- Information on how to manage and enforce encryption for EBS volumes.

Question: 65

A company has migrated several applications to AWS in the past 3 months. The company wants to know the breakdown of costs for each of these applications. The company wants to receive a regular report that includes this information.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Budgets to download data for the past 3 months into a csv file. Look up the desired information.
- B. Load AWS Cost and Usage Reports into an Amazon RDS DB instance. Run SQL queries to get the desired information.
- C. Tag all the AWS resources with a key for cost and a value of the application's name. Activate cost allocation tags. Use Cost Explorer to get the desired information.
- D. Tag all the AWS resources with a key for cost and a value of the application's name. Use the AWS Billing and Cost Management console to download bills for the past 3 months. Look up the desired information.

Answer: C

Explanation:

This solution is the most cost-effective and efficient way to break down costs per application.

Tagging Resources: By tagging all AWS resources with a specific key (e.g., "cost") and a value representing the application's name, you can easily identify and categorize costs associated with each application. This tagging strategy allows for granular tracking of costs within AWS.

Activating Cost Allocation Tags: Once tags are applied to resources, you need to activate cost allocation tags in the AWS Billing and Cost Management console. This ensures that the costs associated with each tag are included in your billing reports and can be used for cost analysis.

AWS Cost Explorer: Cost Explorer is a powerful tool that allows you to visualize, understand, and manage your AWS costs and usage over time. You can filter and group your cost data by the tags you've applied to resources, enabling you to easily see the cost breakdown for each application. Cost Explorer also supports generating regular reports, which can be scheduled and emailed to stakeholders.

Why Not Other Options?:

Option A (AWS Budgets): AWS Budgets is more focused on setting cost and usage thresholds and monitoring them, rather than providing detailed cost breakdowns by application.

Option B (Load Cost and Usage Reports into RDS): This approach is less cost-effective and involves more operational overhead, as it requires setting up and maintaining an RDS instance and running SQL queries.

Option D (AWS Billing and Cost Management Console): While you can download bills, this method is more manual and less dynamic compared to using Cost Explorer with activated tags.

AWS Reference:

AWS Tagging Strategies- Overview of how to use tagging to organize and track AWS resources.

AWS Cost Explorer- Details on how to use Cost Explorer to analyze costs.

Question: 66

A company is migrating its on-premises Oracle database to an Amazon RDS for Oracle database. The company needs to retain data for 90 days to meet regulatory requirements. The company must also be able to restore the database to a specific point in time for up to 14 days.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create Amazon RDS automated backups. Set the retention period to 90 days.
- B. Create an Amazon RDS manual snapshot every day. Delete manual snapshots that are older than 90 days.
- C. Use the Amazon Aurora Clone feature for Oracle to create a point-in-time restore. Delete clones that are older than 90 days.
- D. Create a backup plan that has a retention period of 90 days by using AWS Backup for Amazon RDS.

Answer: D

Explanation:

AWS Backup is the most appropriate solution for managing backups with minimal operational overhead while meeting the regulatory requirement to retain data for 90 days and enabling point-in-time restore for up to 14 days.

AWS Backup: AWS Backup provides a centralized backup management solution that supports automated backup scheduling, retention management, and compliance reporting across AWS services, including Amazon RDS. By creating a backup plan, you can define a retention period (in this case, 90 days) and automate the backup process.

Point-in-Time Restore (PITR): Amazon RDS supports point-in-time restore for up to 35 days with automated backups. By using AWS Backup in conjunction with RDS, you ensure that your backup strategy meets the requirement for restoring data to a specific point in time within the last 14 days.

Why Not Other Options?:

Option A (RDS Automated Backups): While RDS automated backups support PITR, they do not directly support retention beyond 35 days without manual intervention.

Option B (Manual Snapshots): Manually creating and managing snapshots is operationally intensive and less automated compared to AWS Backup.

Option C (Aurora Clones): Aurora Clone is a feature specific to Amazon Aurora and is not applicable to Amazon RDS for Oracle.

AWS Reference:

AWS Backup- Overview of AWS Backup and its capabilities.

Amazon RDS Automated Backups- Information on how RDS automated backups work and their limitations.

Question: 67

A company is building an application in the AWS Cloud. The application is hosted on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses Amazon Route 53 for the DNS.

The company needs a managed solution with proactive engagement to detect against DDoS attacks.

Which solution will meet these requirements?

- A. Enable AWS Config. Configure an AWS Config managed rule that detects DDoS attacks.
- B. Enable AWS WAF on the ALB Create an AWS WAF web ACL with rules to detect and prevent DDoS attacks. Associate the web ACL with the ALB.
- C. Store the ALB access logs in an Amazon S3 bucket. Configure Amazon GuardDuty to detect and take automated preventative actions for DDoS attacks.
- D. Subscribe to AWS Shield Advanced. Configure hosted zones in Route 53 Add ALB resources as protected resources.

Answer: D

Explanation:

AWS Shield Advanced is designed to provide enhanced protection against DDoS attacks with proactive engagement and response capabilities, making it the best solution for this scenario.

AWS Shield Advanced: This service provides advanced protection against DDoS attacks. It includes detailed attack diagnostics, 24/7 access to the AWS DDoS Response Team (DRT), and financial protection against DDoS-related scaling charges. Shield Advanced also integrates with Route 53 and the Application Load Balancer (ALB) to ensure comprehensive protection for your web applications.

Route 53 and ALB Protection: By adding your Route 53 hosted zones and ALB resources to AWS Shield Advanced, you ensure that these components are covered under the enhanced protection plan. Shield Advanced actively monitors traffic and provides real-time attack mitigation, minimizing the impact of DDoS attacks on your application.

Why Not Other Options?:

Option A (AWS Config): AWS Config is a configuration management service and does not provide DDoS protection or detection capabilities.

Option B (AWS WAF): While AWS WAF can help mitigate some types of attacks, it does not provide the comprehensive DDoS protection and proactive engagement offered by Shield Advanced.

Option C (GuardDuty): GuardDuty is a threat detection service that identifies potentially malicious activity within your AWS environment, but it is not specifically designed to provide DDoS protection.

AWS Reference:

AWS Shield Advanced- Overview of AWS Shield Advanced and its DDoS protection capabilities.

Integrating AWS Shield Advanced with Route 53 and ALB- Detailed guidance on how to protect Route 53 and ALB with AWS Shield Advanced.

Question: 68

A company has an employee web portal. Employees log in to the portal to view payroll details. The company is developing a new system to give employees the ability to upload scanned documents for reimbursement. The company runs a program to extract text-based data from the documents and attach the extracted information to each employee's reimbursement IDs for processing.

The employee web portal requires 100% uptime. The document extract program runs infrequently throughout the day on an on-demand basis. The company wants to build a scalable and cost-effective new system that will require minimal changes to the existing web portal.

The company does not want to make any code changes.

Which solution will meet these requirements with the LEAST implementation effort?

- A. Run Amazon EC2 On-Demand Instances in an Auto Scaling group for the web portal. Use an AWS Lambda function to run the document extract program. Invoke the Lambda function when an employee uploads a new reimbursement document.
- B. Run Amazon EC2 Spot Instances in an Auto Scaling group for the web portal. Run the document extract program on EC2 Spot Instances. Start document extract program instances when an employee uploads a new reimbursement document.
- C. Purchase a Savings Plan to run the web portal and the document extract program. Run the web portal and the document extract program in an Auto Scaling group.
- D. Create an Amazon S3 bucket to host the web portal. Use Amazon API Gateway and an AWS Lambda function for the existing functionalities. Use the Lambda function to run the document extract program. Invoke the Lambda function when the API that is associated with a new document upload is called.

Answer: A

Explanation:

This solution offers the most scalable and cost-effective approach with minimal changes to the existing web portal and no code modifications.

Amazon EC2 On-Demand Instances in an Auto Scaling Group: Running the web portal on EC2 On-Demand instances ensures 100% uptime

and scalability. The Auto Scaling group will maintain the desired number of instances, automatically scaling up or down as needed, ensuring high availability for the employee web portal.

AWS Lambda for Document Extraction: Lambda is a serverless compute service that allows you to run code in response to events without provisioning or managing servers. By using Lambda to run the document extraction program, you can trigger the function whenever an employee uploads a document. This approach is cost-effective since you only pay for the compute time used by the Lambda function.

No Code Changes Required: This solution integrates with the existing infrastructure with minimal implementation effort and does not require any modifications to the web portal's code.

Why Not Other Options?:

Option B (Spot Instances): Spot Instances are not suitable for workloads requiring 100% uptime, as they can be terminated by AWS with short notice.

Option C (Savings Plan): A Savings Plan could reduce costs but does not address the requirement for running the document extraction program efficiently or without code changes.

Option D (S3 with API Gateway and Lambda): This would require significant changes to the existing web portal setup, including moving the portal to S3 and reconfiguring its architecture, which contradicts the requirement of minimal implementation effort and no code changes.

AWS Reference:

Amazon EC2 Auto Scaling- Information on how to use Auto Scaling for EC2 instances.

AWS Lambda- Overview of AWS Lambda and its use cases.

Question: 69

An ecommerce company wants a disaster recovery solution for its Amazon RDS DB instances that run Microsoft SQL Server Enterprise Edition. The company's current recovery point objective (RPO) and recovery time objective (RTO) are 24 hours.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a cross-Region read replica and promote the read replica to the primary instance
- B. Use AWS Database Migration Service (AWS DMS) to create RDS cross-Region replication.
- C. Use cross-Region replication every 24 hours to copy native backups to an Amazon S3 bucket
- D. Copy automatic snapshots to another Region every 24 hours.

Answer: D

Explanation:

This solution is the most cost-effective and meets the RPO and RTO requirements of 24 hours.

Automatic Snapshots: Amazon RDS automatically creates snapshots of your DB instance at regular intervals. By copying these snapshots to another AWS Region every 24 hours, you ensure that you have a backup available in a different geographic location, providing disaster recovery capability.

RPO and RTO: Since the company's RPO and RTO are both 24 hours, copying snapshots daily to another Region is sufficient. In the event of a disaster, you can restore the DB instance from the most recent snapshot in the target Region.

Why Not Other Options?:

Option A (Cross-Region Read Replica): This could provide a faster recovery time but is more costly due to the ongoing replication and resource usage in another Region.

Option B (DMS Cross-Region Replication): While effective for continuous replication, it introduces complexity and cost that isn't necessary given the 24-hour RPO/RTO.

Option C (Cross-Region Native Backup Copy): This involves more manual steps and doesn't offer as straightforward a solution as automated snapshot copying.

AWS Reference:

Amazon RDS Automated Backups and Snapshots- Details on automated backups and snapshots in RDS.

Copying an Amazon RDS DB Snapshot- How to copy DB snapshots to another Region.

Question: 70

An ecommerce company runs its application on AWS. The application uses an Amazon Aurora PostgreSQL cluster in Multi-AZ mode for the underlying database. During a recent promotional campaign, the application experienced heavy read load and write load. Users experienced timeout issues when they attempted to access the application.

A solutions architect needs to make the application architecture more scalable and highly available.

Which solution will meet these requirements with the LEAST downtime?

- A. Create an Amazon EventBridge rule that has the Aurora cluster as a source. Create an AWS Lambda function to log the state change events of the Aurora cluster. Add the Lambda function as a target for the EventBridge rule. Add additional reader nodes to fail over to.
- B. Modify the Aurora cluster and activate the zero-downtime restart (ZDR) feature. Use Database Activity Streams on the cluster to track the cluster status.
- C. Add additional reader instances to the Aurora cluster. Create an Amazon RDS Proxy target group for the Aurora cluster.
- D. Create an Amazon ElastiCache for Redis cache. Replicate data from the Aurora cluster to Redis by using AWS Database Migration Service (AWS DMS) with a write-around approach.

Answer: C

Explanation:

This solution directly addresses the scalability and high availability requirements with minimal downtime.

Additional Reader Instances: Adding more reader instances to the Aurora cluster will distribute the read load, improving the performance of the application under heavy read traffic. Aurora reader instances automatically replicate the data from the writer instance, enabling you to scale out read operations.

Amazon RDS Proxy: RDS Proxy improves database availability by managing database connections more efficiently and providing a connection pool. This reduces the overhead on the Aurora cluster during peak loads, further enhancing performance and availability without requiring changes to the application code.

Why Not Other Options?:

Option A (EventBridge and Lambda): This doesn't directly address the performance and availability issues. Logging state changes and adding reader nodes on failure events doesn't provide proactive scalability.

Option B (Zero-Downtime Restart and Activity Streams): Zero-Downtime Restart (ZDR) is useful for minimizing downtime during maintenance but doesn't directly improve scalability. Database Activity Streams are more for security monitoring than for performance enhancement.

Option D (ElastiCache for Redis): While adding a caching layer can help with read performance, it introduces complexity and may not be necessary if additional reader instances can handle the load.

AWS Reference:

Amazon Aurora Scaling- Information on scaling Aurora clusters with reader instances.

Amazon RDS Proxy- Details on how RDS Proxy can improve database performance and availability.

Question: 71

A company is developing a new application that uses a relational database to store user data and application configurations. The company expects the application to have steady user growth. The company expects the database usage to be variable and read-heavy, with occasional writes.

The company wants to cost-optimize the database solution. The company wants to use an AWS managed database solution that will provide the necessary performance.

Which solution will meet these requirements MOST cost-effectively?

A. Deploy the database on Amazon RDS. Use Provisioned IOPS SSD storage to ensure consistent performance for read and write operations.

B. Deploy the database on Amazon Aurora Serverless to automatically scale the database capacity based on actual usage to

accommodate the workload.

C. Deploy the database on Amazon DynamoDB. Use on-demand capacity mode to automatically scale throughput to accommodate the workload.

D. Deploy the database on Amazon RDS Use magnetic storage and use read replicas to accommodate the workload

Answer: B

Explanation:

Amazon Aurora Serverless is a cost-effective, on-demand, autoscaling configuration for Amazon Aurora. It automatically adjusts the database's capacity based on the current demand, which is ideal for workloads with variable and unpredictable usage patterns. Since the application is expected to be read-heavy with occasional writes and steady growth, Aurora Serverless can provide the necessary performance without requiring the management of database instances.

Cost-Optimization: Aurora Serverless only charges for the database capacity you use, making it a more cost-effective solution compared to always running provisioned database instances, especially for workloads with fluctuating demand.

Scalability: It automatically scales database capacity up or down based on actual usage, ensuring that you always have the right amount of resources available.

Performance: Aurora Serverless is built on the same underlying storage as Amazon Aurora, providing high performance and availability.

Why Not Other Options?:

Option A (RDS with Provisioned IOPS SSD): While Provisioned IOPS SSD ensures consistent performance, it is generally more expensive and less flexible compared to the autoscaling nature of Aurora Serverless.

Option C (DynamoDB with On-Demand Capacity): DynamoDB is a NoSQL database and may not be the best fit for applications requiring relational database features.

Option D (RDS with Magnetic Storage and Read Replicas): Magnetic storage is outdated and generally slower. While read replicas help with read-heavy workloads, the overall performance might not be optimal, and magnetic storage doesn't provide the necessary performance.

AWS Reference:

Amazon Aurora Serverless- Information on how Aurora Serverless works and its use cases.

Amazon Aurora Pricing- Details on the cost-effectiveness of Aurora Serverless.

Question: 72

A company is running a media store across multiple Amazon EC2 instances distributed across multiple Availability Zones in a single VPC. The company wants a high-performing solution to share data between all the EC2 instances, and prefers to keep the data within the VPC only.

What should a solutions architect recommend?

- A. Create an Amazon S3 bucket and call the service APIs from each instance's application.
- B. Create an Amazon S3 bucket and configure all instances to access it as a mounted volume.
- C. Configure an Amazon Elastic Block Store (Amazon EBS) volume and mount it across all instances.
- D. Configure an Amazon Elastic File System (Amazon EFS) file system and mount it across all instances.

Answer: D

Explanation:

Amazon Elastic File System (EFS) is a managed file storage service that can be mounted across multiple EC2 instances. It provides a scalable and high-performing solution to share data among instances within a VPC.

High Performance: EFS provides scalable performance for workloads that require high throughput and IOPS. It is particularly well-suited for applications that need to share data across multiple instances.

Ease of Use: EFS can be easily mounted on multiple instances across different Availability Zones, providing a shared file system accessible to all the instances within the VPC.

Security: EFS can be configured to ensure that data remains within the VPC, and it supports encryption at rest and in transit.

Why Not Other Options?:

Option A (Amazon S3 bucket with APIs): While S3 is excellent for object storage, it is not a file system and does not provide the low-latency access required for shared data between instances.

Option B (S3 bucket as a mounted volume): S3 is not designed to be mounted as a file system, and this approach would introduce unnecessary complexity and latency.

Option C (EBS volume shared across instances): EBS volumes cannot be attached to multiple instances simultaneously. It is not designed to be shared across instances like EFS.

AWS Reference:

Amazon EFS- Overview of Amazon EFS and its features.

Best Practices for Amazon EFS- Recommendations for using EFS with multiple instances.

Question: 73

A company currently runs an on-premises stock trading application by using Microsoft Windows Server. The company wants to migrate the application to the AWS Cloud. The company needs to design a highly available solution that provides low-latency access to block storage across multiple Availability Zones. Which solution will meet these requirements with the LEAST implementation effort?

- A. Configure a Windows Server cluster that spans two Availability Zones on Amazon EC2 instances. Install the application on both cluster nodes. Use Amazon FSx for Windows File Server as shared storage between the two cluster nodes.
- B. Configure a Windows Server cluster that spans two Availability Zones on Amazon EC2 instances. Install the application on both cluster nodes. Use Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp3) volumes as storage attached to the EC2 instances. Set up application-level replication to sync data from one EBS volume in one Availability Zone to another EBS volume in the second Availability Zone.
- C. Deploy the application on Amazon EC2 instances in two Availability Zones. Configure one EC2 instance as active and the second EC2 instance in standby mode. Use an Amazon FSx for NetApp ONTAP Multi-AZ file system to access the data by using Internet Small Computer Systems Interface (iSCSI) protocol.
- D. Deploy the application on Amazon EC2 instances in two Availability Zones. Configure one EC2 instance as active and the second EC2 instance in standby mode. Use Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io2) volumes as storage attached to the EC2 instances. Set up Amazon EBS level replication to sync data from one io2 volume in one Availability Zone to another io2 volume in the second Availability Zone.

Answer: A

Explanation:

This solution is designed to provide high availability and low-latency access to block storage across multiple Availability Zones with minimal implementation effort.

Windows Server Cluster Across AZs: Configuring a Windows Server Failover Cluster (WSFC) that spans two Availability Zones ensures that the application can failover from one instance to another in case of a failure, meeting the high availability requirement.

Amazon FSx for Windows File Server: FSx for Windows File Server provides fully managed Windows file storage that is accessible via the SMB protocol, which is suitable for Windows-based applications. It offers high availability and can be used as shared storage between the cluster nodes, ensuring that both nodes have access to the same data with low latency.

Why Not Other Options?:

Option B (EBS with application-level replication): This requires complex configuration and management, as EBS volumes cannot be directly shared across AZs. Application-level replication is more complex and prone to errors.

Option C (FSx for NetApp ONTAP with iSCSI): While this is a viable option, it introduces additional complexity with iSCSI and requires more specialized knowledge for setup and management.

Option D (EBS with EBS-level replication): EBS-level replication is not natively supported across AZs, and setting up a custom replication solution would increase the implementation effort.

AWS Reference:

Amazon FSx for Windows File Server- Overview and benefits of using FSx for Windows File Server.

Windows Server Failover Clustering on AWS- Guide on setting up a Windows Server cluster on AWS.

Question: 74

A company is designing an application on AWS that processes sensitive data. The application stores and processes financial data for multiple customers.

To meet compliance requirements, the data for each customer must be encrypted separately at rest by using a secure, centralized key management solution. The company wants to use AWS Key Management Service (AWS KMS) to implement encryption.

Which solution will meet these requirements with the LEAST operational overhead'?

- A. Generate a unique encryption key for each customer. Store the keys in an Amazon S3 bucket. Enable server-side encryption.
- B. Deploy a hardware security appliance in the AWS environment that securely stores customer- provided encryption keys. Integrate the security appliance with AWS KMS to encrypt the sensitive data in the application.
- C. Create a single AWS KMS key to encrypt all sensitive data across the application.
- D. Create separate AWS KMS keys for each customer's data that have granular access control and logging enabled.

Answer: D

Explanation:

This solution meets the requirement of encrypting each customer's data separately with the least operational overhead by leveraging AWS Key Management Service (KMS).

Separate AWS KMS Keys: By creating separate KMS keys for each customer, you can ensure that each customer's data is encrypted with a unique key. This approach satisfies the compliance requirement for separate encryption and provides fine-grained control over access to the keys.

Granular Access Control: AWS KMS allows you to define key policies and use IAM policies to grant specific permissions to the keys. This ensures that only authorized users or services can access the keys, thereby maintaining the principle of least privilege.

Logging and Monitoring: AWS KMS integrates with AWS CloudTrail, which logs all key usage and management activities. This provides an audit trail that is essential for meeting compliance requirements.

Why Not Other Options?:

Option A (Store keys in S3): Storing keys in S3 is not recommended because it does not provide the same level of security, access control, or integration with AWS services as KMS does.

Option B (Hardware security appliance): Deploying a hardware security appliance adds significant operational overhead and complexity, which is unnecessary given that KMS already provides a secure and centralized key management solution.

Option C (Single KMS key for all data): Using a single KMS key does not meet the requirement of encrypting each customer's data separately.

AWS Reference:

AWS Key Management Service (KMS)- Overview of KMS, its features, and best practices for key management.

Using AWS KMS for Multi-Tenant Applications- Guidance on how to design applications using KMS for multi-tenancy.

Question: 75

A company is building a cloud-based application on AWS that will handle sensitive customer data. The application uses Amazon RDS for the database. Amazon S3 for object storage, and S3 Event Notifications that invoke AWS Lambda for serverless processing.

The company uses AWS IAM Identity Center to manage user credentials. The development, testing, and operations teams need secure access to Amazon RDS and Amazon S3 while ensuring the confidentiality of sensitive customer data. The solution must comply with the principle of least privilege.

Which solution meets these requirements with the LEAST operational overhead?

- A. Use IAM roles with least privilege to grant all the teams access. Assign IAM roles to each team with customized IAM policies defining specific permission for Amazon RDS and S3 object access based on team responsibilities.
- B. Enable IAM Identity Center with an Identity Center directory. Create and configure permission sets with granular access to Amazon RDS and Amazon S3. Assign all the teams to groups that have specific access with the permission sets.
- C. Create individual IAM users for each member in all the teams with role-based permissions. Assign the IAM roles with predefined policies for RDS and S3 access to each user based on user needs. Implement IAM Access Analyzer for periodic credential evaluation.
- D. Use AWS Organizations to create separate accounts for each team. Implement cross-account IAM roles with least privilege Grant specific permission for RDS and S3 access based on team roles and responsibilities.

Answer: B

Explanation:

This solution allows for secure and least-privilege access with minimal operational overhead.

IAM Identity Center: AWS IAM Identity Center (formerly AWS SSO) enables you to centrally manage access to multiple AWS accounts and applications. By using IAM Identity Center, you can assign permission sets that define what users or groups can access, ensuring that only necessary permissions are granted.

Permission Sets: Permission sets in IAM Identity Center allow you to define granular access controls for specific services, such as Amazon RDS and S3. You can tailor these permissions to meet the needs of different teams, adhering to the principle of least privilege.

Group Management: By assigning users to groups and associating those groups with specific permission sets, you reduce the complexity and overhead of managing individual IAM roles and policies. This method also simplifies compliance and audit processes.

Why Not Other Options?:

Option A (IAM roles): While IAM roles can provide least-privilege access, managing multiple roles and policies across teams increases operational overhead compared to using IAM Identity Center.

Option C (Individual IAM users): Managing individual IAM users and roles can be cumbersome and does not scale well compared to group-based management in IAM Identity Center.

Option D (AWS Organizations with cross-account roles): Creating separate accounts and crossaccount roles adds unnecessary complexity and overhead for this use case, where IAM Identity Center provides a more straightforward solution.

AWS Reference:

[AWS IAM Identity Center- Overview and best practices for using IAM Identity Center.](#)

[Managing Access Permissions Using IAM Identity Center- Guide on creating and managing permission sets for secure access.](#)

Question: 76

A company has applications that run in an organization in AWS Organizations. The company outsources operational support of the applications. The company needs to provide access for the external support engineers without compromising security.

The external support engineers need access to the AWS Management Console. The external support engineers also need operating system access to the company's fleet of Amazon EC2 instances that run Amazon Linux in private subnets.

Which solution will meet these requirements MOST securely?

A. Confirm that AWS Systems Manager Agent (SSM Agent) is installed on all instances. Assign an instance profile with the necessary policy to connect to Systems Manager. Use AWS IAM IdentityCenter to provide the external support engineers console access. Use Systems Manager Session Manager to assign the required permissions.

B. Confirm that AWS Systems Manager Agent (SSM Agent) is installed on all instances. Assign an instance profile with the necessary policy to connect to Systems Manager. Use Systems Manager Session Manager to provide local IAM user credentials in each AWS account to the external support engineers for console access.

C. Confirm that all instances have a security group that allows SSH access only from the external support engineers source IP address ranges. Provide local IAM user credentials in each AWS account to the external support engineers for console access. Provide each external support engineer an SSH key pair to log in to the application instances.

D. Create a bastion host in a public subnet. Set up the bastion host security group to allow access from only the external engineers' IP address ranges. Ensure that all instances have a security group that allows SSH access from the bastion host. Provide each external support engineer an SSH key pair to log in to the application instances. Provide local account IAM user credentials to the engineers for console access.

Answer: A

Explanation:

This solution provides the most secure access for external support engineers with the least exposure to potential security risks.

AWS Systems Manager (SSM) and Session Manager: Systems Manager Session Manager allows secure and auditable access to EC2 instances without the need to open inbound SSH ports or manage SSH keys. This reduces the attack surface significantly. The SSM Agent must be installed and configured on all instances, and the instances must have an instance profile with the necessary IAM permissions to connect to Systems Manager.

IAM Identity Center: IAM Identity Center provides centralized management of access to the AWS Management Console for external support engineers. By using IAM Identity Center, you can control console access securely and ensure that external engineers have the appropriate permissions based on their roles.

Why Not Other Options?:

Option B (Local IAM user credentials): This approach is less secure because it involves managing local IAM user credentials and does not leverage the centralized management and security benefits of IAM Identity Center.

Option C (Security group with SSH access): Allowing SSH access opens up the infrastructure to potential security risks, even when restricted by IP addresses. It also requires managing SSH keys, which can be cumbersome and less secure.

Option D (Bastion host): While a bastion host can secure SSH access, it still requires managing SSH keys and opening ports. This approach is less secure and more operationally intensive compared to using Session Manager.

AWS Reference:

AWS Systems Manager Session Manager- Documentation on using Session Manager for secure instance access.

AWS IAM Identity Center- Overview of IAM Identity Center and its capabilities for managing user access.

Question: 77

A solutions architect needs to host a high performance computing (HPC) workload in the AWS Cloud. The workload will run on hundreds of Amazon EC2 instances and will require parallel access to a shared file system to enable distributed processing of large datasets. Datasets will be accessed across multiple instances simultaneously. The workload requires access latency within 1 ms. After processing has

completed, engineers will need access to the dataset for manual postprocessing.

Which solution will meet these requirements?

- A. Use Amazon Elastic File System (Amazon EFS) as a shared file system. Access the dataset from Amazon EFS.
- B. Mount an Amazon S3 bucket to serve as the shared file system. Perform postprocessing directly from the S3 bucket.
- C. Use Amazon FSx for Lustre as a shared file system. Link the file system to an Amazon S3 bucket for postprocessing.
- D. Configure AWS Resource Access Manager to share an Amazon S3 bucket so that it can be mounted to all instances for processing and postprocessing.

Answer: C

Explanation:

Amazon FSx for Lustre is the ideal solution for high-performance computing (HPC) workloads that require parallel access to a shared file system with low latency. FSx for Lustre is designed specifically to meet the needs of such workloads, offering sub-millisecond latencies, which makes it well-suited for the 1 ms latency requirement mentioned in the question.

Here is why FSx for Lustre is the best fit:

Parallel File System: FSx for Lustre is a parallel file system that can scale across hundreds of Amazon EC2 instances, providing high throughput and low-latency access to data. It is optimized for processing large datasets in parallel, which is essential for HPC workloads.

Low Latency: FSx for Lustre is capable of providing access latencies well within 1 ms, making it ideal for performance-sensitive workloads like HPC.

Seamless Integration with Amazon S3: FSx for Lustre can be linked to an Amazon S3 bucket. This integration allows data to be imported from S3 into FSx for Lustre before the workload begins and exported back to S3 after processing. This feature is crucial for manual postprocessing because it enables engineers to access the dataset in S3 after processing.

Performance: FSx for Lustre is built for workloads that require high performance, such as machine learning, analytics, media processing, and financial simulations, which are typical for HPC environments.

In contrast:

Amazon EFS (Option A): While EFS provides shared file storage and scales across multiple EC2 instances, it does not offer the same level of performance or sub-millisecond latencies as FSx for Lustre. EFS is more suited for general-purpose workloads, not high-performance computing.

Mounting S3 as a file system (Option B and D): S3 is object storage, not a file system designed for low-latency access and parallel processing. Mounting S3 buckets directly or using AWS Resource Access Manager to share the bucket would not meet the low-latency (1 ms) or performance requirements needed for HPC workloads.

Therefore, Amazon FSx for Lustre (Option C) is the most appropriate and verified solution for this scenario.

AWS Reference:

Amazon FSx for Lustre

Best Practices for High Performance Computing (HPC)

Amazon FSx and Amazon S3 Integration

Question: 78

A company is building a new furniture inventory application. The company has deployed the application on a fleet of Amazon EC2 instances across multiple Availability Zones. The EC2 instances run behind an Application Load Balancer (ALB) in their VPC.

A solutions architect has observed that incoming traffic seems to favor one EC2 instance, resulting in latency for some requests.

What should the solutions architect do to resolve this issue?

- A. Disable session affinity (sticky sessions) on the ALB.
- B. Replace the ALB with a Network Load Balancer.
- C. Increase the number of EC2 instances in each Availability Zone.
- D. Adjust the frequency of the health checks on the ALB's target group.

Answer: A

Explanation:

The issue described in the question, where incoming traffic seems to favor one EC2 instance, is often caused by session affinity (also known as sticky sessions) being enabled on the Application Load

Balancer (ALB). When session affinity is enabled, the ALB routes requests from the same client to the same EC2 instance. This can cause an imbalance in traffic distribution, leading to performance bottlenecks on certain instances while others remain underutilized.

To resolve this issue, disabling session affinity ensures that the ALB distributes incoming traffic evenly across all EC2 instances, allowing better load distribution and reducing latency. The ALB will rely on its round-robin or least outstanding requests algorithm (depending on the configuration) to distribute traffic more evenly across instances.

Option B (Network Load Balancer): The NLB is designed for Layer 4 (TCP) traffic and low latency use cases, but it is not needed here as the problem is with load balancing logic at the application layer (Layer 7). The ALB is more appropriate for HTTP/HTTPS traffic.

Option C (Increase EC2 Instances): Adding more EC2 instances does not solve the root issue of uneven traffic distribution.

Option D (Health Check Frequency): Adjusting health check frequency won't address the imbalance caused by session affinity.

AWS Reference:

Application Load Balancer Sticky Sessions

Question: 79

A company hosts its multi-tier, public web application in the AWS Cloud. The web application runs on Amazon EC2 instances, and its database runs on Amazon RDS. The company is anticipating a large increase in sales during an upcoming holiday weekend. A solutions architect needs to build a solution to analyze the performance of the web application with a granularity of no more than 2 minutes.

What should the solutions architect do to meet this requirement?

- A. Send Amazon CloudWatch logs to Amazon Redshift. Use Amazon QuickSight to perform further analysis.
- B. Enable detailed monitoring on all EC2 instances. Use Amazon CloudWatch metrics to perform further analysis.
- C. Create an AWS Lambda function to fetch EC2 logs from Amazon CloudWatch Logs. Use Amazon CloudWatch metrics to perform further analysis.
- D. Send EC2 logs to Amazon S3. Use Amazon Redshift to fetch logs from the S3 bucket to process raw data for further analysis with Amazon QuickSight.

Answer: B

Explanation:

To analyze the performance of the web application with granularity of no more than 2 minutes, enabling detailed monitoring on EC2 instances is the best solution. By default, CloudWatch provides metrics at a 5-minute interval. Enabling detailed monitoring allows you to collect metrics at 1-minute intervals, which will give you the level of granularity you need to analyze performance during peak traffic.

Amazon CloudWatch metrics can then be used to analyze CPU utilization, memory usage, disk I/O, and network throughput, among other performance-related metrics, at the desired granularity.

Option A: Sending CloudWatch logs to Redshift for analysis is unnecessary and overcomplicated for simple performance analysis, which can be done using CloudWatch metrics alone.

Option C: Fetching EC2 logs via Lambda adds complexity, and CloudWatch metrics already provide the required data for performance analysis.

Option D: Sending logs to S3 and using Redshift for analysis is also more complex than necessary for simple performance monitoring.

AWS Reference:

Monitoring Amazon EC2 with CloudWatch

Question: 80

A solutions architect is designing an application that helps users fill out and submit registration forms. The solutions architect plans to use a two-tier architecture that includes a web application server tier and a worker tier.

The application needs to process submitted forms quickly. The application needs to process each form exactly once. The solution must ensure that no data is lost.

Which solution will meet these requirements?

- A. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue between the web application server tier and the worker tier to store and forward form data.
- B. Use an Amazon API Gateway HTTP API between the web application server tier and the worker tier to store and forward form data.
- C. Use an Amazon Simple Queue Service (Amazon SQS) standard queue between the web application server tier and the worker tier to store and forward form data.
- D. Use an AWS Step Functions workflow. Create a synchronous workflow between the web application server tier and the worker tier that stores and forwards form data.

Answer: A

Explanation:

To process each form exactly once and ensure no data is lost, using an Amazon SQS FIFO (First-In- First-Out) queue is the most appropriate solution. SQS FIFO queues guarantee that messages are processed in the exact order they are sent and ensure that each message is processed exactly once. This ensures data consistency and reliability, both of which are crucial for processing user-submitted forms without data loss.

SQS acts as a buffer between the web application server and the worker tier, ensuring that submitted forms are stored reliably and forwarded to the worker tier for processing. This also decouples the application, improving its scalability and resilience.

Option B (API Gateway): API Gateway is better suited for API management rather than acting as a message queue for form processing.

Option C (SQS Standard Queue): While SQS Standard queues offer high throughput, they do not guarantee exactly-once processing or the strict ordering needed for this use case.

Option D (Step Functions): Step Functions are useful for orchestrating workflows but add unnecessary complexity for simple message queuing and form processing.

AWS Reference:

Amazon SQS FIFO Queues

Decoupling Application Tiers Using Amazon SQS

Question: 81

A company has a large data workload that runs for 6 hours each day. The company cannot lose any data while the process is running. A solutions architect is designing an Amazon EMR cluster configuration to support this critical data workload.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure a long-running cluster that runs the primary node and core nodes on On-Demand Instances and the task nodes on Spot Instances.
- B. Configure a transient cluster that runs the primary node and core nodes on On-Demand Instances and the task nodes on Spot Instances.
- C. Configure a transient cluster that runs the primary node on an On-Demand Instance and the core nodes and task nodes on Spot Instances.
- D. Configure a long-running cluster that runs the primary node on an On-Demand Instance, the core nodes on Spot Instances, and the task nodes on Spot Instances.

Answer: B

Explanation:

For cost-effectiveness and high availability in Amazon EMR workloads, the best approach is to configure a transient cluster (which runs for the duration of the job and then terminates) with On-Demand Instances for the primary and core nodes, and Spot Instances for the task nodes. Here's why:

Primary and core nodes on On-Demand Instances: These nodes are critical because they manage the cluster and store data on HDFS. Running them on On-Demand Instances ensures stability and that no data is lost, as Spot Instances can be interrupted.

Task nodes on Spot Instances: Task nodes handle additional processing and can be used with Spot Instances to reduce costs. Spot Instances are much cheaper but can be interrupted, which is fine for non-critical tasks as the framework can handle retries.

A transient cluster is more cost-effective than a long-running cluster for workloads that only run for 6 hours a day. Transient clusters automatically terminate after the workload completes, saving costs by not keeping the cluster running when it's not needed.

Option A: A long-running cluster may result in unnecessary costs when the cluster isn't being used.

Option C: Running core nodes on Spot Instances risks data loss if the Spot Instances are interrupted, violating the requirement for

zero data loss.

Option D: Running both core and task nodes on Spot Instances is highly risky for data-critical workloads.

AWS Reference:

Amazon EMR Cluster Management

Using Spot Instances in EMR

Question: 82

A company runs an application that stores and shares photos. Users upload the photos to an Amazon S3 bucket. Every day, users upload approximately 150 photos. The company wants to design a solution that creates a thumbnail of each new photo and stores the thumbnail in a second S3 bucket.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Amazon EventBridge scheduled rule to invoke a script every minute on a long-running Amazon EMR cluster. Configure the script to generate thumbnails for the photos that do not have thumbnails. Configure the script to upload the thumbnails to the second S3 bucket.
- B. Configure an Amazon EventBridge scheduled rule to invoke a script every minute on a memory-optimized Amazon EC2 instance that is always on. Configure the script to generate thumbnails for the photos that do not have thumbnails. Configure the script to upload the thumbnails to the second S3 bucket.
- C. Configure an S3 event notification to invoke an AWS Lambda function each time a user uploads a new photo to the application. Configure the Lambda function to generate a thumbnail and to upload the thumbnail to the second S3 bucket.
- D. Configure S3 Storage Lens to invoke an AWS Lambda function each time a user uploads a new photo to the application. Configure the Lambda function to generate a thumbnail and to upload the thumbnail to a second S3 bucket.

Answer: C

Explanation:

The most cost-effective and scalable solution for generating thumbnails when photos are uploaded to an S3 bucket is to use S3 event notifications to trigger an AWS Lambda function. This approach avoids the need for a long-running EC2 instance or EMR cluster, making it highly cost-effective because Lambda only charges for the time it takes to process each event.

S3 Event Notifications: Automatically triggers the Lambda function when a new photo is uploaded to the S3 bucket.

AWS Lambda: A serverless compute service that scales automatically and only charges for execution time, which makes it the most economical choice when dealing with periodic events like photo uploads.

The Lambda function can generate the thumbnail and upload it to a second S3 bucket, fulfilling the requirement efficiently.

Option A and Option B (EMR or EC2 with scheduled scripts)**: These are less cost-effective as they involve continuously running infrastructure, which incurs unnecessary costs.

Option D (S3 Storage Lens): S3 Storage Lens is a tool for storage analytics and is not designed for event-based photo processing.

AWS Reference:

Amazon S3 Event Notifications

AWS Lambda Pricing

Question: 83

A company uses Amazon RDS (or PostgreSQL) to run its applications in the us-east-1 Region. The company also uses machine learning (ML) models to forecast annual revenue based on near realtime reports. The reports are generated by using the same RDS for PostgreSQL

database. The database performance slows during business hours. The company needs to improve database performance.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a cross-Region read replica. Configure the reports to be generated from the read replica.
- B. Activate Multi-AZ DB instance deployment for RDS for PostgreSQL. Configure the reports to be generated from the standby database.
- C. Use AWS Data Migration Service (AWS DMS) to logically replicate data to a new database. Configure the reports to be generated from the new database.
- D. Create a read replica in us-east-1. Configure the reports to be generated from the read replica.

Answer: D

Explanation:

To improve the performance of the primary RDS PostgreSQL database during business hours and reduce the load, the best solution is to create a read replica in the same region (us-east-1). This will offload the read-heavy operations (like generating reports) to the replica, reducing the burden on the primary instance, which improves overall performance. Additionally, read replicas provide near realtime replication, making them ideal for real-time reporting use cases.

Option A (cross-Region read replica): This adds unnecessary latency for real-time reporting and increased costs due to cross-region data transfer.

Option B (Multi-AZ): Multi-AZ deployments are for high availability and disaster recovery but won't offload the read traffic, as the standby database cannot serve read requests.

Option C (AWS DMS replication): This adds complexity and is not as cost-effective as using an RDS read replica for the same region.

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

AWS Reference:

Amazon RDS Read Replicas

Amazon RDS Performance Best Practices

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Question: 84

A company has developed a non-production application that is composed of multiple microservices for each of the company's business units. A single development team maintains all the microservices.

The current architecture uses a static web frontend and a Java-based backend that contains the application logic. The architecture also uses a MySQL database that the company hosts on an Amazon EC2 instance.

The company needs to ensure that the application is secure and available globally.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon CloudFront and AWS Amplify to host the static web frontend. Refactor the microservices to use AWS Lambda functions that the microservices access by using Amazon API Gateway. Migrate the MySQL database to an Amazon EC2 Reserved Instance.
- B. Use Amazon CloudFront and Amazon S3 to host the static web frontend. Refactor the microservices to use AWS Lambda functions that the microservices access by using Amazon API Gateway. Migrate the MySQL database to Amazon RDS for MySQL.
- C. Use Amazon CloudFront and Amazon S3 to host the static web frontend. Refactor the microservices to use AWS Lambda functions that are in a target group behind a Network Load Balancer. Migrate the MySQL database to Amazon RDS for MySQL.
- D. Use Amazon S3 to host the static web frontend. Refactor the microservices to use AWS Lambda functions that are in a target group behind an Application Load Balancer. Migrate the MySQL database to an Amazon EC2 Reserved Instance.

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

Answer: B

Explanation:

This solution offers the least operational overhead while meeting the security and global availability requirements:

Amazon CloudFront and S3: Hosting the static frontend on S3 and serving it via CloudFront provides low-latency global distribution, high availability, and security. S3 is a cost-effective and serverless option for hosting static assets, and CloudFront ensures that the application is cached closer to the users, reducing latency globally.

AWS Lambda and API Gateway: Refactoring the microservices to use Lambda functions with API Gateway allows for a fully serverless, scalable, and highly available backend. This reduces the need for managing EC2 instances, as Lambda automatically scales to meet demand and only charges for the actual usage.

RDS for MySQL: Migrating the MySQL database from an EC2 instance to Amazon RDS significantly reduces operational overhead. RDS manages backups, patching, and scaling, and it offers high availability options (e.g., Multi-AZ).

Option A and D involve using EC2 Reserved Instances for the database, which requires more operational maintenance than using RDS.

Option C suggests using a Network Load Balancer with Lambda, which adds unnecessary complexity for this use case.

AWS Reference:

Amazon S3 and CloudFront Integration

AWS Lambda with API Gateway

Amazon RDS for MySQL

Question: 85

A company is designing a new internal web application in the AWS Cloud. The new application must securely retrieve and store multiple employee usernames and passwords from an AWS managed service. Which solution will meet these requirements with the LEAST operational overhead?

- A. Store the employee credentials in AWS Systems Manager Parameter Store. Use AWS Cloud Formation and the BatchGetSecretValue API to retrieve usernames and passwords from Parameter Store.
- B. Store the employee credentials in AWS Secrets Manager. Use AWS Cloud Formation and AWS Batch with the BatchGetSecretValue API to retrieve the usernames and passwords from Secrets Manager.
- C. Store the employee credentials in AWS Systems Manager Parameter Store. Use AWS Cloud Formation and AWS Batch with the BatchGetSecretValue API to retrieve the usernames and passwords from Parameter Store.
- D. Store the employee credentials in AWS Secrets Manager. Use AWS Cloud Formation and the BatchGetSecretValue API to retrieve the usernames and passwords from Secrets Manager.

Answer: D

Explanation:

AWS Secrets Manager is the best solution for securely storing and managing sensitive information, such as usernames and passwords. Secrets Manager provides automatic rotation, fine-grained access control, and encryption of credentials. It is designed to integrate easily with other AWS services, such as CloudFormation, to automate the retrieval of secrets via the BatchGetSecretValue API.

Secrets Manager has a lower operational overhead than manually managing credentials, and it offers features like automatic secret rotation that reduce the need for human intervention.

Option A and C (Parameter Store): While Systems Manager Parameter Store can store secrets, Secrets Manager provides more specialized capabilities for securely managing and rotating credentials with less operational overhead.

Option B and C (AWS Batch): Introducing AWS Batch unnecessarily complicates the solution. Secrets Manager already provides simple API calls for retrieving secrets without needing an additional service.

AWS Reference:

AWS Secrets Manager

Secrets Manager with CloudFormation

Question: 86

A company hosts an application on AWS. The application gives users the ability to upload photos and store the photos in an Amazon S3 bucket. The company wants to use Amazon CloudFront and a custom domain name to upload the photo files to the S3 bucket in the eu-west-1 Region.

Which solution will meet these requirements? (Select TWO.)

- A. Use AWS Certificate Manager (ACM) to create a public certificate in the us-east-1 Region. Use the certificate in CloudFront
- B. Use AWS Certificate Manager (ACM) to create a public certificate in eu-west-1. Use the certificate in CloudFront.
- C. Configure Amazon S3 to allow uploads from CloudFront. Configure S3 Transfer Acceleration.
- D. Configure Amazon S3 to allow uploads from CloudFront origin access control (OAC).
- E. Configure Amazon S3 to allow uploads from CloudFront. Configure an Amazon S3 website endpoint.

Answer: B,D

Explanation:

To upload photos to an S3 bucket using Amazon CloudFront with a custom domain name, the following components are required:

ACM in us-east-1 (Option A): When using CloudFront with HTTPS, the SSL/TLS certificate must be created in the us-east-1 Region. AWS Certificate Manager (ACM) handles the provisioning, management, and renewal of public certificates, making this a cost-effective and low-maintenance solution.

S3 Transfer Acceleration (Option C): Transfer Acceleration allows faster uploads to S3 from CloudFront by routing traffic through AWS's edge locations. This significantly speeds up the data upload process, especially for users that are geographically distant from the S3 bucket's region.

Option B (ACM in eu-west-1): CloudFront only supports certificates created in us-east-1.

Option D and E (OAC and website endpoint): These are not ideal for handling secure uploads or efficient data transfers in this case.

AWS Reference:

Using ACM with CloudFront

Amazon S3 Transfer Acceleration

Question: 87

A company maintains its accounting records in a custom application that runs on Amazon EC2 instances. The company needs to migrate the data to an AWS managed service for development and maintenance of the application data. The solution must require minimal operational support and provide immutable, cryptographically verifiable logs of data changes.

Which solution will meet these requirements MOST cost-effectively?

- A. Copy the records from the application into an Amazon Redshift cluster.
- B. Copy the records from the application into an Amazon Neptune cluster.
- C. Copy the records from the application into an Amazon Timestream database.
- D. Copy the records from the application into an Amazon Quantum Ledger Database (Amazon QLDB) ledger.

Answer: D

Explanation:

Amazon QLDB is the most cost-effective and suitable service for maintaining immutable, cryptographically verifiable logs of data changes. QLDB provides a fully managed ledger database with a built-in cryptographic hash chain, making it ideal for recording changes to accounting records, ensuring data integrity and security.

QLDB reduces operational overhead by offering fully managed services, so there's no need for server management, and it's built specifically to ensure immutability and verifiability, making it the best fit for the given requirements.

Option A (Redshift): Redshift is designed for analytics and not for immutable, cryptographically verifiable logs.

Option B (Neptune): Neptune is a graph database, which is not suitable for this use case.

Option C (Timestream): Timestream is a time series database optimized for time-stamped data, but it does not provide immutable or cryptographically verifiable logs.

AWS Reference:

Amazon QLDB

How QLDB Works

Question: 88

A company hosts a video streaming web application in a VPC. The company uses a Network Load Balancer (NLB) to handle TCP traffic for real-time data processing. There have been unauthorized attempts to access the application.

The company wants to improve application security with minimal architectural change to prevent unauthorized attempts to access the application.

Which solution will meet these requirements?

- A. Implement a series of AWS WAF rules directly on the NLB to filter out unauthorized traffic.
- B. Recreate the NLB with a security group to allow only trusted IP addresses.
- C. Deploy a second NLB in parallel with the existing NLB configured with a strict IP address allow list.
- D. Use AWS Shield Advanced to provide enhanced DDoS protection and prevent unauthorized access attempts.

Answer: D

Explanation:

Question: 89

A company runs its workloads on Amazon Elastic Container Service (Amazon ECS). The container images that the ECS task definition uses need to be scanned for Common Vulnerabilities and Exposures (CVEs). New container images that are created also need to be scanned.

Which solution will meet these requirements with the FEWEST changes to the workloads?

- A. Use Amazon Elastic Container Registry (Amazon ECR) as a private image repository to store the container images. Specify scan on push filters for the ECR basic scan.
- B. Store the container images in an Amazon S3 bucket. Use Amazon Macie to scan the images. Use an S3 Event Notification to initiate a Made scan for every event with an s3:ObjectCreated:Put event type
- C. Deploy the workloads to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon Elastic Container Registry (Amazon ECR) as a private image repository. Specify scan on push filters for the ECR enhanced scan.
- D. Store the container images in an Amazon S3 bucket that has versioning enabled. Configure an S3 Event Notification for s3:ObjectCreated:* events to invoke an AWS Lambda function. Configure the Lambda function to initiate an Amazon Inspector scan.

Answer: A

Explanation:

Question: 90

A company has an application that runs on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2 instances. The application has a U1 that uses Amazon DynamoDB and data services that use Amazon S3 as part of the application deployment.

The company must ensure that the EKS Pods for the U1 can access only Amazon DynamoDB and that the EKS Pods for the data services can access only Amazon S3. The company uses AWS Identity and Access Management (IAM).

Which solution meets these requirements?

- A. Create separate IAM policies (or Amazon S3 and DynamoDB access with the required permissions. Attach both IAM policies to the EC2 instance profile. Use role-based access control (RBAC) to control access to Amazon S3 or DynamoDB (or the respective EKS Pods.
- B. Create separate IAM policies (or Amazon S3 and DynamoDB access with the required permissions. Attach the Amazon S3 IAM policy directly to the EKS Pods (or the data services and the DynamoDB policy to the EKS Pods for the U1.
- C. Create separate Kubernetes service accounts for the U1 and data services to assume an IAM role. Attach the Amazon S3 Full Access policy to the data services account and the AmazonDynamoDBFullAccess policy to the U1 service account.
- D. Create separate Kubernetes service accounts for the U1 and data services to assume an IAM role. Use IAM Role for Service Accounts (IRSA) to provide access to the EKS Pods for the U1 to Amazon S3 and the EKS Pods for the data services to DynamoDB.

Answer: A

Explanation:

Question: 91

A company needs to give a globally distributed development team secure access to the company's AWS resources in a way that complies with security policies.

The company currently uses an on-premises Active Directory for internal authentication. The company uses AWS Organizations to manage multiple AWS accounts that support multiple projects.

The company needs a solution to integrate with the existing infrastructure to provide centralized identity management and access control.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Set up AWS Directory Service to create an AWS managed Microsoft Active Directory on AWS. Establish a trust relationship with the on-premises Active Directory. Use IAM roles that are assigned to Active Directory groups to access AWS resources within the company's AWS accounts.
- B. Create an IAM user for each developer. Manually manage permissions for each IAM user based on each user's involvement with each

project. Enforce multi-factor authentication (MFA) as an additional layer of security.

C. Use AD Connector in AWS Directory Service to connect to the on-premises Active Directory. Integrate AD Connector with AWS IAM Identity Center. Configure permissions sets to give each AD group access to specific AWS accounts and resources.

D. Use Amazon Cognito to deploy an identity federation solution. Integrate the identity federation solution with the on-premises Active Directory. Use Amazon Cognito to provide access tokens for developers to access AWS accounts and resources.

Answer: C

Explanation:

Using AD Connector with AWS IAM Identity Center (formerly AWS Single Sign-On) allows the company to leverage its existing on-premises Active Directory for centralized identity management and access control. AD Connector acts as a proxy to the on-premises AD without requiring additional infrastructure or complex setup. This solution integrates seamlessly with AWS, allowing the development team to use their existing AD credentials to access AWS resources across multiple accounts managed by AWS Organizations. The permissions for AWS resources can be managed centrally through IAM Identity Center by configuring permission sets.

This solution provides:

Least operational overhead: AD Connector is fully managed, and IAM Identity Center allows centralized management of permissions across accounts.

Secure access: The solution complies with security policies by using existing AD authentication mechanisms.

Option A (AWS Managed AD): Setting up a fully managed AWS AD and establishing a trust is more complex and involves additional operational overhead.

Option B (IAM Users): Manually managing IAM users and permissions is less scalable and increases operational complexity.

Option D (Cognito): Amazon Cognito is more suited for user-facing applications rather than internal identity management for AWS resources.

AWS Reference:

AD Connector with IAM Identity Center

AWS IAM Identity Center

Question: 92

A company wants to improve the availability and performance of its hybrid application. The application consists of a stateful TCP-based workload hosted on Amazon EC2 instances in different AWS Regions and a stateless UDP-based workload hosted on premises.

Which combination of actions should a solutions architect take to improve availability and performance? (Select TWO.)

- A. Create an accelerator using AWS Global Accelerator. Add the load balancers as endpoints.
- B. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the load balancers.
- C. Configure two Application Load Balancers in each Region. The first will route to the EC2 endpoints, and the second will route to the on-premises endpoints.
- D. Configure a Network Load Balancer in each Region to address the EC2 endpoints. Configure a Network Load Balancer in each Region that routes to the on-premises endpoints.
- E. Configure a Network Load Balancer in each Region to address the EC2 endpoints. Configure an Application Load Balancer in each Region that routes to the on-premises endpoints.

Answer: A,D

Explanation:

For improving availability and performance of the hybrid application, the following solutions are optimal:

AWS Global Accelerator (Option A): Global Accelerator provides high availability and improves performance by using the AWS global network to route user traffic to the nearest healthy endpoint (across AWS Regions). By adding the Network Load Balancers as endpoints, Global Accelerator ensures that traffic is routed efficiently to the closest endpoint, improving both availability and performance.

Network Load Balancer (Option D): The stateful TCP-based workload hosted on Amazon EC2 instances and the stateless UDP-based workload hosted on-premises are best served by Network Load

Balancers (NLBs). NLBs are designed to handle TCP and UDP traffic with ultra-low latency and can route traffic to both EC2 and on-premises endpoints.

Option B (CloudFront and Route 53): CloudFront is better suited for HTTP/HTTPS workloads, not for TCP/UDP-based applications.

Option C (ALB): Application Load Balancers do not support the stateless UDP-based workload, making NLBs the better choice for both TCP and UDP.

AWS Reference:

AWS Global Accelerator

Network Load Balancer

Question: 93

A company runs a production database on Amazon RDS for MySQL. The company wants to upgrade the database version for security compliance reasons. Because the database contains critical data, the company wants a quick solution to upgrade and test functionality without losing any data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an RDS manual snapshot. Upgrade to the new version of Amazon RDS for MySQL.
- B. Use native backup and restore. Restore the data to the upgraded new version of Amazon RDS for MySQL.
- C. Use AWS Database Migration Service (AWS DMS) to replicate the data to the upgraded new version of Amazon RDS for MySQL.
- D. Use Amazon RDS Blue/Green Deployments to deploy and test production changes.

Answer: D

Explanation:

Amazon RDS Blue/Green Deployments is the ideal solution for upgrading the database version with minimal operational overhead and no data loss. Blue/Green Deployments allows you to create a separate, fully managed "green" environment with the upgraded database version. You can test the new version in the green environment while the "blue" environment continues serving production traffic. Once testing is complete, you can seamlessly switch traffic to the green environment without downtime.

This solution provides:

Fast, non-disruptive upgrade: Traffic is only switched to the new environment after testing, ensuring zero data loss.

Minimal operational overhead: AWS handles the infrastructure management, reducing manual intervention.

Option A (Manual snapshot): This requires manual intervention and involves more operational overhead.

Option B (Native backup/restore): This approach is more labor-intensive and slower than Blue/Green Deployments.

Option C (DMS): AWS DMS adds unnecessary complexity for a simple version upgrade when Blue/Green Deployments can handle the task more efficiently.

AWS Reference:

Amazon RDS Blue/Green Deployments

Question: 94

A digital image processing company wants to migrate its on-premises monolithic application to the AWS Cloud. The company processes

thousands of images and generates large files as part of the processing workflow.

The company needs a solution to manage the growing number of image processing jobs. The solution must also reduce the manual tasks in the image processing workflow. The company does not want to manage the underlying infrastructure of the solution.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 Spot Instances to process the images. Configure Amazon Simple Queue Service (Amazon SQS) to orchestrate the workflow. Store the processed files in Amazon Elastic File System (Amazon EFS).
- B. Use AWS Batch jobs to process the images. Use AWS Step Functions to orchestrate the workflow. Store the processed files in an Amazon S3 bucket.
- C. Use AWS Lambda functions and Amazon EC2 Spot Instances to process the images. Store the processed files in Amazon FSx.
- D. Deploy a group of Amazon EC2 instances to process the images. Use AWS Step Functions to orchestrate the workflow. Store the processed files in an Amazon Elastic Block Store (Amazon EBS) volume.

Answer: B

Explanation:

For processing thousands of images and generating large files while minimizing manual tasks and operational overhead, using AWS Batch is the best solution. AWS Batch allows you to run large-scale, parallel, and managed batch computing jobs without needing to manage the underlying infrastructure.

AWS Batch: Automates the image processing jobs, dynamically allocating the necessary resources based on the job requirements, which reduces operational overhead.

AWS Step Functions: Orchestrates the entire image processing workflow, ensuring that tasks are executed in the correct sequence, improving manageability.

Amazon S3: Stores the processed files, providing scalable and cost-effective storage.

Option A (ECS with EC2 Spot Instances): While cost-effective, managing ECS and Spot Instances involves more operational effort.

Option C (Lambda with EC2 Spot): Lambda functions have size and duration limitations, making them less suited for large image processing tasks.

Option D (EC2 with Step Functions): Managing EC2 instances involves more overhead than using AWS Batch.

AWS Reference:

AWS Batch

AWS Step Functions

Question: 95

A company is developing an application in the AWS Cloud. The application's HTTP API contains critical information that is published in Amazon API Gateway. The critical information must be accessible from only a limited set of trusted IP addresses that belong to the company's internal network.

Which solution will meet these requirements?

- A. Set up an API Gateway private integration to restrict access to a predefined set of IP addresses.
- B. Create a resource policy for the API that denies access to any IP address that is not specifically allowed.
- C. Directly deploy the API in a private subnet. Create a network ACL. Set up rules to allow the traffic from specific IP addresses.
- D. Modify the security group that is attached to API Gateway to allow inbound traffic from only the trusted IP addresses.

Answer: B

Explanation:

Amazon API Gateway supports resource policies, which allow you to control access to your API by specifying the IP addresses or ranges that can access the API. By creating a resource policy that explicitly denies access to any IP address outside the allowed set, you can ensure that only trusted IP addresses (such as those from your internal network) can access the critical information in your API. This approach provides fine-grained access control without the need for additional infrastructure or complex configurations.

Option A (Private integration): API Gateway private integrations are for creating private APIs that are only accessible within a VPC, but this solution is about restricting access to certain IP addresses.

Option C (Private subnet and ACLs): Deploying the API in a private subnet and using network ACLs adds unnecessary complexity and isn't the best fit for HTTP APIs.

Option D (Security group): API Gateway doesn't have a security group because it isn't a resource inside a VPC. Instead, resource policies are the correct mechanism for controlling IP-based access.

AWS Reference:

Controlling Access to API Gateway with Resource Policies

Question: 96

A company runs its databases on Amazon RDS for PostgreSQL. The company wants a secure solution to manage the master user password by rotating the password every 30 days. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon EventBridge to schedule a custom AWS Lambda function to rotate the password every 30 days.

- B. Use the modify-db-instance command in the AWS CLI to change the password.
- C. Integrate AWS Secrets Manager with Amazon RDS for PostgreSQL to automate password rotation.
- D. Integrate AWS Systems Manager Parameter Store with Amazon RDS for PostgreSQL to automate password rotation.

Answer: C

Explanation:

AWS Secrets Manager can integrate directly with Amazon RDS for automatic and seamless password rotation. Secrets Manager handles the complexity of password management, including generating strong passwords and rotating them at a defined interval (e.g., every 30 days). It also automatically updates the connection information for RDS, minimizing operational overhead.

Option A (Lambda with EventBridge): While possible, this requires custom coding and operational management of Lambda, which introduces additional complexity.

Option B (Manual password change): Using the modify-db-instance command requires manual intervention and is not automated, leading to more operational effort.

Option D (Parameter Store): Systems Manager Parameter Store is less specialized for password management than Secrets Manager and does not have built-in automated rotation for RDS credentials.

AWS Reference:

AWS Secrets Manager Rotation for RDS

Question: 97

A company wants to implement new security compliance requirements for its development team to limit the use of approved Amazon Machine Images (AMIs).

The company wants to provide access to only the approved operating system and software for all its Amazon EC2 instances. The company wants the solution to have the least amount of lead time for launching EC2 instances.

Which solution will meet these requirements?

- A. Create a portfolio by using AWS Service Catalog that includes only EC2 instances launched with approved AMIs. Ensure that all required software is preinstalled on the AMIs. Create the necessary permissions for developers to use the portfolio.
- B. Create an AMI that contains the approved operating system and software by using EC2 Image Builder. Give developers access to that AMI to launch the EC2 instances.
- C. Create an AMI that contains the approved operating system Tell the developers to use the approved AMI Create an Amazon

EventBridge rule to run an AWS Systems Manager script when a new EC2 instance is launched. Configure the script to install the required software from a repository.

D. Create an AWS Config rule to detect the launch of EC2 instances with an AMI that is not approved. Associate a remediation rule to terminate those instances and launch the instances again with the approved AMI. Use AWS Systems Manager to automatically install the approved software on the launch of an EC2 instance.

Answer: A

Explanation:

AWS Service Catalog is designed to allow organizations to manage a catalog of approved products (including AMIs) that users can deploy. By creating a portfolio that contains only EC2 instances launched with preapproved AMIs, the company can enforce compliance with the approved operating system and software for all EC2 instances. Service Catalog also streamlines the process of launching EC2 instances, reducing the lead time while ensuring that developers use only the approved configurations.

Option B (EC2 Image Builder): While EC2 Image Builder helps in creating and managing AMIs, it doesn't provide the enforcement mechanism that Service Catalog does.

Option C (EventBridge rule and Systems Manager script): This solution is reactive and involves more operational complexity compared to Service Catalog.

Option D (AWS Config rule): This option is reactive (it terminates non-compliant instances after launch) and introduces additional operational overhead.

AWS Reference:

AWS Service Catalog

Question: 98

A company needs a solution to automate email ingestion. The company needs to automatically parse email messages, look for email attachments, and save any attachments to an Amazon S3 bucket in near real time. Email volume varies significantly from day to day.

Which solution will meet these requirements?

A. Set up email receiving in Amazon Simple Email Service (Amazon SES). Create a rule set and a receipt rule. Create an AWS Lambda function that Amazon SES can invoke to process the email bodies and attachments.

B. Set up email content filtering in Amazon Simple Email Service (Amazon SES). Create a content filtering rule based on sender, recipient, message body, and attachments.

C. Set up email receiving in Amazon Simple Email Service (Amazon SES). Configure Amazon SES and S3 Event Notifications to

process the email bodies and attachments.

D. Create an AWS Lambda function to process the email bodies and attachments. Use Amazon EventBridge to invoke the Lambda function. Configure an EventBridge rule to listen for incoming emails.

Answer: A

Explanation:

AmazonSES (Simple Email Service) allows for the automatic ingestion of incoming emails. By setting up email receiving in SES and creating a rule set with a receipt rule, you can configure SES to invoke an AWS Lambda function whenever an email is received. The Lambda function can then process the email body and attachments, saving any attachments to an Amazon S3 bucket. This solution is highly scalable, cost-effective, and provides near real-time processing of emails with minimal operational overhead.

Option B (Content filtering): This only filters emails based on content and does not provide the functionality to save attachments to S3.

Option C (S3 Event Notifications): While SES can store emails in S3, SES with Lambda offers more flexibility for processing attachments in real-time.

Option D (EventBridge rule): EventBridge cannot directly listen for incoming emails, making this solution incorrect.

AWS Reference:

Receiving Email with Amazon SES

Invoking Lambda from SES

Question: 99

A company tracks customer satisfaction by using surveys that the company hosts on its website. The surveys sometimes reach thousands of customers every hour. Survey results are currently sent in email messages to the company so company employees can manually review results and assess customer sentiment.

The company wants to automate the customer survey process. Survey results must be available for the previous 12 months.

Which solution will meet these requirements in the MOST scalable way?

A. Send the survey results data to an Amazon API Gateway endpoint that is connected to an Amazon Simple Queue Service (Amazon SQS) queue. Create an AWS Lambda function to poll the SQS queue, call Amazon Comprehend for sentiment analysis, and save the results to an Amazon DynamoDB table. Set the TTL for all records to 365 days in the future.

B. Send the survey results data to an API that is running on an Amazon EC2 instance. Configure the API to store the survey results as a new record in an Amazon DynamoDB table, call Amazon Comprehend for sentiment analysis, and save the results in a

second DynamoDB table. Set the TTL for all records to 365 days in the future.

C. Write the survey results data to an Amazon S3 bucket. Use S3 Event Notifications to invoke an AWS Lambda function to read the data and call Amazon Rekognition for sentiment analysis. Store the sentiment analysis results in a second S3 bucket. Use S3 Lifecycle policies on each bucket to expire objects after 365 days.

D. Send the survey results data to an Amazon API Gateway endpoint that is connected to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the SQS queue to invoke an AWS Lambda function that calls Amazon Lex for sentiment analysis and saves the results to an Amazon DynamoDB table. Set the TTL for all records to 365 days in the future.

Answer: A

Explanation:

This solution is the most scalable and efficient way to handle large volumes of survey data while automating sentiment analysis:

API Gateway and SQS: The survey results are sent to API Gateway, which forwards the data to an SQS queue. SQS can handle large volumes of messages and ensures that messages are not lost.

AWS Lambda: Lambda is triggered by polling the SQS queue, where it processes the survey data.

Amazon Comprehend: Comprehend is used for sentiment analysis, providing insights into customer satisfaction.

DynamoDB with TTL: Results are stored in DynamoDB with a Time to Live (TTL) attribute set to expire after 365 days, automatically removing old data and reducing storage costs.

Option B (EC2 API): Running an API on EC2 requires more maintenance and scalability management compared to API Gateway.

Option C (S3 and Rekognition): Amazon Rekognition is for image and video analysis, not sentiment analysis.

Option D (Amazon Lex): Amazon Lex is used for building conversational interfaces, not sentiment analysis.

AWS Reference:

Amazon Comprehend for Sentiment Analysis

Amazon SQS

DynamoDB TTL

Question: 100

An online gaming company is transitioning user data storage to Amazon DynamoDB to support the company's growing user base. The current architecture includes DynamoDB tables that contain user profiles, achievements, and in-game transactions.

The company needs to design a robust, continuously available, and resilient DynamoDB architecture to maintain a seamless gaming experience for users.

Which solution will meet these requirements MOST cost-effectively?

- A. Create DynamoDB tables in a single AWS Region. Use on-demand capacity mode. Use global tables to replicate data across multiple Regions.
- B. Use DynamoDB Accelerator (DAX) to cache frequently accessed data. Deploy tables in a single AWS Region and enable auto scaling. Configure Cross-Region Replication manually to additional Regions.
- C. Create DynamoDB tables in multiple AWS Regions. Use on-demand capacity mode. Use DynamoDB Streams for Cross-Region Replication between Regions.
- D. Use DynamoDB global tables for automatic multi-Region replication. Deploy tables in multiple AWS Regions. Use provisioned capacity mode. Enable auto scaling.

Answer: D

Explanation:

DynamoDB Global Tables provide a fully managed, multi-region, and multi-master database solution that allows you to deploy DynamoDB tables in multiple AWS Regions. This ensures high availability and resiliency across different geographical locations, providing a seamless gaming experience for users. Using provisioned capacity mode with auto-scaling ensures cost-efficiency by scaling up or down based on actual demand.

Option A: While on-demand capacity mode is flexible, provisioned capacity with auto-scaling is more cost-effective for predictable workloads.

Option B (DAX): DAX improves read performance, but it doesn't provide the multi-region replication needed for high availability and resiliency.

Option C: DynamoDB Streams with manual cross-region replication adds more complexity and operational overhead compared to Global Tables.

AWS Reference:

DynamoDB Global Tables

Question: 101

A logistics company is creating a data exchange platform to share shipment status information with shippers. The logistics company can see all shipment information and metadata. The company distributes shipment data updates to shippers.

Each shipper should see only shipment updates that are relevant to their company. Shippers should not see the full detail that is visible to

the logistics company. The company creates an Amazon Simple Notification Service (Amazon SNS) topic for each shipper to share data.

Some shippers use a mobile app to submit shipment status updates.

The company needs to create a data exchange platform that provides each shipper specific access to the data that is relevant to their company.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Ingest the shipment updates from the mobile app into Amazon Simple Queue Service (Amazon SQS). Publish the updates to the SNS topic. Apply a filter policy to rewrite the body of each message.
- B. Ingest the shipment updates from the mobile app into Amazon Simple Queue Service (Amazon SQS). Use an AWS Lambda function to consume the updates from Amazon SQS and rewrite the body of each message. Publish the updates to the SNS topic.
- C. Ingest the shipment updates from the mobile app into a second SNS topic. Publish the updates to the shipper SNS topic. Apply a filter policy to rewrite the body of each message.
- D. Ingest the shipment updates from the mobile app into Amazon Simple Queue Service (Amazon SQS). Filter and rewrite the messages in Amazon EventBridge Pipes. Publish the updates to the SNS topic.

Answer: B

Explanation:

The best solution is to use Amazon SQS to receive updates from the mobile app and process them with an AWS Lambda function. The Lambda function can rewrite the message body as necessary for each shipper and then publish the updates to the appropriate SNS topic for distribution. This setup ensures that each shipper receives only the relevant data and minimizes operational overhead by using managed services.

Option A (SNS filter policy): SNS does not have the capability to rewrite message bodies before forwarding.

Option C (Second SNS topic): Using an additional SNS topic adds unnecessary complexity without solving the message rewriting requirement.

Option D (EventBridge Pipes): EventBridge Pipes is more complex than necessary for this use case, and Lambda can handle the logic more efficiently.

AWS Reference:

Amazon SQS

Amazon SNS with Lambda

Question: 102

A company recently launched a new application for its customers. The application runs on multiple Amazon EC2 instances across two Availability Zones. End users use TCP to communicate with the application.

The application must be highly available and must automatically scale as the number of users increases.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Add a Network Load Balancer in front of the EC2 instances.
- B. Configure an Auto Scaling group for the EC2 instances.
- C. Add an Application Load Balancer in front of the EC2 instances.
- D. Manually add more EC2 instances for the application.
- E. Add a Gateway Load Balancer in front of the EC2 instances.

Answer: A,B

Explanation:

For an application requiring TCP communication and high availability:

Network Load Balancer (NLB) is the best choice for load balancing TCP traffic because it is designed for handling high-throughput, low-latency connections.

Auto Scaling group ensures that the application can automatically scale based on demand, adding or removing EC2 instances as needed, which is crucial for handling user growth.

Option C (Application Load Balancer): ALB is primarily for HTTP/HTTPS traffic, not ideal for TCP.

Option D (Manual scaling): Manually adding instances does not provide the automation or scalability required.

Option E (Gateway Load Balancer): GLB is used for third-party virtual appliances, not for direct application load balancing.

AWS Reference:

Network Load Balancer

Auto Scaling Group

Question: 103

A company needs to set up a centralized solution to audit API calls to AWS for workloads that run on AWS services and non AWS services. The company must store logs of the audits for 7 years.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Set up a data lake in Amazon S3. Incorporate AWS CloudTrail logs and logs from non AWS services into the data lake. Use CloudTrail to store the logs for 7 years.
- B. Configure custom integrations for AWS CloudTrail Lake to collect and store CloudTrail events from AWS services and non AWS services. Use CloudTrail to store the logs for 7 years.
- C. Enable AWS CloudTrail for AWS services. Ingest non AWS services into CloudTrail to store the logs for 7 years
- D. Create new Amazon CloudWatch Logs groups. Send the audit data from non AWS services to the CloudWatch Logs groups. Enable AWS CloudTrail for workloads that run on AWS. Use CloudTrail to store the logs for 7 years.

Answer: B

Explanation:

AWS CloudTrail Lake is a fully managed service that allows the collection, storage, and querying of CloudTrail events for both AWS and non-AWS services. CloudTrail Lake can be customized to collect logs from various sources, ensuring a centralized audit solution. It also supports long-term storage, so logs can be retained for 7 years, meeting the compliance requirement.

Option A (Data Lake): Setting up a data lake in S3 introduces unnecessary operational complexity compared to CloudTrail Lake.

Option C (Ingest non-AWS services into CloudTrail): CloudTrail Lake is better suited for this task with less operational overhead.

Option D (CloudWatch Logs): While CloudWatch can store logs, CloudTrail Lake is specifically designed for API auditing and storage.

AWS Reference:

AWS CloudTrail Lake

Question: 104

A company runs its production workload on Amazon EC2 instances with Amazon Elastic Block Store (Amazon EBS) volumes. A solutions architect needs to analyze the current EBS volume cost and to recommend optimizations. The recommendations need to include estimated monthly saving opportunities.

Which solution will meet these requirements?

- A. Use Amazon Inspector reporting to generate EBS volume recommendations for optimization.
- B. Use AWS Systems Manager reporting to determine EBS volume recommendations for optimization.
- C. Use Amazon CloudWatch metrics reporting to determine EBS volume recommendations for optimization.

D. Use AWS Compute Optimizer to generate EBS volume recommendations for optimization.

Answer: D

Explanation:

AWS Compute Optimizer provides detailed recommendations for optimizing Amazon EBS volumes, including insights into underutilized volumes, inefficient configurations, and potential cost savings. It analyzes usage patterns and provides recommendations based on historical data, making it the ideal tool for this use case.

Option A (Amazon Inspector): Amazon Inspector is for security assessments, not for cost optimization.

Option B (Systems Manager): Systems Manager does not specifically provide EBS optimization recommendations.

Option C (CloudWatch): CloudWatch metrics help monitor usage but do not offer optimization recommendations like Compute Optimizer.

AWS Reference:

AWS Compute Optimizer

Question: 105

A company needs to migrate a MySQL database from an on-premises data center to AWS within 2 weeks. The database is 180 TB in size. The company cannot partition the database.

The company wants to minimize downtime during the migration. The company's internet connection speed is 100 Mbps.

Which solution will meet these requirements?

A. Order an AWS Snowball Edge Storage Optimized device. Use AWS Database Migration Service (AWS DMS) and the AWS Schema Conversion Tool (AWS SCT) to migrate the database to Amazon RDS for MySQL and replicate ongoing changes. Send the Snowball Edge device back to AWS to finish the migration. Continue to replicate ongoing changes.

B. Establish an AWS Site-to-Site VPN connection between the data center and AWS. Use AWS Database Migration Service (AWS DMS) and the AWS Schema Conversion Tool (AWS SCT) to migrate the database to Amazon RDS for MySQL and replicate ongoing changes.

C. Establish a 10 Gbps dedicated AWS Direct Connect connection between the data center and AWS. Use AWS DataSync to replicate the database to Amazon S3. Create a script to import the data from Amazon S3 to a new Amazon RDS for MySQL database instance.

D. Use the company's existing internet connection. Use AWS DataSync to replicate the database to Amazon S3. Create a script to import the data from Amazon S3 to a new Amazon RDS for MySQL database instance.

Answer: A

Explanation:

Given the large size (180 TB) of the database and the time constraint, AWS Snowball Edge Storage Optimized is the best solution. Snowball Edge allows for the physical transfer of large datasets to AWS efficiently without relying on slow internet connections. AWS DMS and SCT can be used to perform ongoing replication of any changes made during the migration, ensuring minimal downtime.

Option B (VPN): Using a 100 Mbps internet connection would take far too long to transfer 180 TB.

Option C (Direct Connect): Establishing a 10 Gbps Direct Connect link might not be feasible within the 2-week timeframe.

Option D (DataSync over internet): With the existing internet connection, DataSync would also take too long.

AWS Reference:

AWS Snowball Edge

AWS DMS

Question: 106

A company has multiple Amazon RDS DB instances that run in a development AWS account. All the instances have tags to identify them as development resources. The company needs the development DB instances to run on a schedule only during business hours.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon CloudWatch alarm to identify RDS instances that need to be stopped. Create an AWS Lambda function to start and stop the RDS instances.
- B. Create an AWS Trusted Advisor report to identify RDS instances to be started and stopped. Create an AWS Lambda function to start and stop the RDS instances.
- C. Create AWS Systems Manager State Manager associations to start and stop the RDS instances.
- D. Create an Amazon EventBridge rule that invokes AWS Lambda functions to start and stop the RDS instances.

Answer: D

Explanation:

To run RDS instances only during business hours with the least operational overhead, you can use Amazon EventBridge to schedule events that invoke AWS Lambda functions. The Lambda functions can be configured to start and stop the RDS instances based on the specified

schedule (business hours). EventBridge rules allow you to define recurring events easily, and Lambda functions provide a serverless way to manage RDS instance start and stop operations, reducing administrative overhead.

Option A: While CloudWatch alarms could be used, they are more suited for monitoring, and using Lambda with EventBridge is simpler.

Option B (Trusted Advisor): Trusted Advisor is not ideal for scheduling tasks.

Option C (Systems Manager): Systems Manager could also work, but EventBridge and Lambda offer a more streamlined and lower-overhead solution.

AWS Reference:

Amazon EventBridge Scheduler

AWS Lambda

Question: 107

A company is deploying a new gaming application on Amazon EC2 instances. The gaming application needs to have access to shared storage.

The company requires a high-performance solution to give the application the ability to use an existing custom protocol to access shared storage. The solution must ensure low latency and must be operationally efficient.

Which solution will meet these requirements?

- A. Create an Amazon FSx File Gateway. Create a file share that uses the existing custom protocol. Connect the EC2 instances that host the application to the file share.
- B. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the EC2 instances that host the application to the file share.
- C. Create an Amazon Elastic File System (Amazon EFS) file system. Configure the file system to support Lustre. Connect the EC2 instances that host the application to the file system.
- D. Create an Amazon FSx for Lustre file system. Connect the EC2 instances that host the application to the file system.

Answer: D

Explanation:

Amazon FSx for Lustre is a high-performance, fully managed file system that is ideal for applications requiring low-latency access to shared storage, especially in use cases like gaming where high throughput and low latency are essential. It integrates easily with EC2 instances, providing fast and scalable shared storage, and supports custom protocols for specific application needs.

Option A (FSx File Gateway): FSx File Gateway is designed for hybrid cloud storage and is not suited for high-performance gaming workloads.

Option B (EC2 Windows instance): Setting up a file share on a Windows instance would introduce additional administrative overhead and would not provide the necessary performance.

Option C (EFS with Lustre): While Lustre is integrated with FSx, EFS does not natively support Lustre.

AWS Reference:

Amazon FSx for Lustre

Question: 108

A company is developing a rating system for its ecommerce web application. The company needs a solution to save ratings that users submit in an Amazon DynamoDB table.

The company wants to ensure that developers do not need to interact directly with the DynamoDB table. The solution must be scalable and reusable.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Application Load Balancer (ALB). Create an AWS Lambda function, and set the function as a target group in the ALB. Invoke the Lambda function by using the `put_item` method through the ALB.
- B. Create an AWS Lambda function. Configure the Lambda function to interact with the DynamoDB table by using the `put-item` method from Boto3. Invoke the Lambda function from the web application.
- C. Create an Amazon Simple Queue Service (Amazon SQS) queue and an AWS Lambda function that has an SQS trigger type. Instruct the developers to add customer ratings to the SQS queue as JSON messages. Configure the Lambda function to fetch the ratings from the queue and store the ratings in DynamoDB.
- D. Create an Amazon API Gateway REST API. Define a resource and create a new POST method. Choose AWS as the integration type, and select DynamoDB as the service. Set the action to `PutItem`.

Answer: D

Explanation:

Amazon API Gateway provides a scalable and reusable solution for interacting with DynamoDB without requiring direct access by developers. By setting up a REST API with a POST method that integrates with DynamoDB's `PutItem` action, developers can submit data (such as user ratings) to the DynamoDB table through API Gateway, without having to directly interact with the database. This solution is serverless and minimizes operational overhead.

Option A: Using ALB with Lambda adds complexity and is less efficient for this use case.

Option B: While using Lambda is possible, API Gateway provides a more scalable, reusable interface.

Option C: SQS with Lambda introduces unnecessary components for a simple put operation.

AWS Reference:

Amazon API Gateway with DynamoDB

Question: 109

A marketing team wants to build a campaign for an upcoming multi-sport event. The team has news reports from the past five years in PDF format. The team needs a solution to extract insights about the content and the sentiment of the news reports. The solution must use Amazon Textract to process the news reports.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Provide the extracted insights to Amazon Athena for analysis. Store the extracted insights and analysis in an Amazon S3 bucket.
- B. Store the extracted insights in an Amazon DynamoDB table. Use Amazon SageMaker to build a sentiment model.
- C. Provide the extracted insights to Amazon Comprehend for analysis. Save the analysis to an Amazon S3 bucket.
- D. Store the extracted insights in an Amazon S3 bucket. Use Amazon QuickSight to visualize and analyze the data.

Answer: C

Explanation:

Amazon Textract can extract text from the PDFs, and Amazon Comprehend is the most suitable service to analyze the extracted text for sentiment and insights. Comprehend offers a fully managed, low-operational overhead solution for analyzing text data. The results can then be stored in an Amazon S3 bucket, ensuring scalability and easy access.

Option A: Athena is for querying structured data and is not suitable for sentiment analysis.

Option B: SageMaker adds complexity and is not necessary when Comprehend can handle sentiment analysis natively.

Option D: QuickSight is used for visualization and analytics, but it does not provide sentiment analysis.

AWS Reference:

Amazon Comprehend

Amazon Textract

Question: 110

An ecommerce company is migrating its on-premises workload to the AWS Cloud. The workload currently consists of a web application and a backend Microsoft SQL database for storage.

The company expects a high volume of customers during a promotional event. The new infrastructure in the AWS Cloud must be highly available and scalable.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Migrate the web application to two Amazon EC2 instances across two Availability Zones behind an Application Load Balancer. Migrate the database to Amazon RDS for Microsoft SQL Server with read replicas in both Availability Zones.
- B. Migrate the web application to an Amazon EC2 instance that runs in an Auto Scaling group across two Availability Zones behind an Application Load Balancer. Migrate the database to two EC2 instances across separate AWS Regions with database replication.
- C. Migrate the web application to Amazon EC2 instances that run in an Auto Scaling group across two Availability Zones behind an Application Load Balancer. Migrate the database to Amazon RDS with Multi-AZ deployment.
- D. Migrate the web application to three Amazon EC2 instances across three Availability Zones behind an Application Load Balancer. Migrate the database to three EC2 instances across three Availability Zones.

Answer: C

Explanation:

To ensure high availability and scalability, the web application should run in an Auto Scaling group across two Availability Zones behind an Application Load Balancer (ALB). The database should be migrated to Amazon RDS with Multi-AZ deployment, which ensures fault tolerance and automatic failover in case of an AZ failure. This setup minimizes administrative overhead while meeting the company's requirements for high availability and scalability.

Option A: Read replicas are typically used for scaling read operations, and Multi-AZ provides better availability for a transactional database.

Option B: Replicating across AWS Regions adds unnecessary complexity for a single web application.

Option D: EC2 instances across three Availability Zones add unnecessary complexity for this scenario.

AWS Reference:

Auto Scaling Groups

Amazon RDS Multi-AZ

Question: 111

A company wants to publish a private website for its on-premises employees. The website consists of several HTML pages and image files. The website must be available only through HTTPS and must be available only to on-premises employees. A solutions architect plans to store the website files in an Amazon S3 bucket.

Which solution will meet these requirements?

A. Create an S3 bucket policy to deny access when the source IP address is not the public IP address of the on-premises environment. Set up an Amazon Route 53 alias record to point to the S3 bucket.

Provide the alias record to the on-premises employees to grant the employees access to the website.

B. Create an S3 access point to provide website access. Attach an access point policy to deny access when the source IP address is not the public IP address of the on-premises environment. Provide the S3 access point alias to the on-premises employees to grant the employees access to the website.

C. Create an Amazon CloudFront distribution that includes an origin access control (OAC) that is configured for the S3 bucket. Use AWS Certificate Manager for SSL. Use AWS WAF with an IP set rule that allows access for the on-premises IP address. Set up an Amazon Route 53 alias record to point to the CloudFront distribution.

D. Create an Amazon CloudFront distribution that includes an origin access control (OAC) that is configured for the S3 bucket. Create a CloudFront signed URL for the objects in the bucket. Set up an Amazon Route 53 alias record to point to the CloudFront distribution. Provide the signed URL to the on-premises employees to grant the employees access to the website.

Answer: C

Explanation:

This solution uses CloudFront to serve the website securely over HTTPS using AWS Certificate Manager (ACM) for SSL certificates. Origin Access Control (OAC) ensures that only CloudFront can access the S3 bucket directly. AWS WAF with an IP set rule restricts access to the website, allowing only the on-premises IP address. Route 53 is used to create an alias record pointing to the CloudFront distribution. This setup ensures secure, private access to the website with low administrative overhead.

Option A and B: S3 bucket policies and access points do not provide HTTPS support, nor do they offer the same level of security as CloudFront with WAF.

Option D: Signed URLs are more suitable for temporary, expiring access rather than a permanent solution for on-premises employees.

AWS Reference:

Amazon CloudFront with Origin Access Control

Question: 112

A company needs a solution to enforce data encryption at rest on Amazon EC2 instances. The solution must automatically identify noncompliant resources and enforce compliance policies on findings.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use an IAM policy that allows users to create only encrypted Amazon Elastic Block Store (Amazon EBS) volumes. Use AWS Config and AWS Systems Manager to automate the detection and remediation of unencrypted EBS volumes.
- B. Use AWS Key Management Service (AWS KMS) to manage access to encrypted Amazon Elastic Block Store (Amazon EBS) volumes. Use AWS Lambda and Amazon EventBridge to automate the detection and remediation of unencrypted EBS volumes.
- C. Use Amazon Macie to detect unencrypted Amazon Elastic Block Store (Amazon EBS) volumes. Use AWS Systems Manager Automation rules to automatically encrypt existing and new EBS volumes.
- D. Use Amazon Inspector to detect unencrypted Amazon Elastic Block Store (Amazon EBS) volumes. Use AWS Systems Manager Automation rules to automatically encrypt existing and new EBS volumes.

Answer: A

Explanation:

The best solution to enforce encryption at rest for Amazon EBS volumes is to use an IAM policy to restrict the creation of unencrypted volumes. To automatically identify and remediate unencrypted volumes, you can use AWS Config rules, which continuously monitor the compliance of resources, and AWS Systems Manager to automate the remediation by encrypting existing unencrypted volumes. This setup requires minimal administrative overhead while ensuring compliance.

Option B (KMS): KMS is for managing encryption keys, but Config and Systems Manager provide a better solution for automatic detection and enforcement.

Option C (Macie): Macie is for data classification and is not suitable for this use case.

Option D (Inspector): Inspector is used for security vulnerabilities, not encryption compliance.

AWS Reference:

AWS Config Rules

AWS Systems Manager

Question: 113

A company deploys its applications on Amazon Elastic Kubernetes Service (Amazon EKS) behind an Application Load Balancer in an AWS Region. The application needs to store data in a PostgreSQL database engine. The company wants the data in the database to be highly available. The company also needs increased capacity for read workloads.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Create an Amazon DynamoDB database table configured with global tables.
- B. Create an Amazon RDS database with Multi-AZ deployments
- C. Create an Amazon RDS database with Multi-AZ DB cluster deployment.
- D. Create an Amazon RDS database configured with cross-Region read replicas.

Answer: C

Explanation:

Amazon RDS Multi-AZ DB cluster deployment ensures high availability by automatically replicating data across multiple Availability Zones (AZs), and it supports failover in case of a failure in one AZ. This setup also provides increased capacity for read workloads by allowing read scaling with reader instances in different AZs. This solution offers the most operational efficiency with minimal manual intervention.

Option A (DynamoDB): DynamoDB is not suitable for a relational database workload, which requires a PostgreSQL engine.

Option B (RDS with Multi-AZ): While this provides high availability, it doesn't offer read scaling capabilities.

Option D (Cross-Region Read Replicas): This adds complexity and is not necessary if the requirement is high availability within a single region.

AWS Reference:

Amazon RDS Multi-AZ DB Cluster

Question: 114

A manufacturing company runs an order processing application in its VPC. The company wants to securely send messages from the application to an external Salesforce system that uses Open Authorization (OAuth).

A solutions architect needs to integrate the company's order processing application with the external Salesforce system.

Which solution will meet these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic in a fanout configuration that pushes data to an HTTPS endpoint. Configure the order processing application to publish messages to the SNS topic.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic in a fanout configuration that pushes data to an Amazon Data Firehose delivery stream that has a HTTP destination. Configure the order processing application to publish messages to the SNS topic.
- C. Create an Amazon EventBridge rule and configure an Amazon EventBridge API destination partner. Configure the order processing

application to publish messages to Amazon EventBridge.

D. Create an Amazon Managed Streaming for Apache Kafka (Amazon MSK) topic that has an outbound MSK Connect connector. Configure the order processing application to publish messages to the MSK topic.

Answer: C

Explanation:

AmazonEventBridgeAPI destinations allow you to send data from AWS to external systems, like Salesforce, using HTTP APIs, including those secured with OAuth. This provides a secure and scalable solution for sending messages from the order processing application to Salesforce.

Option A and B (SNS): SNS is not ideal for OAuth-secured external APIs and lacks the necessary OAuth integration.

Option D (MSK): Amazon MSK is a Kafka-based streaming solution, which is overkill for simple message forwarding to Salesforce.

AWS Reference:

Amazon EventBridge API Destinations

Question: 115

A company uses an Amazon EC2 Auto Scaling group to host an API. The EC2 instances are in a target group that is associated with an Application Load Balancer (ALB). The company stores data in an Amazon Aurora PostgreSQL database.

The API has a weekly maintenance window. The company must ensure that the API returns a static maintenance response during the weekly maintenance window.

Which solution will meet this requirement with the LEAST operational overhead?

A. Create a table in Aurora PostgreSQL that has fields to contain keys and values. Create a key for a maintenance flag. Set the flag when the maintenance window starts. Configure the API to query the table for the maintenance flag and to return a maintenance response if the flag is set. Reset the flag when the maintenance window is finished.

B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Subscribe the EC2 instances to the queue. Publish a message to the queue when the maintenance window starts. Configure the API to return a maintenance message if the instances receive a maintenance start message from the queue. Publish another message to the queue when the maintenance window is finished to restore normal operation.

C. Create a listener rule on the ALB to return a maintenance response when the path on a request matches a wildcard. Set the rule priority to one. Perform the maintenance. When the maintenance window is finished, delete the listener rule.

D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the EC2 instances to the topic. Publish a message to the topic when the maintenance window starts. Configure the API to return a maintenance response if the instances receive the maintenance start message from the topic. Publish another message to the topic when the maintenance window finishes to restore normal

operation.

Answer: C

Explanation:

Creating a listener rule on the Application Load Balancer (ALB) to return a maintenance response during the maintenance window is the most straightforward solution with the least operational overhead. The rule can be configured to match all incoming requests and return a custom response, and it can be easily removed once maintenance is complete.

Option A (Aurora table flag): This adds unnecessary complexity for a temporary maintenance response.

Option B and D (SQS or SNS): These options introduce more components than needed for a simple maintenance message.

AWS Reference:

ALB Listener Rules

Question: 116

An online education platform experiences lag and buffering during peak usage hours, when thousands of students access video lessons concurrently. A solutions architect needs to improve the performance of the education platform.

The platform needs to handle unpredictable traffic surges without losing responsiveness. The platform must provide smooth video playback performance at all times. The platform must create multiple copies of each video lesson and store the copies in various bitrates to serve users who have different internet speeds. The smallest video size is 7 GB.

Which solution will meet these requirements MOST cost-effectively?

- A. Use Amazon ElastiCache to cache videos in all the required bitrates. Use AWS Lambda functions to process the videos and to convert the videos to the required bitrates.
- B. Create an Auto Scaling group that includes Amazon EC2 instances that are sized to meet peak loads. Use the Auto Scaling group to serve videos. Use the Auto Scaling group to convert the videos to the required bitrates.
- C. Store a copy of every video in every required bitrate in an Amazon S3 bucket. Use a single Amazon EC2 instance to serve the videos.
- D. Use Amazon Kinesis Video Streams to store and serve the videos. Use AWS Lambda functions to process the videos and to convert the videos to the required bitrates.

Answer: C

Explanation:

The most cost-effective solution for serving video content with different bitrates is to store multiple versions of each video in Amazon S3. S3 provides scalable and cost-effective storage for large media files. Serving the videos from a single Amazon EC2 instance ensures low-latency delivery, and S3 storage helps minimize costs.

Option A (ElastiCache): Caching large video files in memory would be prohibitively expensive and unnecessary.

Option B (Auto Scaling group): Using Auto Scaling groups to serve video is less cost-effective compared to leveraging S3 for static storage.

Option D (Kinesis Video Streams): Kinesis Video Streams is designed for real-time video streaming and is not suitable for storing and serving pre-recorded videos.

AWS Reference:

Amazon S3 for Media Storage

Question: 117

A company has Amazon EC2 instances in multiple AWS Regions. The instances all store and retrieve confidential data from the same Amazon S3 bucket. The company wants to improve the security of its current architecture.

The company wants to ensure that only the Amazon EC2 instances within its VPC can access the S3 bucket. The company must block all other access to the bucket.

Which solution will meet this requirement?

- A. Use IAM policies to restrict access to the S3 bucket.
- B. Use server-side encryption (SSE) to encrypt data in the S3 bucket at rest. Store the encryption key on the EC2 instances.
- C. Create a VPC endpoint for Amazon S3. Configure an S3 bucket policy to allow connections only from the endpoint.
- D. Use AWS Key Management Service (AWS KMS) with customer-managed keys to encrypt the data before sending the data to the S3 bucket.

Answer: C

Explanation:

Creating a VPC endpoint for S3 and configuring a bucket policy to allow access only from the endpoint ensures that only EC2 instances within the VPC can access the S3 bucket. This solution improves security by restricting access at the network level without the need for public internet access.

Option A (IAM policies): IAM policies alone cannot restrict access based on the network location.

Option B and D (Encryption): Encryption secures data at rest but does not restrict network access to the bucket.

AWS Reference:

Amazon S3 VPC Endpoints

Question: 118

A company recently launched a new product that is highly available in one AWS Region. The product consists of an application that runs on Amazon Elastic Container Service (Amazon ECS), a public Application Load Balancer (ALB), and an Amazon DynamoDB table. The company wants a solution that will make the application highly available across Regions.

Which combination of steps will meet these requirements? (Select THREE.)

- A. In a different Region, deploy the application to a new ECS cluster that is accessible through a new ALB.
- B. Create an Amazon Route 53 failover record.
- C. Modify the DynamoDB table to create a DynamoDB global table.
- D. In the same Region, deploy the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is accessible through a new ALB.
- E. Modify the DynamoDB table to create global secondary indexes (GSIs).
- F. Create an AWS PrivateLink endpoint for the application.

Answer: A,B,C

Explanation:

To make the application highly available across regions:

Deploy the application in a different region using a new ECS cluster and ALB to ensure regional redundancy.

Use Route 53 failover routing to automatically direct traffic to the healthy region in case of failure.

Use DynamoDB Global Tables to ensure the database is replicated and available across multiple regions, supporting read and write operations in each region.

Option D (EKS cluster in the same region): This does not provide regional redundancy.

Option E (Global Secondary Indexes): GSIs improve query performance but do not provide multi-region availability.

Option F (PrivateLink): PrivateLink is for secure communication, not for cross-region high availability.

AWS Reference:

DynamoDB Global Tables

Amazon ECS with ALB

Question: 119

A company wants to restrict access to the content of its web application. The company needs to protect the content by using authorization techniques that are available on AWS. The company also wants to implement a serverless architecture for authorization and authentication that has low login latency.

The solution must integrate with the web application and serve web content globally. The application currently has a small user base, but the company expects the application's user base to increase

Which solution will meet these requirements?

- A. Configure Amazon Cognito for authentication. Implement Lambda@Edge for authorization. Configure Amazon CloudFront to serve the web application globally
- B. Configure AWS Directory Service for Microsoft Active Directory for authentication. Implement AWS Lambda for authorization. Use an Application Load Balancer to serve the web application globally.
- C. Configure Amazon Cognito for authentication. Implement AWS Lambda for authorization Use Amazon S3 Transfer Acceleration to serve the web application globally.
- D. Configure AWS Directory Service for Microsoft Active Directory for authentication. Implement Lambda@Edge for authorization. Use AWS Elastic Beanstalk to serve the web application globally.

Answer: A

Explanation:

Amazon Cognito provides scalable, serverless authentication, and Lambda@Edge is used for authorization, providing low-latency access control at the edge. Amazon CloudFront serves the web application globally with reduced latency and ensures secure access for users around the world. This solution minimizes operational overhead while providing scalability and security.

Option B (Directory Service): Directory Service is more suitable for enterprise use cases involving Active Directory, not for web-based applications.

Option C (S3 Transfer Acceleration): S3 Transfer Acceleration helps with file transfers but does not provide authorization features.

Option D (Elastic Beanstalk): Elastic Beanstalk adds unnecessary overhead when CloudFront can handle global delivery efficiently.

AWS Reference:

Amazon Cognito

Lambda@Edge

Question: 120

A company runs a payment processing system in the AWS Cloud. Sometimes when a payment fails because of insufficient funds or technical issues, users attempt to resubmit the payment. Sometimes payment resubmissions invoke multiple payment messages for the same payment ID.

A solutions architect needs to ensure that the payment processing system receives payment messages that have the same payment ID sequentially, according to when the messages were generated. The processing system must process the messages in the order in which the messages are received. The solution must retain all payment messages for 10 days for analytics.

Which solutions will meet these requirements? (Select TWO.)

- A. Write the payment messages to an Amazon DynamoDB table that uses the payment ID as the partition key.
- B. Write the payment messages to an Amazon Kinesis data stream that uses the payment ID as the partition key.
- C. Write the payment messages to an Amazon ElastiCache for Memcached cluster that uses the payment ID as the key.
- D. Write the payment messages to an Amazon Simple Queue Service (Amazon SQS) queue. Set the message attribute to use the payment ID.
- E. Write the payment messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the message group to use the payment ID.

Answer: B,E

Explanation:

Both Amazon Kinesis and SQS FIFO queues ensure the sequential processing of messages. By using the payment ID as the partition key in Kinesis or as the message group in the SQS FIFO queue, messages are processed in order. Both solutions also allow for long-term retention (up to 10 days) of messages, making them suitable for this payment processing use case.

Option A (DynamoDB): DynamoDB does not guarantee message ordering for real-time processing.

Option C (ElastiCache): ElastiCache is for caching, not suitable for sequential message processing.

Option D (Standard SQS queue): A standard SQS queue does not guarantee ordering of messages.

AWS Reference:

Amazon Kinesis

Amazon SQS FIFO Queues

Question: 121

A consulting company provides professional services to customers worldwide. The company provides solutions and tools for customers to expedite gathering and analyzing data on AWS. The company needs to centrally manage and deploy a common set of solutions and tools for customers to use for self-service purposes.

Which solution will meet these requirements?

- A. Create AWS Cloud Formation templates for the customers.
- B. Create AWS Service Catalog products for the customers.
- C. Create AWS Systems Manager templates for the customers.
- D. Create AWS Config items for the customers.

Answer: B

Explanation:

AWS Service Catalog allows organizations to centrally manage commonly deployed IT services and offers self-service deployment capabilities to customers. By creating Service Catalog products, the consulting company can package their solutions and tools for easy reuse by customers while maintaining central control over configuration and access. This provides a standardized and automated solution with the least operational overhead for managing and deploying solutions across different customers.

Option A (CloudFormation): CloudFormation templates are useful but don't provide the same level of management and user-friendly self-service capabilities as Service Catalog.

Option C (Systems Manager): Systems Manager is more focused on managing infrastructure and doesn't offer the same self-service capabilities.

Option D (AWS Config): AWS Config is used for tracking resource configurations, not for deploying solutions.

AWS Reference:

AWS Service Catalog

Question: 122

A company stores customer data in a multitenant Amazon S3 bucket. Each customer's data is stored in a prefix that is unique to the

customer. The company needs to migrate data for specific customers to a new, dedicated S3 bucket that is in the same AWS Region as the source bucket. The company must preserve object metadata such as creation date and version IDs.

After the migration is finished, the company must delete the source data for the migrated customers from the original multitenant S3 bucket.

Which combination of solutions will meet these requirements with the LEAST overhead? (Select THREE.)

- A. Create a new S3 bucket as a destination bucket. Enable versioning on the new bucket.
- B. Use S3 batch operations to copy objects from the specified prefixes to the destination bucket.
- C. Use the S3 CopyObject API, and create a script to copy data to the destination S3 bucket.
- D. Configure S3 Same-Region Replication (SRR) to replicate existing data from the specified prefixes in the source bucket to the destination bucket.
- E. Configure AWS DataSync to migrate data from the specified prefixes in the source bucket to the destination bucket.
- F. Use an S3 Lifecycle policy to delete objects from the source bucket after the data is migrated to the destination bucket.

Answer: A,B,F

Explanation:

The combination of these solutions provides an efficient and automated way to migrate data while preserving metadata and ensuring cleanup:

Create a new S3 bucket with versioning enabled (Option A) to preserve object metadata like version IDs during migration.

Use S3 batch operations (Option B) to efficiently copy data from specific prefixes in the source bucket to the destination bucket, ensuring minimal overhead.

Use an S3 Lifecycle policy (Option F) to automatically delete the data from the source bucket after it has been migrated, reducing manual intervention.

Option C (CopyObject API): This approach would require more manual scripting and effort.

Option D (Same-Region Replication): SRR is designed for ongoing replication, not for one-time migrations.

Option E (DataSync): DataSync adds more complexity than necessary for this task.

AWS Reference:

S3 Batch Operations

S3 Lifecycle Policies

Question: 123

A media company is using video conversion tools that run on Amazon EC2 instances. The video conversion tools run on a combination of Windows EC2 instances and Linux EC2 instances. Each video file is tens of gigabytes in size. The video conversion tools must process the video files in the shortest possible amount of time. The company needs a single, centralized file storage solution that can be mounted on all the EC2 instances that host the video conversion tools.

Which solution will meet these requirements?

- A. Deploy Amazon FSx for Windows File Server with hard disk drive (HDD) storage.
- B. Deploy Amazon FSx for Windows File Server with solid state drive (SSD) storage.
- C. Deploy Amazon Elastic File System (Amazon EFS) with Max I/O performance mode.
- D. Deploy Amazon Elastic File System (Amazon EFS) with General Purpose performance mode.

Answer: C

Explanation:

Amazon EFS with Max I/O performance mode is designed for workloads that require high levels of parallelism, such as video processing across multiple EC2 instances. EFS provides shared file storage that can be mounted on both Windows and Linux EC2 instances, and the Max I/O mode ensures the best performance for handling large files and concurrent access across multiple instances.

Option A and B (FSx for Windows File Server): FSx for Windows File Server is optimized for Windows workloads and would not be ideal for Linux instances or high-throughput, parallel workloads.

Option D (EFS General Purpose mode): General Purpose mode offers lower latency but doesn't support the high throughput needed for large, concurrent workloads.

AWS Reference:

Amazon EFS Performance Modes

Question: 124

A company has customers located across the world. The company wants to use automation to secure its systems and network infrastructure. The company's security team must be able to track and audit all incremental changes to the infrastructure.

Which solution will meet these requirements?

- A. Use AWS Organizations to set up the infrastructure. Use AWS Config to track changes.

- B. Use AWS Cloud Formation to set up the infrastructure. Use AWS Config to track changes.
- C. Use AWS Organizations to set up the infrastructure. Use AWS Service Catalog to track changes.
- D. Use AWS Cloud Formation to set up the infrastructure. Use AWS Service Catalog to track changes.

Answer: B

Explanation:

AWS CloudFormation allows for the automated, repeatable setup of infrastructure, reducing human error and ensuring consistency. AWS Config provides the ability to track changes in the infrastructure, ensuring that all changes are logged and auditable, which satisfies the requirement for tracking incremental changes.

Option A and C (AWS Organizations): AWS Organizations manage multiple accounts, but they are not designed for infrastructure setup or change tracking.

Option D (Service Catalog): Service Catalog is used for deploying products, not for setting up infrastructure or tracking changes.

AWS Reference:

AWS Config

AWS CloudFormation

Question: 125

A company has a static website that is hosted on Amazon CloudFront in front of Amazon S3. The static website uses a database backend. The company notices that the website does not reflect updates that have been made in the website's Git repository. The company checks the continuous integration and continuous delivery (CI/CD) pipeline between the Git repository and Amazon S3. The company verifies that the webhooks are configured properly and that the CI/CD pipeline is sending messages that indicate successful deployments.

A solutions architect needs to implement a solution that displays the updates on the website.

Which solution will meet these requirements?

- A. Add an Application Load Balancer.
- B. Add Amazon ElastiCache for Redis or Memcached to the database layer of the web application.
- C. Invalidate the CloudFront cache.
- D. Use AWS Certificate Manager (ACM) to validate the website's SSL certificate.

Explanation:

Amazon CloudFront is a content delivery network (CDN) service that caches copies of your content at edge locations around the world. This helps improve performance by serving content from the edge nearest to the user. However, when the content in Amazon S3 (your origin) is updated, those updates may not immediately reflect on the website if they are cached at the CloudFront edge locations.

The issue described in the question suggests that the CI/CD pipeline is functioning correctly, and updates are being deployed to S3. However, since CloudFront caches this content, the edge locations may still be serving outdated content, causing the updates to not be reflected on the website.

To resolve this issue, you need to invalidate the CloudFront cache. By invalidating the cache, CloudFront will remove the outdated content and retrieve the latest version from the S3 origin.

AWS documentation on this process:

CloudFront cache invalidation allows you to clear items from the cache so that CloudFront retrieves the latest version from the origin. You can create invalidation requests via the AWS Management Console, AWS CLI, or SDKs.

AWS CloudFront Documentation

Why the other options are incorrect:

A . Add an Application Load Balancer: ALBs are used to distribute incoming application traffic and are not relevant to caching or serving content from CloudFront.

B . Add Amazon ElastiCache for Redis or Memcached: This would help in caching database queries but has no relation to static website content hosted on CloudFront and S3.

D . Use AWS Certificate Manager (ACM): ACM is used for managing SSL/TLS certificates and is unrelated to the issue of content not being updated on CloudFront.

Question: 126

A company is migrating applications from an on-premises Microsoft Active Directory that the company manages to AWS. The company deploys the applications in multiple AWS accounts. The company uses AWS Organizations to manage the accounts centrally.

The company's security team needs a single sign-on solution across all the company's AWS accounts. The company must continue to manage users and groups that are in the on-premises Active Directory

Which solution will meet these requirements?

- A. Create an Enterprise Edition Active Directory in AWS Directory Service for Microsoft Active Directory. Configure the Active Directory to be the identity source for AWS IAM Identity Center
- B. Enable AWS IAM Identity Center. Configure a two-way forest trust relationship to connect the company's self-managed Active

Directory with IAM Identity Center by using AWS Directory Service for Microsoft Active Directory.

C. Use AWS Directory Service and create a two-way trust relationship with the company's selfmanaged Active Directory.

D. Deploy an identity provider (IdP) on Amazon EC2. Link the IdP as an identity source within AWS IAM Identity Center.

Answer: B

Explanation:

The company is looking for a solution that provides single sign-on (SSO) across multiple AWS accounts while continuing to manage users and groups in their on-premises Active Directory (AD). AWS IAM Identity Center (formerly AWS SSO) is the recommended solution for this type of requirement.

AWS IAM Identity Center provides a centralized identity management solution, enabling single sign-on across multiple AWS accounts and other cloud applications. It can integrate with on-premises Active Directory to leverage existing users and groups.

By configuring a two-way forest trust relationship between AWS Directory Service for Microsoft Active Directory and the company's on-premises Active Directory, users can be authenticated by their on-premises AD and still access AWS resources through IAM Identity Center. This solution allows centralized management of AWS accounts within AWS Organizations.

The two-way trust allows mutual access between the on-premises AD and the AWS Directory Service. This means that users and groups in the on-premises AD can be used for authentication in AWS IAM Identity Center while maintaining the existing identity management system.

AWS Reference:

[AWS IAM Identity Center Documentation](#)

[AWS Directory Service for Microsoft Active Directory Trust Relationships](#)

[AWS Directory Service Integration with IAM Identity Center](#)

Why the other options are incorrect:

A . Create an Enterprise Edition Active Directory in AWS Directory Service: This would require setting up a new directory and managing it in AWS, which adds unnecessary overhead. The requirement is to continue using the existing on-premises AD, making this option unsuitable.

C . Use AWS Directory Service and create a two-way trust relationship: While this approach establishes a trust between on-premises AD and AWS Directory Service, it does not address the single sign-on (SSO) requirements across multiple AWS accounts through IAM Identity Center.

D . Deploy an identity provider (IdP) on Amazon EC2: This is more complex than necessary and introduces more management overhead.

AWS IAM Identity Center natively supports integration with on-premises Active Directory without requiring a custom IdP.

Question: 127

A company is designing a microservice-based architecture for a new application on AWS. Each microservice will run on its own set of Amazon EC2 instances. Each microservice will need to interact with multiple AWS services such as Amazon S3 and Amazon Simple Queue Service (Amazon SQS).

The company wants to manage permissions for each EC2 instance based on the principle of least privilege.

Which solution will meet this requirement?

- A. Assign an IAM user to each micro-service. Use access keys stored within the application code to authenticate AWS service requests.
- B. Create a single IAM role that has permission to access all AWS services. Associate the IAM role with all EC2 instances that run the microservices.
- C. Use AWS Organizations to create a separate account for each microservice. Manage permissions at the account level.
- D. Create individual IAM roles based on the specific needs of each microservice. Associate the IAM roles with the appropriate EC2 instances.

Answer: D

Explanation:

When designing a microservice architecture where each microservice interacts with different AWS services, it's essential to follow the principle of least privilege. This means granting each microservice only the permissions it needs to perform its tasks, reducing the risk of unauthorized access or accidental actions.

The recommended approach is to create individual IAM roles with policies that grant each microservice the specific permissions it requires. Then, these roles should be associated with the EC2 instances that run the corresponding microservice. By doing so, each EC2 instance will assume its specific IAM role, and permissions will be automatically managed by AWS.

IAM roles provide temporary credentials via the instance metadata service, eliminating the need to hard-code credentials in your application code, which enhances security.

AWS Reference:

[IAM Roles for Amazon EC2](#) explains how EC2 instances can use IAM roles to securely access AWS services without managing long-term credentials.

[Best Practices for IAM](#) includes recommendations for implementing the least privilege principle and using IAM roles effectively.

Why the other options are incorrect:

- A. Assign an IAM user to each microservice: This requires managing long-term credentials (access keys), which should be avoided. Storing

keys in application code is insecure and creates a maintenance burden.

B . Create a single IAM role: This violates the principle of least privilege, as a single role with broad permissions across all services is less secure.

C . Use AWS Organizations: This approach adds unnecessary complexity. Managing permissions at the account level for each microservice is excessive for this use case and doesn't adhere to the principle of least privilege.

Question: 128

A media company hosts its video processing workload on AWS. The workload uses Amazon EC2 instances in an Auto Scaling group to handle varying levels of demand. The workload stores the original videos and the processed videos in an Amazon S3 bucket.

The company wants to ensure that the video processing workload is scalable. The company wants to prevent failed processing attempts because of resource constraints. The architecture must be able to handle sudden spikes in video uploads without impacting the processing capability.

Which solution will meet these requirements with the LEAST overhead?

A. Migrate the workload from Amazon EC2 instances to AWS Lambda functions. Configure an Amazon S3 event notification to invoke the Lambda functions when a new video is uploaded. Configure the Lambda functions to process videos directly and to save processed videos back to the S3 bucket.

B. Migrate the workload from Amazon EC2 instances to AWS Lambda functions. Use Amazon S3 to invoke an Amazon Simple Notification Service (Amazon SNS) topic when a new video is uploaded. Subscribe the Lambda functions to the SNS topic. Configure the Lambda functions to process the videos asynchronously and to save processed videos back to the S3 bucket.

C. Configure an Amazon S3 event notification to send a message to an Amazon Simple Queue Service (Amazon SQS) queue when a new video is uploaded. Configure the existing Auto Scaling group to poll the SQS queue, process the videos, and save processed videos back to the S3 bucket.

D. Configure an Amazon S3 upload trigger to invoke an AWS Step Functions state machine when a new video is uploaded. Configure the state machine to orchestrate the video processing workflow by placing a job message in the Amazon SQS queue. Configure the job message to invoke the EC2 instances to process the videos. Save processed videos back to the S3 bucket.

Answer: C

Explanation:

This solution addresses the scalability needs of the workload while preventing failed processing attempts due to resource constraints.

Amazon S3 event notifications can be used to trigger a message to an SQS queue whenever a new video is uploaded.

The existing Auto Scaling group of EC2 instances can poll the SQS queue, ensuring that the EC2 instances only process videos when

there is a job in the queue.

SQS decouples the video upload and processing steps, allowing the system to handle sudden spikes in video uploads without overloading EC2 instances.

The use of Auto Scaling ensures that the EC2 instances can scale in or out based on the demand, maintaining cost efficiency while avoiding processing failures due to insufficient resources.

AWS Reference:

S3 Event Notifications details how to configure notifications for S3 events.

Amazon SQS is a fully managed message queuing service that decouples components of the system.

Auto Scaling EC2 explains how to manage automatic scaling of EC2 instances based on demand.

Why the other options are incorrect:

A . AWS Lambda functions: While Lambda can handle some workloads, video processing is often resource-intensive and long-running, making EC2 a more suitable solution.

B . Using SNS with Lambda: Similar to A, Lambda is not ideal for large-scale video processing due to its time and memory limitations.

D . AWS Step Functions: While a valid orchestration solution, this introduces more complexity and overhead compared to the simpler SQS-based solution.

Question: 129

A company uses a set of Amazon EC2 instances to host a website. The website uses an Amazon S3 bucket to store images and media files.

The company wants to automate website infrastructure creation to deploy the website to multiple AWS Regions. The company also wants to provide the EC2 instances access to the S3 bucket so the instances can store and access data by using AWS Identity and Access Management (IAM).

Which solution will meet these requirements MOST securely?

A. Create an AWS CloudFormation template for the web server EC2 instances. Save an IAM access key in the UserData section of the AWS::EC2::Instance entity in the CloudFormation template.

B. Create a file that contains an IAM secret access key and access key ID. Store the file in a new S3 bucket. Create an AWS CloudFormation template. In the template, create a parameter to specify the location of the S3 object that contains the access key and access key ID.

C. Create an IAM role and an IAM access policy that allows the web server EC2 instances to access the S3 bucket. Create an AWS CloudFormation template for the web server EC2 instances that contains an IAM instance profile entity that references the IAM role and the IAM access policy.

D. Create a script that retrieves an IAM secret access key and access key ID from IAM and stores them on the web server EC2 instances. Include the script in the UserData section of the AWS::EC2::Instance entity in an AWS CloudFormation template.

Answer: C

Explanation:

The most secure solution for allowing EC2 instances to access an S3 bucket is by using IAM roles. An IAM role can be created with an access policy that grants the required permissions (e.g., to read and write to the S3 bucket). The IAM role is then associated with the EC2 instances through an IAM instance profile.

By associating the role with the instances, the EC2 instances can securely assume the role and receive temporary credentials via the instance metadata service. This avoids the need to store credentials (such as access keys) on the instances or within the application, enhancing security and reducing the risk of credentials being exposed.

AWS CloudFormation can be used to automate the creation of the entire infrastructure, including EC2 instances, IAM roles, and associated policies.

AWS Reference:

IAM Roles for EC2 Instances outlines the use of IAM roles for secure access to AWS services.

AWS CloudFormation User Guide details how to create and manage resources using CloudFormation templates.

Why the other options are incorrect:

A . Save IAM access key in UserData: This is insecure because it involves storing long-term credentials in the instance user data, which can be exposed.

B . Store access keys in S3: This is also insecure, as it involves managing and distributing long-term credentials, which should be avoided.

D . Retrieve access keys via a script: This approach is unnecessarily complex and less secure than using IAM roles, which provide temporary credentials automatically.

Question: 130

A company creates operations data and stores the data in an Amazon S3 bucket for the company's annual audit, an external consultant needs to access an annual report that is stored in the S3 bucket. The external consultant needs to access the report for 7 days.

The company must implement a solution to allow the external consultant access to only the report.

Which solution will meet these requirements with the MOST operational efficiency?

A. Create a new S3 bucket that is configured to host a public static website. Migrate the operations data to the new S3 bucket.

Share the S3 website URL with the external consultant.

B. Enable public access to the S3 bucket for 7 days. Remove access to the S3 bucket when the external consultant completes the audit.

C. Create a new IAM user that has access to the report in the S3 bucket. Provide the access keys to the external consultant. Revoke the access keys after 7 days.

D. Generate a presigned URL that has the required access to the location of the report on the S3 bucket. Share the presigned URL with the external consultant.

Answer: D

Explanation:

A presigned URL allows temporary access to a specific object in an S3 bucket without needing to make the bucket public or creating and managing additional IAM users. The URL is time-limited, and permissions are granted only to the specific object (in this case, the annual report), making it a highly secure and operationally efficient solution.

With a presigned URL, the consultant can access the report for the specified duration (7 days), after which the URL will expire automatically, removing the need for manual intervention to revoke access.

AWS Reference:

Amazon S3 Presigned URLs explain how to generate a presigned URL to grant temporary access to S3 objects.

Best Practices for S3 Security emphasize using presigned URLs for sharing temporary access to S3 objects securely.

Why the other options are incorrect:

A . Public static website: This approach involves making the S3 bucket publicly accessible, which is unnecessary and insecure for sensitive data.

B . Enable public access: Granting public access to the entire bucket, even temporarily, is a security risk and violates best practices.

C . Create an IAM user: Creating an IAM user and managing credentials is unnecessary overhead and less secure compared to a presigned URL for this short-term need.

Question: 131

A finance company is migrating its trading platform to AWS. The trading platform processes a high volume of market data and processes

stock trades. The company needs to establish a consistent, low-latency network connection from its on-premises data center to AWS.

The company will host resources in a VPC. The solution must not use the public internet.

Which solution will meet these requirements?

- A. Use AWS Client VPN to connect the on-premises data center to AWS.
- B. Use AWS Direct Connect to set up a connection from the on-premises data center to AWS.
- C. Use AWS PrivateLink to set up a connection from the on-premises data center to AWS.
- D. Use AWS Site-to-Site VPN to connect the on-premises data center to AWS.

Answer: B

Explanation:

AWS Direct Connect is the best solution for establishing a consistent, low-latency connection from an on-premises data center to AWS without using the public internet. Direct Connect offers dedicated, high-throughput, and low-latency network connections, which are ideal for performance-sensitive applications like a trading platform that processes high volumes of market data and stock trades.

Direct Connect provides a private connection to your AWS VPC, ensuring that data doesn't traverse the public internet, which enhances both security and performance consistency.

AWS Reference:

AWS Direct Connect provides a dedicated network connection to AWS services with consistent, low-latency performance.

Best Practices for High Performance on AWS for performance-sensitive workloads like trading platforms.

Why the other options are incorrect:

A . AWS Client VPN: While this offers secure connectivity, it's over the public internet and is not designed for the low-latency, high-performance needs of a trading platform.

C . AWS PrivateLink: PrivateLink is used for connecting VPCs and services within AWS, but it is not designed for connecting on-premises data centers to AWS.

D . AWS Site-to-Site VPN: Although this provides secure connectivity, it uses the public internet, which can introduce latency and doesn't meet the low-latency requirements of the use case.

Question: 132

A solutions architect is designing the architecture for a company website that is composed of static content. The company's target

customers are located in the United States and Europe.

Which architecture should the solutions architect recommend to MINIMIZE cost?

- A. Store the website files on Amazon S3 in the us-east-2 Region. Use an Amazon CloudFront distribution with the price class configured to limit the edge locations in use.
- B. Store the website files on Amazon S3 in the us-east-2 Region. Use an Amazon CloudFront distribution with the price class configured to maximize the use of edge locations.
- C. Store the website files on Amazon S3 in the us-east-2 Region and the eu-west-1 Region. Use an Amazon CloudFront geolocation routing policy to route requests to the closest Region to the user.
- D. Store the website files on Amazon S3 in the us-east-2 Region and the eu-west-1 Region. Use an Amazon CloudFront distribution with an Amazon Route 53 latency routing policy to route requests to the closest Region to the user.

Answer: A

Explanation:

The question focuses on minimizing costs while serving static content to users in the US and Europe.

Option A uses a single S3 bucket and configures CloudFront to limit edge locations, reducing costs by using fewer edge locations while still improving performance.

Option B maximizes edge locations, which increases costs unnecessarily.

Options C and D involve storing data in multiple regions, which increases storage and operational costs. Thus, Option A is the most cost-effective solution.

Question: 133

A company hosts a database that runs on an Amazon RDS instance deployed to multiple Availability Zones. A periodic script negatively affects a critical application by querying the database. How can application performance be improved with minimal costs?

- A. Add functionality to the script to identify the instance with the fewest active connections and query that instance.
- B. Create a read replica of the database. Configure the script to query only the read replica.
- C. Instruct the development team to manually export new entries at the end of the day.
- D. Use Amazon ElastiCache to cache the common queries the script runs.

Answer: B

Explanation:

Option A introduces complexity and does not scale well.

Option B creates a read replica, offloading read traffic from the primary RDS instance without impacting the critical application.

Option C is manual and inefficient.

Option D might help for caching frequently queried data but is not ideal for ad-hoc reporting. Therefore, Option B is the best choice.

Question: 134

A company has developed an API using Amazon API Gateway REST API and AWS Lambda. How can latency be reduced for users worldwide?

- A. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Enable content encoding to compress data in transit.
- B. Deploy the REST API as a Regional API endpoint. Enable caching. Enable content encoding to compress data in transit.
- C. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Configure reserved concurrency for Lambda functions.
- D. Deploy the REST API as a Regional API endpoint. Enable caching. Configure reserved concurrency for Lambda functions.

Answer: A

Explanation:

Edge-optimized API endpoints route requests through CloudFront, reducing latency for global users. Option A correctly implements edge-optimization, caching, and compression to minimize latency. Options B and D do not use edge optimization, leading to higher latency for global users. Reserved concurrency in Options C and D improves backend scaling but does not address global latency directly.

Question: 135

How can a law firm make files publicly readable while preventing modifications or deletions until a specific future date?

- A. Upload files to an Amazon S3 bucket configured for static website hosting. Grant read-only IAM permissions to any AWS principals.
- B. Create an S3 bucket. Enable S3 Versioning. Use S3 Object Lock with a retention period. Create a CloudFront distribution. Use a bucket policy to restrict access.
- C. Create an S3 bucket. Enable S3 Versioning. Configure an event trigger with AWS Lambda to restore modified objects from a private S3 bucket.
- D. Upload files to an S3 bucket for static website hosting. Use S3 Object Lock with a retention period.

Grant read-only IAM permissions.

Answer: B

Explanation:

Option B ensures the use of S3 Object Lock and Versioning to meet compliance for immutability. CloudFront enhances performance while a bucket policy ensures secure access.

Option A lacks immutability safeguards.

Option C introduces unnecessary complexity.

Option D misses out on additional security benefits offered by CloudFront.

Question: 136

A media company hosts a web application on AWS for uploading videos. Only authenticated users should upload within a specified time frame after authentication.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure the application to generate IAM temporary security credentials for authenticated users.
- B. Create an AWS Lambda function that generates pre-signed URLs when a user authenticates.
- C. Develop a custom authentication service that integrates with Amazon Cognito to control and log direct S3 bucket access through the application.
- D. Use AWS Security Token Service (AWS STS) to assume a pre-defined IAM role that grants authenticated users temporary permissions to upload videos directly to the S3 bucket.

Answer: B

Explanation:

Option B: Pre-signed URLs provide temporary, authenticated access to S3, limiting uploads to the time frame specified. This solution is lightweight, efficient, and easy to implement.

Option A requires the management of IAM temporary credentials, adding complexity.

Option C involves unnecessary development effort.

Option D introduces more complexity with STS and roles than pre-signed URLs.

Question: 137

A company needs to ingest and analyze telemetry data from vehicles at scale for machine learning and reporting.

Which solution will meet these requirements?

- A. Use Amazon Timestream for LiveAnalytics to store data points. Grant Amazon SageMaker permission to access the data. Use Amazon QuickSight to visualize the data.
- B. Use Amazon DynamoDB to store data points. Use DynamoDB Connector to ingest data into Amazon EMR for processing. Use Amazon QuickSight to visualize the data.
- C. Use Amazon Neptune to store data points. Use Amazon Kinesis Data Streams to ingest data into a Lambda function for processing. Use Amazon QuickSight to visualize the data.
- D. Use Amazon Timestream for LiveAnalytics to store data points. Grant Amazon SageMaker permission to access the data. Use Amazon Athena to visualize the data.

Answer: A

Explanation:

Amazon Timestream is purpose-built for storing and analyzing time-series data like telemetry.

Option A leverages Timestream, SageMaker for ML, and QuickSight for visualization, meeting all requirements with minimal complexity.

Option B involves more complex DynamoDB-EMR integration.

Option C uses Neptune, which is designed for graph databases, not telemetry data.

Option D incorrectly uses Athena for visualization instead of QuickSight.

Question: 138

A company runs an application on EC2 instances that need access to RDS credentials stored in AWS Secrets Manager.

Which solution meets this requirement?

- A. Create an IAM role, and attach the role to each EC2 instance profile. Use an identity-based policy to grant the role access to the secret.
- B. Create an IAM user, and attach the user to each EC2 instance profile. Use a resource-based policy to grant the user access to the

secret.

C. Create a resource-based policy for the secret. Use EC2 Instance Connect to access the secret.

D. Create an identity-based policy for the secret. Grant direct access to the EC2 instances.

Answer: A

Explanation:

Option A uses an IAM role attached to the EC2 instance profile, enabling secure and automated access to Secrets Manager. This is the recommended approach.

Option B uses IAM users, which is less secure and harder to manage.

Option C is not practical for accessing secrets programmatically.

Option D violates best practices by granting direct access to the EC2 instance.

Question: 139

A company needs a cloud-based solution for backup, recovery, and archiving while retaining encryption key material control.

Which combination of solutions will meet these requirements? (Select TWO)

A. Create an AWS Key Management Service (AWS KMS) key without key material. Import the company's key material into the KMS key.

B. Create an AWS KMS encryption key that contains key material generated by AWS KMS.

C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Use S3 Bucket Keys with AWS KMS keys.

D. Store the data in an Amazon S3 Glacier storage class. Use server-side encryption with customer-provided keys (SSE-C).

E. Store the data in AWS Snowball devices. Use server-side encryption with AWS KMS keys (SSE-KMS).

Answer: A,D

Explanation:

Option A allows importing your own encryption keys into AWS KMS, ensuring control over key material.

Option D uses S3 Glacier with SSE-C, where the customer controls the encryption keys, meeting compliance needs.

Option B uses AWS-managed key material, violating the requirement for key material control.

Option C and E are not fully compliant with the control requirement.

Question: 140

A website uses EC2 instances with Auto Scaling and EFS. How can the company optimize costs?

- A. Reconfigure the Auto Scaling group to set a desired number of instances. Turn off scheduled scaling.
- B. Create a new launch template version that uses larger EC2 instances.
- C. Reconfigure the Auto Scaling group to use a target tracking scaling policy.
- D. Replace the EFS volume with instance store volumes.

Answer: C

Explanation:

Option C ensures dynamic scaling based on demand using a target tracking scaling policy, optimizing costs.

Option A results in over-provisioning, leading to higher costs.

Option B increases costs by using larger instances.

Option D is not feasible as instance store volumes are ephemeral and unsuitable for shared storage like EFS.

Question: 141 A company is developing a social media application that must scale to meet demand spikes and handle ordered processes.

Which AWS services meet these requirements?

- A. ECS with Fargate, RDS, and SQS for decoupling.
- B. ECS with Fargate, RDS, and SNS for decoupling.
- C. DynamoDB, Lambda, DynamoDB Streams, and Step Functions.
- D. Elastic Beanstalk, RDS, and SNS for decoupling.

Answer: A

Explanation:

Option A combines ECS with Fargate for scalability, RDS for relational data, and SQS for decoupling with message ordering (FIFO queues).

Option B uses SNS, which does not maintain message order.

Option C is suitable for serverless workflows but not relational data.

Option D relies on Elastic Beanstalk, which offers less flexibility for scaling.

Question: 142

How can a company detect and notify security teams about PII in S3 buckets?

- A. Use Amazon Macie. Create an EventBridge rule for SensitiveData findings and send an SNS notification.
- B. Use Amazon GuardDuty. Create an EventBridge rule for CRITICAL findings and send an SNS notification.
- C. Use Amazon Macie. Create an EventBridge rule for SensitiveData:S3Object/Personal findings and send an SQS notification.
- D. Use Amazon GuardDuty. Create an EventBridge rule for CRITICAL findings and send an SQS notification.

Answer: A

Explanation:

Amazon Macie is purpose-built for detecting PII in S3.

Option A uses EventBridge to filter SensitiveData findings and notify via SNS, meeting the requirements.

Options B and D involve GuardDuty, which is not designed for PII detection.

Option C uses SQS, which is less suitable for immediate notifications.

Question: 143

A company runs HPC workloads requiring high IOPS.

Which combination of steps will meet these requirements? (Select TWO)

- A. Use Amazon EFS as a high-performance file system.
- B. Use Amazon FSx for Lustre as a high-performance file system.
- C. Create an Auto Scaling group of EC2 instances. Use Reserved Instances. Configure a spread placement group. Use AWS Batch for analytics.

- D. Use Mountpoint for Amazon S3 as a high-performance file system.
- E. Create an Auto Scaling group of EC2 instances. Use mixed instance types and a cluster placement group. Use Amazon EMR for analytics.

Answer: B,E

Explanation:

Option B: FSx for Lustre is designed for HPC workloads with high IOPS.

Option E: A cluster placement group ensures low-latency networking for HPC analytics workloads.

Option A: Amazon EFS is not optimized for HPC.

Option D: Mountpoint for S3 does not meet high IOPS needs.

Question: 144

How can trade data from DynamoDB be ingested into an S3 data lake for near real-time analysis?

- A. Use DynamoDB Streams to invoke a Lambda function that writes to S3.
- B. Use DynamoDB Streams to invoke a Lambda function that writes to Data Firehose, which writes to S3.
- C. Enable Kinesis Data Streams on DynamoDB. Configure it to invoke a Lambda function that writes to S3.
- D. Enable Kinesis Data Streams on DynamoDB. Use Data Firehose to write to S3.

Answer: A

Explanation:

Option A is the simplest solution, using DynamoDB Streams and Lambda for real-time ingestion into S3.

Options B, C, and D add unnecessary complexity with Data Firehose or Kinesis.

Question: 145

How can DynamoDB data be made available for long-term analytics with minimal operational overhead?

- A. Configure DynamoDB incremental exports to S3.
- B. Configure DynamoDB Streams to write records to S3.
- C. Configure EMR to copy DynamoDB data to S3.
- D. Configure EMR to copy DynamoDB data to HDFS.

Answer: A

Explanation:

Option A is the most automated and cost-efficient solution for exporting data to S3 for analytics.

Option B involves manual setup of Streams to S3.

Options C and D introduce complexity with EMR.

Question: 146

A company runs a Microsoft Windows SMB file share on-premises to support an application. The company wants to migrate the application to AWS. The company wants to share storage across multiple Amazon EC2 instances.

Which solutions will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. Create an Amazon Elastic File System (Amazon EFS) file system with elastic throughput.
- B. Create an Amazon FSx for NetApp ONTAP file system.
- C. Use Amazon Elastic Block Store (Amazon EBS) to create a self-managed Windows file share on the instances.
- D. Create an Amazon FSx for Windows File Server file system.
- E. Create an Amazon FSx for OpenZFS file system.

Answer: A,D

Explanation:

A . Amazon EFS: Provides a scalable, shared file storage solution with minimal operational overhead. It's ideal for Linux-based workloads.

B . Amazon FSx for NetApp ONTAP: More suited for workloads requiring NetApp-specific features.

- C . Amazon EBS:Requires manual management of file shares, which increases operational overhead.
- D . Amazon FSx for Windows File Server:Best suited for Windows SMB workloads with low operational overhead.
- E . Amazon FSx for OpenZFS:Better for Linux and Unix-based workloads.

Reference:Amazon EFS,Amazon FSx

Question: 147

A solutions architect needs to implement a solution that can handle up to 5,000 messages per second. The solution must publish messages as events to multiple consumers. The messages are upto 500 KB in size. The message consumers need to have the ability to use multiple programming languages to consume the messages with minimal latency. The solution must retain published messages for more than 3 months. The solution must enforce strict ordering of the messages.

Which solution will meet these requirements?

- A. Publish messages to an Amazon Kinesis Data Streams data stream. Enable enhanced fan-out. Ensure that consumers ingest the data stream by using dedicated throughput.
- B. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to subscribe to the topic.
- C. Publish messages to Amazon EventBridge. Allow each consumer to create rules to deliver messages to the consumer's own target.
- D. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use Amazon Data Firehose to subscribe to the topic.

Answer: A

Explanation:

- A . Kinesis Data Streams:Supports high throughput, strict ordering, multiple consumers, and data retention for 365 days.
- B . SNS + SQS FIFO:Can enforce ordering but lacks native support for 500 KB messages and retention requirements.
- C . EventBridge:Lacks strict ordering and message size compatibility.
- D . SNS + Firehose:Not designed for strict ordering or large message sizes.

Reference:Amazon Kinesis Data Streams

Question: 148

A company is planning to migrate an on-premises online transaction processing (OLTP) database that uses MySQL to an AWS managed database management system. Several reporting and analytics applications use the on-premises database heavily on weekends and at the end of each month. The cloud-based solution must be able to handle read-heavy surges during weekends and at the end of each month.

Which solution will meet these requirements?

- A. Migrate the database to an Amazon Aurora MySQL cluster. Configure Aurora Auto Scaling to use replicas to handle surges.
- B. Migrate the database to an Amazon EC2 instance that runs MySQL. Use an EC2 instance type that has ephemeral storage. Attach Amazon EBS Provisioned IOPS SSD (io2) volumes to the instance.
- C. Migrate the database to an Amazon RDS for MySQL database. Configure the RDS for MySQL database for a Multi-AZ deployment, and set up auto scaling.
- D. Migrate from the database to Amazon Redshift. Use Amazon Redshift as the database for both OLTP and analytics applications.

Answer: A

Explanation:

- A . Aurora MySQL:Handles OLTP workloads efficiently with built-in replication and auto-scaling capabilities.
- B . EC2 with MySQL:Requires heavy manual maintenance and does not scale seamlessly.
- C . RDS for MySQL:Limited in auto-scaling compared to Aurora.
- D . Redshift:Primarily for OLAP, not suitable for OLTP workloads.

Reference:Amazon Aurora

Question: 149

A company recently migrated a data warehouse to AWS. The company has an AWS Direct Connect connection to AWS. Company users query the data warehouse by using a visualization tool. The average size of the queries that the data warehouse returns is 50 MB. The average visualization that the visualization tool produces is 500 KB in size. The result sets that the data warehouse returns are not cached.

The company wants to optimize costs for data transfers between the data warehouse and the company.

Which solution will meet this requirement?

- A. Host the visualization tool on premises. Connect to the data warehouse directly through the internet.
- B. Host the visualization tool in the same AWS Region as the data warehouse. Access the visualization tool through the internet.

C. Host the visualization tool on premises. Connect to the data warehouse through the Direct Connect connection.

D. Host the visualization tool in the same AWS Region as the data warehouse. Access the visualization tool through the Direct Connect connection.

Answer: D

Explanation:

A . On-premises tool via internet:Incurs high costs due to large data transfers over the internet.

B . AWS Region tool via internet:Does not utilize Direct Connect, leading to potential latency and higher costs.

C . On-premises tool via Direct Connect:Adds latency for querying and visualization.

D . AWS Region tool via Direct Connect:Reduces latency and leverages Direct Connect for optimized data transfer costs.

Reference:AWS Direct Connect

Question: 150

A company runs an order management application on AWS. The application allows customers to place orders and pay with a credit card. The company uses an Amazon CloudFront distribution to deliver the application. A security team has set up logging for all incoming requests. The security team needs a solution to generate an alert if any user modifies the logging configuration.

Which combination of solutions will meet these requirements? (Select TWO.)

A. Configure an Amazon EventBridge rule that is invoked when a user creates or modifies a CloudFront distribution. Add the AWS Lambda function as a target of the EventBridge rule.

B. Create an Application Load Balancer (ALB). Enable AWS WAF rules for the ALB. Configure an AWS Config rule to detect security violations.

C. Create an AWS Lambda function to detect changes in CloudFront distribution logging. Configure the Lambda function to use Amazon Simple Notification Service (Amazon SNS) to send notifications to the security team.

D. Set up Amazon GuardDuty. Configure GuardDuty to monitor findings from the CloudFront distribution. Create an AWS Lambda function to address the findings.

E. Create a private API in Amazon API Gateway. Use AWS WAF rules to protect the private API from common security problems.

Answer: A,C

Explanation:

- A . EventBridge rule:Triggers an event whenever there is a change in CloudFront distribution, ensuring real-time monitoring.
- B . ALB with WAF:Focuses on application-level security, not CloudFront logging.
- C . Lambda + SNS:Provides notifications upon detection of changes in logging configuration.
- D . GuardDuty:Monitors anomalies but does not specifically address CloudFront logging changes.
- E . Private API + WAF:Irrelevant to CloudFront logging changes.

Reference:Amazon EventBridge,AWS Lambda

Question: 151

A company has a website that handles dynamic traffic loads. The website architecture is based on Amazon EC2 instances in an Auto Scaling group that is configured to use scheduled scaling. Each EC2 instance runs code from an Amazon Elastic File System (Amazon EFS) volume and stores shared data back to the same volume.

The company wants to optimize costs for the website.

Which solution will meet this requirement?

- A. Reconfigure the Auto Scaling group to set a desired number of instances. Turn off scheduled scaling.
- B. Create a new launch template version for the Auto Scaling group that uses larger EC2 instances.
- C. Reconfigure the Auto Scaling group to use a target tracking scaling policy.
- D. Replace the EFS volume with instance store volumes.

Answer: C

Explanation:

- A . Fixed desired instances:Does not adapt to traffic load fluctuations, leading to inefficiencies.
- B . Larger EC2 instances:Increases costs unnecessarily.
- C . Target tracking scaling policy:Adjusts capacity based on actual demand, optimizing costs.
- D . Instance store volumes:Not persistent and unsuitable for shared data across instances.

Reference:Auto Scaling

Question: 152

A company wants to provide a third-party system that runs in a private data center with access to its AWS account. The company wants to call AWS APIs directly from the third-party system. The company has an existing process for managing digital certificates. The company does not want to use SAML or OpenID Connect (OIDC) capabilities and does not want to store long-term AWS credentials.

Which solution will meet these requirements?

- A. Configure mutual TLS to allow authentication of the client and server sides of the communication channel.
- B. Configure AWS Signature Version 4 to authenticate incoming HTTPS requests to AWS APIs.
- C. Configure Kerberos to exchange tickets for assertions that can be validated by AWS APIs.
- D. Configure AWS Identity and Access Management (IAM) Roles Anywhere to exchange X.509 certificates for AWS credentials to interact with AWS APIs.

Answer: D

Explanation:

- A . Mutual TLS:Provides secure communication but does not integrate with AWS credential exchange.
- B . AWS Signature v4:Requires direct integration with AWS and is less secure for external systems.
- C . Kerberos:Not natively supported for AWS API authentication.
- D . IAM Roles Anywhere:Enables AWS API access using X.509 certificates without long-term credentials.

Reference:IAM Roles Anywhere

Question: 153

A company is migrating a new application from an on-premises data center to a new VPC in the AWS Cloud. The company has multiple AWS accounts and VPCs that share many subnets and applications. The company wants to have fine-grained access control for the new application. The company wants to ensure that all network resources across accounts and VPCs that are granted permission to access the new application can access the application.

Which solution will meet these requirements?

- A. Set up a VPC peering connection for each VPC that needs access to the new application VPC. Update route tables in each VPC to enable connectivity.
- B. Deploy a transit gateway in the account that hosts the new application. Share the transit gateway with each account that needs to connect to the application. Update route tables in the VPC that hosts the new application and in the transit gateway to enable connectivity.

- C. Use an AWS PrivateLink endpoint service to make the new application accessible to other VPCs. Control access to the application by using an endpoint policy.
- D. Use an Application Load Balancer (ALB) to expose the new application to the internet. Configure authentication and authorization processes to ensure that only specified VPCs can access the application.

Answer: B

Explanation:

- A . VPC peering:Creates a fully meshed architecture, which is complex to manage for multiple VPCs.
- B . Transit gateway:Simplifies network management by connecting multiple VPCs and on-premises networks via a central hub.
- C . PrivateLink:Restricts communication to the application endpoint but may not allow full VPC connectivity.
- D . ALB with internet exposure:Not secure or specific to private network communication.

Reference:AWS Transit Gateway

Question: 154

A company is enhancing the security of its AWS environment, where the company stores a significant amount of sensitive customer data. The company needs a solution that automatically identifies and classifies sensitive data that is stored in multiple Amazon S3 buckets. The solution must automatically respond to data breaches and alert the company's security team through email immediately when noncompliant data is found.

Which solution will meet these requirements?

- A. Use Amazon GuardDuty. Configure an AWS Lambda function to route alerts to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team to the SNS topic.
- B. Use Amazon GuardDuty. Configure an AWS Lambda function to route alerts to an Amazon Simple Queue Service (Amazon SQS) queue. Configure a second Lambda function to periodically poll the SQS queue and to send emails to the security team by using Amazon Simple Email Service (Amazon SES).
- C. Use Amazon Macie. Integrate Amazon EventBridge with Macie, and configure EventBridge to send alerts to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team to the SNS topic.
- D. Use Amazon Macie. Integrate Amazon EventBridge with Macie, and configure EventBridge to route alerts to an Amazon Simple Queue Service (Amazon SQS) queue. Configure an AWS Lambda function to periodically poll the SQS queue and to send alerts to the security team by using Amazon Simple Email Service (Amazon SES).

Answer: C

Explanation:

A & B. GuardDuty: Designed for threat detection, not for identifying or classifying sensitive data in S3 buckets.

C . Macie with EventBridge + SNS: Automatically identifies sensitive data, triggers alerts, and uses SNS for immediate notification via email.

D . Macie with EventBridge + SQS: Introduces latency due to periodic polling and adds unnecessary complexity.

Reference: Amazon Macie, Amazon EventBridge

Question: 155

An ecommerce company is planning to migrate an on-premises Microsoft SQL Server database to the AWS Cloud. The company needs to migrate the database to SQL Server Always On availability groups. The cloud-based solution must be highly available.

Options:

A. Deploy three Amazon EC2 instances with SQL Server across three Availability Zones. Attach one Amazon Elastic Block Store (Amazon EBS) volume to the EC2 instances.

B. Migrate the database to Amazon RDS for SQL Server. Configure a Multi-AZ deployment and read replicas.

C. Deploy three Amazon EC2 instances with SQL Server across three Availability Zones. Use Amazon FSx for Windows File Server as the storage tier.

D. Deploy three Amazon EC2 instances with SQL Server across three Availability Zones. Use Amazon S3 as the storage tier.

Answer: C

Explanation:

A . EC2 with EBS: Does not support SQL Server Always On availability groups effectively.

B . RDS Multi-AZ: Provides high availability but does not support SQL Server Always On availability groups.

C . EC2 with FSx for Windows: Best solution for SQL Server Always On as FSx provides shared storage compatible with SQL Server clustering.

D . EC2 with S3: S3 is not suitable for SQL Server storage.

Reference: Amazon FSx for Windows

Question: 156

A company is developing a highly available natural language processing (NLP) application. The application handles large volumes of concurrent requests. The application performs NLP tasks such as entity recognition, sentiment analysis, and key phrase extraction on text data.

The company needs to store data that the application processes in a highly available and scalable database.

Options:

A. Create an Amazon API Gateway REST API endpoint to handle incoming requests. Configure the REST API to invoke an AWS Lambda function for each request. Configure the Lambda function to call Amazon Comprehend to perform NLP tasks on the text data. Store the processed data in Amazon

DynamoDB.

B. Create an Amazon API Gateway HTTP API endpoint to handle incoming requests. Configure the HTTP API to invoke an AWS Lambda function for each request. Configure the Lambda function to call Amazon Translate to perform NLP tasks on the text data. Store the processed data in Amazon ElastiCache.

C. Create an Amazon SQS queue to buffer incoming requests. Deploy the NLP application on Amazon EC2 instances in an Auto Scaling group. Use Amazon Comprehend to perform NLP tasks. Store the processed data in an Amazon RDS database.

D. Create an Amazon API Gateway WebSocket API endpoint to handle incoming requests. Configure the WebSocket API to invoke an AWS Lambda function for each request. Configure the Lambda function to call Amazon Textract to perform NLP tasks on the text data. Store the processed data in Amazon ElastiCache.

Answer: A

Explanation:

A . API Gateway + DynamoDB:Provides high scalability, low latency, and seamless integration with Amazon Comprehend for NLP tasks.

B . HTTP API + Translate + ElastiCache:Translate is not relevant for NLP tasks like sentiment analysis or entity recognition. ElastiCache is unsuitable for permanent storage.

C . SQS + EC2 + RDS:Increases complexity and operational overhead. RDS may not scale effectively for high concurrent loads.

D. WebSocket API + Textract:Textract is irrelevant for NLP tasks. WebSocket API is not the optimal choice for this use case.

Reference:Amazon Comprehend,Amazon DynamoDB

Question: 157

A company has developed an API using an Amazon API Gateway REST API and AWS Lambda functions. The API serves static and dynamic content to users worldwide. The company wants to decrease the latency of transferring content for API requests.

Options:

- A. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.
- B. Deploy the REST API as a Regional API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.
- C. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.
- D. Deploy the REST API as a Regional API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.

Answer: A

Explanation:

- A . Edge-optimized API + Caching: Reduces latency by using Amazon CloudFront for edge locations and enables caching for dynamic content. Compression reduces data transfer latency.
- B . Regional API + Caching: Increases latency for global users due to the lack of edge locations.
- C . Edge-optimized API + Reserved Concurrency: Reserved concurrency ensures Lambda availability but does not address latency for dynamic content.
- D . Regional API + Reserved Concurrency: Lacks edge optimization, increasing latency for global users.

Reference: Amazon API Gateway

Question: 158

A gaming company is building an application that uses a database to store user data. The company wants the database to have an active-active configuration that allows data writes to a secondary AWS Region. The database must achieve a sub-second recovery point objective (RPO).

Options:

- A. Deploy an Amazon ElastiCache (Redis OSS) cluster. Configure a global data store for disaster recovery. Configure the ElastiCache cluster to cache data from an Amazon RDS database that is deployed in the primary Region.
- B. Deploy an Amazon DynamoDB table in the primary Region and the secondary Region. Configure Amazon DynamoDB Streams to invoke

an AWS Lambda function to write changes from the table in the primary Region to the table in the secondary Region.

C. Deploy an Amazon Aurora MySQL database in the primary Region. Configure a global database for the secondary Region.

D. Deploy an Amazon DynamoDB table in the primary Region. Configure global tables for the secondary Region.

Answer: D

Explanation:

A . ElastiCache:Provides in-memory caching, not suitable for persistent, scalable databases.

B . DynamoDB Streams + Lambda:Manages replication manually, increasing latency and operational complexity.

C . Aurora Global Database:Provides high availability but does not support active-active configuration.

D . DynamoDB Global Tables:Provides active-active configuration and sub-second RPO.

Reference:Amazon DynamoDB Global Tables

Question: 159

A company wants to implement a data lake in the AWS Cloud. The company must ensure that only specific teams have access to sensitive data in the data lake. The company must have row-level access control for the data lake.

Options:

A. Use Amazon RDS to store the data. Use IAM roles and permissions for data governance and access control.

B. Use Amazon Redshift to store the data. Use IAM roles and permissions for data governance and access control.

C. Use Amazon S3 to store the data. Use AWS Lake Formation for data governance and access control.

D. Use AWS Glue Catalog to store the data. Use AWS Glue DataBrew for data governance and access control.

Answer: C

Explanation:

Detailed

A . RDS:Suitable for relational databases but does not provide native support for data lakes or rowlevel access.

- B . Redshift:Primarily for analytics, not designed for large-scale data lake governance.
- C . S3 + Lake Formation:Provides native support for data lakes with granular access control, including row-level permissions.
- D . Glue Catalog + DataBrew:Focused on data preparation and metadata management, not row-level access control.

Reference:AWS Lake Formation

Question: 160

A company hosts a multi-tier inventory reporting application on AWS. The company needs a costeffective solution to generate inventory reports on demand. Admin users need to have the ability to generate new reports. Reports take approximately 5-10 minutes to finish. The application must send reports to the email address of the admin user who generates each report.

Options:

- A. Use Amazon Elastic Container Service (Amazon ECS) to host the report generation code. Use an Amazon API Gateway HTTP API to invoke the code. Use Amazon Simple Email Service (Amazon SES) to send the reports to admin users.
- B. Use Amazon EventBridge to invoke a scheduled AWS Lambda function to generate the reports. Use Amazon Simple Notification Service (Amazon SNS) to send the reports to admin users.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS) to host the report generation code. Use an Amazon API Gateway REST API to invoke the code. Use Amazon Simple Notification Service (Amazon SNS) to send the reports to admin users.
- D. Create an AWS Lambda function to generate the reports. Use a function URL to invoke the function. Use Amazon Simple Email Service (Amazon SES) to send the reports to admin users.

Answer: D

Explanation:

Detailed

- A . ECS + API Gateway:Overly complex and costly for an on-demand, intermittent workload.
- B . EventBridge + SNS:EventBridge schedules are unnecessary for on-demand generation.
- C . EKS + API Gateway:Overkill for this use case, with high operational overhead.
- D . Lambda + SES:Most cost-effective and efficient solution for generating and emailing reports on demand.

Reference:AWS Lambda,Amazon SES

Question: 161

A company that has multiple AWS accounts maintains an on-premises Microsoft Active Directory. The company needs a solution to

implement Single Sign-On for its employees. The company wants to use AWS IAM Identity Center.

The solution must meet the following requirements:

Allow users to access AWS accounts and third-party applications by using existing Active Directory credentials.

Enforce multi-factor authentication (MFA) to access AWS accounts.

Centrally manage permissions to access AWS accounts and applications.

Options:

- A. Create an IAM identity provider for Active Directory in each AWS account. Ensure that Active Directory users and groups access AWS accounts directly through IAM roles. Use IAM Identity Center to enforce MFA in each account for all users.
- B. Use AWS Directory Service to create a new AWS Managed Microsoft AD Active Directory. Configure IAM Identity Center in each account to use the new AWS Managed Microsoft AD Active Directory as the identity source. Use IAM Identity Center to enforce MFA for all users.
- C. Use IAM Identity Center with the existing Active Directory as the identity source. Enforce MFA for all users. Use AWS Organizations and Active Directory groups to manage access permissions for AWS accounts and application access.
- D. Use AWS Lambda functions to periodically synchronize Active Directory users and groups with IAM users and groups in each AWS account. Use IAM roles and policies to manage application access. Create a second Lambda function to enforce MFA.

Answer: C

Explanation:

Detailed

- A . IAM identity provider:Does not support centralized management across multiple accounts.
- B . AWS Managed AD:Unnecessary if an on-premises Active Directory already exists.
- C . IAM Identity Center + Existing AD:Best approach to integrate existing Active Directory for SSO, with MFA and centralized permissions.
- D . Lambda for synchronization:Adds complexity and does not leverage IAM Identity Center capabilities.

Reference:AWS IAM Identity Center

Question: 162

A company runs an order management application on AWS. The application allows customers to place orders and pay with a credit card. The company uses an Amazon CloudFront distribution to deliver the application.

A security team has set up logging for all incoming requests. The security team needs a solution to generate an alert if any user modifies the logging configuration.

Options (Select TWO):

- A. Configure an Amazon EventBridge rule that is invoked when a user creates or modifies a CloudFront distribution. Add the AWS Lambda function as a target of the EventBridge rule.
- B. Create an Application Load Balancer (ALB). Enable AWS WAF rules for the ALB. Configure an AWS Config rule to detect security violations.
- C. Create an AWS Lambda function to detect changes in CloudFront distribution logging. Configure the Lambda function to use Amazon Simple Notification Service (Amazon SNS) to send notifications to the security team.
- D. Set up Amazon GuardDuty. Configure GuardDuty to monitor findings from the CloudFront distribution. Create an AWS Lambda function to address the findings.
- E. Create a private API in Amazon API Gateway. Use AWS WAF rules to protect the private API from common security problems.

Answer: A,C

Explanation:

Detailed

A . EventBridge Rule: Detects modifications to CloudFront distributions in real time and triggers the Lambda function for further action.

B . ALB + Config: Focuses on ALB security violations, not relevant for CloudFront logging changes.

C . Lambda + SNS: Provides real-time notifications about changes in logging configuration.

D . GuardDuty: Focuses on threat detection, not logging configuration changes.

E . API Gateway + WAF: Unrelated to CloudFront logging changes.

Reference: Amazon EventBridge, Amazon SNS

Question: 163

A company has developed an API by using an Amazon API Gateway REST API and AWS Lambda functions. The API serves static content and dynamic content to users worldwide. The company wants to decrease the latency of transferring the content for API requests. Which solution will meet these requirements?

- A. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.

- B. Deploy the REST API as a Regional API endpoint. Enable caching. Enable content encoding in the API definition to compress the application data in transit.
- C. Deploy the REST API as an edge-optimized API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.
- D. Deploy the REST API as a Regional API endpoint. Enable caching. Configure reserved concurrency for the Lambda functions.

Answer: A

Explanation:

Edge-Optimized API: Designed for global users by routing requests through CloudFront's edge locations, reducing latency.

Content Encoding: Enabling content encoding compresses data, further optimizing performance by decreasing payload size.

Caching: Adding API Gateway caching reduces the number of calls to Lambda and database backends, improving latency.

Reserved Concurrency: Although useful, this does not directly affect latency for transferring static and dynamic content.

AWS API Gateway Edge-Optimized APIs Documentation

Question: 164

A law firm needs to make hundreds of files readable for the general public. The law firm must prevent members of the public from modifying or deleting the files before a specified future date. Which solution will meet these requirements MOST securely?

- A. Upload the files to an Amazon S3 bucket that is configured for static website hosting. Grant readonly IAM permissions to any AWS principals that access the S3 bucket until the specified date.
- B. Create a new Amazon S3 bucket. Enable S3 Versioning. Use S3 Object Lock and set a retention period based on the specified date. Create an Amazon CloudFront distribution to serve content from the bucket. Use an S3 bucket policy to restrict access to the CloudFront origin access control (OAC).
- C. Create a new Amazon S3 bucket. Enable S3 Versioning. Configure an event trigger to run an AWS Lambda function if a user modifies or deletes an object. Configure the Lambda function to replace the modified or deleted objects with the original versions of the objects from a private S3 bucket.
- D. Upload the files to an Amazon S3 bucket that is configured for static website hosting. Select the folder that contains the files. Use S3 Object Lock with a retention period based on the specified date. Grant read-only IAM permissions to any AWS principals that access the S3 bucket.

Answer: B

Explanation:

S3 Object Lock: Enables Write Once Read Many (WORM) protection for data, preventing objects from being deleted or modified for a set retention period.

S3 Versioning: Helps maintain object versions and ensures a recovery path for accidental overwrites.

CloudFront Distribution: Ensures secure and efficient public access by serving content through an edge-optimized delivery network while protecting S3 data with OAC.

Bucket Policy for OAC: Restricts public access to only the CloudFront origin, ensuring maximum security.

Amazon S3 Object Lock Documentation

Question: 165

A media company hosts a web application on AWS. The application gives users the ability to upload and view videos. The application stores the videos in an Amazon S3 bucket. The company wants to ensure that only authenticated users can upload videos. Authenticated users must have the ability to upload videos only within a specified time frame after authentication. Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure the application to generate IAM temporary security credentials for authenticated users.
- B. Create an AWS Lambda function that generates pre-signed URLs when a user authenticates.
- C. Develop a custom authentication service that integrates with Amazon Cognito to control and log direct S3 bucket access through the application.
- D. Use AWS Security Token Service (AWS STS) to assume a pre-defined IAM role that grants authenticated users temporary permissions to upload videos directly to the S3 bucket.

Answer: B

Explanation:

Pre-Signed URLs: Allow temporary access to S3 buckets, making it easy to manage time-limited upload permissions without complex operational overhead.

Lambda for Automation: Automatically generates and provides pre-signed URLs when users authenticate, minimizing manual steps and code complexity.

Least Operational Overhead: Requires no custom authentication service or deep integration with STS or Cognito.

Amazon S3 Pre-Signed URLs Documentation

Question: 166

A company has a large fleet of vehicles that are equipped with internet connectivity to send telemetry to the company. The company receives over 1 million data points every 5 minutes from the vehicles. The company uses the data in machine learning (ML) applications to predict vehicle maintenance needs and to preorder parts. The company produces visual reports based on the captured data. The company wants to migrate the telemetry ingestion, processing, and visualization workloads to AWS. Which solution will meet these requirements?

- A. Use Amazon Timestream for LiveAnalytics to store the data points. Grant Amazon SageMaker permission to access the data for processing. Use Amazon QuickSight to visualize the data.
- B. Use Amazon DynamoDB to store the data points. Use DynamoDB Connector to ingest data from DynamoDB into Amazon EMR for processing. Use Amazon QuickSight to visualize the data.
- C. Use Amazon Neptune to store the data points. Use Amazon Kinesis Data Streams to ingest data from Neptune into an AWS Lambda function for processing. Use Amazon QuickSight to visualize the data.
- D. Use Amazon Timestream to for LiveAnalytics to store the data points. Grant Amazon SageMaker permission to access the data for processing. Use Amazon Athena to visualize the data.

Answer: A

Explanation:

Amazon Timestream: Purpose-built time series database optimized for telemetry and IoT data ingestion and analytics.

Amazon SageMaker: Provides ML capabilities for predictive maintenance workflows.

Amazon QuickSight: Efficiently generates interactive, real-time visual reports from Timestream data.

Optimized for Scale: Timestream efficiently handles large-scale telemetry data with time-series indexing and queries.

Amazon Timestream Documentation

Question: 167

A company runs an application on Amazon EC2 instances. The instances need to access an Amazon RDS database by using specific credentials. The company uses AWS Secrets Manager to contain the credentials the EC2 instances must use. Which solution will meet this requirement?

- A. Create an IAM role, and attach the role to each EC2 instance profile. Use an identity-based policy to grant the new IAM role access to the secret that contains the database credentials.
- B. Create an IAM user, and attach the user to each EC2 instance profile. Use a resource-based policy to grant the new IAM user access to the secret that contains the database credentials.
- C. Create a resource-based policy for the secret that contains the database credentials. Use EC2 Instance Connect to access the secret.
- D. Create an identity-based policy for the secret that contains the database credentials. Grant direct access to the EC2 instances.

Answer: A

Explanation:

IAM Role: Attaching an IAM role to an EC2 instance profile is a secure way to manage permissions without embedding credentials.

AWS Secrets Manager: Grants controlled access to database credentials and automatically rotates secrets if configured.

Identity-Based Policy: Ensures the IAM role only has access to specific secrets, enhancing security.

AWS Secrets Manager Documentation

Question: 168

A company stores petabytes of historical medical information on premises. The company has a process to manage encryption of the data to comply with regulations. The company needs a cloudbased solution for data backup, recovery, and archiving. The company must retain control over the encryption key material. Which combination of solutions will meet these requirements? (Select TWO.)

- A. Create an AWS Key Management Service (AWS KMS) key without key material. Import the company's key material into the KMS key.
- B. Create an AWS Key Management Service (AWS KMS) encryption key that contains key material generated by AWS KMS.
- C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) storage. Use S3 Bucket Keys with AWS Key Management Service (AWS KMS) keys.
- D. Store the data in an Amazon S3 Glacier storage class. Use server-side encryption with customer- provided keys (SSE-C).
- E. Store the data in AWS Snowball devices. Use server-side encryption with AWS KMS keys (SSE- KMS).

Answer: A,D

Explanation:

Option A: Importing customer-managed keys into AWS KMS ensures that encryption key material remains under the company's control.

Option D: S3 Glacier with server-side encryption using customer-provided keys (SSE-C) complies with the need for controlled encryption and provides cost-effective storage for backups.

[AWS Key Management Service Importing Keys Documentation](#), [S3 Encryption Documentation](#)

Question: 169

A company is developing a social media application. The company anticipates rapid and unpredictable growth in users and data volume. The application needs to handle a continuous high volume of user requests. User requests include long-running processes that store large amounts of user-generated content and user profiles in a relational format. The processes must run in a specific order. The company requires an architecture that can scale resources to meet demand spikes without downtime or performance degradation. The company must ensure that the components of the application can evolve independently without affecting other parts of the system. Which combination of AWS services will meet these requirements?

- A. Deploy the application on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Use Amazon RDS as the database. Use Amazon Simple Queue Service (Amazon SQS) to decouple message processing between components.
- B. Deploy the application on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Use Amazon RDS as the database. Use Amazon Simple Notification Service (Amazon SNS) to decouple message processing between components.
- C. Use Amazon DynamoDB as the database. Use AWS Lambda functions to implement the application. Configure Amazon DynamoDB Streams to invoke the Lambda functions. Use AWS Step Functions to manage workflows between services.
- D. Use an AWS Elastic Beanstalk environment with auto scaling to deploy the application. Use Amazon RDS as the database. Use Amazon Simple Notification Service (Amazon SNS) to decouple message processing between components.

Answer: A

Explanation:

ECS with Fargate: Allows containerized workloads to scale rapidly without managing underlying servers, handling unpredictable growth effectively.

RDS for Relational Data: Manages large relational datasets efficiently while supporting high availability.

SQS for Decoupling: Ensures message processing occurs in a specific order, decoupling application components and allowing independent evolution.

[AWS ECS with Fargate Documentation](#), [AWS SQS Documentation](#)

Question: 170

A company stores data in Amazon S3. According to regulations, the data must not contain personally identifiable information (PII). The company recently discovered that S3 buckets have some objects that contain PII. The company needs to automatically detect PII in S3 buckets and to notify the company's security team. Which solution will meet these requirements?

- A. Use Amazon Macie. Create an Amazon EventBridge rule to filter the SensitiveData event type from Macie findings and to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- B. Use Amazon GuardDuty. Create an Amazon EventBridge rule to filter the CRITICAL event type from GuardDuty findings and to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- C. Use Amazon Macie. Create an Amazon EventBridge rule to filter the SensitiveData:S3Object/Personal event type from Macie findings and to send an Amazon Simple Queue Service (Amazon SQS) notification to the security team.
- D. Use Amazon GuardDuty. Create an Amazon EventBridge rule to filter the CRITICAL event type from GuardDuty findings and to send an Amazon Simple Queue Service (Amazon SQS) notification to the security team.

Answer: A

Explanation:

Amazon Macie: Detects sensitive data such as PII in S3 buckets using machine learning. EventBridge Rule: Filters Macie findings for specific sensitive data events (e.g., SensitiveData). SNS Notification: Provides real-time alerts to the security team for immediate action.

[Amazon Macie Documentation](#), [Amazon EventBridge Documentation](#)

Question: 171

A company plans to use AWS to run high-performance computing (HPC) workloads and analytics workloads. The company will run HPC workloads on Amazon EC2 instances. The workloads require a high-performance file system that can scale to millions of input/output operations per second (IOPS). Which combination of steps will meet these requirements? (Select TWO.)

- A. Use Amazon Elastic File System (Amazon EFS) as a high-performance file system.
- B. Use Amazon FSx for Lustre as a high-performance file system.
- C. Create an Auto Scaling group of Amazon EC2 instances. Use Reserved Instances. Configure a spread placement group. Use AWS Batch to run the analytics workloads.
- D. Use Mountpoint for Amazon S3 as a high-performance file system.
- E. Create an Auto Scaling group of Amazon EC2 instances. Use a mix of On-Demand Instances, Reserved Instances, and Spot Instances. Configure a cluster placement group. Use Amazon EMR to run the analytics workloads.

Answer: B,E

Explanation:

Option B (Amazon FSx for Lustre): FSx for Lustre is optimized for high-performance file systems required by HPC workloads, scaling to millions of IOPS and supporting parallelized data access.

Option E (Cluster Placement Group with Auto Scaling): A cluster placement group ensures low-latency communication between EC2 instances, critical for HPC workloads. Amazon EMR simplifies running large-scale analytics jobs.

[Amazon FSx for Lustre Documentation](#), [AWS Placement Groups Documentation](#)

Question: 172

A company provides a trading platform to customers. The platform uses an Amazon API Gateway REST API, AWS Lambda functions, and an Amazon DynamoDB table. Each trade that the platform processes invokes a Lambda function that stores the trade data in Amazon DynamoDB. The company wants to ingest trade data into a data lake in Amazon S3 for near real-time analysis. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon DynamoDB Streams to capture the trade data changes. Configure DynamoDB Streams to invoke a Lambda function that writes the data to Amazon S3.
- B. Use Amazon DynamoDB Streams to capture the trade data changes. Configure DynamoDB Streams to invoke a Lambda function that writes the data to Amazon Data Firehose. Write the data from Data Firehose to Amazon S3.
- C. Enable Amazon Kinesis Data Streams on the DynamoDB table to capture the trade data changes. Configure Kinesis Data Streams to invoke a Lambda function that writes the data to Amazon S3.
- D. Enable Amazon Kinesis Data Streams on the DynamoDB table to capture the trade data changes. Configure a data stream to be the input for Amazon Data Firehose. Write the data from Data Firehose to Amazon S3.

Answer: A

Explanation:

DynamoDB Streams: Captures real-time changes in DynamoDB tables and allows integration with Lambda for processing the changes.

Minimal Operational Overhead: Using a Lambda function directly to write data to S3 ensures simplicity and reduces the complexity of the pipeline.

[Amazon DynamoDB Streams Documentation](#)

Question: 173

A company has a large amount of data in an Amazon DynamoDB table. A large batch of data is appended to the table once each day. The company wants a solution that will make all the existing and future data in DynamoDB available for analytics on a long-term basis. Which solution meets these requirements with the LEAST operational overhead?

- A. Configure DynamoDB incremental exports to Amazon S3.
- B. Configure Amazon DynamoDB Streams to write records to Amazon S3.
- C. Configure Amazon EMR to copy DynamoDB data to Amazon S3.
- D. Configure Amazon EMR to copy DynamoDB data to Hadoop Distributed File System (HDFS).

Answer: A

Explanation:

Incremental Exports: Exporting DynamoDB data directly to Amazon S3 provides an automated, serverless way to make data available for analytics without operational overhead.

Analytics-Friendly Storage: Amazon S3 supports long-term analytics workloads and can integrate with tools like Athena or QuickSight.

[DynamoDB Export to S3 Documentation](#)

Question: 174

A company is building a serverless application to process large video files that users upload. The application performs multiple tasks to process each video file. Processing can take up to 30 minutes for the largest files.

The company needs a scalable architecture to support the processing application.

Which solution will meet these requirements?

- A. Store the uploaded video files in Amazon Elastic File System (Amazon EFS). Configure a schedule in Amazon EventBridge Scheduler to invoke an AWS Lambda function periodically to check for new files. Configure the Lambda function to perform all the processing tasks.
- B. Store the uploaded video files in Amazon Elastic File System (Amazon EFS). Configure an Amazon EFS event notification to start an AWS Step Functions workflow that uses AWS Fargate tasks to perform the processing tasks.
- C. Store the uploaded video files in Amazon S3. Configure an Amazon S3 event notification to send an event to Amazon EventBridge when a user uploads a new video file. Configure an AWS Step Functions workflow as a target for an EventBridge rule. Use the workflow to manage AWS Fargate tasks to perform the processing tasks.

D. Store the uploaded video files in Amazon S3. Configure an Amazon S3 event notification to invoke an AWS Lambda function when a user uploads a new video file. Configure the Lambda function to perform all the processing tasks.

Answer: C

Explanation:

The requirements include:

Scalability: The solution must scale as video files are uploaded.

Long-running tasks: Processing tasks can take up to 30 minutes. AWS Lambda has a maximum execution time of 15 minutes, which rules out options that involve Lambda performing all the processing.

Serverless and event-driven architecture: Ensures cost-effectiveness and high availability.

Analysis of Options:

Option A:

AWS Lambda has a 15-minute timeout, which cannot support tasks that take up to 30 minutes.

EventBridge Scheduler is unnecessary for monitoring files when native event notifications are available. Not a valid choice.

Option B:

AWS Step Functions and AWS Fargate can handle long-running processes, but Amazon EFS is not the ideal storage for uploaded video files in a serverless architecture.

Processing tasks triggered by EFS events are not a common pattern and may introduce complexities. Not the best practice.

Option C:

Amazon S3 is used for storing uploaded files, which integrates natively with event-driven services like EventBridge and Step Functions.

Amazon S3 event notifications trigger a Step Functions workflow, which can orchestrate Fargate tasks to process large video files, meeting the scalability and execution time requirements. Correct choice.

Option D:

Similar to Option A, AWS Lambda cannot handle long-running processes due to its 15-minute timeout.

Invoking Lambda for processing directly is not feasible for tasks that take up to 30 minutes. Not a valid choice.

AWS Reference:

Amazon S3 Event Notifications: [AWS Documentation - S3 Event Notifications](#)

AWS Step Functions:AWS Documentation - Step Functions

AWS Fargate:AWS Documentation - Fargate

Comparison of Storage Services:AWS Storage Options

By leveragingAmazon S3,Step Functions, andFargate, this solution provides a scalable, efficient, and serverless approach to handling video processing tasks.

Question: 175

A company is migrating a new application from an on-premises data center to a new VPC in the AWS Cloud. The company has multiple AWS accounts and VPCs that share many subnets and applications.

The company wants to have fine-grained access control for the new application. The company wants to ensure that all network resources across accounts and VPCs that are granted permission to access the new application can access the application.

- A. Set up a VPC peering connection for each VPC that needs access to the new application VPC. Update route tables in each VPC to enable connectivity.
- B. Deploy a transit gateway in the account that hosts the new application. Share the transit gateway with each account that needs to connect to the application. Update route tables in the VPC that hosts the new application and in the transit gateway to enable connectivity.
- C. Use an AWS PrivateLink endpoint service to make the new application accessible to other VPCs. Control access to the application by using an endpoint policy.
- D. Use an Application Load Balancer (ALB) to expose the new application to the internet. Configure authentication and authorization processes to ensure that only specified VPCs can access the application.

Answer: C

Explanation:

AWS PrivateLinkis the most suitable solution for providing fine-grained access control while allowing multiple VPCs, potentially across multiple accounts, to access the new application. This approach offers the following advantages:

Fine-grained control: Endpoint policies can restrict access to specific services or principals.

No need for route table updates: Unlike VPC peering or transit gateways, AWS PrivateLink does not require complex route table management.

Scalable architecture: PrivateLink scales to support traffic from multiple VPCs.

Secure connectivity: Ensures private connectivity over the AWS network, without exposing resources to the internet.

Why Other Options Are Not Ideal:

Option A:

VPC peering is not scalable when connecting multiple VPCs or accounts.

Route table management becomes complex as the number of VPCs increases. Not scalable.

Option B:

While transit gateways provide scalable VPC connectivity, they are not ideal for fine-grained access control.

Transit gateways allow connectivity but do not inherently restrict access to specific applications. Not ideal for fine-grained access control.

Option D:

Exposing the application through an ALB over the internet is not secure and does not align with the requirement to use private network resources. Security risk.

AWS Reference:

AWS PrivateLink: AWS Documentation - PrivateLink

AWS Networking Services Comparison: AWS Whitepaper - Networking Services

Question: 176

A solutions architect needs to implement a solution that can handle up to 5,000 messages per second. The solution must publish messages as events to multiple consumers. The messages are up to 500 KB in size. The message consumers need to have the ability to use multiple programming languages to consume the messages with minimal latency. The solution must retain published messages for more than 3 months. The solution must enforce strict ordering of the messages.

- A. Publish messages to an Amazon Kinesis Data Streams data stream. Enable enhanced fan-out. Ensure that consumers ingest the data stream by using dedicated throughput.
- B. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to subscribe to the topic.
- C. Publish messages to Amazon EventBridge. Allow each consumer to create rules to deliver messages to the consumer's own target.
- D. Publish messages to an Amazon Simple Notification Service (Amazon SNS) topic. Ensure that consumers use Amazon Data Firehose to subscribe to the topic.

Answer: A

Explanation:

Amazon Kinesis Data Streams is the best choice for this scenario:

Message throughput: Kinesis Data Streams supports high throughput with enhanced fan-out and dedicated throughput for consumers.

Large message size: Supports message sizes up to 1 MB, meeting the 500 KB requirement.

Message retention: Data streams can retain messages for up to 365 days.

Strict ordering: Guarantees message ordering within shards.

Why Other Options Are Not Ideal:

Option B:

While SQS FIFO supports strict ordering, SNS topics do not. SNS also does not natively support message retention or strict ordering across consumers. Does not meet requirements.

Option C:

EventBridge does not provide strict ordering guarantees or message retention beyond 24 hours. Does not meet requirements.

Option D:

SNS topics with Data Firehose are not designed for use cases requiring strict ordering or long message retention. Does not meet requirements.

AWS Reference:

Amazon Kinesis Data Streams: [AWS Documentation - Kinesis Data Streams](#)

AWS Messaging Services Comparison: [AWS Documentation - Messaging Services](#)

Question: 177

A healthcare provider is planning to store patient data on AWS as PDF files. To comply with regulations, the company must encrypt the data and store the files in multiple locations. The data must be available for immediate access from any environment.

- A. Store the files in an Amazon S3 bucket. Use the Standard storage class. Enable server-side encryption with Amazon S3 managed keys (SSE-S3) on the bucket. Configure cross-Region replication on the bucket.
- B. Store the files in an Amazon Elastic File System (Amazon EFS) volume. Use an AWS KMS managed key to encrypt the EFS volume. Use AWS DataSync to replicate the EFS volume to a second AWS Region.
- C. Store the files in an Amazon Elastic Block Store (Amazon EBS) volume. Configure AWS Backup to back up the volume on a

regular schedule. Use an AWS KMS key to encrypt the backups.

D. Store the files in an Amazon S3 bucket. Use the S3 Glacier Flexible Retrieval storage class. Ensure that all PDF files are encrypted by using client-side encryption before the files are uploaded. Configure cross-Region replication on the bucket.

Answer: A

Explanation:

AmazonS3 with the Standard storage class is the best solution:

Encryption: SSE-S3 ensures server-side encryption of the data, meeting compliance requirements.

Immediate access: The Standard storage class provides low-latency and high-throughput access to data.

Multi-location storage: Cross-Region replication ensures data is stored in multiple AWS Regions for redundancy.

Why Other Options Are Not Ideal:

Option B:

Amazon EFS is more costly and suited for file systems rather than object storage. Not cost-effective.

Option C:

Amazon EBS is block storage and not optimized for object storage like PDFs. Backup schedules do not ensure immediate availability. Not suitable.

Option D:

S3 Glacier Flexible Retrieval is designed for archival, not immediate access. Does not meet access requirements.

AWS Reference:

Amazon S3 Standard Storage: [AWS Documentation - S3 Storage Classes](#)

Amazon S3 Cross-Region Replication: [AWS Documentation - Cross-Region Replication](#)

AWS Encryption Options: [AWS Documentation - S3 Encryption](#)

Question: 178

A company wants to use an API to translate text from one language to another. The API must receive an HTTP header value and pass the value to an embedded library. The API translates documents in 6 minutes. The API requires a custom authorization mechanism.

- A. Configure an Amazon API Gateway REST API with AWS_PROXY integration to synchronously call an AWS Lambda function to perform translations.
- B. Configure an AWS Lambda function with a Lambda function URL to synchronously call a second function to perform translations.
- C. Configure an Amazon API Gateway REST API with AWS_PROXY integration to asynchronously call an AWS Lambda function to perform translations.
- D. Configure an Amazon API Gateway REST API with HTTP_PROXY integration to synchronously call a web endpoint that is hosted on an EC2 instance.

Answer: A

Explanation:

The AWS_PROXY integration with Amazon API Gateway allows the API to invoke a Lambda function synchronously, making it a suitable solution for the custom authorization mechanism and text translation use case.

Synchronous Invocation: The API Gateway REST API with AWS_PROXY integration enables synchronous processing of HTTP requests and responses, which is required for document translation.

Custom Authorization: API Gateway supports custom authorizers for fine-grained access control.

Lambda Function Execution: Although Lambda's execution time limit is 15 minutes, this is sufficient for the 6-minute document translation requirement.

Why Other Options Are Not Ideal:

Option B:

Introducing a Lambda function URL to invoke another Lambda function unnecessarily complicates the architecture. Not efficient.

Option C:

Asynchronous invocation cannot guarantee real-time response delivery for document translation tasks. Not suitable.

Option D:

Hosting the API on an EC2 instance increases operational overhead. HTTP_PROXY integration does not add significant benefits here. Not cost-effective or efficient.

AWS Reference:

API Gateway Lambda Proxy Integration: [AWS Documentation - Proxy Integration](#)

Question: 179

A company uses Amazon S3 to store customer data that contains personally identifiable information (PII) attributes. The company needs to make the customer information available to company resources through an AWS Glue Catalog. The company needs to have fine-grained access control for the data so that only specific IAM roles can access the PII data.

- A. Create one IAM policy that grants access to PII. Create a second IAM policy that grants access to non-PII data. Assign the PII policy to the specified IAM roles.
- B. Create one IAM role that grants access to PII. Create a second IAM role that grants access to non- PII data. Assign the PII policy to the specified IAM roles.
- C. Use AWS Lake Formation to provide the specified IAM roles access to the PII data.
- D. Use AWS Glue to create one view for PII data. Create a second view for non-PII data. Provide the specified IAM roles access to the PII view.

Answer: C

Explanation:

AWS Lake Formation is designed for managing fine-grained access control to data in an efficient manner:

Granular Permissions: Lake Formation allows column-level, row-level, and table-level access controls, which can precisely define access to PII data.

Integration with AWS Glue Catalog: Lake Formation natively integrates with AWS Glue for seamless data cataloging and access control.

Operational Efficiency: Centralized access control policies minimize the need for separate IAM roles or policies.

Why Other Options Are Not Ideal:

Option A:

Creating multiple IAM policies introduces complexity and lacks column-level access control. Not efficient.

Option B:

Managing multiple IAM roles for granular access is operationally complex. Not efficient.

Option D:

Creating views in Glue adds unnecessary complexity and may not provide the level of granularity that Lake Formation offers. Not the best choice.

AWS Reference:

AWS Lake Formation: AWS Documentation - Lake Formation

Fine-Grained Permissions with Lake Formation: AWS Documentation - Fine-Grained Permissions

Question: 180

A company stores 5 PB of archived data on physical tapes. The company needs to preserve the data for another 10 years. The data center that stores the tapes has a 10 Gbps Direct Connect connection to an AWS Region. The company wants to migrate the data to AWS within the next 6 months.

- A. Read the data from the tapes on premises. Use local storage to stage the data. Use AWS DataSync to migrate the data to Amazon S3 Glacier Flexible Retrieval storage.
- B. Use an on-premises backup application to read the data from the tapes. Use the backup application to write directly to Amazon S3 Glacier Deep Archive storage.
- C. Order multiple AWS Snowball Edge devices. Copy the physical tapes to virtual tapes on the Snowball Edge devices. Ship the Snowball Edge devices to AWS. Create an S3 Lifecycle policy to move the tapes to Amazon S3 Glacier Instant Retrieval storage.
- D. Configure an on-premises AWS Storage Gateway Tape Gateway. Create virtual tapes in the AWS Cloud. Use backup software to copy the physical tapes to the virtual tapes. Move the virtual tapes to Amazon S3 Glacier Deep Archive storage.

Answer: D

Explanation:

Analysis:

The company's requirements are to migrate 5 PB of data from physical tapes to AWS within 6 months, preserve the data for 10 years, and ensure cost efficiency. AWS Storage Gateway Tape Gateway is purpose-built for such use cases, as it seamlessly integrates with backup applications and provides virtual tape storage in Amazon S3 Glacier Deep Archive.

Why Option D is Correct:

Tape Gateway: Enables the migration of physical tapes to virtual tapes. Virtual tapes are stored in Amazon S3 and can later be archived in Amazon S3 Glacier Deep Archive for long-term storage.

Cost Efficiency: Amazon S3 Glacier Deep Archive is the lowest-cost storage class for long-term data preservation.

Operational Simplicity: Tape Gateway integrates with existing on-premises backup software, reducing the need for additional tools or manual processes.

Scalability: Can handle the migration of large datasets, such as 5 PB, within the required timeframe.

Why Other Options Are Not Ideal:

Option A:

AWS DataSync is not designed for reading data directly from physical tapes. Staging the data on local storage adds unnecessary complexity and cost. Not suitable.

Option B:

Using backup applications to write directly to S3 Glacier Deep Archive may not leverage AWS-native services optimally. Tape Gateway simplifies the workflow significantly. Less efficient.

Option C:

Snowball Edge is ideal for environments without high-bandwidth connectivity. However, the company already has a 10 Gbps Direct Connect, making Tape Gateway a better choice. Not cost effective.

AWS Reference:

AWS Storage Gateway - Tape Gateway: [AWS Documentation - Tape Gateway](#)

Amazon S3 Glacier Deep Archive: [AWS Documentation - Glacier Deep Archive](#)

Question: 181

A company is using microservices to build an ecommerce application on AWS. The company wants to preserve customer transaction information after customers submit orders. The company wants to store transaction data in an Amazon Aurora database. The company expects sales volumes to vary throughout each year.

- A. Use an Amazon API Gateway REST API to invoke an AWS Lambda function to send transaction data to the Aurora database. Send transaction data to an Amazon Simple Queue Service (Amazon SQS) queue that has a dead-letter queue. Use a second Lambda function to read from the SQS queue and to update the Aurora database.
- B. Use an Amazon API Gateway HTTP API to send transaction data to an Application Load Balancer (ALB). Use the ALB to send the transaction data to Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use ECS tasks to store the data in Aurora database.
- C. Use an Application Load Balancer (ALB) to route transaction data to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon EKS to send the data to the Aurora database.
- D. Use Amazon Data Firehose to send transaction data to Amazon S3. Use AWS Database Migration Service (AWS DMS) to migrate the data from Amazon S3 to the Aurora database.

Answer: A

Explanation:

Analysis:

The solution must handle variable sales volumes, preserve transaction information, and store data in an Amazon Aurora database with minimal operational overhead. Using API Gateway, AWS Lambda, and SQS is the best option because it provides scalability, reliability, and resilience.

Why Option A is Correct:

API Gateway: Serves as an entry point for transaction data in a serverless, scalable manner.

AWS Lambda: Processes the transactions and sends them to Amazon SQS for queuing.

Amazon SQS: Buffers the transaction data, ensuring durability and resilience against spikes in transaction volume.

Second Lambda Function: Processes messages from the SQS queue and updates the Aurora database, decoupling the workflow for better scalability.

Dead-Letter Queue (DLQ): Ensures failed transactions are logged for later debugging or reprocessing.

Why Other Options Are Not Ideal:

Option B:

Using an ALB with ECS on EC2 introduces operational overhead, such as managing EC2 instances and scaling ECS tasks. Not cost-effective.

Option C:

EKS is highly operationally intensive and requires Kubernetes cluster management, which is unnecessary for this use case. Too complex.

Option D:

Amazon Data Firehose and DMS are not designed for real-time transactional workflows. They are better suited for data analytics pipelines. Not suitable.

AWS Reference:

Amazon API Gateway: [AWS Documentation - API Gateway](#)

AWS Lambda: [AWS Documentation - Lambda](#)

Amazon SQS: [AWS Documentation - SQS](#)

Amazon Aurora: [AWS Documentation - Aurora](#)

Question: 182

A company has an application that receives and processes purchase orders. The application supports only XML data. The company needs to configure the application to accept orders in JSON format. The company does not want to modify the application.

A solutions architect is using an Amazon API Gateway HTTP API to create a new purchase order API.

The solutions architect needs to modify the application DNS record to point to the new HTTP API.

- A. Use an HTTP proxy integration to pass XML requests to the application. For JSON requests, use API Gateway mappings to convert the purchase orders to XML. Use an AWS Lambda function that is integrated with API Gateway to call the application.
- B. Use an HTTP proxy integration to pass XML requests to the application. For JSON requests, use an AWS Lambda function that is integrated with API Gateway to convert the purchase orders from JSON to XML and to call the application.
- C. Use an HTTP custom integration to pass XML requests to the application. For JSON requests, use API Gateway mappings to convert the purchase orders to XML. Use an AWS Lambda function that is integrated with API Gateway to call the application.
- D. Use an HTTP custom integration to pass XML requests to the application. For JSON requests, use an AWS Lambda function that is integrated with API Gateway to convert the purchase orders to JSON and to call the application.

Answer: B

Explanation:

Why Option B is Correct:

HTTP Proxy Integration: Passes XML requests directly to the application, which already supports XML.

JSON Conversion: An AWS Lambda function converts JSON requests to XML and calls the application. API Gateway: Acts as a front end to handle JSON requests and integrates seamlessly with Lambda for the transformation process.

Why Other Options Are Not Ideal:

Option A: Suggests using API Gateway mappings to convert JSON to XML. API Gateway mapping templates are limited in functionality and are not ideal for complex transformations.

Option C and D: Use HTTP custom integration unnecessarily, which adds complexity without additional benefits.

AWS Reference:

Amazon API Gateway Integration: [AWS Documentation - API Gateway Integration](#)

AWS Lambda: [AWS Documentation - Lambda](#)

Question: 183

A company stores data for multiple business units in a single Amazon S3 bucket that is in the company's payer AWS account. To maintain

data isolation, the business units store data in separate prefixes in the S3 bucket by using an S3 bucket policy.

The company plans to add a large number of dynamic prefixes. The company does not want to rely on a single S3 bucket policy to manage data access at scale. The company wants to develop a secure access management solution in addition to the bucket policy to enforce prefix-level data isolation.

- A. Configure the S3 bucket policy to deny s3:GetObject permissions for all users. Configure the bucket policy to allow s3:* access to individual business units.
- B. Enable default encryption on the S3 bucket by using server-side encryption with Amazon S3 managed keys (SSE-S3).
- C. Configure resource-based permissions on the S3 bucket by creating an S3 access point for each business unit.
- D. Use pre-signed URLs to provide access to the S3 bucket.

Answer: C

Explanation:

Why Option C is Correct:

S3 Access Points: Provide scalable management of access to large datasets with specific permissions for individual prefixes.

Dynamic Prefixes: Access points simplify managing access to a growing number of prefixes without relying solely on a single bucket policy.

Fine-Grained Control: Resource-based permissions on access points enforce prefix-level isolation effectively.

Why Other Options Are Not Ideal:

Option A: Using deny/allow bucket policies introduces complexity and is less scalable for dynamic prefixes.

Option B: Encryption ensures data security but does not address access management.

Option D: Pre-signed URLs are temporary and not suitable for managing access at scale.

AWS Reference:

Amazon S3 Access Points: [AWS Documentation - S3 Access Points](#)

Question: 184

A company is migrating an online marketplace application from a mainframe system to an Auto Scaling group of Amazon EC2 instances. The EC2 instances access an Amazon Aurora cluster. The application requires a scalable, persistent caching solution to store the results of in-progress transactions and SQL queries.

- A. Use an Amazon ElastiCache (Redis OSS) cluster to serve transaction and query results.
- B. Use an Amazon CloudFront distribution with an Amazon S3 bucket as the origin to cache the transactions. Add an Amazon EC2 instance store volume to the EC2 instances for query result caching.
- C. Use an Amazon ElastiCache (Memcached) cluster to serve transaction and query results.
- D. Use an Amazon ElastiCache (Redis OSS) cluster to cache the transactions. Add an Amazon EC2 instance store volume to the EC2 instances for query result caching.

Answer: A

Explanation:

Why Option A is Correct:

ElastiCache for Redis: Provides persistent, scalable caching for in-progress transactions and SQL queries. Redis supports data durability and advanced features, making it suitable for transactional workloads.

Integration with Aurora: Easily integrates with the Aurora cluster to improve query performance.

Why Other Options Are Not Ideal:

Option B: CloudFront and S3 are unsuitable for transactional caching. EC2 instance store volumes are ephemeral and lack persistence.

Option C: Memcached does not offer persistence or advanced transactional support, unlike Redis.

Option D: Combining Redis with EC2 instance store is unnecessary; Redis alone meets all caching requirements.

AWS Reference:

Amazon ElastiCache:AWS Documentation - ElastiCache

Question: 185

A company is launching a new application that will be hosted on Amazon EC2 instances. A solutions architect needs to design a solution that does not allow public IPv4 access that originates from the internet. However, the solution must allow the EC2 instances to make outbound IPv4 internet requests.

- A. Deploy a NAT gateway in public subnets in both Availability Zones. Create and configure one route table for each private subnet.
- B. Deploy an internet gateway in public subnets in both Availability Zones. Create and configure a shared route table for the private subnets.
- C. Deploy a NAT gateway in public subnets in both Availability Zones. Create and configure a shared route table for the private

subnets.

D. Deploy an egress-only internet gateway in public subnets in both Availability Zones. Create and configure one route table for each private subnet.

Answer: C

Explanation:

Why Option C is Correct:

NAT Gateway: Allows private subnets to access the internet for outbound requests while preventing inbound connections.

High Availability: Deploying NAT gateways in both AZs ensures fault tolerance.

Shared Route Table: Simplifies routing configuration for private subnets.

Why Other Options Are Not Ideal:

Option A: Creating separate route tables for each subnet adds unnecessary complexity.

Option B: Internet gateways allow inbound access, violating the requirement to block public IPv4 access.

Option D: Egress-only internet gateways are designed for IPv6, not IPv4.

AWS Reference:

Amazon VPC NAT Gateway: [AWS Documentation - NAT Gateway](#)

Question: 186

A company hosts an application on Amazon EC2 instances that are part of a target group behind an Application Load Balancer (ALB). The company has attached a security group to the ALB.

During a recent review of application logs, the company found many unauthorized login attempts from IP addresses that belong to countries outside the company's normal user base. The company wants to allow traffic only from the United States and Australia.

A. Edit the default network ACL to block IP addresses from outside of the allowed countries.

B. Create a geographic match rule in AWS WAF. Attach the rule to the ALB.

C. Configure the ALB security group to allow the IP addresses of company employees. Edit the default network ACL to block IP addresses from outside of the allowed countries.

D. Use a host-based firewall on the EC2 instances to block IP addresses from outside of the allowed countries. Configure the ALB security

group to allow the IP addresses of company employees.

Answer: B

Explanation:

Why Option B is Correct:

AWS WAF: Provides a simple way to create geographic match rules to block or allow traffic based on country IP ranges.

Least Operational Overhead: Attaching the WAF rule to the ALB ensures centralized control without modifying ACLs or instance firewalls.

Why Other Options Are Not Ideal:

Option A: Network ACLs operate at the subnet level and can become complex to manage for dynamic or evolving IP ranges.

Option C: Managing IP-based rules in security groups and ACLs lacks scalability and does not provide country-based filtering.

Option D: Configuring host-based firewalls increases operational overhead and does not leverage AWS-managed solutions.

AWS Reference:

AWS WAF Geomatch: [AWS Documentation - WAF Geomatch](#)

Question: 187

A company has an ecommerce application that users access through multiple mobile apps and web applications. The company needs a solution that will receive requests from the mobile apps and web applications through an API.

Request traffic volume varies significantly throughout each day. Traffic spikes during sales events. The solution must be loosely coupled and ensure that no requests are lost.

- A. Create an Application Load Balancer (ALB). Create an AWS Elastic Beanstalk endpoint to process the requests. Add the Elastic Beanstalk endpoint to the target group of the ALB.
- B. Set up an Amazon API Gateway REST API with an integration to an Amazon Simple Queue Service (Amazon SQS) queue. Configure a dead-letter queue. Create an AWS Lambda function to poll the queue to process the requests.
- C. Create an Application Load Balancer (ALB). Create an AWS Lambda function to process the requests. Add the Lambda function as a target of the ALB.
- D. Set up an Amazon API Gateway HTTP API with an integration to an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function to process the requests. Subscribe the function to the SNS topic to process the requests.

Answer: B

Explanation:

Why Option B is Correct:

Amazon SQS: Ensures no requests are lost, even during traffic spikes.

API Gateway: Handles dynamic traffic patterns efficiently, integrating with SQS for asynchronous processing.

Lambda: Polls the queue and processes requests in a serverless and scalable manner.

Dead-Letter Queue (DLQ): Ensures failed messages are retried or logged for debugging.

Why Other Options Are Not Ideal:

Option A: Elastic Beanstalk cannot handle queue-based decoupling, making it unsuitable for spiky traffic.

Option C: ALB to Lambda does not provide buffering for traffic spikes, risking request loss.

Option D: SNS is better suited for notifications, not reliable for ensuring message durability.

AWS Reference:

Amazon SQS: AWS Documentation - SQS

Question: 188

A company is developing a new application that will run on Amazon EC2 instances. The application needs to access multiple AWS services.

The company needs to ensure that the application will not use long-term access keys to access AWS services.

- A. Create an IAM user. Assign the IAM user to the application. Create programmatic access keys for the IAM user. Embed the access keys in the application code.
- B. Create an IAM user that has programmatic access keys. Store the access keys in AWS Secrets Manager. Configure the application to retrieve the keys from Secrets Manager when the application runs.
- C. Create an IAM role that can access AWS Systems Manager Parameter Store. Associate the role with each EC2 instance profile. Create IAM access keys for the AWS services, and store the keys in Parameter Store. Configure the application to retrieve the keys from Parameter Store when the application runs.
- D. Create an IAM role that has permissions to access the required AWS services. Associate the IAM role with each EC2 instance profile.

Answer: D

Explanation:

Why Option D is Correct:

IAM Roles with Instance Profiles: Allow applications to access AWS services securely without hardcoding long-term access keys.

Short-Term Credentials: IAM roles issue short-term credentials dynamically managed by AWS.

Why Other Options Are Not Ideal:

Option A and B: Embedding or retrieving long-term access keys introduces security risks and operational overhead.

Option C: Combining IAM roles with Parameter Store adds unnecessary complexity.

AWS Reference:

IAM Roles and Instance Profiles: [AWS Documentation - IAM Roles](#)

Question: 189

A company is developing a containerized web application that needs to be highly available and scalable. The application requires access to GPU resources.

- A. Package the application as an AWS Lambda function in a container image. Use Lambda to run the containerized application on a runtime with GPU access.
- B. Deploy the application container to Amazon Elastic Kubernetes Service (Amazon EKS). Use AWS Fargate to manage compute resources and access to GPU resources.
- C. Deploy the application container to Amazon Elastic Container Registry (Amazon ECR). Use Amazon ECR to run the containerized application with an attached GPU.
- D. Run the application on Amazon EC2 instances from a GPU instance family by using Amazon Elastic Container Service (Amazon ECS) for orchestration.

Answer: D

Explanation:

Why Option D is Correct:

GPU Access: Only EC2 instances in the GPU family (e.g., P2, P3) can provide GPU resources.

ECS Orchestration: Simplifies container deployment and management.

Why Other Options Are Not Ideal:

Option A: Lambda does not support GPU-based runtimes.

Option B: AWS Fargate does not support GPU-based workloads.

Option C: ECR is a container registry, not an orchestration or execution service.

AWS Reference:

Amazon ECS with GPU Instances: [AWS Documentation - ECS GPU Instances](#)

Question: 190

A solutions architect needs to secure an Amazon API Gateway REST API. Users need to be able to log in to the API by using common external social identity providers (IdPs). The social IdPs must use standard authentication protocols such as SAML or OpenID Connect (OIDC). The solutions architect needs to protect the API against attempts to exploit application vulnerabilities.

Which combination of steps will meet these security requirements? (Select TWO.)

- A. Create an AWS WAF web ACL that is associated with the REST API. Add the appropriate managed rules to the ACL.
- B. Subscribe to AWS Shield Advanced. Enable DDoS protection. Associate Shield Advanced with the REST API.
- C. Create an Amazon Cognito user pool with a federation to the social IdPs. Integrate the user pool with the REST API.
- D. Create an API key in API Gateway. Associate the API key with the REST API.
- E. Create an IP address filter in AWS WAF that allows only the social IdPs. Associate the filter with the web ACL and the API.

Answer: A,C

Explanation:

Step A: AWS WAF with managed rules protects the API against application-layer attacks, such as SQL injection and cross-site scripting (XSS).

Step C: Amazon Cognito provides secure authentication and supports federation with social IdPs using OIDC or SAML. It integrates seamlessly with API Gateway.

Option B: AWS Shield Advanced provides DDoS protection, which is not explicitly required in this scenario.

Option D: API keys provide identification, not authentication, and are insufficient for this use case.

Option E: IP filters in WAF are overly restrictive for federated authentication scenarios.

AWS Documentation Reference:

Amazon Cognito Federation

AWS WAF Managed Rules

Question: 191

A finance company uses backup software to back up its data to physical tape storage on-premises. To comply with regulations, the company needs to store the data for 7 years. The company must be able to restore archived data within one week when necessary.

The company wants to migrate the backup data to AWS to reduce costs. The company does not want to change the current backup software.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Storage Gateway Tape Gateway to copy the data to virtual tapes. Use AWS DataSync to migrate the virtual tapes to the Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Change the target of the backup software to S3 Standard-IA.
- B. Convert the physical tapes to virtual tapes. Use AWS DataSync to migrate the virtual tapes to Amazon S3 Glacier Flexible Retrieval. Change the target of the backup software to the S3 Glacier Flexible Retrieval.
- C. Use AWS Storage Gateway Tape Gateway to copy the data to virtual tapes. Migrate the virtual tapes to Amazon S3 Glacier Deep Archive. Change the target of the backup software to the virtual tapes.
- D. Convert the physical tapes to virtual tapes. Use AWS Snowball Edge storage-optimized devices to migrate the virtual tapes to Amazon S3 Glacier Flexible Retrieval. Change the target of the backup software to S3 Glacier Flexible Retrieval.

Answer: C

Explanation:

AWS Storage Gateway Tape Gateway provides a seamless way to migrate backup data to AWS without requiring changes to the backup software. Migrating to S3 Glacier Deep Archive ensures long-term, cost-effective storage for data that rarely needs retrieval.

Option A: S3 Standard-IA is more expensive than Glacier for long-term storage.

Option B and D: Glacier Flexible Retrieval is costlier than Glacier Deep Archive for archival use cases with low retrieval frequency.

AWS Documentation Reference:

AWS Storage Gateway Tape Gateway

S3 Glacier Storage Classes

Question: 192

A company is designing a new application that uploads files to an Amazon S3 bucket. The uploaded files are processed to extract metadata.

Processing must take less than 5 seconds. The volume and frequency of the uploads vary from a few files each hour to hundreds of concurrent uploads.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure AWS CloudTrail trails to log Amazon S3 API calls. Use AWS AppSync to process the files.
- B. Configure a new object created S3 event notification within the bucket to invoke an AWS Lambda function to process the files.
- C. Configure Amazon Kinesis Data Streams to deliver the files to the S3 bucket. Invoke an AWS Lambda function to process the files.
- D. Deploy an Amazon EC2 instance. Create a script that lists all files in the S3 bucket and processes new files. Use a cron job that runs every minute to run the script.

Answer: B

Explanation:

Using S3 event notifications to trigger AWS Lambda for file processing is a cost-effective and serverless solution. Lambda scales automatically with upload volume, and processing each file takes less than 5 seconds, fitting within Lambda's execution time.

Option A: AWS AppSync is designed for GraphQL APIs and is not suitable for file processing.

Option C: Kinesis is overkill and more expensive for this use case.

Option D: Running an EC2 instance incurs ongoing costs and is less flexible compared to Lambda.

AWS Documentation Reference:

[Amazon S3 Event Notifications](#)

[AWS Lambda Overview](#)

Question: 193

A solutions architect is designing the network architecture for an application that runs on Amazon EC2 instances in an Auto Scaling group.

The application needs to access data that is in Amazon S3 buckets.

Traffic to the S3 buckets must not use public IP addresses. The solutions architect will deploy the application in a VPC that has public and private subnets.

Which solutions will meet these requirements? (Select TWO.)

- A. Deploy the EC2 instances in a private subnet. Configure a default route to an egress-only internet gateway.
- B. Deploy the EC2 instances in a public subnet. Create a gateway endpoint for Amazon S3. Associate the endpoint with the subnet's route table.
- C. Deploy the EC2 instances in a public subnet. Create an interface endpoint for Amazon S3. Configure DNS hostnames and DNS resolution for the VPC.
- D. Deploy the EC2 instances in a private subnet. Configure a default route to a NAT gateway in a public subnet.
- E. Deploy the EC2 instances in a private subnet. Configure a default route to a customer gateway.

Answer: B,D

Explanation:

Option B: A gateway endpoint for S3 allows traffic to S3 without using public IPs and integrates with route tables.

Option D: Deploying EC2 instances in a private subnet with a NAT gateway enables outbound internet connectivity for other requirements without public IPs.

Option A: Egress-only internet gateways are for IPv6 traffic and do not work for IPv4 in this context.

Option C: Interface endpoints are not required for S3 as gateway endpoints are more suitable and cost-effective.

Option E: A customer gateway is for hybrid connectivity (e.g., on-premises), not suitable for this case.

AWS Documentation Reference:

VPC Endpoints

Amazon S3 Gateway Endpoints

Question: 194

A company is building a serverless application to process orders from an ecommerce site. The application needs to handle bursts of traffic during peak usage hours and to maintain high availability. The orders must be processed asynchronously in the order the application receives them.

Which solution will meet these requirements?

- A. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use an AWS Lambda function to process the orders.
- B. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to receive orders. Use an AWS Lambda function to process the orders.
- C. Use an Amazon Simple Queue Service (Amazon SQS) standard queue to receive orders. Use AWS Batch jobs to process the orders.
- D. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use AWS Batch jobs to process the orders.

Answer: B

Explanation:

Amazon SQS FIFO queues ensure that orders are processed in the exact order received and maintain message deduplication.

AWS Lambda scales automatically, handling bursts and maintaining high availability in a cost-effective manner.

Option A and D: Amazon SNS does not guarantee ordered processing.

Option C: Standard SQS queues do not guarantee order.

AWS Documentation Reference:

Amazon SQS FIFO Queues

Question: 195

A company's reporting system delivers hundreds of .csv files to an Amazon S3 bucket each day. The company must convert these files to Apache Parquet format and must store the files in a transformed data bucket.

Which solution will meet these requirements with the LEAST development effort?

- A. Create an Amazon EMR cluster with Apache Spark installed. Write a Spark application to transform the data. Use EMR File System (EMRFS) to write files to the transformed data bucket.
- B. Create an AWS Glue crawler to discover the data. Create an AWS Glue extract, transform, and load (ETL) job to transform the data. Specify the transformed data bucket in the output step.
- C. Use AWS Batch to create a job definition with Bash syntax to transform the data and output the data to the transformed data bucket. Use the job definition to submit a job. Specify an array job as the job type.
- D. Create an AWS Lambda function to transform the data and output the data to the transformed data bucket. Configure an event notification for the S3 bucket. Specify the Lambda function as the destination for the event notification.

Answer: B

Explanation:

AWS Glue provides a serverless ETL solution requiring minimal development. Glue supports conversion to Parquet with managed jobs and integrates with S3 for output.

AWS Documentation Reference:

[AWS Glue Overview](#)

Question: 196

A company is designing a new Amazon Elastic Kubernetes Service (Amazon EKS) deployment to host multi-tenant applications that use a single cluster. The company wants to ensure that each pod has its own hosted environment. The environments must not share CPU, memory, storage, or elastic network interfaces.

Which solution will meet these requirements?

- A. Use Amazon EC2 instances to host self-managed Kubernetes clusters. Use taints and tolerations to enforce isolation boundaries.
- B. Use Amazon EKS with AWS Fargate. Use Fargate to manage resources and to enforce isolation boundaries.
- C. Use Amazon EKS and self-managed node groups. Use taints and tolerations to enforce isolation boundaries.
- D. Use Amazon EKS and managed node groups. Use taints and tolerations to enforce isolation boundaries.

Answer: B

Explanation:

AWS Fargate provides per-pod isolation for CPU, memory, storage, and networking, making it ideal for multi-tenant use cases.

AWS Documentation Reference:

[EKS with Fargate](#)

Question: 197

A company is developing a public web application that needs to access multiple AWS services. The application will have hundreds of users who must log in to the application first before using the services.

The company needs to implement a secure and scalable method to grant the web application temporary access to the AWS

resources.

Which solution will meet these requirements?

- A. Create an IAM role for each AWS service that the application needs to access. Assign the roles directly to the instances that the web application runs on.
- B. Create an IAM role that has the access permissions the web application requires. Configure the web application to use AWS Security Token Service (AWS STS) to assume the IAM role. Use STS tokens to access the required AWS services.
- C. Use AWS IAM Identity Center to create a user pool that includes the application users. Assign access credentials to the web application users. Use the credentials to access the required AWS services.
- D. Create an IAM user that has programmatic access keys for the AWS services. Store the access keys in AWS Systems Manager Parameter Store. Retrieve the access keys from Parameter Store. Use the keys in the web application.

Answer: B

Explanation:

Option B is the correct solution because:

AWS Security Token Service (STS) allows the web application to request temporary security credentials that grant access to AWS resources. These temporary credentials are secure and shortlived, reducing the risk of misuse.

Using STS and IAM roles ensures scalability by enabling the application to dynamically assume roles with the required permissions for each AWS service.

Option A: Assigning IAM roles directly to instances is less flexible and would grant the same permissions to all applications on the instance, which is not ideal for a multi-service web application. Option C: AWS IAM Identity Center is used for managing single sign-on (SSO) for workforce users and is not designed for granting programmatic access to web applications. Option D: Storing long-term access keys, even in AWS Systems Manager Parameter Store, is less secure and does not scale well compared to temporary credentials from STS.

AWS Documentation Reference:

AWS Security Token Service (STS)

IAM Roles for Temporary Credentials

Question: 198

A company wants to run big data workloads on Amazon EMR. The workloads need to process terabytes of data in memory.

A solutions architect needs to identify the appropriate EMR cluster instance configuration for the workloads.

Which solution will meet these requirements?

- A. Use a storage optimized instance for the primary node. Use compute optimized instances for core nodes and task nodes.
- B. Use a memory optimized instance for the primary node. Use storage optimized instances for core nodes and task nodes.
- C. Use a general purpose instance for the primary node. Use memory optimized instances for core nodes and task nodes.
- D. Use general purpose instances for the primary, core, and task nodes.

Answer: C

Explanation:

Big data workloads that need to process terabytes of data in memory require memory-optimized instances for the core and task nodes to ensure sufficient memory for processing data efficiently.

Primary Node: A general purpose instance is suitable because it manages cluster operations, including coordination and monitoring, and does not process data directly.

Core and Task Nodes: These nodes handle data storage and processing. Memory-optimized instances are ideal because they provide high memory-to-CPU ratios, which is critical for in-memory big data workloads.

Why Other Options Are Incorrect:

Option A: Storage optimized and compute optimized instances are not suitable for workloads that rely heavily on in-memory processing.

Option B: A memory-optimized primary node is unnecessary because the primary node does not process data.

Option D: General purpose instances for all nodes will not provide sufficient memory for processing terabytes of data in memory.

AWS Documentation Reference:

Amazon EMR Instance Types

Memory-Optimized Instances

Question: 199

A solutions architect needs to optimize a large data analytics job that runs on an Amazon EMR cluster. The job takes 13 hours to finish. The cluster has multiple core nodes and worker nodes deployed on large, compute-optimized instances.

After reviewing EMR logs, the solutions architect discovers that several nodes are idle for more than 5 hours while the job is running. The solutions architect needs to optimize cluster performance.

Which solution will meet this requirement MOST cost-effectively?

- A. Increase the number of core nodes to ensure there is enough processing power to handle the analytics job without any idle time.
- B. Use the EMR managed scaling feature to automatically resize the cluster based on workload.
- C. Migrate the analytics job to a set of AWS Lambda functions. Configure reserved concurrency for the functions.
- D. Migrate the analytics job core nodes to a memory-optimized instance type to reduce the total job runtime.

Answer: B

Explanation:

EMR managed scaling dynamically resizes the cluster by adding or removing nodes based on the workload. This feature helps minimize idle time and reduces costs by scaling the cluster to meet processing demands efficiently.

Option A: Increasing the number of core nodes might increase idle time further, as it does not address the root cause of underutilization.

Option C: Migrating the job to Lambda is infeasible for large analytics jobs due to resource and runtime constraints.

Option D: Changing to memory-optimized instances may not necessarily reduce idle time or optimize costs.

AWS Documentation Reference:

EMR Managed Scaling

Question: 200

A solutions architect is building an Amazon S3 data lake for a company. The company uses Amazon Kinesis Data Firehose to ingest customer personally identifiable information (PII) and transactional data in near real-time to an S3 bucket. The company needs to mask all PII data before storing the data in the data lake.

Which solution will meet these requirements?

- A. Create an AWS Lambda function to detect and mask PII. Invoke the function from Kinesis Data Firehose.
- B. Use Amazon Macie to scan the S3 bucket. Configure Macie to detect and mask PII.
- C. Enable server-side encryption (SSE) on the S3 bucket.
- D. Create an AWS Lambda function that integrates with AWS CloudHSM. Configure the function to detect and mask PII.

Answer: A

Explanation:

Using a Lambda function as part of the Kinesis Data Firehose pipeline allows for real-time detection and masking of PII before data is written to S3. This ensures that PII is never stored in its raw form in the data lake.

Option B: Amazon Macie can scan and classify data but does not provide in-line PII masking for data ingestion.

Option C: Server-side encryption secures data but does not mask PII.

Option D: CloudHSM is unnecessary for PII masking and adds complexity without addressing the requirements.

AWS Documentation Reference:

Using Lambda with Kinesis Data Firehose

Question: 201

An ecommerce company runs an application that uses an Amazon DynamoDB table in a single AWS Region. The company wants to deploy the application to a second Region. The company needs to support multi-active replication with low latency reads and writes to the existing DynamoDB table in both Regions.

Which solution will meet these requirements in the MOST operationally efficient way?

A. Create a DynamoDB global secondary index (GSI) for the existing table. Create a new table in the second Region. Convert the existing DynamoDB table to a global table. Specify the new table as the secondary table.

B. Enable Amazon DynamoDB Streams for the existing table. Create a new table in the second Region. Create a new application that uses the DynamoDB Streams Kinesis Adapter and the Amazon Kinesis Client Library (KCL). Configure the new application to read data from the DynamoDB table in the first Region and to write the data to the new table in the second Region.

C. Convert the existing DynamoDB table to a global table. Choose the appropriate second Region to achieve active-active write capabilities in both Regions.

D. Enable Amazon DynamoDB Streams for the existing table. Create a new table in the second Region. Create an AWS Lambda function in the first Region that reads data from the table in the first Region and writes the data to the new table in the second Region. Set a DynamoDB stream as the input trigger for the Lambda function.

Answer: C

Explanation:

Converting the existing DynamoDB table to a global table provides active-active replication and low-latency reads and writes in both

Regions. DynamoDB global tables are specifically designed for multi-Region and multi-active use cases.

Option A: GSIs do not provide multi-Region replication or active-active capabilities.

Option B and D: Using DynamoDB Streams and custom replication is less operationally efficient than global tables and introduces additional complexity.

AWS Documentation Reference:

DynamoDB Global Tables

Question: 202

A company runs a critical public application on Amazon Elastic Kubernetes Service (Amazon EKS) clusters. The application has a microservices architecture. The company needs to implement a solution that collects, aggregates, and summarizes metrics and logs from the application in a centralized location.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Run the Amazon CloudWatch agent in the existing EKS cluster. Use a CloudWatch dashboard to view the metrics and logs.
- B. Configure a data stream in Amazon Kinesis Data Streams. Use Amazon Kinesis Data Firehose to read events and to deliver the events to an Amazon S3 bucket. Use Amazon Athena to view the events.
- C. Configure AWS CloudTrail to capture data events. Use Amazon OpenSearch Service to query CloudTrail.
- D. Configure Amazon CloudWatch Container Insights in the existing EKS cluster. Use a CloudWatch dashboard to view the metrics and logs.

Answer: D

Explanation:

Amazon CloudWatch Container Insights is designed for monitoring containerized environments like EKS. It provides native support for collecting and visualizing metrics and logs in a centralized location through CloudWatch dashboards, offering the most operationally efficient solution.

Option A: Using the CloudWatch agent provides basic metrics but lacks the specific insights required for containerized applications.

Option B: Kinesis Data Streams and Firehose add unnecessary complexity for this use case.

Option C: CloudTrail is for auditing API activity and is not designed for application metrics and log aggregation.

AWS Documentation Reference:

Question: 203

A company hosts an application that processes highly sensitive customer transactions on AWS. The application uses Amazon RDS as its database. The company manages its own encryption keys to secure the data in Amazon RDS.

The company needs to update the customer-managed encryption keys at least once each year.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Set up automatic key rotation in AWS Key Management Service (AWS KMS) for the encryption keys.
- B. Configure AWS Key Management Service (AWS KMS) to alert the company to rotate the encryption keys annually.
- C. Schedule an AWS Lambda function to rotate the encryption keys annually.
- D. Create an AWS CloudFormation stack to run an AWS Lambda function that deploys new encryption keys once each year.

Answer: A

Explanation:

AWS KMS automatic key rotation is the simplest and most operationally efficient solution. Enabling automatic key rotation ensures that KMS automatically generates new key material for the key every year without requiring manual intervention.

Option B: Configuring alerts to rotate keys introduces operational overhead as the actual rotation must still be managed manually.

Option C: Scheduling a Lambda function to rotate keys adds unnecessary complexity compared to enabling automatic key rotation.

Option D: Using a CloudFormation stack to run a Lambda function for key rotation increases operational overhead and complexity unnecessarily.

AWS Documentation Reference:

AWS KMS Key Rotation

Using Customer-Managed Keys with Amazon RDS

Question: 204

A company is designing a web application with an internet-facing Application Load Balancer (ALB).

The company needs the ALB to receive HTTPS web traffic from the public internet. The ALB must send only HTTPS traffic to the web application servers hosted on the Amazon EC2 instances on port 443. The ALB must perform a health check of the web application servers

over HTTPS on port 8443.

Which combination of configurations of the security group that is associated with the ALB will meet these requirements? (Select THREE.)

- A. Allow HTTPS inbound traffic from 0.0.0.0/0 for port 443.
- B. Allow all outbound traffic to 0.0.0.0/0 for port 443.
- C. Allow HTTPS outbound traffic to the web application instances for port 443.
- D. Allow HTTPS inbound traffic from the web application instances for port 443.
- E. Allow HTTPS outbound traffic to the web application instances for the health check on port 8443.
- F. Allow HTTPS inbound traffic from the web application instances for the health check on port 8443.

Answer: A,C,E

Explanation:

Option A: The ALB must accept HTTPS traffic from the public internet. Allowing inbound traffic on port 443 from 0.0.0.0/0 enables this functionality.

Option C: The ALB must forward HTTPS traffic to the web application servers on port 443. Outbound traffic for port 443 must be allowed for this communication.

Option E: The ALB must perform health checks on the web application servers over HTTPS on port 8443. Outbound traffic for port 8443 must be allowed for this purpose.

Option B: Allowing all outbound traffic is overly permissive and does not align with the specific requirements.

Option D and F: Inbound traffic to the ALB from the web application instances is unnecessary because the flow of traffic is from the ALB to the web application instances, not vice versa.

AWS Documentation Reference:

Application Load Balancer Security Groups

Health Checks for ALBs

Question: 205

A finance company has a web application that generates credit reports for customers. The company hosts the frontend of the web application on a fleet of Amazon EC2 instances that is associated with an Application Load Balancer (ALB). The application generates reports by running queries on an Amazon RDS for SQL Server database.

The company recently discovered that malicious traffic from around the world is abusing the application by submitting unnecessary requests. The malicious traffic is consuming significant compute resources. The company needs to address the malicious traffic. Which solution will meet this requirement?

- A. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Update the web ACL to block IP addresses that are associated with malicious traffic.
- B. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Use the AWS WAF Bot Control managed rule feature.
- C. Set up AWS Shield to protect the ALB and the database.
- D. Use AWS WAF to create a web ACL. Associate the web ACL with the ALB. Configure the AWS WAF IP reputation rule.

Answer: B

Explanation:

The AWS WAF Bot Control managed rule is designed to automatically detect and mitigate bot traffic. This feature is particularly useful for addressing malicious traffic and conserving compute resources by filtering unnecessary requests at the ALB level.

Option A: Blocking IP addresses manually introduces significant operational overhead and is not scalable against dynamic, worldwide malicious traffic.

Option C: AWS Shield provides DDoS protection, but the scenario does not describe a DDoS attack. WAF is better suited for managing application-layer threats like bot traffic.

Option D: The AWS WAF IP reputation rule helps block traffic from known bad IPs but may not address bot traffic effectively.

AWS Documentation Reference:

AWS WAF Bot Control

AWS WAF Managed Rules

Question: 206

A media company is launching a new product platform that artists from around the world can use to upload videos and images directly to an Amazon S3 bucket. The company owns and maintains the S3 bucket. The artists must be able to upload files from personal devices without the need for AWS credentials or an AWS account.

Which solution will meet these requirements MOST securely?

- A. Enable cross-origin resource sharing (CORS) on the S3 bucket.
- B. Turn off block public access for the S3 bucket. Share the bucket URL to the artists to enable uploads without credentials.
- C. Use an IAM role that has upload permissions for the S3 bucket to generate presigned URLs for S3 prefixes that are specific to each artist. Share the URLs to the artists.
- D. Create a web interface that uses an IAM role that has permission to upload and view objects in the S3 bucket. Share the web interface URL to the artists.

Answer: C

Explanation:

Option C is the most secure and practical solution. Presigned URLs allow temporary, limited access to upload files to specific S3 prefixes. This ensures that artists can upload files securely without needing AWS credentials or accounts. Each artist receives a unique URL with permissions tied to the intended S3 location, and the URL can be configured to expire after a certain time, minimizing security risks.

Why Other Options Are Incorrect:

Option A: Enabling CORS allows cross-origin access but does not provide authentication or authorization for uploads. This does not secure the uploads or restrict access to specific artists.

Option B: Turning off block public access and sharing the bucket URL exposes the bucket to potential misuse and unauthorized uploads. This approach is highly insecure.

Option D: While a web interface might work, it introduces additional complexity and potential security risks by exposing upload functionality through a public-facing application.

AWS Documentation Reference:

Using Amazon S3 Presigned URLs

AWS Best Practices for Secure S3 Access

Question: 207

A company is building a serverless application to process clickstream data from its website. The clickstream data is sent to an Amazon Kinesis Data Streams data stream from the application web servers.

The company wants to enrich the clickstream data by joining the clickstream data with customer profile data from an Amazon Aurora Multi-AZ database. The company wants to use Amazon Redshift to analyze the enriched data. The solution must be highly

available.

Which solution will meet these requirements?

- A. Use an AWS Lambda function to process and enrich the clickstream data. Use the same Lambda function to write the clickstream data to Amazon S3. Use Amazon Redshift Spectrum to query the enriched data in Amazon S3.
- B. Use an Amazon EC2 Spot Instance to poll the data stream and enrich the clickstream data. Configure the EC2 instance to use the COPY command to send the enriched results to Amazon Redshift.
- C. Use an Amazon Elastic Container Service (Amazon ECS) task with AWS Fargate Spot capacity to poll the data stream and enrich the clickstream data. Configure an Amazon EC2 instance to use the COPY command to send the enriched results to Amazon Redshift.
- D. Use Amazon Kinesis Data Firehose to load the clickstream data from Kinesis Data Streams to Amazon S3. Use AWS Glue crawlers to infer the schema and populate the AWS Glue Data Catalog. Use Amazon Athena to query the raw data in Amazon S3.

Answer: A

Explanation:

Option A is the best solution as it leverages AWS Lambda for serverless, scalable, and highly available processing and enrichment of clickstream data. Lambda can process the data in real-time, join it with the Aurora database data, and write the enriched results to Amazon S3. From S3, Amazon Redshift Spectrum can directly query the enriched data without needing to load the data into Redshift, enabling cost efficiency and high availability.

Why Other Options Are Incorrect:

Option B: EC2 Spot Instances are not guaranteed to be highly available, as Spot Instances can be interrupted at any time. This does not align with the requirement for high availability.

Option C: While ECS with AWS Fargate provides scalability, using EC2 for the COPY command introduces operational overhead and compromises high availability.

Option D: Kinesis Data Firehose and Athena are suitable for querying raw data, but they do not directly support enriching the data by joining with Aurora. This solution fails to meet the requirement for data enrichment.

Key AWS Features Used:

AWS Lambda: Real-time serverless processing with integration capabilities for Aurora and S3.

Amazon S3: Cost-effective storage for enriched data.

Amazon Redshift Spectrum: Direct querying of data stored in S3 without loading it into Redshift.

AWS Documentation Reference:

AWS Lambda Function Overview

Amazon Redshift Spectrum

Processing Streaming Data with Kinesis Data Streams

Question: 208

A company wants to create an API to authorize users by using JSON Web Tokens (JWTs). The company needs to support dynamic access to multiple AWS services by using path-based routing.

Which solution will meet these requirements?

- A. Deploy an Application Load Balancer behind an Amazon API Gateway REST API. Configure IAM authorization.
- B. Deploy an Application Load Balancer behind an Amazon API Gateway HTTP API. Use Amazon Cognito for authorization.
- C. Deploy a Network Load Balancer behind an Amazon API Gateway REST API. Use an AWS Lambda function as a custom authorizer.
- D. Deploy a Network Load Balancer behind an Amazon API Gateway HTTP API. Use Amazon Cognito for authorization.

Answer: C

Explanation:

Question: 209

A company is performing a security review of its Amazon EMR API usage. The company's developers use an integrated development environment (IDE) that is hosted on Amazon EC2 instances. The IDE is configured to authenticate users to AWS by using access keys. Traffic between the company's EC2 instances and EMR cluster uses public IP addresses.

A solutions architect needs to improve the company's overall security posture. The solutions architect needs to reduce the company's use of long-term credentials and to limit the amount of communication that uses public IP addresses.

Which combination of steps will MOST improve the security of the company's architecture? (Select TWO.)

- A. Set up a gateway endpoint to the EMR cluster.
- B. Set up interface VPC endpoints to connect to the EMR cluster.
- C. Set up a private NAT gateway to connect to the EMR cluster.
- D. Set up IAM roles for the developers to use to connect to the Amazon EMR API.

E. Set up AWS Systems Manager Parameter Store to store access keys for each developer.

Answer: B,D

Explanation:

Question: 210

A company is implementing a new policy to enhance the security of its AWS environment. The policy requires all administrative actions that users perform on the AWS Management Console to be secured by multi-factor authentication (MFA).

Which solution will allow the company to enforce this policy in the MOST operationally efficient way?

- A. Enable MFA on the root account. Ensure that all administrators use the root account to perform administrative actions.
- B. Create an IAM policy that requires MFA to be enabled for the IAM roles that administrators assume to perform administrative actions.
- C. Configure an Amazon CloudWatch alarm that sends an email notification when an administrator performs an administrative action without MFA.
- D. Use AWS Config to periodically audit IAM users and to automatically attach an IAM policy that requires MFA when AWS Config detects administrative actions.

Answer: B

Explanation:

Question: 211

A company is building a new application that uses multiple serverless architecture components. The application architecture includes an Amazon API Gateway REST API and AWS Lambda functions to manage incoming requests.

The company needs a service to send messages that the REST API receives to multiple target Lambda functions for processing. The service must filter messages so each target Lambda function receives only the messages the function needs.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Send the requests from the REST API to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe multiple Amazon Simple Queue Service (Amazon SQS) queues to the SNS topic. Configure the target Lambda functions to poll the SQS queues.
- B. Send the requests from the REST API to a set of Amazon EC2 instances that are configured to process messages. Configure the instances to filter messages and to invoke the target Lambda functions.

- C. Send the requests from the REST API to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Configure Amazon MSK to publish the messages to the target Lambda functions.
- D. Send the requests from the REST API to multiple Amazon Simple Queue Service (Amazon SQS) queues. Configure the target Lambda functions to poll the SQS queues.

Answer: C

Explanation:

Question: 212

A solutions architect is investigating compute options for a critical analytics application. The application uses long-running processes to prepare and aggregate data. The processes cannot be interrupted. The application has a known baseline load. The application needs to handle occasional usage surges.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon EC2 Auto Scaling group. Set the Min capacity and Desired capacity parameters to the number of instances required to handle the baseline load. Purchase Reserved Instances for the Auto Scaling group.
- B. Create an Amazon EC2 Auto Scaling group. Set the Min capacity, Max capacity, and Desired capacity parameters to the number of instances required to handle the baseline load. Use On-Demand Instances to address occasional usage surges.
- C. Create an Amazon EC2 Auto Scaling group. Set the Min capacity and Desired capacity parameters to the number of instances required to handle the baseline load. Purchase Reserved Instances for the Auto Scaling group. Use the OnDemandPercentageAboveBaseCapacity parameter to configure the launch template to launch Spot Instances.
- D. Re-architect the application to use AWS Lambda functions instead of Amazon EC2 instances. Purchase a one-year Compute Savings Plan to reduce the cost of Lambda usage.

Answer: C

Explanation:

Question: 213

A company wants to deploy an AWS Lambda function that will read and write objects to Amazon S3 bucket. The Lambda function must be connected to the company's VPC. The company must deploy the Lambda function only to private subnets in the VPC. The Lambda function must not be allowed to access the internet.

Which solutions will meet these requirements? (Select TWO.)

- A. Create a private NAT gateway to access the S3 bucket.

- B. Attach an Elastic IP address to the NAT gateway.
- C. Create a gateway VPC endpoint for the S3 bucket.
- D. Create an interface VPC endpoint for the S3 bucket.
- E. Create a public NAT gateway to access the S3 bucket.

Answer: C,D

Explanation:

Question: 214

A company is creating a low-latency payment processing application that supports TLS connections from IPv4 clients. The application requires outbound access to the public internet. Users must access the application from a single entry point.

The bank wants to use Amazon Elastic Container Service (Amazon ECS) tasks to deploy the application. The company wants to enable AWSVPC network mode.

Which solution will meet these requirements MOST securely?

- A. Create a VPC that has an internet gateway, public subnets, and private subnets. Deploy a Network Load Balancer and a NAT gateway in the public subnets. Deploy the ECS tasks in the private subnets.
- B. Create a VPC that has an outbound-only internet gateway, public subnets, and private subnets. Deploy an Application Load Balancer and a NAT gateway in the public subnets. Deploy the ECS tasks in the private subnets.
- C. Create a VPC that has an internet gateway, public subnets, and private subnets. Deploy an Application Load Balancer in the public subnets. Deploy the ECS tasks in the public subnets.
- D. Create a VPC that has an outbound-only internet gateway, public subnets, and private subnets. Deploy a Network Load Balancer in the public subnets. Deploy the ECS tasks in the public subnets.

Answer: A

Explanation:

Question: 215

A company hosts an Amazon EC2 instance in a private subnet in a new VPC. The VPC also has a public subnet that has the default route set to an internet gateway. The private subnet does not have outbound internet access.

The EC2 instance needs to have the ability to download monthly security updates from an outside vendor. However, the company must block any connections that are initiated from the internet.

Which solution will meet these requirements?

- A. Configure the private subnet route table to use the internet gateway as the default route.
- B. Create a NAT gateway in the public subnet. Configure the private subnet route table to use the NAT gateway as the default route.
- C. Create a NAT instance in the private subnet. Configure the private subnet route table to use the NAT instance as the default route.
- D. Create a NAT instance in the private subnet. Configure the private subnet route table to use the internet gateway as the default route.

Answer: B

Explanation:

Question: 216

A company wants to create a payment processing application. The application must run when a payment record arrives in an existing Amazon S3 bucket. The application must process each payment record exactly once. The company wants to use an AWS Lambda function to process the payments.

Which solution will meet these requirements?

- A. Configure the existing S3 bucket to send object creation events to Amazon EventBridge. Configure EventBridge to route events to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Configure the Lambda function to run when a new event arrives in the SQS queue.
- B. Configure the existing S3 bucket to send object creation events to an Amazon Simple Notification Service (Amazon SNS) topic. Configure the Lambda function to run when a new event arrives in the SNS topic.
- C. Configure the existing S3 bucket to send object creation events to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the Lambda function to run when a new event arrives in the SQS queue.
- D. Configure the existing S3 bucket to send object creation events directly to the Lambda function. Configure the Lambda function to handle object creation events and to process the payments.

Answer: B

Explanation:

Question: 217

A company is redesigning a static website. The company needs a solution to host the new website in the company's AWS account. The solution must be secure and scalable.

Which combination of solutions will meet these requirements? (Select THREE.)

- A. Configure an Amazon CloudFront distribution. Set the Amazon S3 bucket as the origin.
- B. Associate an AWS Certificate Manager (ACM) TLS certificate to the Amazon CloudFront distribution.
- C. Enable static website hosting for the Amazon S3 bucket.
- D. Create an Amazon S3 bucket to store the static website content.
- E. Export the website's SSL/TLS certificate from AWS Certificate Manager (ACM) to the root of the Amazon S3 bucket.
- F. Turn off Block Public Access for the Amazon S3 bucket.

Answer: A,B,D

Explanation:

Question: 218

A global company is migrating its workloads from an on-premises data center to AWS. The AWS environment includes multiple AWS accounts. IAM roles. AWS Config rules, and a VPC.

The company wants an automated process to provision new accounts on demand when the company's business units require new accounts.

Which solution will meet these requirements with LEAST effort?

- A. Use AWS Control Tower to set up an organization in AWS Organizations. Use AWS Control Tower Account Factory for Terraform (AFT) to provision new AWS accounts.
- B. Create an organization in AWS Organizations. Use the AWS CLI CreateAccount API action to provision new AWS accounts. Organize the business units with organizational units (OUs).
- C. Create an AWS Lambda function that uses the AWS Organizations API to create new accounts. Invoke the Lambda function from an AWS CloudFormation template in AWS Service Catalog.
- D. Create an organization in AWS Organizations. Use AWS Step Functions to orchestrate the account creation process. Send account creation requests to an Amazon API Gateway API endpoint to invoke an AWS Lambda function that creates new accounts.

Answer: A

Explanation:

Question: 219

A company wants to design a microservices architecture for an application. Each microservice must perform operations that can be completed within 30 seconds.

The microservices need to expose RESTful APIs and must automatically scale in response to varying loads. The APIs must also provide client access control and rate limiting to maintain equitable usage and service availability.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 to host each microservice. Use Amazon API Gateway to manage the RESTful API requests.
- B. Deploy each microservice as a set of AWS Lambda functions. Use Amazon API Gateway to manage the RESTful API requests.
- C. Host each microservice on Amazon EC2 instances in Auto Scaling groups behind an Elastic Load Balancing (ELB) load balancer. Use the ELB to manage the RESTful API requests.
- D. Deploy each microservice on Amazon Elastic Beanstalk. Use Amazon CloudFront to manage the RESTful API requests.

Answer: C

Explanation:

Question: 220

A company is launching a new gaming application. The company will use Amazon EC2 Auto Scaling groups to deploy the application. The application stores user data in a relational database.

The company has office locations around the world that need to run analytics on the user data in the database. The company needs a cost-effective database solution that provides cross-Region disaster recovery with low-latency read performance across AWS Regions.

Which solution will meet these requirements?

- A. Create an Amazon ElastiCache for Redis cluster in the Region where the application is deployed. Create read replicas in Regions where the company offices are located. Ensure the company offices read from the read replica instances.
- B. Create Amazon DynamoDB global tables. Deploy the tables to the Regions where the company offices are located and to the Region where the application is deployed. Ensure that each company office reads from the tables that are in the same Region as the office.
- C. Create an Amazon Aurora global database. Configure the primary cluster to be in the Region where the application is deployed. Configure the secondary Aurora replicas to be in the Regions where the company offices are located. Ensure the company offices read from the Aurora replicas.
- D. Create an Amazon RDS Multi-AZ DB cluster deployment in the Region where the application is deployed. Ensure the company offices read from read replica instances.

Answer: A

Explanation:

Question: 221

An international company needs to share data from an Amazon S3 bucket to employees who are located around the world. The company needs a secure solution to provide employees with access to the S3 bucket. The employees are already enrolled in AWS IAM Identity Center.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a help desk application to generate an Amazon S3 presigned URL for each employee. Configure the presigned URLs to have short expirations. Instruct employees to contact the company help desk to receive a presigned URL to access the S3 bucket.
- B. Create a group for Amazon S3 access in IAM Identity Center. Add the employees who require access to the S3 bucket to the group. Create an IAM policy to allow Amazon S3 access from the group. Instruct employees to use the AWS access portal to access the AWS Management Console and navigate to the S3 bucket.
- C. Create an Amazon S3 File Gateway. Create one share for data uploads and a second share for data downloads. Set up an SFTP service on an Amazon EC2 instance. Mount the shares to the EC2 instance. Instruct employees to use the SFTP server.
- D. Configure AWS Transfer Family SFTP endpoints. Select the custom identity provider option. Use AWS Secrets Manager to manage the user credentials. Instruct employees to use Transfer Family SFTP.

Answer: C

Explanation:

Question: 222

A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics application is highly resilient and is designed to run in stateless mode.

The company notices that the application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon Machine Image (AMI) of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use an Application Load Balancer to distribute the load across the two EC2 instances.
- B. Create an Amazon Machine Image (AMI) of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.

- C. Create an AWS Lambda function to stop the EC2 instance and change the instance type. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization is more than 75%.
- D. Create an Amazon Machine Image (AMI) of the web application. Apply the AMI to a launch template. Create an Auto Scaling group that includes the launch template. Configure the launch template to use a Spot Fleet. Attach an Application Load Balancer to the Auto Scaling group.

Answer: D

Explanation:

Auto Scaling is the most effective solution for ensuring seamless scalability of a stateless application. Key points:

D leverages Auto Scaling with a Spot Fleet for cost efficiency and attaches an ALB to distribute traffic.

A and B do not provide automated scaling and would require manual intervention to add more instances.

C changes the instance type but does not scale out horizontally, which is required here.

AWS Documentation Reference:

Amazon EC2 Auto Scaling

Question: 223

A company hosts its applications in multiple private and public subnets in a VPC. The applications in the private subnets need to access an API. The API is available on the internet and is hosted in the company's on-premises data center. A solutions architect needs to establish connectivity for applications in the private subnets.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a transit gateway to connect the VPC to the on-premises network. Use the transit gateway to route API calls from the private subnets to the on-premises data center.
- B. Create a NAT gateway in the public subnet of the VPC. Use the NAT gateway to allow the private subnets to access the API over the internet.
- C. Establish an AWS PrivateLink connection to connect the VPC to the on-premises network. Use PrivateLink to make API calls from the private subnets to the on-premises data center.
- D. Implement an AWS Site-to-Site VPN connection between the VPC and the on-premises data center. Use the VPN connection to make API calls from the private subnets to the on-premises data center.

Answer: D

Explanation:

AWS Site-to-Site VPN is a cost-effective way to securely connect your on-premises data center with AWS resources. In this scenario:

Applications in private subnets require access to the API hosted in the on-premises data center.

A Site-to-Site VPN connection is a secure and cost-efficient option to route traffic between the VPC and on-premises resources.

Transit Gateway and PrivateLink are not cost-effective for this use case.

NAT Gateway only provides internet access for private subnets, which is not suitable for reaching an on-premises resource.

AWS Documentation Reference:

AWS Site-to-Site VPN

Question: 224

A healthcare company is developing an AWS Lambda function that publishes notifications to an encrypted Amazon Simple Notification Service (Amazon SNS) topic. The notifications contain protected health information (PHI).

The SNS topic uses AWS Key Management Service (AWS KMS) customer-managed keys for encryption. The company must ensure that the application has the necessary permissions to publish messages securely to the SNS topic.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Create a resource policy for the SNS topic that allows the Lambda function to publish messages to the topic.
- B. Use server-side encryption with AWS KMS keys (SSE-KMS) for the SNS topic instead of customer managed keys.
- C. Create a resource policy for the encryption key that the SNS topic uses that has the necessary AWS KMS permissions.
- D. Specify the Lambda function's Amazon Resource Name (ARN) in the SNS topic's resource policy.
- E. Associate an Amazon API Gateway HTTP API with the SNS topic to control access to the topic by using API Gateway resource policies.
- F. Configure a Lambda execution role that has the necessary IAM permissions to use a customer managed key in AWS KMS.

Answer: A,C,F

Explanation:

To securely publish messages to an encrypted Amazon SNS topic, the following steps are required:

A . Resource policy for SNS topic:Ensures that the Lambda function is explicitly allowed to publish messages to the topic.

C . Resource policy for KMS key:Provides the necessary permissions to use the customer-managed key for encryption.

F . Lambda execution role:Grants the Lambda function the necessary IAM permissions to use the encryption key.

Other options:

Bis invalid because using SSE-KMS does not eliminate the need for resource policies.

D overlaps with A, but specifying the ARN in the topic policy is covered by creating the resource policy.

E is unrelated as API Gateway is not required for this setup.

AWS Documentation Reference:

Amazon SNS and KMS Permissions

Question: 225

A company stores sensitive customer data in an Amazon DynamoDB table. The company frequently updates the data. The company wants to use the data to personalize offers for customers.

The company's analytics team has its own AWS account. The analytics team runs an application on Amazon EC2 instances that needs to process data from the DynamoDB tables. The company needs to follow security best practices to create a process to regularly share data from DynamoDB to the analytics team.

Which solution will meet these requirements?

- A. Export the required data from the DynamoDB table to an Amazon S3 bucket as multiple JSON files. Provide the analytics team with the necessary IAM permissions to access the S3 bucket.
- B. Allow public access to the DynamoDB table. Create an IAM user that has permission to access DynamoDB. Share the IAM user with the analytics team.
- C. Allow public access to the DynamoDB table. Create an IAM user that has read-only permission for DynamoDB. Share the IAM user with the analytics team.
- D. Create a cross-account IAM role. Create an IAM policy that allows the AWS account ID of the analytics team to access the DynamoDB table. Attach the IAM policy to the IAM role. Establish a trust relationship between accounts.

Answer: D

Explanation:

Using cross-account IAM roles is the most secure and scalable way to share data between AWS accounts.

A trust relationship allows the analytics team's account to assume the role in the main account and access the DynamoDB table directly.

A is feasible but involves data duplication and additional costs for storing the JSON files in S3.

B and C violate security best practices by allowing public access to sensitive data and sharing credentials, which is highly discouraged.

AWS Documentation Reference:

Cross-Account Access with Roles

Best Practices for Amazon DynamoDB Security

Question: 226

A financial service company has a two-tier consumer banking application. The frontend serves static web content. The backend consists of APIs. The company needs to migrate the frontend component to AWS. The backend of the application will remain on premises. The company must protect the application from common web vulnerabilities and attacks.

Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate the frontend to Amazon EC2 instances. Deploy an Application Load Balancer (ALB) in front of the instances. Use the instances to invoke the on-premises APIs. Associate AWS WAF rules with the instances.

B. Deploy the frontend as an Amazon CloudFront distribution that has multiple origins. Configure one origin to be an Amazon S3 bucket that serves the static web content. Configure a second origin to route traffic to the on-premises APIs based on the URL pattern. Associate AWS WAF rules with the distribution.

C. Migrate the frontend to Amazon EC2 instances. Deploy a Network Load Balancer (NLB) in front of the instances. Use the instances to invoke the on-premises APIs. Create an AWS Network Firewall instance. Route all traffic through the Network Firewall instance.

D. Deploy the frontend as a static website based on an Amazon S3 bucket. Use an Amazon API Gateway REST API and a set of Amazon EC2 instances to invoke the on-premises APIs. Associate AWS WAF rules with the REST API and the S3 bucket.

Answer: B

Explanation:

Comprehensive Deploying the frontend as a CloudFront distribution with multiple origins provides an efficient and scalable solution. Using WAF rules with CloudFront protects against web vulnerabilities, while the multi-origin configuration allows traffic routing to the on-

premises backend APIs. This approach minimizes operational overhead compared to managing EC2 instances.

Reference:

Amazon CloudFront Features

AWS WAF Integration with CloudFront

Question: 227

A company is deploying a new application to a VPC on existing Amazon EC2 instances. The application has a presentation tier that uses an Auto Scaling group of EC2 instances. The application also has a database tier that uses an Amazon RDS Multi-AZ database.

The VPC has two public subnets that are split between two Availability Zones. A solutions architect adds one private subnet to each Availability Zone for the RDS database. The solutions architect wants to restrict network access to the RDS database to block access from EC2 instances that do not host the new application.

Which solution will meet this requirement?

- A. Modify the RDS database security group to allow traffic from a CIDR range that includes IP addresses of the EC2 instances that host the new application.
- B. Associate a new ACL with the private subnets. Deny all incoming traffic from IP addresses that belong to any EC2 instance that does not host the new application.
- C. Modify the RDS database security group to allow traffic from the security group that is associated with the EC2 instances that host the new application.
- D. Associate a new ACL with the private subnets. Deny all incoming traffic except for traffic from a CIDR range that includes IP addresses of the EC2 instances that host the new application.

Answer: C

Explanation:

Correct Approach:

AWS Security Groups:

Security groups operate at the instance level, making them the ideal tool for controlling access to specific resources such as an Amazon RDS database.

By default, security groups deny all incoming traffic. You can allow access by explicitly specifying another security group.

Associating an RDS database security group with the EC2 instances' security group ensures only the specified EC2 instances can

access the RDS database.

Incorrect Options Analysis:

Option A: Using CIDR blocks for IP-based access is less secure and more difficult to manage. Additionally, Auto Scaling groups dynamically allocate IP addresses, making this approach impractical.

Option B: Network ACLs (NACLs) operate at the subnet level and are stateless. While NACLs can deny or allow traffic, they are not suited to application-specific access control.

Option D: Similar to Option B, using a NACL with CIDR ranges for EC2 IPs is difficult to manage and not application-specific.

Reference:

Amazon RDS Security Groups

Security Group Best Practices

Differences Between Security Groups and NACLs

Question: 228

A developer is creating a serverless application that performs video encoding. The encoding process runs as background jobs and takes several minutes to encode each video. The process must not send an immediate result to users.

The developer is using Amazon API Gateway to manage an API for the application. The developer needs to run test invocations and request validations. The developer must distribute API keys to control access to the API.

Which solution will meet these requirements?

- A. Create an HTTP API. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the HTTP API. Use the Event invocation type to call the Lambda function.
- B. Create a REST API with the default endpoint type. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the REST API. Use the Event invocation type to call the Lambda function.
- C. Create an HTTP API. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the HTTP API. Use the RequestResponse invocation type to call the Lambda function.
- D. Create a REST API with the default endpoint type. Create an AWS Lambda function to handle the encoding jobs. Integrate the function with the REST API. Use the RequestResponse invocation type to call the Lambda function.

Answer: B

Explanation:

Background Jobs with Event Invocation Type:

The Event invocation type is asynchronous, meaning the Lambda function does not send an immediate result to the API Gateway and processes the request in the background. This is ideal for video encoding tasks that take time.

REST API vs. HTTP API:

REST APIs support advanced features like API keys, request validation, and throttling that HTTP APIs do not support fully.

Since the developer needs API keys and request validations, a REST API is the correct choice.

Integration with Lambda:

AWS Lambda integration is seamless with REST APIs, and using the Event invocation ensures asynchronous processing.

Incorrect Options Analysis:

Option A: HTTP API lacks full support for API keys and validation.

Option C and D: RequestResponse invocation type requires immediate responses, unsuitable for background jobs.

Reference:

AWS Lambda Invocation Types

Amazon API Gateway REST APIs

Question: 229

A developer needs to export the contents of several Amazon DynamoDB tables into Amazon S3 buckets to comply with company data regulations. The developer uses the AWS CLI to run commands to export from each table to the proper S3 bucket. The developer sets up AWS credentials correctly and grants resources appropriate permissions. However, the exports of some tables fail.

What should the developer do to resolve this issue?

- A. Ensure that point-in-time recovery is enabled on the DynamoDB tables.
- B. Ensure that the target S3 bucket is in the same AWS Region as the DynamoDB table.
- C. Ensure that DynamoDB streaming is enabled for the tables.
- D. Ensure that DynamoDB Accelerator (DAX) is enabled.

Answer: A

Explanation:

Export Requirements:

To export data from DynamoDB to Amazon S3, point-in-time recovery (PITR) must be enabled for the tables. This feature creates a snapshot of the data, which is essential for exports.

Incorrect Options Analysis:

Option B: S3 buckets and DynamoDB tables do not need to be in the same region for exports.

Option C: DynamoDB streams are unrelated to the export functionality.

Option D: DAX accelerates reads but has no role in exports.

Reference:

Exporting DynamoDB Data to Amazon S3

Question: 230

A developer is creating an ecommerce workflow in an AWS Step Functions state machine that includes an HTTP Task state. The task passes shipping information and order details to an endpoint.

The developer needs to test the workflow to confirm that the HTTP headers and body are correct and that the responses meet expectations.

Which solution will meet these requirements?

- A. Use the TestState API to invoke only the HTTP Task. Set the inspection level to TRACE.
- B. Use the TestState API to invoke the state machine. Set the inspection level to DEBUG.
- C. Use the data flow simulator to invoke only the HTTP Task. View the request and response data.
- D. Change the log level of the state machine to ALL. Run the state machine.

Answer: D

Explanation:

State Machine Testing with Logs:

Changing the log level to ALL enables capturing detailed request and response data. This helps verify HTTP headers, body, and responses.

Incorrect Options Analysis:

Option A and B: The TestState API is not a valid option for Step Functions.

Option C: A data flow simulator does not exist for AWS Step Functions.

Reference:

Step Functions Logging and Monitoring

Question: 231

A developer used the AWS SDK to create an application that aggregates and produces log records for 10 services. The application delivers data to an Amazon Kinesis Data Streams stream.

Each record contains a log message with a service name, creation timestamp, and other log information. The stream has 15 shards in provisioned capacity mode. The stream uses service name as the partition key.

The developer notices that when all the services are producing logs, `ProvisionedThroughputExceededException` errors occur during `PutRecord` requests. The stream metrics show that the write capacity the applications use is below the provisioned capacity.

How should the developer resolve this issue?

- A. Change the capacity mode from provisioned to on-demand.
- B. Double the number of shards until the throttling errors stop occurring.
- C. Change the partition key from service name to creation timestamp.
- D. Use a separate Kinesis stream for each service to generate the logs.

Answer: C

Explanation:

Partition Key Issue:

Using "service name" as the partition key results in uneven data distribution. Some shards may become hot due to excessive logs from certain services, leading to throttling errors.

Changing the partition key to "creation timestamp" ensures a more even distribution of records across shards.

Incorrect Options Analysis:

Option A: On-demand capacity mode eliminates throughput management but is more expensive and does not address the root cause.

Option B: Adding more shards does not solve the issue if the partition key still creates hot shards.

Option D: Using separate streams increases complexity and is unnecessary.

Reference:

Kinesis Data Streams Partition Key Best Practices

Question: 232

A company has a serverless web application that is comprised of AWS Lambda functions. The application experiences spikes in traffic that cause increased latency because of cold starts. The company wants to improve the application's ability to handle traffic spikes and to minimize latency. The solution must optimize costs during periods when traffic is low.

Which solution will meet these requirements?

- A. Configure provisioned concurrency for the Lambda functions. Use AWS Application Auto Scaling to adjust the provisioned concurrency.
- B. Launch Amazon EC2 instances in an Auto Scaling group. Add a scheduled scaling policy to launch additional EC2 instances during peak traffic periods.
- C. Configure provisioned concurrency for the Lambda functions. Set a fixed concurrency level to handle the maximum expected traffic.
- D. Create a recurring schedule in Amazon EventBridge Scheduler. Use the schedule to invoke the Lambda functions periodically to warm the functions.

Answer: A

Explanation:

Provisioned Concurrency:

AWS Lambda's provisioned concurrency ensures that a predefined number of execution environments are pre-warmed and ready to handle requests, reducing latency during traffic spikes.

This solution optimizes costs during low-traffic periods when combined with AWS Application Auto Scaling to dynamically adjust the provisioned concurrency based on demand.

Incorrect Options Analysis:

Option B: Switching to EC2 would increase complexity and cost for a serverless application.

Option C: A fixed concurrency level may result in over-provisioning during low-traffic periods, leading to higher costs.

Option D: Periodically warming functions does not effectively handle sudden spikes in traffic.

Reference:

AWS Lambda Provisioned Concurrency

Question: 233

A solutions architect is designing the storage architecture for a new web application used for storing and viewing engineering drawings. All application components will be deployed on the AWS infrastructure. The application design must support caching to minimize the amount of time that users wait for the engineering drawings to load. The application must be able to store petabytes of data.

Which combination of storage and caching should the solutions architect use?

- A. Amazon S3 with Amazon CloudFront
- B. Amazon S3 Glacier Deep Archive with Amazon ElastiCache
- C. Amazon Elastic Block Store (Amazon EBS) volumes with Amazon CloudFront
- D. AWS Storage Gateway with Amazon ElastiCache

Answer: A

Explanation:

Amazon S3 with Amazon CloudFront:

Amazon S3 provides highly scalable and durable storage for petabytes of data.

Amazon CloudFront, as a content delivery network (CDN), caches frequently accessed data at edge locations to reduce latency. This combination is ideal for storing and accessing engineering drawings.

Incorrect Options Analysis:

Option B: Amazon S3 Glacier Deep Archive is for long-term archival storage, not frequent access.

Option C: Amazon EBS is unsuitable for large-scale, multi-user data access and does not support caching directly.

Option D: AWS Storage Gateway is for hybrid cloud storage, which is unnecessary for a fully cloudbased architecture.

Reference:

Amazon S3

Amazon CloudFront

Question: 234

A media company has an ecommerce website to sell music. Each music file is stored as an MP3 file. Premium users of the website purchase music files and download the files. The company wants to store music files on AWS. The company wants to provide access only to the premium users. The company wants to use the same URL for all premium users.

Which solution will meet these requirements?

- A. Store the MP3 files on a set of Amazon EC2 instances that have Amazon Elastic Block Store (Amazon EBS) volumes attached. Manage

access to the files by creating an IAM user and an IAM policy for each premium user.

B. Store all the MP3 files in an Amazon S3 bucket. Create a presigned URL for each MP3 file. Share the presigned URLs with the premium users.

C. Store all the MP3 files in an Amazon S3 bucket. Create an Amazon CloudFront distribution that uses the S3 bucket as the origin. Generate CloudFront signed cookies for the music files. Share the signed cookies with the premium users.

D. Store all the MP3 files in an Amazon S3 bucket. Create an Amazon CloudFront distribution that uses the S3 bucket as the origin. Use a CloudFront signed URL for each music file. Share the signed URLs with the premium users.

Answer: C

Explanation:

CloudFront Signed Cookies:

CloudFront signed cookies allow the company to provide access to premium users while maintaining a single, consistent URL.

This approach is simpler and more scalable than managing presigned URLs for each file.

Incorrect Options Analysis:

Option A: Using EC2 and EBS increases complexity and cost.

Option B: Managing presigned URLs for each file is not scalable.

Option D: CloudFront signed URLs require unique URLs for each file, which does not meet the requirement for a single URL.

Reference:

Serving Private Content with CloudFront

Question: 235

A company is building a serverless application to process orders from an e-commerce site. The application needs to handle bursts of traffic during peak usage hours and to maintain high availability. The orders must be processed asynchronously in the order the application receives them.

A. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use an AWS Lambda function to process the orders.

B. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to receive orders. Use an AWS Lambda function to process the orders.

C. Use an Amazon Simple Queue Service (Amazon SQS) standard queue to receive orders. Use AWS Batch jobs to process the orders.

D. Use an Amazon Simple Notification Service (Amazon SNS) topic to receive orders. Use AWS Batch jobs to process the orders.

Answer: B

Explanation:

Key Requirements:

Serverless architecture.

Handle traffic bursts with high availability.

Process orders asynchronously in the order they are received.

Analysis of Options:

Option A: Amazon SNS delivers messages to subscribers. However, SNS does not ensure ordering, making it unsuitable for FIFO (First In, First Out) requirements.

Option B: Amazon SQS FIFO queues support ordering and ensure messages are delivered exactly once. AWS Lambda functions can be triggered by SQS to process messages asynchronously and efficiently. This satisfies all requirements.

Option C: Amazon SQS standard queues do not guarantee message order and have "at-least-once" delivery, making them unsuitable for the FIFO requirement.

Option D: Similar to Option A, SNS does not ensure message ordering, and using AWS Batch adds complexity without directly addressing the requirements.

AWS Reference:

Amazon SQS FIFO Queues

AWS Lambda and SQS Integration

Question: 236

An e-commerce company has an application that uses Amazon DynamoDB tables configured with provisioned capacity. Order data is stored in a table named Orders. The Orders table has a primary key of order-ID and a sort key of product-ID. The company configured an AWS Lambda function to receive DynamoDB streams from the Orders table and update a table named Inventory. The company has noticed that during peak sales periods, updates to the Inventory table take longer than the company can tolerate. Which solutions will resolve the slow table updates? (Select TWO.)

- A. Add a global secondary index to the Orders table. Include the product-ID attribute.
- B. Set the batch size attribute of the DynamoDB streams to be based on the size of items in the Orders table.
- C. Increase the DynamoDB table provisioned capacity by 1,000 write capacity units (WCUs).

- D. Increase the DynamoDB table provisioned capacity by 1,000 read capacity units (RCUs).
- E. Increase the timeout of the Lambda function to 15 minutes.

Answer: B,C

Explanation:

Key Problem:

Delayed Inventory table updates during peak sales.

DynamoDB Streams and Lambda processing require optimization.

Analysis of Options:

Option A: Adding a GSI is unrelated to the issue. It does not address stream processing delays or capacity issues.

Option B: Optimizing batch size reduces latency and allows the Lambda function to process larger chunks of data at once, improving performance during peak load.

Option C: Increasing write capacity for the Inventory table ensures that it can handle the increased volume of updates during peak times.

Option D: Increasing read capacity for the Orders table does not directly resolve the issue since the problem is with updates to the Inventory table.

Option E: Increasing Lambda timeout only addresses longer processing times but does not solve the underlying throughput problem.

AWS Reference:

DynamoDB Streams Best Practices

Provisioned Throughput in DynamoDB

Question: 237

A company has an e-commerce site. The site is designed as a distributed web application hosted in multiple AWS accounts under one AWS Organizations organization. The web application is comprised of multiple microservices. All microservices expose their AWS services either through Amazon CloudFront distributions or public Application Load Balancers (ALBs). The company wants to protect public endpoints from malicious attacks and monitor security configurations. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS WAF to protect the public endpoints. Use AWS Firewall Manager from a dedicated security account to manage rules in AWS WAF. Use AWS Config rules to monitor the Regional and global WAF configurations.

- B. Use AWS WAF to protect the public endpoints. Apply AWS WAF rules in each account. Use AWS Config rules and AWS Security Hub to monitor the WAF configurations of the ALBs and the CloudFront distributions.
- C. Use AWS WAF to protect the public endpoints. Use AWS Firewall Manager from a dedicated security account to manage the rules in AWS WAF. Use Amazon Inspector and AWS Security Hub to monitor the WAF configurations of the ALBs and the CloudFront distributions.
- D. Use AWS Shield Advanced to protect the public endpoints. Use AWS Config rules to monitor the Shield Advanced configuration for each account.

Answer: A

Explanation:

Key Requirements:

Protect public endpoints (CloudFront distributions and ALBs) from malicious attacks.

Centralized management across multiple accounts in an organization.

Ability to monitor security configuration effectively.

Minimize operational overhead.

Analysis of Options

Option A:

AWS WAF: Protects web applications by filtering and blocking malicious requests. Rules can be applied to both ALBs and CloudFront distributions.

AWS Firewall Manager: Enables centralized management of WAF rules across multiple accounts in an AWS Organizations organization. It simplifies rule deployment, avoiding the need to configure rules individually in each account.

AWS Config: Monitors compliance by using rules that check Regional and global WAF configurations.

Ensures that security configurations align with organizational policies.

Operational Overhead: Centralized management and automated monitoring reduce the operational burden.

Correct Approach: Meets all requirements with the least overhead.

Option B:

This approach involves applying WAF rules in each account manually.

While AWS Config and AWS Security Hub provide monitoring capabilities, managing individual WAF configurations in multiple accounts

introduces significant operational overhead.

Incorrect Approach: Higher overhead compared to centralized management with AWS Firewall Manager.

Option C:

Similar to Option A but includes Amazon Inspector, which is not designed for monitoring WAF configurations.

AWS Security Hub is appropriate for monitoring but is redundant when Firewall Manager and Config are already in use.

Incorrect Approach: Adds unnecessary complexity and does not focus on monitoring WAF specifically.

Option D:

AWS Shield Advanced: Focuses on mitigating large-scale DDoS attacks but does not provide the fine-grained web application protection offered by WAF.

AWS Config: Can monitor Shield Advanced configurations but does not fulfill the WAF monitoring requirements.

Incorrect Approach: Does not address the need for WAF or centralized rule management.

Why Option A is Correct

Protection:

AWS WAF provides fine-grained filtering and protection against SQL injection, cross-site scripting, and other web vulnerabilities.

Rules can be applied at both ALBs and CloudFront distributions, covering all public endpoints.

Centralized Management:

AWS Firewall Manager enables security teams to centrally define and manage WAF rules across all accounts in the organization.

Monitoring:

AWS Config ensures compliance with WAF configurations by checking rules and generating alerts for misconfigurations.

Operational Overhead:

Centralized management via Firewall Manager and automated compliance monitoring via AWS Config greatly reduce manual effort.

AWS Solution Architect Reference

AWS WAF Documentation

AWS Firewall Manager Documentation

Question: 238

A company has a serverless web application that is comprised of AWS Lambda functions. The application experiences spikes in traffic that cause increased latency because of cold starts. The company wants to improve the application's ability to handle traffic spikes and to minimize latency. The solution must optimize costs during periods when traffic is low.

- A. Configure provisioned concurrency for the Lambda functions. Use AWS Application Auto Scaling to adjust the provisioned concurrency.
- B. Launch Amazon EC2 instances in an Auto Scaling group. Add a scheduled scaling policy to launch additional EC2 instances during peak traffic periods.
- C. Configure provisioned concurrency for the Lambda functions. Set a fixed concurrency level to handle the maximum expected traffic.
- D. Create a recurring schedule in Amazon EventBridge Scheduler. Use the schedule to invoke the Lambda functions periodically to warm the functions.

Answer: A

Explanation:

Key Requirements:

Handle traffic spikes efficiently and reduce latency caused by cold starts.

Optimize costs during low traffic periods.

Analysis of Options:

Option A:

Provisioned Concurrency: Reduces cold start latency by pre-warming Lambda environments for the required number of concurrent executions.

AWS Application Auto Scaling: Automatically adjusts provisioned concurrency based on demand, ensuring cost optimization by scaling down during low traffic.

Correct Approach:Provides a balance between performance during traffic spikes and cost optimization during idle periods.

Option B:

Using EC2 instances with Auto Scaling introduces unnecessary complexity for a serverless architecture. It requires additional management and does not address the issue of cold starts for Lambda.

Incorrect Approach:Contradicts the serverless design philosophy and increases operational overhead.

Option C:

Setting a fixed concurrency level ensures performance during spikes but does not optimize costs during low traffic. This approach would maintain provisioned instances unnecessarily.

Incorrect Approach:Lacks cost optimization.

Option D:

Using EventBridge Scheduler for periodic invocations may reduce cold starts but does not dynamically scale based on traffic demand. It also leads to unnecessary invocations during idle times.

Incorrect Approach:Suboptimal for high traffic fluctuations and cost control.

AWS Solution Architect Reference:

AWS Lambda Provisioned Concurrency

AWS Application Auto Scaling with Lambda

Question: 239

A financial services company has a two-tier consumer banking application. The frontend serves static web content. The backend consists of APIs. The company needs to migrate the frontend component to AWS. The backend of the application will remain on-premises. The company must protect the application from common web vulnerabilities and attacks.

A. Migrate the frontend to Amazon EC2 instances. Deploy an Application Load Balancer (ALB) in front of the instances. Use the instances to invoke the on-premises APIs. Associate AWS WAF rules with the instances.

B. Deploy the frontend as an Amazon CloudFront distribution that has multiple origins. Configure one origin to be an Amazon S3 bucket that serves the static web content. Configure a second origin to route traffic to the on-premises APIs based on the URL pattern. Associate AWS WAF rules with the distribution.

C. Migrate the frontend to Amazon EC2 instances. Deploy a Network Load Balancer (NLB) in front of the instances. Use the instances to invoke the on-premises APIs. Create an AWS Network Firewall instance. Route all traffic through the Network Firewall instance.

D. Deploy the frontend as a static website based on an Amazon S3 bucket. Use an Amazon API Gateway REST API and a set of Amazon EC2 instances to invoke the on-premises APIs. Associate AWS WAF rules with the REST API and the S3 bucket.

Answer: B

Explanation:

Key Requirements:

Host the frontend on AWS as a static website.

Protect the application from common web vulnerabilities.

Minimal operational overhead.

Analysis of Options:

Option A:

Hosting the frontend on EC2 with an ALB introduces unnecessary complexity for serving static content.

AWS WAF rules can protect the ALB, but managing EC2 instances adds operational overhead.

Incorrect Approach: High operational complexity for a simple static website.

Option B:

Amazon CloudFront: Acts as a global CDN, reducing latency and protecting against DDoS attacks.

Multiple Origins: Allows static content to be served from S3 while routing API traffic to the on-premises backend.

AWS WAF: Integrates with CloudFront to provide web application protection.

Correct Approach: Offers low operational overhead with optimal security and performance.

Option C:

Using NLB and Network Firewall is unnecessary for a static website. This approach increases cost and complexity without addressing the frontend requirements effectively.

Incorrect Approach: Over-engineered solution.

Option D:

Hosting the frontend on S3 and using API Gateway is a viable option, but managing AWS WAF rules separately for both the S3 bucket and the REST API increases complexity.

Incorrect Approach: Less efficient than using CloudFront with multiple origins.

AWS Solution Architect Reference:

Question: 240

A company is deploying a critical application by using Amazon RDS for MySQL. The application must be highly available and must recover automatically. The company needs to support interactive users (transactional queries) and batch reporting (analytical queries) with no more than a 4-hour lag. The analytical queries must not affect the performance of the transactional queries.

- A. Configure Amazon RDS for MySQL in a Multi-AZ DB instance deployment with one standby instance. Point the transactional queries to the primary DB instance. Point the analytical queries to a secondary DB instance that runs in a different Availability Zone.
- B. Configure Amazon RDS for MySQL in a Multi-AZ DB cluster deployment with two standby instances. Point the transactional queries to the primary DB instance. Point the analytical queries to the reader endpoint.
- C. Configure Amazon RDS for MySQL to use multiple read replicas across multiple Availability Zones. Point the transactional queries to the primary DB instance. Point the analytical queries to one of the replicas in a different Availability Zone.
- D. Configure Amazon RDS for MySQL as the primary database for the transactional queries with automated backups enabled. Configure automated backups. Each night, create a read-only database from the most recent snapshot to support the analytical queries. Terminate the previously created database.

Answer: C

Explanation:

Key Requirements:

High availability and automatic recovery.

Separate transactional and analytical queries with minimal performance impact.

Allow up to a 4-hour lag for analytical queries.

Analysis of Options:

Option A:

Multi-AZ deployments provide high availability but do not include read replicas for separating transactional and analytical queries.

Analytical queries on the secondary DB instance would impact the transactional workload.

Incorrect Approach: Does not meet the requirement of query separation.

Option B:

Multi-AZ DB clusters provide high availability and include a reader endpoint. However, these are better suited for Aurora and not RDS for MySQL.

Incorrect Approach: Not applicable to standard RDS for MySQL.

Option C:

Multiple read replicas allow separation of transactional and analytical workloads.

Queries can be pointed to a replica in a different AZ, ensuring no impact on transactional queries.

Correct Approach: Meets all requirements with high availability and query separation.

Option D:

Creating nightly snapshots and read-only databases adds significant operational overhead and does not support the 4-hour lag requirement.

Incorrect Approach: Not practical for dynamic query separation.

AWS Solution Architect Reference:

Amazon RDS Read Replicas

Multi-AZ Deployments

Question: 241

A company plans to use an Amazon S3 bucket to archive backup data. Regulations require the company to retain the backup data for 7 years.

During the retention period, the company must prevent users, including administrators, from deleting the data. The company can delete the data after 7 years.

Which solution will meet these requirements?

- A. Create an S3 bucket policy that denies delete operations for 7 years. Create an S3 Lifecycle policy to delete the data after 7 years.
- B. Create an S3 Object Lock default retention policy that retains data for 7 years in governance mode. Create an S3 Lifecycle policy to delete the data after 7 years.
- C. Create an S3 Object Lock default retention policy that retains data for 7 years in compliance mode. Create an S3 Lifecycle policy to delete the data after 7 years.

D. Create an S3 Batch Operations job to set a legal hold on each object for 7 years. Create an S3

Lifecycle policy to delete the data after 7 years.

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step

The requirement is to prevent data deletion by any user, including administrators, for 7 years while allowing automatic deletion afterward.

S3 Object Lock in Compliance Mode (Correct Choice - C)

Compliance mode ensures that even the root user cannot delete or modify the objects during the retention period.

After 7 years, the S3 Lifecycle policy automatically deletes the objects.

This meets both immutability and automatic deletion requirements.

Governance Mode (Option B - Incorrect)

Governance mode prevents deletion, but administrators can override it.

The requirement explicitly states that even administrators must not be able to delete the data.

S3 Bucket Policy (Option A - Incorrect)

An S3 bucket policy can deny deletes, but policies can be modified at any time by administrators.

It does not enforce strict retention like Object Lock.

S3 Batch Operations Job (Option D - Incorrect)

A legal hold does not have an automatic expiration.

Legal holds must be manually removed, which is not efficient.

Why Option C is Correct:

S3 Object Lock in Compliance Mode prevents deletion by all users, including administrators.

The S3 Lifecycle policy deletes the data automatically after 7 years, reducing operational overhead.

Reference:

S3 Object Lock Compliance Mode

Question: 242

A company runs multiple applications on Amazon EC2 instances in a VPC.

Application A runs in a private subnet that has a custom route table and network ACL.

Application B runs in a second private subnet in the same VPC.

The company needs to prevent Application A from sending traffic to Application B.

Which solution will meet this requirement?

- A. Add a deny outbound rule to a security group associated with Application B. Configure the rule to prevent Application B from sending traffic to Application A.
- B. Add a deny outbound rule to a security group associated with Application A. Configure the rule to prevent Application A from sending traffic to Application B.
- C. Add a deny outbound rule to the custom network ACL for the Application B subnet. Configure the rule to prevent Application B from sending traffic to the IP addresses associated with Application A.
- D. Add a deny outbound rule to the custom network ACL for the Application A subnet. Configure the rule to prevent Application A from sending traffic to the IP addresses associated with Application B.

Answer: D

Explanation:

Comprehensive and Detailed Step-by-Step

The requirement is to prevent Application A from sending traffic to Application B.

Understanding AWS Network Security Components:

Security Groups

Stateful (if traffic is allowed in one direction, it is automatically allowed in the reverse).

Do not support explicit deny rules, only allow rules.

Not suitable for blocking traffic in this scenario.

Network ACLs (NACLs)

Stateless(must define explicit rules for both inbound and outbound traffic).

Support explicit DENY rules.

Best suited for blocking traffic between subnets.

Analysis of the Options:

Option A: Deny Outbound Rule in Security Group for Application B (Incorrect)

Security Groups do not support explicit deny rules.

Does not block traffic from Application A to Application B.

Option B: Deny Outbound Rule in Security Group for Application A (Incorrect)

Security Groups do not support explicit deny rules.

Cannot effectively prevent Application A from sending traffic to Application B.

Option C: Deny Outbound Rule in NACL for Application B Subnet (Incorrect)

This would prevent Application B from sending traffic, but the requirement is to block traffic from Application A to Application B.

Incorrect subnet is being modified.

Option D: Deny Outbound Rule in NACL for Application A Subnet (Correct Choice)

Prevents Application A from sending traffic to Application B by blocking outbound requests at the network level.

Effectively stops communication from A to B at the subnet level.

Why Option D is the Best Choice?

NACLs support explicit deny rules, unlike security groups. Blocks outbound traffic from Application A before it reaches Application B.

Works at the subnet level, making it scalable.

Reference:

AWS Network ACLs

Question: 243

A company generates approximately 20 GB of data multiple times each day. The company uses AWS DataSync to copy all data from on-premises storage to Amazon S3 every 6 hours for further processing. The analytics team wants to modify the copy process to copy only data relevant to the analytics team and ignore the rest of the data. The team wants to copy data as soon as possible and receive a notification when the copy process is finished. Which combination of steps will meet these requirements MOST cost-effectively?

(Select THREE.)

- A. Modify the data generation process on-premises to create a manifest file at the end of the copy process with the names of the objects to be copied to Amazon S3. Create a custom script to upload the manifest file to an S3 bucket.
- B. Modify the data generation process on-premises to create a manifest file at the end of the copy process with the names of the objects to be copied to Amazon S3. Create an AWS Lambda function to load the manifest file data into an Amazon DynamoDB table.
- C. Create an AWS Lambda function that Amazon EventBridge invokes when the manifest file is loaded into Amazon DynamoDB. Configure the Lambda function to copy the data from on-premises storage to the S3 bucket that uses the manifest file.
- D. Create an AWS Lambda function that an S3 Event Notification invokes when the manifest file is uploaded. Configure the Lambda function to invoke the DataSync task by calling the StartTaskExecution API action with a manifest.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an Amazon EventBridge rule to send an email notification to the SNS topic when the DataSync task execution status changes to SUCCESS or to ERROR.
- F. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function to send an email notification to the SNS topic when the DataSync task execution status changes to SUCCESS or to ERROR.

Answer: A,D,E

Explanation:

Comprehensive and Detailed Step-by-Step

To meet the requirements of copying only relevant data as soon as possible and receiving notifications upon completion, the following steps are recommended:

Generate a Manifest File (Option A):

Action: Modify the on-premises data generation process to create a manifest file at the end of each data generation cycle. This manifest should list the names of the objects that need to be copied to Amazon S3.

Implementation: Develop a custom script that runs after data generation. This script compiles the list of relevant data files into a manifest file and uploads it to a designated S3 bucket.

Justification: Using a manifest allows AWS DataSync to transfer only the specified files, reducing unnecessary data transfer and associated costs.

docs.aws.amazon.com

Automate DataSync Task Execution (Option D):

Action: Set up an S3 Event Notification to trigger an AWS Lambda function whenever a new manifest file is uploaded to the S3 bucket.

Implementation: Configure the Lambda function to invoke the DataSync task by calling the StartTaskExecution API action, specifying the manifest file. This ensures that only the files listed in the manifest are copied from on-premises storage to Amazon S3.

Justification: This automation ensures timely data transfer as soon as relevant data is available, minimizing delays and manual intervention.

docs.aws.amazon.com

Set Up Completion Notifications (Option E):

Action: Create an Amazon SNS topic to handle notifications. Then, establish an Amazon EventBridge rule that monitors the DataSync task execution status and sends an email notification to the SNS topic when the status changes to SUCCESS or ERROR.

Implementation: Configure EventBridge to capture state changes of the DataSync task. When a task completes successfully or encounters an error, EventBridge triggers a notification to the SNS topic, which then sends an email to the subscribed recipients.

Justification: This setup provides immediate feedback on the data transfer process, allowing the analytics team to act promptly based on the success or failure of the data copy operation.

docs.aws.amazon.com

Reference:

AWS DataSync User Guide: Transferring specific files or objects by using a manifest

AWS DataSync API Reference: StartTaskExecution

Amazon EventBridge User Guide: Creating an EventBridge rule that triggers on an AWS API call

Amazon SNS User Guide: Sending Amazon SNS messages to HTTP/HTTPS endpoints

Question: 244

A company uses an AWS Transfer for SFTP public server endpoint and Amazon S3 storage to host large datasets for its customers. The company provides customers SSH private keys to authenticate and download their datasets. The Transfer for SFTP server is configured with structured logging that is saved to an S3 bucket. The company wants to charge customers based on their monthly data download usage.

Which solution will meet these requirements?

- A. Configure VPC Flow Logs to write to a new S3 bucket. Run monthly queries on the flow logs to identify customer usage and calculate cost. Add the charges to the customers' monthly bills.
- B. Each month, use AWS Cost Explorer to examine the costs for Transfer for SFTP and obtain a breakdown by customer. Add the charges to the customers' monthly bills.
- C. Enable requester pays on the S3 bucket that hosts the software. Allocate the charges to each customer based on the customer's requests.
- D. Run Amazon Athena queries on the logging S3 bucket monthly to identify customer usage and calculate costs. Add the charges to the customers' monthly bills.

Answer: D

Explanation:

Comprehensive and Detailed Step-by-Step

To accurately charge customers based on their monthly data download usage, the following solution is recommended:

Structured Logging Configuration:

Action: Ensure that the AWS Transfer for SFTP server is configured to log user activity, including details about file downloads, to Amazon S3 in a structured format.

Implementation: Utilize AWS Transfer Family's structured logging feature to capture detailed information about user sessions, including actions performed and data transferred.

docs.aws.amazon.com

Justification: Structured logs provide comprehensive data necessary for analyzing customer-specific download activities.

Data Analysis with Amazon Athena:

Action: Use Amazon Athena to run SQL queries on the structured log data stored in the S3 bucket to calculate the amount of data each customer has downloaded.

Implementation:

a. **Define a Schema:** Create a table in Athena that maps to the structure of your log files. This involves specifying the format of the logs and the location in S3.

b. **Query Data:** Write SQL queries to sum the total bytes downloaded by each customer over the billing period. This can be achieved by filtering logs based on user identifiers and summing the data transfer amounts.

Justification: Athena allows for efficient querying

Question: 245

A company collects data for temperature, humidity, and atmospheric pressure in cities across multiple continents. The average volume of data that the company collects from each site daily is 500 GB. Each site has a high-speed internet connection.

The company wants to aggregate the data from all these global sites as quickly as possible in a single Amazon S3 bucket. The solution must minimize operational complexity.

Which solution meets these requirements?

A. Turn on S3 Transfer Acceleration on the destination S3 bucket. Use multipart uploads to directly upload site data to the destination S3 bucket.

- B. Upload the data from each site to an S3 bucket in the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket. Then remove the data from the origin S3 bucket.
- C. Schedule AWS Snowball Edge Storage Optimized device jobs daily to transfer data from each site to the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket.
- D. Upload the data from each site to an Amazon EC2 instance in the closest Region. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. At regular intervals, take an EBS snapshot and copy it to the Region that contains the destination S3 bucket. Restore the EBS volume in that Region.

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step The goal is to transfer 500 GB daily from multiple global locations quickly into a single S3 bucket while keeping operational complexity low.

Option A:

Turn on S3 Transfer Acceleration on the destination S3 bucket. Use multipart uploads to directly upload site data to the destination S3 bucket.

S3 Transfer Acceleration (S3-TA) allows faster global uploads by routing traffic through Amazon CloudFront's globally distributed edge locations.

Multipart uploads improve efficiency by breaking large files into smaller parts, transferring them in parallel.

Low operational complexity: No need for additional resources or manual replication.

Why is this best? It ensures high-speed transfers while minimizing complexity.

Reference: Amazon S3 Transfer Acceleration

Option B:

Upload the data from each site to an S3 bucket in the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket. Then remove the data from the origin S3 bucket.

While S3 Cross-Region Replication (CRR) can copy objects, it adds latency due to sequential replication rather than a direct fast transfer.

Why not? S3 Transfer Acceleration is faster and avoids extra steps.

Reference:Cross-Region Replication

Option C:

Use AWS Snowball Edge for daily transfers.

AWS Snowball is for bulk offline transfers, not daily high-speed internet transfers.

Why not? Unnecessary physical devices add operational overhead.

Reference:AWS Snowball Edge

Option D:

Upload to EC2, store in EBS, snapshot, and restore in the destination Region.

This approach is overly complex and not optimized for direct S3 ingestion. Why not? Too many steps and higher costs.

Reference:Amazon EBS Snapshots

Question: 246

An e-commerce company stores inventory, order, and user information in multiple Amazon Redshift clusters. The Redshift clusters must comply with the company's security policies. The company must receive notifications about any security configuration violations.

Which solution will meet these requirements?

- A. Create an Amazon EventBridge rule that uses the Redshift clusters as the source. Create an AWS Lambda function to evaluate the Redshift cluster security configuration. Configure the Lambda function to notify the company of any violations of the security policies. Add the Lambda function as a target of the EventBridge rule.
- B. Create an AWS Lambda function to check the validity of the Redshift cluster security configurations. Create an Amazon EventBridge rule that invokes the Lambda function when Redshift clusters are created. Notify the company of any violations of security policies.
- C. Set up Amazon Redshift Advisor in the company's AWS account to monitor cluster configurations. Configure Redshift Advisor to generate notifications for security items that the company must address.
- D. Create an AWS Lambda function to check the Redshift clusters for any violation of the security configurations. Create an AWS Config custom rule to invoke the Lambda function when Redshift cluster security configurations are modified. Provide the compliance state of each Redshift cluster to AWS Config. Configure AWS Config to notify the company of any violations of the security policies.

Answer: D

Explanation:

Comprehensive and Detailed Step-by-Step

The company needs automatic monitoring and notifications for security violations in Amazon Redshift clusters.

Option A:

Create an Amazon EventBridge rule that uses the Redshift clusters as the source. Create an AWS Lambda function to evaluate the Redshift cluster security configuration. Configure the Lambda function to notify the company of any violations of the security policies. Add the Lambda function as a target of the EventBridge rule.

EventBridge can trigger actions based on events.

However, it does not track changes in Redshift security configurations.

Why not? AWS Config is better suited for compliance monitoring.

Reference: Amazon EventBridge

Option B:

Create an AWS Lambda function to check the validity of the Redshift cluster security configurations. Create an Amazon EventBridge rule that invokes the Lambda function when Redshift clusters are created. Notify the company of any violations of security policies.

This solution only checks security at cluster creation.

It does not detect changes that happen after the cluster is created.

Why not? Security policies need continuous monitoring.

Reference: EventBridge Rules

Option C:

Set up Amazon Redshift Advisor in the company's AWS account to monitor cluster configurations. Configure Redshift Advisor to generate notifications for security items that the company must address.

Redshift Advisor provides performance tuning recommendations but does not monitor security settings.

Why not? It does not provide compliance checks.

Reference: Amazon Redshift Advisor

Option D:

Create an AWS Lambda function to check the Redshift clusters for any violation of the security configurations. Create an AWS Config custom rule to invoke the Lambda function when Redshift cluster security configurations are modified. Provide the compliance state of each Redshift cluster to AWS Config. Configure AWS Config to notify the company of any violations of the security policies.

AWS Config tracks configuration changes in Redshift clusters.

Custom AWS Config rules allow security compliance enforcement.

Lambda can perform custom security checks and notify the company.

Why is this best? AWS Config is designed for compliance monitoring.

Reference: AWS Config for Security Compliance

Question: 247

A company uses AWS Cost Explorer to monitor its AWS costs. The company notices that Amazon Elastic Block Store (Amazon EBS) storage and snapshot costs increase every month. However, the company does not purchase additional EBS storage every month. The company wants to optimize monthly costs for its current storage usage.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use logs in Amazon CloudWatch Logs to monitor the storage utilization of Amazon EBS. Use Amazon EBS Elastic Volumes to reduce the size of the EBS volumes.
- B. Use a custom script to monitor space usage. Use Amazon EBS Elastic Volumes to reduce the size of the EBS volumes.
- C. Delete all expired and unused snapshots to reduce snapshot costs.
- D. Delete all nonessential snapshots. Use Amazon Data Lifecycle Manager to create and manage the snapshots according to the company's snapshot policy requirements.

Answer: D

Explanation:

Amazon Data Lifecycle Manager (DLM) automates the creation, retention, and deletion of EBS snapshots. This allows organizations to define policies that ensure snapshots are only kept as long as needed, reducing costs automatically and minimizing manual effort. AWS recommends using DLM for optimizing storage and managing backup lifecycle with minimal overhead.

Reference: AWS Documentation - Amazon Data Lifecycle Manager

Question: 248

A company uses Amazon Elastic Container Service (Amazon ECS) to run workloads that belong to service teams. Each service team uses an owner tag to specify the ECS containers that the team owns. The company wants to generate an AWS Cost Explorer report that shows how much each service team spends on ECS containers on a monthly basis.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select TWO.)

- A. Create a custom report in Cost Explorer. Apply a filter for Amazon ECS.
- B. Create a custom report in Cost Explorer. Apply a filter for the owner resource tag.
- C. Set up AWS Compute Optimizer. Review the rightsizing recommendations.
- D. Activate the owner tag as a cost allocation tag. Group the Cost Explorer report by linked account.
- E. Activate the owner tag as a cost allocation tag. Group the Cost Explorer report by the owner cost allocation tag.

Answer: A,E

Explanation:

To allocate costs based on team ownership, AWS recommends tagging resources using cost allocation tags. Activating the “owner” tag as a cost allocation tag allows AWS Cost Explorer to categorize and group spending. Additionally, filtering for Amazon ECS services enables the visibility into servicespecific usage. This approach is both scalable and operationally efficient.

Reference: AWS Documentation - AWS Billing and Cost Management, Cost Allocation Tags

Question: 249

A company is planning to run an AI/ML workload on AWS. The company needs to train a model on a dataset that is in Amazon S3 Standard. A model training application requires multiple compute nodes and single-digit millisecond access to the data.

Which solution will meet these requirements in the MOST cost-effective way?

- A. Move the data to S3 Intelligent-Tiering. Point the model training application to S3 Intelligent-Tiering as the data source.
- B. Add partitions to the S3 bucket by adding random prefixes. Reconfigure the model training application to point to the new prefixes as the data source.
- C. Move the data to S3 Express One Zone. Point the model training application to S3 Express One Zone as the data source.
- D. Move the data to a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instance. Point the model training application to the gp3 volume as the data source.

Answer: C

Explanation:

Amazon S3 Express One Zone provides single-digit millisecond latency and high throughput, making it ideal for ML workloads that require multiple compute nodes and fast access. It is also more cost-effective than traditional file or block storage for temporary, high-speed needs.

Reference: AWS Documentation - Amazon S3 Storage Classes

Question: 250

A company is building a mobile gaming app. The company wants to serve users from around the world with low latency. The company needs a scalable solution to host the application and to route user requests to the location that is nearest to each user.

Which solution will meet these requirements?

- A. Use an Application Load Balancer to route requests to Amazon EC2 instances that are deployed across multiple Availability Zones.
- B. Use a Regional Amazon API Gateway REST API to route requests to AWS Lambda functions.
- C. Use an edge-optimized Amazon API Gateway REST API to route requests to AWS Lambda functions.
- D. Use an Application Load Balancer to route requests to containers in an Amazon ECS cluster.

Answer: C

Explanation:

Edge-optimized API Gateway endpoints utilize the Amazon CloudFront global network to decrease latency for clients globally. This setup ensures that the request is routed to the closest edge location, significantly reducing response time and improving performance for worldwide users.

Reference: AWS Documentation - Amazon API Gateway Endpoint Types

Question: 251

A solutions architect is designing a web application that will run on Amazon EC2 instances behind an Application Load Balancer (ALB). The company strictly requires that the application be resilient against malicious internet activity and attacks, and protect against new common

vulnerabilities and exposures.

What should the solutions architect recommend?

- A. Leverage Amazon CloudFront with the ALB endpoint as the origin.
- B. Deploy an appropriate managed rule for AWS WAF and associate it with the ALB.
- C. Subscribe to AWS Shield Advanced and ensure common vulnerabilities and exposures are blocked.
- D. Configure network ACLs and security groups to allow only ports 80 and 443 to access the EC2 instances.

Answer: B

Explanation:

AWS WAF allows web applications to protect themselves from common web exploits and vulnerabilities. Using AWS managed rule groups ensures protection against known attack patterns, such as SQL injection and cross-site scripting. Associating AWS WAF with the ALB provides application-layer security and real-time threat mitigation.

Reference: AWS Documentation - AWS WAF and Shield Developer Guide

Question: 252

A company is planning to deploy a business-critical application in the AWS Cloud. The application requires durable storage with consistent, low-latency performance.

Which type of storage should a solutions architect recommend to meet these requirements?

- A. Instance store volume
- B. Amazon ElastiCache (Memcached) cluster
- C. Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume
- D. Throughput Optimized HDD Amazon Elastic Block Store (Amazon EBS) volume

Answer: C

Explanation:

Provisioned IOPS SSD (io1 or io2) volumes are designed to deliver predictable, high performance for I/O-intensive workloads such as databases. They offer consistent, low-latency performance and durability, making them ideal for business-critical applications that cannot tolerate latency or data loss.

Reference: AWS Documentation - Amazon EBS Volume Types

Question: 253

A company is creating an application. The company stores data from tests of the application in multiple on-premises locations.

The company needs to connect the on-premises locations to VPCs in an AWS Region in the AWS Cloud. The number of accounts and VPCs will increase during the next year. The network architecture must simplify the administration of new connections and must provide the ability to scale.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Create a peering connection between the VPCs. Create a VPN connection between the VPCs and the on-premises locations.
- B. Launch an Amazon EC2 instance. On the instance, include VPN software that uses a VPN connection to connect all VPCs and on-premises locations.
- C. Create a transit gateway. Create VPC attachments for the VPC connections. Create VPNAttachments for the on-premises connections.
- D. Create an AWS Direct Connect connection between the on-premises locations and a central VPC. Connect the central VPC to other VPCs by using peering connections.

Answer: C

Explanation:

AWS Transit Gateway simplifies network connectivity by acting as a hub that can connect VPCs and on-premises networks through VPN or Direct Connect. It provides scalability and reduces administrative overhead by eliminating the need to manage complex peering relationships as the number of accounts and VPCs grows.

Reference: AWS Documentation - Transit Gateway

Question: 254

A company has an organization in AWS Organizations that has all features enabled. The company has multiple Amazon S3 buckets in multiple AWS Regions around the world. The S3 buckets contain sensitive data.

The company needs to ensure that no personally identifiable information (PII) is stored in the S3 buckets. The company also needs a scalable solution to identify PII.

Which solution will meet these requirements?

- A. In the Organizations management account, configure an Amazon Macie administrator IAM user as the delegated administrator for the global organization. Use the Macie administrator user to configure Macie settings to scan for PII.

- B. For each Region in the Organizations management account, designate a delegated Amazon Macie administrator account. In the Macie administrator account, add all accounts in the organization. Use the Macie administrator account to enable Macie. Configure automated sensitive data discovery for all accounts in the organization.
- C. For each Region in the Organizations management account, configure a service control policy (SCP) to identify PII. Apply the SCP to the organization root.
- D. In the Organizations management account, configure AWS Lambda functions to scan for PII in each Region.

Answer: B

Explanation:

Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. To scale across Regions and accounts in AWS Organizations, Macie supports delegated administration, automated sensitive data discovery, and multi-account aggregation through a centralized admin account.

Reference: AWS Documentation - Amazon Macie Multi-Account Configuration

Question: 255

A company is planning to deploy a data processing platform on AWS. The data processing platform is based on PostgreSQL. The company stores the data that the platform must process on premises.

To comply with regulations, the company must not migrate the data to the cloud. However, the company wants to use AWS managed data analytics solutions.

Which solution will meet these requirements?

- A. Create an Amazon RDS for PostgreSQL database in a VPC. Create an interface VPC endpoint to connect the on-premises PostgreSQL database to the RDS for PostgreSQL database.
- B. Create Amazon EC2 instances in an Auto Scaling group on AWS Outposts. Install PostgreSQL data analytics software on the instances.
- C. Create an Amazon EMR cluster on AWS Outposts. Connect the EMR cluster to the on-premises PostgreSQL database to perform data processing locally.
- D. Create an Amazon EMR cluster in a VPC. Connect the EMR cluster to Amazon RDS for SQL Server with a linked server to connect to the company's data processing platform.

Answer: C

Explanation:

AWS Outposts extends AWS infrastructure and services to on-premises locations. Running Amazon EMR on Outposts allows for processing data that resides locally while benefiting from the managed services of EMR. This enables compliance with data residency requirements and provides scalability and manageability for analytics.

Reference: AWS Documentation - Amazon EMR on Outposts

Question: 256

A company is testing an application that runs on an Amazon EC2 Linux instance. A single 500 GB Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume is attached to the EC2 instance.

The company will deploy the application on multiple EC2 instances in an Auto Scaling group. All instances require access to the data that is stored in the EBS volume. The company needs a highly available and resilient solution that does not introduce significant changes to the application's code.

Which solution will meet these requirements?

- A. Provision an EC2 instance that uses NFS server software. Attach a single 500 GB gp2 EBS volume to the instance.
- B. Provision an Amazon FSx for Windows File Server file system. Configure the file system as an SMB file store within a single Availability Zone.

- C. Provision an EC2 instance with two 250 GB Provisioned IOPS SSD EBS volumes.
- D. Provision an Amazon Elastic File System (Amazon EFS) file system. Configure the file system to use General Purpose performance mode.

Answer: D

Explanation:

Amazon EFS is a fully managed, scalable file storage service that can be accessed concurrently by thousands of EC2 instances. It supports General Purpose performance mode for latency-sensitive use cases and provides high availability and durability across multiple Availability Zones with minimal changes to application code.

Reference: AWS Documentation - Amazon Elastic File System (EFS)

Question: 257

A company runs a critical three-tier web application that consists of multiple virtual machines (VMs) and virtual databases in an on-premises environment. The company wants to set up a disaster recovery (DR) environment in AWS.

The company requires a 15-minute recovery time objective (RTO). The company must be able to test the failover solution to validate the recovery. The solution must provide an automated failover mechanism.

Which solution will meet these requirements?

- A. Use AWS Backup to create backups of the on-premises VMs and to restore the backups in AWS. Configure recovery to Amazon EC2 instances to meet the RTO requirement.
- B. Use AWS Database Migration Service (AWS DMS) to replicate the on-premises databases to Amazon RDS. Set up AWS Storage Gateway for baseline and incremental data replication to AWS to meet the RTO requirement.
- C. Use AWS DataSync and AWS Storage Gateway to migrate the baseline and incremental data to AWS. Use Amazon EC2, Amazon S3, and an Application Load Balancer to set up the DR environment.
- D. Use AWS Elastic Disaster Recovery to replicate the VMs incrementally to AWS. Configure Elastic Disaster Recovery to automate the DR process.

Answer: D

Explanation:

AWS Elastic Disaster Recovery (AWS DRS) enables fast, reliable, and cost-effective disaster recovery. It replicates on-premises machines to AWS using continuous block-level replication. It supports automated testing and failover, meeting aggressive RTO targets such as 15 minutes.

Reference: AWS Documentation - Elastic Disaster Recovery Overview

Question: 258

A company is launching a new application that requires a structured database to store user profiles, application settings, and transactional data. The database must be scalable with application traffic and must offer backups.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy a self-managed database on Amazon EC2 instances by using open-source software. Use Spot Instances for cost optimization. Configure automated backups to Amazon S3.
- B. Use Amazon RDS. Use on-demand capacity mode for the database with General Purpose SSD storage. Configure automatic backups with a retention period of 7 days.
- C. Use Amazon Aurora Serverless for the database. Use serverless capacity scaling. Configure automated backups to Amazon S3.
- D. Deploy a self-managed NoSQL database on Amazon EC2 instances. Use Reserved Instances for cost optimization. Configure automated

backups directly to Amazon S3 Glacier Flexible Retrieval.

Answer: C

Explanation:

Amazon Aurora Serverless v2 is ideal for applications with unpredictable or intermittent workloads. It automatically scales capacity up or down based on demand, significantly reducing costs. It supports automated backups to Amazon S3, making it suitable and cost-effective for new applications with variable traffic.

Reference: AWS Documentation - Amazon Aurora Serverless

Question: 259

A company collects data from sensors. The company needs a cloud-based solution to store and transform the sensor data to make critical decisions. The solution must store the data for up to 2 days. After 2 days, the solution must delete the data. The company needs to use the transformed data in an automated workflow that has manual approval steps.

Which solution will meet these requirements?

- A. Load the data into an Amazon Simple Queue Service (Amazon SQS) queue that has a retention period of 2 days. Use an Amazon EventBridge pipe to retrieve data from the queue, transform the data, and pass the data to an AWS Step Functions workflow.
- B. Load the data into AWS DataSync. Delete the DataSync task after 2 days. Invoke an AWS Lambda function to retrieve the data, transform the data, and invoke a second Lambda function that performs the remaining workflow steps.
- C. Load the data into an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge pipe to retrieve the data from the topic, transform the data, and send the data to Amazon EC2 instances to perform the remaining workflow steps.
- D. Load the data into an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge pipe to retrieve the data from the topic and transform the data into an appropriate format for an Amazon SQS queue. Use an AWS Lambda function to poll the queue to perform the remaining workflow steps.

Answer: A

Explanation:

Amazon SQS with a 2-day retention ensures the data lives just as long as needed. EventBridge Pipes allow direct integration between event producers and consumers, with optional filtering and transformation. AWS Step Functions supports manual approval steps, which fits the workflow requirement perfectly.

Reference: AWS Documentation - Amazon EventBridge Pipes, AWS Step Functions

Question: 260

A company hosts an application on AWS that gives users the ability to download photos. The company stores all photos in an Amazon S3 bucket that is located in the us-east-1 Region. The company wants to provide the photo download application to global customers with low latency.

Which solution will meet these requirements?

- A. Find the public IP addresses that Amazon S3 uses in us-east-1. Configure an Amazon Route 53 latency-based routing policy that routes to all the public IP addresses.
- B. Configure an Amazon CloudFront distribution in front of the S3 bucket. Use the distribution endpoint to access the photos that are in the S3 bucket.
- C. Configure an Amazon Route 53 geoproximity routing policy to route the traffic to the S3 bucket that is closest to each customer's location.
- D. Create a new S3 bucket in the us-west-1 Region. Configure an S3 Cross-Region Replication rule to copy the photos to the new S3 bucket.

Answer: B

Explanation:

Amazon CloudFront is a content delivery network (CDN) service that distributes content with low latency and high transfer speeds. Placing CloudFront in front of the S3 bucket ensures global users download content from the nearest edge location, reducing latency significantly.

Reference: AWS Documentation - Amazon CloudFront with S3 Origin

Question: 261

A company QUESTION NO: 241

A company has launched an Amazon RDS for MySQL DB instance. Most of the connections to the database come from serverless applications. Application traffic to the database changes significantly at random intervals. At times of high demand, users report that their applications experience database connection rejection errors.

Which solution will resolve this issue with the LEAST operational overhead?

- A. Create a proxy in RDS Proxy. Configure the users' applications to use the DB instance through RDS Proxy.
- B. Deploy Amazon ElastiCache (Memcached) between the users' applications and the DB instance.
- C. Migrate the DB instance to a different instance class that has higher I/O capacity. Configure the users' applications to use the new DB instance.
- D. Configure Multi-AZ for the DB instance. Configure the users' applications to switch between the DB instances.

Answer: A

Explanation:

Amazon RDS Proxy is designed to manage a large number of database connections from applications, especially serverless applications that can scale quickly. It improves application availability and scalability by pooling and sharing established database connections. This reduces the overhead of database connections and prevents overload during traffic spikes.

Reference: AWS Documentation - Amazon RDS Proxy

Question: 262

A company is building a serverless web application with multiple interdependent workflows that millions of users worldwide will access. The application needs to handle bursts of traffic.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy an Amazon API Gateway HTTP API with a usage plan and throttle settings. Use AWS Step Functions with a Standard Workflow.
- B. Deploy an Amazon API Gateway HTTP API with a usage plan and throttle settings. Use AWS Step Functions with an Express Workflow.
- C. Deploy an Amazon API Gateway HTTP API without a usage plan. Use AWS Step Functions with an Express Workflow.
- D. Deploy an Amazon API Gateway HTTP API without a usage plan. Use AWS Step Functions and multiple AWS Lambda functions with reserved concurrency.

Answer: B

Explanation:

Express Workflows in AWS Step Functions are optimized for high-throughput, short-duration, and low-cost workflows. They are suitable for

applications with large volumes of parallel and interdependent tasks. When paired with HTTP APIs from API Gateway, which are more cost-efficient than REST APIs, this setup offers scalability and cost-effectiveness.

Reference: AWS Documentation - Step Functions Express Workflows vs. Standard Workflows

Question: 263

A solutions architect is designing the cloud architecture for a new stateless application that will be deployed on AWS. The solutions architect created an Amazon Machine Image (AMI) and launch template for the application.

Based on the number of jobs that need to be processed, the processing must run in parallel while adding and removing application Amazon EC2 instances as needed. The application must be loosely coupled. The job items must be durably stored.

Which solution will meet these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic to send the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on CPU usage.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue to hold the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on network usage.
- C. Create an Amazon Simple Queue Service (Amazon SQS) queue to hold the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on the number of items in the SQS queue.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic to send the jobs that need to be processed. Create an Auto Scaling group by using the launch template with the scaling policy set to add and remove EC2 instances based on the number of messages published to the SNS topic.

Answer: C

Explanation:

Amazon SQS provides durable, decoupled message storage for distributed systems. Using SQS as a job queue enables each EC2 instance to process a message independently. Scaling the Auto Scaling group based on the SQS queue length ensures parallelism and elasticity, aligning compute resources with workload volume.

Reference: AWS Documentation - Using Amazon SQS with Auto Scaling

Question: 264

A company is hosting multiple websites for several lines of business under its registered parent domain. Users accessing these websites will be routed to appropriate backend Amazon EC2 instances based on the subdomain. The websites host static webpages, images, and server-side scripts like PHP and JavaScript.

Some of the websites experience peak access during the first two hours of business with constant usage throughout the rest of the day. A solutions architect needs to design a solution that will automatically adjust capacity to these traffic patterns while keeping costs low.

Which combination of AWS services or features will meet these requirements? (Select TWO.)

- A. AWS Batch
- B. Network Load Balancer
- C. Application Load Balancer
- D. Amazon EC2 Auto Scaling
- E. Amazon S3 website hosting

Answer: C,D

Explanation:

An Application Load Balancer supports path- and host-based routing, which makes it ideal for routing requests based on subdomains. EC2 Auto Scaling ensures that the number of instances adjusts dynamically based on traffic, which helps manage cost and performance during predictable peak hours.

Reference: AWS Documentation - ALB with Auto Scaling for Web Applications

Question: 265

An internal product team is deploying a new application to a private VPC in a company's AWS account. The application runs on Amazon EC2 instances that are in a security group named App1. The EC2 instances store application data in an Amazon S3 bucket and use AWS Secrets Manager to store application service credentials. The company's security policy prohibits applications in a private VPC from using public IP addresses to communicate.

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Configure gateway endpoints for Amazon S3 and AWS Secrets Manager.
- B. Configure interface VPC endpoints for Amazon S3 and AWS Secrets Manager.
- C. Add routes to the endpoints in the VPC route table.

D. Associate the App1 security group with the interface VPC endpoints. Configure a self-referencing security group rule to allow inbound traffic.

E. Associate the App1 security group with the gateway endpoints. Configure a self-referencing security group rule to allow inbound traffic.

Answer: B,C

Explanation:

To securely access AWS services like S3 and Secrets Manager from a private VPC without using public IPs, interface VPC endpoints are required. These endpoints are accessible via private IP addresses. For the application to reach these endpoints, appropriate routes must be configured in the route table.

Reference: AWS Documentation - Interface VPC Endpoints and PrivateLink

Question: 266

A company wants to optimize costs for its AWS infrastructure. The company wants to receive notifications when actual costs or forecasted costs exceed a specified budget. The company does not want to develop a custom solution.

Which solution will meet these requirements?

A. Use AWS Trusted Advisor to set up budget notifications. Configure Amazon CloudWatch to monitor costs. Export CloudWatch data to Amazon S3. Use machine learning (ML) to estimate future trends based on the CloudWatch data.

B. Create a budget in AWS Budgets that has a specified cost threshold. Create an AWS Lambda function that sends a notification to the company when costs reach the specified threshold. Use AWS Billing and Cost Management reports to monitor costs.

C. Use AWS Cost Explorer to set a specified budget threshold. Create an AWS Lambda function to calculate cost estimates. Configure the Lambda function to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic if estimated costs exceed the specified threshold.

D. Create a budget in AWS Budgets that has a specified cost threshold. Configure AWS Budgets to send budget alerts to an Amazon Simple Notification Service (Amazon SNS) topic. Use AWS Cost

Explorer to monitor costs.

Answer: D

Explanation:

AWS Budgets allows you to set custom cost and usage budgets. When actual or forecasted usage exceeds the threshold, AWS Budgets can automatically send alerts to an Amazon SNS topic. You can use AWS Cost Explorer in parallel for visual tracking of spending. This solution requires no code and has minimal operational overhead.

Reference: AWS Documentation - AWS Budgets and Cost Explorer

Question: 267

A company launches a new web application that uses an Amazon Aurora PostgreSQL database. The company wants to add new features to the application that rely on AI. The company requires vector storage capability to use AI tools.

Which solution will meet this requirement MOST cost-effectively?

- A. Use Amazon OpenSearch Service to create an OpenSearch service. Configure the application to write vector embeddings to a vector index.
- B. Create an Amazon DocumentDB cluster. Configure the application to write vector embeddings to a vector index.
- C. Create an Amazon Neptune ML cluster. Configure the application to write vector embeddings to a vector graph.
- D. Install the pgvector extension on the Aurora PostgreSQL database. Configure the application to write vector embeddings to a vector table.

Answer: D

Explanation:

Aurora PostgreSQL supports the pgvector extension, which allows storage and querying of vector embeddings directly inside the database. This eliminates the need for external vector databases and provides cost-effective and performant integration for AI workloads.

Reference: AWS Documentation - Amazon Aurora PostgreSQL and pgvector Support

Question: 268

A company manages multiple AWS accounts in an organization in AWS Organizations. The company's applications run on Amazon EC2 instances in multiple AWS Regions. The company needs a solution to simplify the management of security rules across the accounts in its organization. The solution must apply shared security group rules, audit security groups, and detect unused and redundant rules in VPC security groups across all AWS environments.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Use AWS Firewall Manager to create a set of rules based on the security requirements. Replicate the rules to all the AWS accounts and Regions.
- B. Use AWS CloudFormation StackSets to provision VPC security groups based on the specifications across multiple accounts and Regions. Deploy AWS Network Firewall to define the firewall rules to control network traffic across multiple accounts and Regions.
- C. Use AWS CloudFormation StackSets to provision VPC security groups based on the specifications across multiple accounts and Regions. Configure AWS Config and AWS Lambda to evaluate compliance information and to automate enforcement across all accounts and Regions.
- D. Use AWS Network Firewall to build policies based on the security requirements. Centrally apply the new policies to all the VPCs and accounts.

Answer: A

Explanation:

AWS Firewall Manager integrates with AWS Organizations to centrally manage and apply security group policies, AWS WAF rules, and AWS Shield Advanced protections. It automates the propagation of rules across accounts and Regions and can also audit and remediate noncompliant configurations.

Reference: AWS Documentation - AWS Firewall Manager for Centralized Security Group Management

Question: 269

A company uses Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to run an application. The company creates one snapshot of each EBS volume every day.

The company needs to prevent users from accidentally deleting the EBS volume snapshots. The solution must not change the administrative rights of a storage administrator user.

Which solution will meet these requirements with the LEAST administrative effort?

- A. Create an IAM role that has permission to delete snapshots. Attach the role to a new EC2 instance. Use the AWS CLI from the new EC2 instance to delete snapshots.
- B. Create an IAM policy that denies snapshot deletion. Attach the policy to the storage administrator user.
- C. Add tags to the snapshots. Create tag-level retention rules in the Recycle Bin for EBS snapshots. Configure rule lock settings for the retention rules.
- D. Take EBS snapshots by using the EBS direct APIs. Copy the snapshots to an Amazon S3 bucket. Configure S3 Versioning and Object Lock on the bucket.

Answer: C

Explanation:

Amazon EBS Snapshots Recycle Bin enables you to specify retention rules for EBS snapshots based on tags. When snapshots are deleted, they are retained in the Recycle Bin for a specified duration, preventing accidental deletion. Tag-level rules allow selective protection without changing IAM roles or user permissions.

Reference: AWS Documentation - Amazon EBS Snapshots and Recycle Bin

Question: 270

A software company needs to upgrade a critical web application. The application is hosted in a public subnet. The EC2 instance runs a MySQL database. The application's DNS records are published in an Amazon Route 53 zone.

A solutions architect must reconfigure the application to be scalable and highly available. The solutions architect must also reduce MySQL read latency.

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Launch a second EC2 instance in a second AWS Region. Use a Route 53 failover routing policy to redirect the traffic to the second EC2 instance.
- B. Create and configure an Auto Scaling group to launch private EC2 instances in multiple Availability Zones. Add the instances to a target group behind a new Application Load Balancer.
- C. Migrate the database to an Amazon Aurora MySQL cluster. Create the primary DB instance and reader DB instance in separate Availability Zones.
- D. Create and configure an Auto Scaling group to launch private EC2 instances in multiple AWS Regions. Add the instances to a target group behind a new Application Load Balancer.
- E. Migrate the database to an Amazon Aurora MySQL cluster with cross-Region read replicas.

Answer: B,C

Explanation:

To improve scalability and availability, EC2 Auto Scaling across multiple Availability Zones with an Application Load Balancer ensures resilient infrastructure. Migrating to Amazon Aurora MySQL with reader endpoints reduces read latency by offloading read traffic to replicas in other AZs, while also increasing high availability.

Reference: AWS Documentation - Aurora Multi-AZ and EC2 Auto Scaling with ALB

Question: 271

A company is designing an application to connect AWS Lambda functions to an Amazon RDS for MySQL DB instance. The DB instance manages many connections. The company needs to modify the application to improve connectivity and recovery.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon RDS Proxy for connection pooling. Modify the application to use the RDS Proxy for connections to the DB instance.
- B. Create a new RDS instance for connection pooling. Modify the application to use the new RDS instance for connectivity.
- C. Create read replicas to distribute the load of the DB instance. Create a Network Load Balancer to distribute the load across the read replicas.
- D. Migrate the RDS for MySQL DB instance to Amazon Aurora MySQL to increase DB instance performance.

Answer: A

Explanation:

Amazon RDS Proxy helps manage thousands of concurrent database connections by pooling and reusing them efficiently. It is especially useful for serverless applications like AWS Lambda that can open numerous connections quickly, potentially overwhelming the database. Using RDS Proxy reduces connection management overhead and improves fault tolerance.

Reference: AWS Documentation - Amazon RDS Proxy

Question: 272

A company has a multi-tier web application. The application's internal service components are deployed on Amazon EC2 instances. The internal service components need to access third-party software as a service (SaaS) APIs that are hosted on AWS.

The company needs to provide secure and private connectivity from the application's internal services to the third-party SaaS application. The company needs to ensure that there is minimal public internet exposure.

Which solution will meet these requirements?

- A. Implement an AWS Site-to-Site VPN to establish a secure connection with the third-party SaaS provider.
- B. Deploy AWS Transit Gateway to manage and route traffic between the application's VPC and the third-party SaaS provider.
- C. Configure AWS PrivateLink to allow only outbound traffic from the VPC without enabling the third-party SaaS provider to establish a return path to the network.

D. Use AWS PrivateLink to create a private connection between the application's VPC and the third- party SaaS provider.

Answer: D

Explanation:

AWS PrivateLink enables private connectivity between VPCs and supported AWS or third-party services without exposing traffic to the public internet. It ensures secure and private communications, making it ideal for connecting internal services to SaaS applications hosted in AWS.

Reference: AWS Documentation - AWS PrivateLink for SaaS Connectivity

Question: 273

A company wants to create an Amazon EMR cluster that multiple teams will use. The company wants to ensure that each team's big data workloads can access only the AWS services that each team needs to interact with. The company does not want the workloads to have access to Instance Metadata Service Version 2 (IMDSv2) on the cluster's underlying EC2 instances.

Which solution will meet these requirements?

- A. Configure interface VPC endpoints for each AWS service that the teams need. Use the required interface VPC endpoints to submit the big data workloads.
- B. Create EMR runtime roles. Configure the cluster to use the runtime roles. Use the runtime roles to submit the big data workloads.
- C. Create an EC2 IAM instance profile that has the required permissions for each team. Use the instance profile to submit the big data workloads.
- D. Create an EMR security configuration that has the EnableApplicationScoped IAM Role option set to false. Use the security configuration to submit the big data workloads.

Answer: B

Explanation:

EMR runtime roles allow fine-grained permissions per job, letting each team access only the services they are authorized to use. This isolates IAM permissions per workload and avoids exposing instancelevel credentials through IMDSv2. Runtime roles improve security posture in multi-tenant EMR environments.

Reference: AWS Documentation - EMR Runtime Roles and Access Isolation

Question: 274

A company has an application that processes information from documents that users upload. When a user uploads a new document to an Amazon S3 bucket, an AWS Lambda function is invoked. The Lambda function processes information from the documents.

The company discovers that the application did not process many recently uploaded documents. The company wants to ensure that the application processes each document with retries if there is an error during the first attempt to process the document.

Which solution will meet these requirements?

- A. Create an Amazon API Gateway REST API that has a proxy integration to the Lambda function. Update the application to send requests to the REST API.
- B. Configure a replication policy on the S3 bucket to stage the documents in another S3 bucket that an AWS Batch job processes on a daily schedule.
- C. Deploy an Application Load Balancer in front of the Lambda function that processes the documents.
- D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source for the Lambda function. Configure an S3 event notification on the S3 bucket to send new document upload events to the SQS queue.

Answer: D

Explanation:

Using SQS as a buffer between S3 and the Lambda function ensures durability and allows for retries in case of processing failures. Messages in the queue can be retried by Lambda, and failed processing can be directed to a dead-letter queue for further inspection. This guarantees reliable and scalable message-driven processing.

Reference: AWS Documentation - Using Amazon SQS as Lambda Event Source with S3 Trigger

Question: 275

A solutions architect is designing the architecture for a web application that has a frontend and a backend. The backend services must receive data from the frontend services for processing. The frontend must manage access to the application by using API keys. The backend must scale without affecting the frontend.

Which solution will meet these requirements?

- A. Deploy an Amazon API Gateway HTTP API as the frontend to direct traffic to an Amazon Simple Queue Service (Amazon SQS) queue. Use AWS Lambda functions as the backend to read from the queue.
- B. Deploy an Amazon API Gateway REST API as the frontend to direct traffic to an Amazon Simple Queue Service (Amazon SQS) queue.

Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate as the backend to read from the queue.

C. Deploy an Amazon API Gateway REST API as the frontend to direct traffic to an Amazon Simple Notification Service (Amazon SNS) topic. Use AWS Lambda functions as the backend. Subscribe the Lambda functions to the topic.

D. Deploy an Amazon API Gateway HTTP API as the frontend to direct traffic to an Amazon Simple Notification Service (Amazon SNS) topic. Use Amazon Elastic Kubernetes Service (Amazon EKS) on AWS Fargate as the backend. Subscribe Amazon EKS to the topic.

Answer: A

Explanation:

Using API Gateway with API keys provides secure access control. Amazon SQS allows asynchronous decoupling between frontend and backend, ensuring that backend processing can scale independently. AWS Lambda reading from SQS ensures scalable, event-driven processing with minimal operational management. This architecture is resilient and decoupled.

Reference: AWS Well-Architected - Decoupling and Microservices, API Gateway with SQS and Lambda Integration

Question: 276

A company is migrating a daily Microsoft Windows batch job from the company's on-premises environment to AWS. The current batch job runs for up to 1 hour. The company wants to modernize the batch job process for the cloud environment.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create a fleet of Amazon EC2 instances in an Auto Scaling group to handle the Windows batch job processing.

B. Implement an AWS Lambda function to process the Windows batch job. Use an Amazon EventBridge rule to invoke the Lambda function.

C. Use AWS Fargate to deploy the Windows batch job as a container. Use AWS Batch to manage the batch job processing.

D. Use Amazon Elastic Kubernetes Service (Amazon EKS) on Amazon EC2 instances to orchestrate Windows containers for the batch job processing.

Answer: C

Explanation:

AWS Batch supports Windows-based jobs and automates provisioning and scaling of compute environments. Paired with AWS Fargate, it removes the need to manage infrastructure. This solution requires the least operational overhead and is cloud-native, providing flexibility and scalability.

Reference: AWS Documentation - AWS Batch with Fargate for Windows Workloads

Question: 277

A company is building a serverless application that processes large volumes of data from a mobile app. The application uses an AWS Lambda function to process the data and store the data in an Amazon DynamoDB table.

The company needs to ensure that the application can recover from failures and continue processing data without losing any records.

Which solution will meet these requirements?

- A. Configure the Lambda function to use a dead-letter queue with an Amazon Simple Queue Service (Amazon SQS) queue. Configure Lambda to retry failed records from the dead-letter queue. Use a retry mechanism by implementing an exponential backoff algorithm.
- B. Configure the Lambda function to read records from Amazon Data Firehose. Replay the Firehose records in case of any failures.
- C. Use Amazon OpenSearch Service to store failed records. Configure AWS Lambda to retry failed records from OpenSearch Service. Use Amazon EventBridge to orchestrate the retry logic.
- D. Use Amazon Simple Notification Service (Amazon SNS) to store the failed records. Configure Lambda to retry failed records from the SNS topic. Use Amazon API Gateway to orchestrate the retry calls.

Answer: A

Explanation:

Dead-letter queues (DLQs) with Amazon SQS allow Lambda functions to offload failed events for later inspection or retry. Using retry logic with exponential backoff ensures resilience and compliance with best practices for fault-tolerant serverless architectures. This guarantees no data is lost due to transient errors.

Reference: AWS Documentation - Lambda Error Handling and Dead-Letter Queues

Question: 278

A company discovers that an Amazon DynamoDB Accelerator (DAX) cluster for the company's web application workload is not encrypting data at rest. The company needs to resolve this security issue.

Which solution will meet this requirement?

- A. Stop the existing DAX cluster. Enable encryption at rest for the existing DAX cluster, and start the cluster again.
- B. Delete the existing DAX cluster. Recreate the DAX cluster, and configure the new cluster to encrypt the data at rest.
- C. Update the configuration of the existing DAX cluster to encrypt the data at rest.
- D. Integrate the existing DAX cluster with AWS Security Hub to automatically enable encryption at rest.

Answer: B

Explanation:

DAX does not support enabling encryption at rest on an existing cluster. To use encryption at rest, you must create a new DAX cluster with encryption enabled at creation time and migrate workloads accordingly.

Reference: AWS Documentation - Amazon DAX Encryption at Rest

Question: 279

A company is developing a serverless web application that gives users the ability to interact with real-time analytics from online games. The data from the games must be streamed in real time. The company needs a durable, low-latency database option for user data. The company does not know how many users will use the application. Any design considerations must provide response times of single-digit milliseconds as the application scales.

Which combination of AWS services will meet these requirements? (Select TWO.)

- A. Amazon CloudFront
- B. Amazon DynamoDB
- C. Amazon Kinesis
- D. Amazon RDS
- E. AWS Global Accelerator

Answer: B,C

Explanation:

Amazon Kinesis allows real-time ingestion of game events at scale, while Amazon DynamoDB provides millisecond-latency access to user data, automatically scaling with demand. This combination ensures real-time processing and fast data retrieval without managing infrastructure.

Reference: AWS Documentation - Real-Time Processing with Kinesis and Low-Latency Databases with DynamoDB

Question: 280

A company is migrating a production environment application to the AWS Cloud. The company uses Amazon RDS for Oracle for the database layer. The company needs to configure the database to meet the needs of high I/O intensive workloads that require low latency and consistent throughput. The database workloads are read intensive and write intensive.

Which solution will meet these requirements?

- A. Use a Multi-AZ DB instance deployment for the RDS for Oracle database.
- B. Configure the RDS for Oracle database to use the Provisioned IOPS SSD storage type.
- C. Configure the RDS for Oracle database to use the General Purpose SSD storage type.
- D. Enable RDS read replicas for RDS for Oracle.

Answer: B

Explanation:

Provisioned IOPS SSD (io1 or io2) is designed for I/O-intensive workloads that require low latency and consistent throughput, which is critical for transactional and production databases. It provides predictable performance, unlike General Purpose SSD, which is burst-based.

Reference: AWS RDS Storage - Amazon RDS Storage Types (Provisioned IOPS SSD for latency-sensitive workloads)

Question: 281

A company hosts a public web application on AWS. The website has a three-tier architecture. The frontend web tier is comprised of Amazon EC2 instances in an Auto Scaling group. The application tier is a second Auto Scaling group. The database tier is an Amazon RDS database.

The company has configured the Auto Scaling groups to handle the application's normal level of demand. During an unexpected spike in demand, the company notices a long delay in the startup time when the frontend and application layers scale out. The company needs to improve the scaling performance of the application without negatively affecting the user experience.

Which solution will meet these requirements MOST cost-effectively?

- A. Decrease the minimum number of EC2 instances for both Auto Scaling groups. Increase the desired number of instances to meet the peak demand requirement.
- B. Configure the maximum number of instances for both Auto Scaling groups to be the number required to meet the peak demand. Create a warm pool.
- C. Increase the maximum number of EC2 instances for both Auto Scaling groups to meet the normal demand requirement. Create a warm pool.
- D. Reconfigure both Auto Scaling groups to use a scheduled scaling policy. Increase the size of the EC2 instance types and the RDS instance types.

Answer: B

Explanation:

EC2 Auto Scaling warm pools allow you to pre-initialize instances, reducing the delay in scale-out events. This results in significantly faster response times during demand surges while remaining cost effective compared to always running at peak capacity.

Reference: AWS Documentation - EC2 Auto Scaling Warm Pools for Faster Scaling

Question: 282

A company wants to enhance its ecommerce order-processing application that is deployed on AWS. The application must process each order exactly once without affecting the customer experience during unpredictable traffic surges.

Which solution will meet these requirements?

- A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Put all the orders in the SQS queue. Configure an AWS Lambda function as the target to process the orders.
- B. Create an Amazon Simple Notification Service (Amazon SNS) standard topic. Publish all the orders to the SNS standard topic. Configure the application as a notification target.
- C. Create a flow by using Amazon AppFlow. Send the orders to the flow. Configure an AWS Lambda function as the target to

process the orders.

D. Configure AWS X-Ray in the application to track the order requests. Configure the application to process the orders by pulling the orders from Amazon CloudWatch.

Answer: A

Explanation:

Amazon SQS FIFO queues guarantee the order of message delivery and ensure that each message is delivered exactly once. Paired with AWS Lambda, this creates a scalable, fault-tolerant architecture that processes each order in order and prevents duplicates, which is critical for ecommerce

workflows.

Reference: AWS Documentation - Amazon SQS FIFO Queues and Exactly-Once Processing

Question: 283

A telemarketing company is designing its customer call center functionality on AWS. The company needs a solution that provides multiple speaker recognition and generates transcript files. The company wants to query the transcript files to analyze the business patterns.

Which solution will meet these requirements?

- A. Use Amazon Rekognition for multiple speaker recognition. Store the transcript files in Amazon S3. Use machine learning (ML) models to analyze the transcript files.
- B. Use Amazon Transcribe for multiple speaker recognition. Use Amazon Athena to analyze the transcript files.
- C. Use Amazon Translate for multiple speaker recognition. Store the transcript files in Amazon Redshift. Use SQL queries to analyze the transcript files.
- D. Use Amazon Rekognition for multiple speaker recognition. Store the transcript files in Amazon S3. Use Amazon Textract to analyze the transcript files.

Answer: B

Explanation:

Amazon Transcribe supports automatic speech recognition (ASR) with speaker diarization (i.e., multiple speaker identification). The transcripts can be stored in Amazon S3 and queried using Amazon Athena, which provides a serverless, pay-as-you-go interactive querying model.

Reference: AWS Documentation - Amazon Transcribe + Amazon Athena for Speech-to-Insights Pipelines

Question: 284

A company has a three-tier web application. An Application Load Balancer (ALB) is in front of Amazon EC2 instances that are in the ALB target group. An Amazon S3 bucket stores documents.

The company requires the application to meet a recovery time objective (RTO) of 60 seconds.

Which solution will meet this requirement?

- A. Replicate S3 objects to a second AWS Region. Create a second ALB and a minimum set of EC2 instances in the second Region. Ensure that the EC2 instances are shut down until they are needed. Configure Amazon Route 53 to fail over to the second Region by using an IP-based routing policy.
- B. Use AWS Backup to take hourly backups of the EC2 instances. Back up the S3 data to a second AWS Region. Use AWS CloudFormation to deploy the entire infrastructure in the second Region when needed.
- C. Create daily snapshots of the EC2 instances in a second AWS Region. Use the snapshots to recreate the instances in the second Region. Back up the S3 data to the second Region. Perform a failover by modifying the application DNS record when needed.
- D. Replicate S3 objects to a second AWS Region. Create a second ALB and a minimum set of EC2 instances in the second Region. Ensure that the EC2 instances in the second Region are running. Configure Amazon Route 53 to fail over to the secondary Region based on health checks.

Answer: D

Explanation:

To achieve a 60-second RTO, pre-warming the DR environment (including running EC2 instances and Route 53 health checks) is essential. Active/passive failover using Route 53 with health checks ensures fast redirection when the primary Region becomes unavailable. S3 cross-region replication ensures document availability.

Reference: AWS Disaster Recovery - Active-Passive Strategy with Route 53 and Health Checks

Question: 285

A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images

that are stored in Amazon S3. This content is the same for all users.

The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.

Which solution meets these requirements MOST cost-effectively?

- A. Deploy an AWS Global Accelerator accelerator in front of the web servers.
- B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.
- C. Deploy an Amazon ElastiCache (Redis OSS) instance in front of the web servers.
- D. Deploy an Amazon ElastiCache (Memcached) instance in front of the web servers.

Answer: B

Explanation:

Amazon CloudFront is a highly cost-effective CDN that caches content like images and videos at edge locations globally. This reduces latency and the load on the origin S3 bucket. It is ideal for static content that is accessed by many users.

Reference: AWS Documentation - Amazon CloudFront with S3 Integration

Question: 286

A media company runs an application on multiple Amazon EC2 instances that requires high storage input/output operations per second (IOPS).

To achieve the necessary performance, a solutions architect wants to stripe multiple Amazon EBS volumes together and attach the volumes to EC2 instances. The solutions architect wants to receive a notification when IOPS are over-provisioned.

Which solution will meet these requirements?

- A. Configure auto scaling for the EBS volumes to automatically increase or decrease IOPS based on the EC2 instance CPU utilization metric.
- B. Deploy the application on an EC2 instance type that supports the highest possible IOPS.
- C. Create a custom AWS Config rule to monitor the provisioned IOPS for the EBS volumes that are attached to the EC2 instances and to send notifications.
- D. Adjust the IOPS of each EBS volume daily based on Amazon CloudWatch metrics for IOPS utilization.

Answer: C

Explanation:

AWS Config allows for creation of custom rules to monitor EBS configurations. Combined with CloudWatch metrics and Amazon SNS, custom rules can track over-provisioned IOPS and send alerts when thresholds are breached, allowing proactive cost and performance management.

Reference: AWS Documentation - AWS Config with Custom Rules for EBS Monitoring

Question: 287

An insurance company runs an application on premises to process contracts. The application processes jobs that are comprised of many tasks. The individual tasks run for up to 5 minutes. Some jobs can take up to 24 hours in total to finish. If a task fails, the task must be reprocessed.

The company wants to migrate the application to AWS. The company will use Amazon S3 as part of the solution. The company wants to configure jobs to start automatically when a contract is uploaded to an S3 bucket.

Which solution will meet these requirements?

- A. Use AWS Lambda functions to process individual tasks. Create a primary Lambda function to handle the overall job processing by calling individual Lambda functions in sequence. Configure the S3 bucket to send an event notification to invoke the primary Lambda function to begin processing.
- B. Use a state machine in AWS Step Functions to handle the overall contract processing job. Configure the S3 bucket to send an event notification to Amazon EventBridge. Create a rule in Amazon EventBridge to target the state machine.
- C. Use an AWS Batch job to handle the overall contract processing job. Configure the S3 bucket to send an event notification to initiate the Batch job.
- D. Use an S3 event notification to notify an Amazon Simple Queue Service (Amazon SQS) queue when a contract is uploaded. Configure an AWS Lambda function to read messages from the queue and to run the contract processing job.

Answer: B

Explanation:

AWS Step Functions supports long-running workflows and error retries, making it ideal for a job composed of many tasks. Integration with EventBridge allows automatic triggering from S3 events. This setup is resilient and supports up to 1-year execution duration.

Reference: AWS Documentation - AWS Step Functions with Amazon EventBridge for Long-Running Workflows

Question: 288

A company wants to visualize its AWS spend and resource usage. The company wants to use an AWS managed service to provide visual dashboards.

Which solution will meet these requirements?

- A. Configure an export in AWS Data Exports. Use Amazon QuickSight to create a cost and usage dashboard. View the data in QuickSight.
- B. Configure one custom budget in AWS Budgets for costs. Configure a second custom budget for usage. Schedule daily AWS Budgets reports by using the two budgets as sources.
- C. Configure AWS Cost Explorer to use user-defined cost allocation tags with hourly granularity to generate detailed data.
- D. Configure an export in AWS Data Exports. Use the standard export option. View the data in Amazon Athena.

Answer: A

Explanation:

By exporting AWS Cost and Usage Reports (CUR) to Amazon S3 and analyzing them with Amazon QuickSight, companies can generate interactive visual dashboards. This solution is fully AWS- managed, requires no third-party tools, and integrates deeply with AWS cost data.

Reference: AWS Documentation - Using QuickSight with CUR for AWS Cost Visualization

Question: 289

A company hosts an application on AWS. The application has generated approximately 2.5 TB of data over the previous 12 years. The company currently stores the data on Amazon EBS volumes.

The company wants a cost-effective backup solution for long-term storage. The company must be able to retrieve the data within minutes when required for audits.

Which solution will meet these requirements?

- A. Create EBS snapshots to back up the data.
- B. Create an Amazon S3 bucket. Use the S3 Glacier Deep Archive storage class to back up the data.
- C. Create an Amazon S3 bucket. Use the S3 Glacier Flexible Retrieval storage class to back up the data.
- D. Create an Amazon Elastic File System (Amazon EFS) file system to back up the data.

Answer: C

Explanation:

Amazon S3 Glacier Flexible Retrieval is a low-cost archival storage class that supports retrieval of data within minutes (expedited access ~1-5 minutes), making it ideal for audit scenarios where occasional, quick access to archived data is required. In contrast, Glacier Deep Archive takes hours to retrieve.

Reference: AWS Documentation - Amazon S3 Glacier Flexible Retrieval

Question: 290

A company is developing a new online gaming application. The application will run on Amazon EC2 instances in multiple AWS Regions and will have a high number of globally distributed users. A solutions architect must design the application to optimize network latency for the users.

Which actions should the solutions architect take to meet these requirements? (Select TWO.)

- A. Configure AWS Global Accelerator. Create Regional endpoint groups in each Region where an EC2 fleet is hosted.
- B. Create a content delivery network (CDN) by using Amazon CloudFront. Enable caching for static and dynamic content, and specify a high expiration period.
- C. Integrate AWS Client VPN into the application. Instruct users to select which Region is closest to them after they launch the application. Establish a VPN connection to that Region.
- D. Create an Amazon Route 53 weighted routing policy. Configure the routing policy to give the highest weight to the EC2 instances in the Region that has the largest number of users.
- E. Configure an Amazon API Gateway endpoint in each Region where an EC2 fleet is hosted. Instruct users to select which Region is closest to them after they launch the application. Use the API Gateway endpoint that is closest to them.

Answer: A,B

Explanation:

AWS Global Accelerator reduces latency by directing users to the optimal Regional endpoint based on global network health and proximity. Amazon CloudFront caches static and dynamic content at edge locations for ultra-low latency access worldwide, improving performance and reducing server load.

Reference: AWS Documentation - AWS Global Accelerator and Amazon CloudFront for Latency Optimization

Question: 291

A company runs a web application in a single AWS Region. A solutions architect wants to ensure that the web application can continue to operate if the application becomes unavailable in the Region.

Which solution will meet this requirement?

- A. Deploy the application in multiple Regions. Use Amazon Route 53 DNS health checks to route traffic to a healthy Region.
- B. Deploy the application in multiple Availability Zones within a single Region. Use Amazon Route 53 DNS health checks to route traffic to healthy application resources.
- C. Deploy the application in multiple Regions. Use an Amazon Route 53 simple routing record to route traffic to a healthy Region.
- D. Deploy the application in multiple Availability Zones within a single Region. Use an Amazon Route 53 latency record in each Availability Zone to route traffic to a healthy Availability Zone.

Answer: A

Explanation:

To protect against a Regional failure, the application must be deployed in multiple Regions. Amazon Route 53 DNS failover with health checks allows traffic to be automatically routed to a healthy Region when the primary becomes unavailable, meeting high availability and disaster recovery requirements.

Reference: AWS Well-Architected Framework - Multi-Region High Availability with Route 53 Health Checks

Question: 292

A company uses Amazon S3 to host its static website. The company wants to add a contact form to the webpage. The contact form will have dynamic server-side components for users to input their name, email address, phone number, and user message.

The company expects fewer than 100 site visits each month. The contact form must notify the company by email when a customer fills out the form.

Which solution will meet these requirements MOST cost-effectively?

- A. Host the dynamic contact form in Amazon Elastic Container Service (Amazon ECS). Set up Amazon Simple Email Service (Amazon SES) to connect to a third-party email provider.
- B. Create an Amazon API Gateway endpoint that returns the contact form from an AWS Lambda function. Configure another Lambda function on the API Gateway to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Host the website by using AWS Amplify Hosting for static content and dynamic content. Use server-side scripting to build the contact form. Configure Amazon Simple Queue Service (Amazon SQS) to deliver the message to the company.

D. Migrate the website from Amazon S3 to Amazon EC2 instances that run Windows Server. Use Internet Information Services (IIS) for Windows Server to host the webpage. Use client-side scripting to build the contact form. Integrate the form with Amazon WorkMail.

Answer: B

Explanation:

Using API Gateway and Lambda enables serverless handling of form submissions with minimal cost and infrastructure. When coupled with Amazon SNS, it allows instant email notifications without running servers, making it ideal for low-traffic workloads.

Reference: AWS Documentation - Serverless Contact Form with API Gateway, Lambda, and SNS

Question: 293

A company is migrating a large amount of data from on-premises storage to AWS. Windows, Mac, and Linux based Amazon EC2 instances in the same AWS Region will access the data by using SMB and NFS storage protocols. The company will access a portion of the data routinely. The company will access the remaining data infrequently.

The company needs to design a solution to host the data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon Elastic File System (Amazon EFS) volume that uses EFS Intelligent-Tiering. Use AWS DataSync to migrate the data to the EFS volume.
- B. Create an Amazon FSx for ONTAP instance. Create an FSx for ONTAP file system with a root volume that uses the auto tiering policy. Migrate the data to the FSx for ONTAP volume.
- C. Create an Amazon S3 bucket that uses S3 Intelligent-Tiering. Migrate the data to the S3 bucket by using an AWS Storage Gateway Amazon S3 File Gateway.
- D. Create an Amazon FSx for OpenZFS file system. Migrate the data to the new volume.

Answer: B

Explanation:

Amazon FSx for ONTAP supports both NFS and SMB protocols and includes automated tiering between SSD and capacity pool storage, optimizing cost and performance. It is ideal for mixed operating systems and varied access patterns with minimal administrative overhead.

Question: 294

An application uses an Amazon SQS queue and two AWS Lambda functions. One of the Lambda functions pushes messages to the queue, and the other function polls the queue and receives **queued messages**.

A solutions architect needs to ensure that only the two Lambda functions can write to or read from the queue.

Which solution will meet these requirements?

- A. Attach an IAM policy to the SQS queue that grants the Lambda function principals read and write access. Attach an IAM policy to the execution role of each Lambda function that denies all access to the SQS queue except for the principal of each function.
- B. Attach a resource-based policy to the SQS queue to deny read and write access to the queue for any entity except the principal of each Lambda function. Attach an IAM policy to the execution role of each Lambda function that allows read and write access to the queue.
- C. Attach a resource-based policy to the SQS queue that grants the Lambda function principals read and write access to the queue. Attach an IAM policy to the execution role of each Lambda function that allows read and write access to the queue.
- D. Attach a resource-based policy to the SQS queue to deny all access to the queue. Attach an IAM policy to the execution role of each Lambda function that grants read and write access to the queue.

Answer: C

Explanation:

To ensure that only specific AWS Lambda functions can read from or write to an Amazon SQS queue, use resource-based policies attached directly to the SQS queue. These policies explicitly grant permissions to the IAM roles used by the Lambda functions. Additionally, the Lambda execution roles must also have IAM policies that permit SQS access. This dual-layer approach follows the AWS security best practice of granting least privilege access and ensures that no other service or entity can interact with the queue.

This is a common and supported pattern documented in the Amazon SQS Developer Guide, where resource-based policies restrict access at the queue level while IAM roles control permissions at the function level.

Reference: AWS Documentation - Amazon SQS Access Control, Lambda Permissions, and ResourceBased Policies

Question: 295

A company wants to relocate its on-premises MySQL database to AWS. The database accepts regular imports from a client-facing application, which causes a high volume of write operations. The company is concerned that the amount of traffic might be causing performance issues within the application.

- A. Provision an Amazon RDS for MySQL DB instance with Provisioned IOPS SSD storage. Monitor write operation metrics by using Amazon CloudWatch. Adjust the provisioned IOPS if necessary.
- B. Provision an Amazon RDS for MySQL DB instance with General Purpose SSD storage. Place an Amazon ElastiCache cluster in front of the DB instance. Configure the application to query ElastiCache instead.
- C. Provision an Amazon DocumentDB (with MongoDB compatibility) instance with a memory- optimized instance type. Monitor Amazon CloudWatch for performance-related issues. Change the instance class if necessary.
- D. Provision an Amazon Elastic File System (Amazon EFS) file system in General Purpose performance mode. Monitor Amazon CloudWatch for IOPS bottlenecks. Change to Provisioned Throughput performance mode if necessary.

Answer: A

Explanation:

For a MySQL database experiencing high write operations, using Amazon RDS with Provisioned IOPS (io1 or io2) SSD storage is recommended to achieve consistent and low-latency performance. Provisioned IOPS allows you to specify a desired IOPS rate, which is crucial for write-intensive workloads.

Monitoring write operation metrics through Amazon CloudWatch enables you to observe performance and adjust the provisioned IOPS as needed to meet application demands.

Reference:

Modifying settings for Provisioned IOPS SSD storage [AWS Documentation](#)

Amazon CloudWatch metrics for Amazon RDS [AWS Documentation](#)

Question: 296

A company wants to migrate its on-premises Oracle database to Amazon Aurora. The company wants to use a secure and encrypted network to transfer the data. Which combination of steps will meet these requirements? (Select TWO.)

- A. Use AWS Application Migration Service to migrate the data.
- B. Use AWS Schema Conversion Tool (AWS SCT) and AWS Database Migration Service (AWS DMS) to migrate the data.
- C. Use AWS Direct Connect SiteLink to transfer data from the on-premises environment to AWS.
- D. Use AWS Site-to-Site VPN to establish a connection to transfer the data from the on-premises environment to AWS.
- E. Use AWS App2Container to migrate the data.

Answer: B,D

Explanation:

To securely migrate an on-premises Oracle database to Amazon Aurora, the following steps are recommended:

Use AWS Schema Conversion Tool (AWS SCT) and AWS Database Migration Service (AWS DMS): AWS SCT helps convert the source database schema to a format compatible with the target database (Aurora). AWS DMS facilitates the actual data migration, ensuring minimal downtime and data integrity.

Use AWS Site-to-Site VPN: Establishing a Site-to-Site VPN connection provides a secure and encrypted tunnel between the on-premises environment and AWS. This ensures that data transferred during the migration is protected against interception and unauthorized access.

Reference:

Migrate an Oracle database to Aurora PostgreSQL using AWS DMS and AWS SCT

Step 2: Configure Your Source Database - AWS Documentation

Question: 297

A company has a single AWS account that contains resources belonging to several teams. The company needs to identify the costs associated with each team. The company wants to use a tag named CostCenter to identify resources that belong to each team.

- A. Tag all resources that belong to each team with the user-defined CostCenter tag.
- B. Create a tag for each team, and set the value to CostCenter.
- C. Activate the CostCenter tag to track cost allocation.
- D. Configure AWS Billing and Cost Management to send monthly invoices to the company through email messages.
- E. Set up consolidated billing in the existing AWS account.

Answer: A,C

Explanation:

To track costs associated with different teams within a single AWS account, the company should implement user-defined cost allocation tags.

Tagging Resources with CostCenter: Assign the CostCenter tag to all resources, specifying the appropriate value for each team. This practice enables the organization to categorize and track AWS costs effectively.

Activating the CostCenter Tag: After tagging resources, activate the CostCenter tag in the AWS Billing and Cost Management console. Activation is necessary for the tags to appear in cost allocation reports and AWS Cost Explorer.

By following these steps, the company can generate detailed billing reports that break down costs by team, facilitating better cost management and accountability.

Reference:

Organizing and tracking costs using AWS cost allocation tags

Activating user-defined cost allocation tags

Question: 298

A company recently migrated its application to AWS. The application runs on Amazon EC2 Linux instances in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon Elastic File System (Amazon EFS) file system that uses EFS Standard-Infrequent Access storage. The application indexes the company's files, and the index is stored in an Amazon RDS database.

The company needs to optimize storage costs with some application and services changes.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon S3 bucket that uses an Intelligent-Tiering lifecycle policy. Copy all files to the S3 bucket. Update the application to use Amazon S3 API to store and retrieve files.
- B. Deploy Amazon FSx for Windows File Server file shares. Update the application to use CIFS protocol to store and retrieve files.
- C. Deploy Amazon FSx for OpenZFS file system shares. Update the application to use the new mount point to store and retrieve files.
- D. Create an Amazon S3 bucket that uses S3 Glacier Flexible Retrieval. Copy all files to the S3 bucket. Update the application to use Amazon S3 API to store and retrieve files as standard retrievals.

Answer: A

Explanation:

To optimize storage costs, migrating data from Amazon EFS to Amazon S3 with the Intelligent-Tiering storage class is a cost-effective solution.

Amazon S3 Intelligent-Tiering: This storage class automatically moves data between frequent, infrequent, and archive access tiers based on access patterns, reducing storage costs without impacting performance.

Cost Savings: By leveraging Intelligent-Tiering, the company can achieve significant cost savings, especially for data with

unpredictable access patterns.

Application Update: The application must be updated to interact with Amazon S3 APIs for storing and retrieving files, ensuring seamless integration with the new storage solution.

This approach provides a scalable, durable, and cost-effective storage solution, aligning with the company's optimization goals.

Reference:

Amazon S3 Intelligent-Tiering Storage Class

Managing storage costs with Amazon S3 Intelligent-Tiering

Question: 299

A company uses AWS to run its e-commerce platform, which is critical to its operations and experiences a high volume of traffic and transactions. The company has configured a multi-factor authentication (MFA) device to secure its AWS account root user credentials. The company wants to ensure that it will not lose access to the root user account if the MFA device is lost.

Which solution will meet these requirements?

- A. Set up a backup administrator account that the company can use to log in if the company loses the MFA device.
- B. Add multiple MFA devices for the root user account to handle the disaster scenario.
- C. Create a new administrator account when the company cannot access the root account.
- D. Attach the administrator policy to another IAM user when the company cannot access the root account.

Answer: B

Explanation:

To prevent loss of access to the AWS root user account due to a lost MFA device, AWS recommends enabling multiple MFA devices for the root user.

Multiple MFA Devices: AWS allows up to eight MFA devices (virtual, hardware, or security keys) to be associated with a single root user account. This redundancy ensures that if one device is lost, others can be used to authenticate and access the account.

Security Best Practices: Implementing multiple MFA devices enhances account security and provides resilience against potential access issues.

By proactively setting up multiple MFA devices, the company can maintain uninterrupted access to its critical AWS resources, even in the event of a lost or malfunctioning MFA device.

Reference:

Multi-factor authentication for AWS account root user

AWS Multi-factor authentication in IAM

Question: 300

A financial services company must retain log data for 1 year. The company stores log files in an Amazon S3 bucket and wants to prevent any user from deleting or overwriting the log files during this period. The data must remain available for read-only requests.

- A. Enable S3 Versioning on the bucket. Use Object Lock in compliance mode with a 1-year retention period.
- B. Enable S3 Transfer Acceleration on the bucket. Create an S3 Lifecycle Configuration rule to move objects to Amazon S3 Glacier Flexible Retrieval after 1 year.
- C. Enable S3 Versioning on the bucket. Create an S3 Lifecycle Configuration rule to move objects to Amazon S3 Glacier Flexible Retrieval after 1 year.
- D. Create an AWS Lambda function to programmatically check the timestamp of S3 data and to move the data to Amazon S3 Glacier Deep Archive if the data is older than 1 year.

Answer: A

Explanation:

To ensure that log files are immutable and cannot be deleted or overwritten for a specified retention period, Amazon S3 offers Object Lock in Compliance Mode. When enabled:

Compliance Mode: Ensures that a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account, for the duration of the retention period. [AWS Documentation](#)

S3 Versioning: Must be enabled on the bucket to use Object Lock. This allows multiple versions of an object to be stored, ensuring that previous versions are preserved and protected. [Amazon Web Services, Inc.](#)

By enabling S3 Versioning and applying Object Lock in Compliance Mode with a 1-year retention period, the company can meet regulatory requirements to retain log data securely and prevent any modifications or deletions during that time.

Question: 301

A company has an ordering application that stores customer information in Amazon RDS for MySQL. During regular business hours, employees run one-time queries for reporting purposes. Timeouts are occurring during order processing because the reporting queries are taking a long time to run. The company needs to eliminate the timeouts without preventing employees from performing queries.

- A. Create a read replica. Move reporting queries to the read replica.
- B. Create a read replica. Distribute the ordering application to the primary DB instance and the read replica.
- C. Migrate the ordering application to Amazon DynamoDB with on-demand capacity.
- D. Schedule the reporting queries for non-peak hours.

Answer: A

Explanation:

Amazon RDS for MySQL supports the creation of read replicas, which are read-only copies of the primary database instance. By offloading read-heavy operations, such as reporting queries, to a read replica:

Performance Improvement: The primary DB instance is relieved from the additional load, reducing the likelihood of timeouts during order processing.

Data Consistency: Read replicas use asynchronous replication, ensuring that they have up-to-date data for accurate reporting.

Scalability: Multiple read replicas can be created to handle increased read traffic.

This approach allows employees to continue running necessary reports without impacting the performance of the ordering application.

Question: 302

A company has an industrial application that controls a process in real time. The company plans to rearchitect the application to distribute jobs across several Amazon EC2 instances in a VPC. The solution needs to maximize the network throughput and minimize the network latency between the instances.

- A. Place the instances in a host-level partition placement group. Choose instance types that support enhanced networking.
- B. Place the instances in several dedicated hosts in the same partition of a partition placement group. Choose dedicated hosts that support enhanced networking.
- C. Place the instances in several dedicated hosts in the same rack of a rack-level placement group. Choose dedicated hosts that support enhanced networking.
- D. Place the instances in a cluster placement group. Choose instance types that support enhanced networking.

Answer: D

Explanation:

For applications requiring high network throughput and low latency between EC2 instances, AWS recommends using a Cluster

Placement Group:

Cluster Placement Group: Packs instances close together inside an Availability Zone, enabling workloads to achieve the low-latency network performance necessary for tightly-coupled node-to- node communication.

Enhanced Networking: Selecting instance types that support enhanced networking provides higher bandwidth, higher packet-per-second (PPS) performance, and consistently lower inter-instance latencies.

By combining a Cluster Placement Group with enhanced networking-enabled instance types, the company can meet the stringent performance requirements of its real-time industrial application.

Question: 303

A company is planning to migrate customer records to an Amazon S3 bucket. The company needs to ensure that customer records are protected against unauthorized access and are encrypted in transit and at rest. The company must monitor all access to the S3 bucket.

- A. Use AWS Key Management Service (AWS KMS) to encrypt customer records at rest. Create an S3 bucket policy that includes the aws:SecureTransport condition. Use an IAM policy to control access to the records. Use AWS CloudTrail to monitor access to the records.
- B. Use AWS Nitro Enclaves to encrypt customer records at rest. Use AWS Key Management Service (AWS KMS) to encrypt the records in transit. Use an IAM policy to control access to the records. Use AWS CloudTrail and AWS Security Hub to monitor access to the records.
- C. Use AWS Key Management Service (AWS KMS) to encrypt customer records at rest. Create an Amazon Cognito user pool to control access to the records. Use AWS CloudTrail to monitor access to the records. Use Amazon GuardDuty to detect threats.
- D. Use server-side encryption with Amazon S3 managed keys (SSE-S3) with default settings to encrypt the records at rest. Access the records by using an Amazon CloudFront distribution that uses the S3 bucket as the origin. Use IAM roles to control access to the records. Use Amazon CloudWatch to monitor access to the records.

Answer: A

Explanation:

Encryption at Rest: AWS Key Management Service (AWS KMS) provides centralized control over the cryptographic keys used to protect data. By using AWS KMS with Amazon S3, you can manage encryption keys and define policies to control access to them.

Encryption in Transit: By enforcing the aws:SecureTransport condition in the S3 bucket policy, you ensure that all data is transmitted over HTTPS, protecting data in transit.

Access Control: IAM policies allow you to define fine-grained permissions for users and roles, ensuring that only authorized entities

can access the customer records.

Monitoring: AWS CloudTrail provides a record of actions taken by a user, role, or AWS service in Amazon S3, enabling you to monitor access to the records.

Reference:

Protecting data using encryption

Using IAM policies to control access to Amazon S3 resources

Logging Amazon S3 API calls using AWS CloudTrail

Question: 304

A company plans to deploy containerized microservices in the AWS Cloud. The containers must mount a persistent file store that the company can manage by using OS-level permissions. The company requires fully managed services to host the containers and file store.

- A. Use AWS Lambda functions and an Amazon API Gateway REST API to handle the microservices. Use Amazon S3 buckets for storage.
- B. Use Amazon EC2 instances to host the microservices. Use Amazon Elastic Block Store (Amazon EBS) volumes for storage.
- C. Use Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate to handle the microservices. Use an Amazon Elastic File System (Amazon EFS) file system for storage.
- D. Use Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate to handle the microservices. Use an Amazon EC2 instance that runs a dedicated file store for storage.

Answer: C

Explanation:

Amazon ECS on AWS Fargate: AWS Fargate is a serverless compute engine for containers that works with Amazon ECS. It allows you to run containers without managing servers or clusters.

Amazon EFS: Amazon Elastic File System (EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It can be mounted to ECS tasks running on Fargate, allowing containers to access a shared file system with standard file system semantics, including OS-level permissions.

Reference:

Using Amazon EFS with Amazon ECS

Amazon EFS: How it works

Question: 305

A company is migrating a data processing application to AWS. The application processes several short-lived batch jobs that cannot be disrupted. The process generates data after each batch job finishes running. The company accesses the data for 30 days following data generation. After 30 days, the company stores the data for 2 years.

The company wants to optimize costs for the application and data storage. Which solution will meet these requirements?

- A. Use Amazon EC2 Spot Instances to run the application. Store the data in Amazon S3 Standard. Move the data to S3 Glacier Instant Retrieval after 30 days. Configure a bucket policy to delete the data after 2 years.
- B. Use Amazon EC2 On-Demand Instances to run the application. Store the data in Amazon S3 Glacier Instant Retrieval. Move the data to S3 Glacier Deep Archive after 30 days. Configure an S3 Lifecycle configuration to delete the data after 2 years.
- C. Use Amazon EC2 Spot Instances to run the application. Store the data in Amazon S3 Standard. Move the data to S3 Glacier Flexible Retrieval after 30 days. Configure a bucket policy to delete the data after 2 years.
- D. Use Amazon EC2 On-Demand Instances to run the application. Store the data in Amazon S3 Standard. Move the data to S3 Glacier Deep Archive after 30 days. Configure an S3 Lifecycle configuration to delete the data after 2 years.

Answer: D

Explanation:

Amazon EC2 On-Demand Instances: Since the batch jobs cannot be disrupted, On-Demand Instances provide the necessary reliability and availability.

Amazon S3 Standard: Storing data in S3 Standard for the first 30 days ensures quick and frequent access.

S3 Glacier Deep Archive: After 30 days, moving data to S3 Glacier Deep Archive significantly reduces storage costs for data that is rarely accessed.

S3 Lifecycle Configuration: Automating the transition and deletion of objects using lifecycle policies ensures cost optimization and compliance with data retention requirements.

Reference:

Amazon S3 Storage Classes

Managing your storage lifecycleAWS Documentation

Question: 306

A company is using an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The company must ensure that Kubernetes service accounts in the EKS cluster have secure and granular access to specific AWS resources by using IAM roles for service accounts

(IRSA).

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Create an IAM policy that defines the required permissions. Attach the policy directly to the IAM role of the EKS nodes.
- B. Implement network policies within the EKS cluster to prevent Kubernetes service accounts from accessing specific AWS services.
- C. Modify the EKS cluster's IAM role to include permissions for each Kubernetes service account. Ensure a one-to-one mapping between IAM roles and Kubernetes roles.
- D. Define an IAM role that includes the necessary permissions. Annotate the Kubernetes service accounts with the Amazon Resource Name (ARN) of the IAM role.
- E. Set up a trust relationship between the IAM roles for the service accounts and an OpenID Connect (OIDC) identity provider.

Answer: D,E

Explanation:

IAM Roles for Service Accounts (IRSA): IRSA allows you to associate an IAM role with a Kubernetes service account. This enables pods to assume the IAM role and access AWS resources securely.

Annotating Service Accounts: By annotating Kubernetes service accounts with the ARN of the IAM role, you establish the association required for IRSA.

OIDC Identity Provider: EKS clusters use OpenID Connect (OIDC) to authenticate service accounts. Setting up a trust relationship between the IAM role and the OIDC provider allows the Kubernetes service account to assume the IAM role.

Question: 307

A genomics research company is designing a scalable architecture for a loosely coupled workload. Tasks in the workload are independent and can be processed in parallel. The architecture needs to minimize management overhead and provide automatic scaling based on demand.

Options:

- A. Use a cluster of Amazon EC2 instances. Use AWS Systems Manager to manage the workload.
- B. Implement a serverless architecture that uses AWS Lambda functions.
- C. Use AWS ParallelCluster to deploy a dedicated high-performance cluster.

D. Implement vertical scaling for each workload task.

Answer: B

Explanation:

For workloads where tasks are independent and can be processed in parallel, and where minimizing management overhead is a priority, a serverless architecture using AWS Lambda is ideal.

AWS Lambda allows you to run code without provisioning or managing servers. It automatically scales your application by running code in response to each trigger.

Parallel Processing: Lambda functions can process multiple tasks concurrently, making it suitable for parallel workloads.

Automatic Scaling: Lambda automatically scales by running code in response to each event, scaling precisely with the size of the workload.

Minimal Management Overhead: With Lambda, there's no need to manage the underlying infrastructure, reducing operational complexity. Wikipedia

Reference:

AWS Lambda - Run Code Without Thinking About Servers

Best Practices for Designing and Architecting with AWS Lambda

Question: 308

A company uses AWS Organizations to manage multiple AWS accounts. Each department in the company has its own AWS account. A security team needs to implement centralized governance and control to enforce security best practices across all accounts. The team wants to have control over which AWS services each account can use. The team needs to restrict access to sensitive resources based on IP addresses or geographic regions. The root user must be protected with multi-factor authentication (MFA) across all accounts.

Options:

A. Use AWS Identity and Access Management (IAM) to manage IAM users and IAM roles in each account. Implement MFA for the root user in each account. Enforce service restrictions by using AWS managed prefix lists.

B. Use AWS Control Tower to establish a multi-account environment. Use service control policies (SCPs) to enforce service restrictions in AWS Organizations. Configure MFA for the root user across all accounts.

C. Use AWS Systems Manager to enforce service restrictions across multiple accounts. Use IAM policies to enforce MFA for the

root user across all accounts.

D. Use AWS IAM Identity Center to manage user access and to enforce service restrictions by using permissions boundaries in each account.

Answer: B

Explanation:

AWS Control Tower provides a straightforward way to set up and govern a secure, multi-account AWS environment based on AWS best practices. It automates the setup of a baseline environment, or **landing zone**, that includes:

Service Control Policies (SCPs): These are used to manage permissions across AWS Organizations, allowing you to set permission guardrails. SCPs can restrict access to specific AWS services and actions, helping enforce security best practices.

Multi-Factor Authentication (MFA): AWS Control Tower can enforce MFA for the root user across all accounts, enhancing security.

Centralized Governance: It offers centralized logging and monitoring, making it easier to manage and audit multiple AWS accounts.

Reference:

AWS Control Tower User Guide

Service Control Policies

Root user best practices for your AWS account

Question: 309

An ecommerce company hosts an API that handles sales requests. The company hosts the API frontend on Amazon EC2 instances that run behind an Application Load Balancer (ALB). The company hosts the API backend on EC2 instances that perform the transactions. The backend tiers are loosely coupled by an Amazon Simple Queue Service (Amazon SQS) queue.

The company anticipates a significant increase in request volume during a new product launch event. The company wants to ensure that the API can handle increased loads successfully.

Options:

A. Double the number of frontend and backend EC2 instances to handle the increased traffic during the product launch event. Create a dead-letter queue to retain unprocessed sales requests when the demand exceeds the system capacity.

B. Place the frontend EC2 instances into an Auto Scaling group. Create an Auto Scaling policy to launch new instances to handle

the incoming network traffic.

C. Place the frontend EC2 instances into an Auto Scaling group. Add an Amazon ElastiCache cluster in front of the ALB to reduce the amount of traffic the API needs to handle.

D. Place the frontend and backend EC2 instances into separate Auto Scaling groups. Create a policy for the frontend Auto Scaling group to launch instances based on incoming network traffic. Create a policy for the backend Auto Scaling group to launch instances based on the SQS queue backlog.

Answer: D

Explanation:

To handle increased loads effectively, it's essential to implement Auto Scaling for both frontend and backend tiers:

Frontend Auto Scaling Group: Scaling based on incoming network traffic ensures that the application can handle increased user requests.

Backend Auto Scaling Group: Scaling based on the Amazon SQS queue backlog ensures that the backend can process messages as they arrive, preventing delays.

This approach allows each tier to scale independently based on its specific load, ensuring optimal resource utilization and performance.

Reference:

Tutorial: Set up a scaled and load-balanced application

Scaling policy based on Amazon SQSAWS Documentation

Question: 310

A machine learning (ML) team is building an application that uses data that is in an Amazon S3 bucket. The ML team needs a storage solution for its model training workflow on AWS. The ML team requires high-performance storage that supports frequent access to training datasets. The storage solution must integrate natively with Amazon S3. Which solution will meet these requirements with the LEAST operational overhead?

Options:

A. Use Amazon Elastic Block Store (Amazon EBS) volumes to provide high-performance storage. Use AWS DataSync to migrate data from the S3 bucket to EBS volumes.

- B. Use Amazon EC2 ML instances to provide high-performance storage. Store training data on Amazon EBS volumes. Use the S3 Copy API to copy data from the S3 bucket to EBS volumes.
- C. Use Amazon FSx for Lustre to provide high-performance storage. Store training datasets in Amazon S3 Standard storage.
- D. Use Amazon EMR to provide high-performance storage. Store training datasets in Amazon S3 Glacier Instant Retrieval storage.

Answer: C

Explanation:

Amazon FSx for Lustre is a high-performance file system optimized for fast processing of workloads such as machine learning, high-performance computing (HPC), and video processing. It integrates natively with Amazon S3, allowing you to:

Access S3 Data: FSx for Lustre can be linked to an S3 bucket, presenting S3 objects as files in the file system.

High Performance: It provides sub-millisecond latencies, high throughput, and millions of IOPS, which are ideal for ML workloads.

Amazon Web Services, Inc.

Minimal Operational Overhead: Being a fully managed service, it reduces the complexity of setting up and managing high-performance file systems.

Reference:

Amazon FSx for Lustre - High-Performance File System Integrated with S3 Amazon Web Services, Inc.

What is Amazon FSx for Lustre?

Question: 311

A company operates an online photo-sharing service and stores data in AWS Account A in a centralized Amazon S3 bucket. The company wants to grant a second AWS account named Account B access to the centralized S3 bucket. The company owns Account B.

Options:

- A. Enable S3 Transfer Acceleration to provide Account B access to the centralized S3 bucket in Account A.
- B. Enable cross-Region replication between Account A and Account B to share the S3 bucket data.
- C. Use Amazon CloudFront to distribute the S3 bucket contents. Grant Account B access to the bucket contents through a signed URL.
- D. Create a bucket policy that grants Account B permission to access the centralized S3 bucket in Account A.

Answer: D

Explanation:

To grant cross-account access to an Amazon S3 bucket, you can use a bucket policy that specifies the AWS account ID of the account you want to grant access to. This method allows Account B to access the S3 bucket in Account A without the need for additional services or configurations.

Reference: Example 2: Bucket owner granting cross-account bucket permissions

Question: 312

A company wants to migrate an application that uses a microservice architecture to AWS. The services currently run on Docker containers on-premises. The application has an event-driven architecture that uses Apache Kafka. The company configured Kafka to use multiple queues to send and receive messages. Some messages must be processed by multiple services. Which solution will meet these requirements with the LEAST management overhead?

Options:

- A. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Deploy a Kafka cluster on EC2 instances to handle service-to-service communication.
- B. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Create multiple Amazon Simple Queue Service (Amazon SQS) queues to handle service-to-service communication.
- C. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Deploy an Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster to handle service-to-service communication.
- D. Migrate the services to Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Use Amazon EventBridge to handle service-to-service communication.

Answer: C

Explanation:

Amazon Managed Streaming for Apache Kafka (Amazon MSK) is a fully managed service that makes it easy to build and run applications that use Apache Kafka to process streaming data. By using

Amazon ECS with the AWS Fargate launch type, you can run containers without managing servers or clusters. This combination reduces operational overhead and provides scalability.

Reference: Power your Kafka Streams application with Amazon MSK and AWS FargateAmazon Web Services, Inc.

Question: 313

A company wants to deploy an internal web application on AWS. The web application must be accessible only from the company's office. The company needs to download security patches for the web application from the internet. The company has created a VPC and has configured an AWS Site- to-Site VPN connection to the company's office. A solutions architect must design a secure architecture for the web application. Which solution will meet these requirements?

Options:

- A. Deploy the web application on Amazon EC2 instances in public subnets behind a public Application Load Balancer (ALB). Attach an internet gateway to the VPC. Set the inbound source of the ALB's security group to 0.0.0.0/0.
- B. Deploy the web application on Amazon EC2 instances in private subnets behind an internal Application Load Balancer (ALB). Deploy NAT gateways in public subnets. Attach an internet gateway to the VPC. Set the inbound source of the ALB's security group to the company's office network CIDR block.
- C. Deploy the web application on Amazon EC2 instances in public subnets behind an internal Application Load Balancer (ALB). Deploy NAT gateways in private subnets. Attach an internet gateway to the VPC. Set the outbound destination of the ALB's security group to the company's office network CIDR block.
- D. Deploy the web application on Amazon EC2 instances in private subnets behind a public Application Load Balancer (ALB). Attach an internet gateway to the VPC. Set the outbound destination of the ALB's security group to 0.0.0.0/0.

Explanation:

Deploying the web application on EC2 instances in private subnets behind an internal ALB ensures that the application is not directly accessible from the internet. By setting up a Site-to-Site VPN connection, the application can be accessed securely from the company's office network. Deploying NAT gateways in public subnets allows instances in private subnets to initiate outbound connections to the internet for downloading security patches.

Reference: Example: VPC with servers in private subnets and NATAWS Documentation

Question: 314

A company wants to migrate an application to AWS. The application runs on Docker containers behind an Application Load Balancer (ALB). The application stores data in a PostgreSQL database. The cloud-based solution must use AWS WAF to inspect all application traffic. The application experiences most traffic on weekdays. There is significantly less traffic on weekends. Which solution will meet these requirements in the MOST cost-effective way?

Options:

- A. Use a Network Load Balancer (NLB). Create a web access control list (web ACL) in AWS WAF that includes the necessary rules. Attach the web ACL to the NLB. Run the application on Amazon Elastic Container Service (Amazon ECS). Use Amazon RDS for PostgreSQL as the database.
- B. Create a web access control list (web ACL) in AWS WAF that includes the necessary rules. Attach the web ACL to the ALB. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon RDS for PostgreSQL as the database.
- C. Create a web access control list (web ACL) in AWS WAF that includes the necessary rules. Attach the web ACL to the ALB. Run the application on Amazon Elastic Container Service (Amazon ECS). Use Amazon Aurora Serverless as the database.
- D. Use a Network Load Balancer (NLB). Create a web access control list (web ACL) in AWS WAF that has the necessary rules. Attach the web ACL to the NLB. Run the application on Amazon Elastic Container Service (Amazon ECS). Use Amazon Aurora Serverless as the database.

Answer: C

Explanation:

Using an Application Load Balancer (ALB) allows for integration with AWS WAF to inspect all incoming traffic. Running the application on Amazon ECS provides a scalable and managed container orchestration service. Utilizing Amazon Aurora Serverless for the PostgreSQL database offers automatic scaling based on application demand, which is cost-effective for workloads with variable traffic patterns, such as higher traffic on weekdays and lower traffic on weekends.

Reference: Optimizing cost savings: The advantage of Amazon Aurora over self-managed opensource databases

Question: 315

A company hosts a public application on AWS. The company uses an Application Load Balancer (ALB) to distribute application traffic to multiple Amazon EC2 instances that are hosted in private subnets.

The company wants to authenticate all the requests by using an on-premises Active Directory Federation Service (AD FS). The company

uses AWS Direct Connect to connect its on-premises data center to AWS.

Which solution will meet this requirement?

- A. Configure an Amazon Cognito user pool. Integrate the user pool with the ALB for AD FS authentication.
- B. Configure an AWS Directory Service directory. Integrate the directory with the ALB for AD FS authentication.
- C. Replace the ALB with a Network Load Balancer (NLB). Use Amazon Connect Agent Workspace to integrate an agent workspace with the NLB.
- D. Configure an AWS Directory Service AD Connector. Integrate the AD Connector with the ALB for AD FS authentication.

Answer: D

Explanation:

To authenticate users using an on-premises Active Directory Federation Service (AD FS), AWS provides the AWS Directory Service AD Connector. AD Connector is a proxy that connects AWS applications to your on-premises Microsoft Active Directory without requiring complex directory synchronization or the need to set up a separate directory in the cloud.

By integrating AD Connector with the Application Load Balancer (ALB), you can authenticate and authorize users using your existing on-premises credentials. This setup allows the ALB to leverage the authentication capabilities of your on-premises AD FS, providing a seamless and secure user experience.

Reference:

How to connect your on-premises Active Directory to AWS using AD Connector

AWS Directory Service AD Connector

Question: 316

A company runs an online order management system on AWS. The company stores order and inventory data for the previous 5 years in an Amazon Aurora MySQL database. The company deletes inventory data after 5 years.

The company wants to optimize costs to archive data.

Options:

- A. Create an AWS Glue crawler to export data to Amazon S3. Create an AWS Lambda function to compress the data.
- B. Use the SELECT INTO OUTFILE S3 query on the Aurora database to export the data to Amazon S3. Configure S3 Lifecycle rules on

the S3 bucket.

- C. Create an AWS Glue DataBrew Job to migrate data from Aurora to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.
- D. Use the AWS Schema Conversion Tool (AWS SCT) to replicate data from Aurora to Amazon S3. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.

Answer: B

Explanation:

The SELECT INTO OUTFILE S3 feature allows you to export Amazon Aurora MySQL data directly to Amazon S3 with minimal operational overhead. This method is efficient and cost-effective for archiving historical data.

You can configure S3 Lifecycle rules to transition the exported data to lower-cost storage (e.g., S3 Glacier or S3 Standard-IA) and eventually delete it after 5 years.

No need for additional ETL tools like Glue or DataBrew unless complex transformations are required.

Reference:

Exporting data from Aurora MySQL to S3

Question: 317

A company runs a mobile game app that stores session data (up to 256 KB) for up to 48 hours. The data updates frequently and must be deleted automatically after expiration. Restorability is also required.

Options:

- A. Use an Amazon DynamoDB table to store the session data. Enable point-in-time recovery (PITR) and TTL.
- B. Use Amazon MemoryDB and enable PITR and TTL.
- C. Store session data in S3 Standard. Enable Versioning and a Lifecycle rule to expire objects after 48 hours.
- D. Store data in S3 Intelligent-Tiering with Versioning and a Lifecycle rule to expire after 48 hours.

Answer: A

Explanation:

Amazon DynamoDB supports both TTL (Time to Live) for automatic deletion and Point-in-Time Recovery (PITR) for backup and restore of table data. It's ideal for frequently updated data with short retention requirements and minimal operational overhead.

TTL attribute ensures expired items are automatically removed.

PITR enables recovery from accidental deletes or application errors.

Reference:

DynamoDB TTL

DynamoDB PITR

Question: 318

A finance company collects streaming data for a real-time search and visualization system. They want to migrate to AWS using a native solution for ingest, search, and visualization.

Options:

- A. Use EC2 to ingest/process data to S3 Athena + Managed Grafana
- B. Use EMR to ingest/process to Redshift Redshift Spectrum + QuickSight
- C. Use EKS to ingest/process to DynamoDB CloudWatch Dashboards
- D. Use Kinesis Data Streams Amazon OpenSearch Service Amazon QuickSight

Answer: D

Explanation:

This is a classic use case for Amazon Kinesis Data Streams + OpenSearch + QuickSight:

Kinesis Data Streams: For real-time ingestion and processing

OpenSearch Service: For fast full-text search, indexing, and analysis

QuickSight: For rich dashboard visualizations

This stack is fully managed, scalable, and native to AWS, minimizing operational overhead.

Reference:

Kinesis + OpenSearch architecture

Question: 319

A company uses Apache Hadoop and Spark on-prem. The infrastructure is complex and not scalable. They want to reduce operational complexity but keep data processing on-premises.

Options:

- A. Use Site-to-Site VPN to access on-prem HDFS. Use Amazon EMR to process the data.
- B. Use AWS DataSync to connect to on-prem HDFS. Use Amazon EMR to process the data.
- C. Migrate to Amazon EMR on AWS Outposts.
- D. Use AWS Snowball to migrate data to S3. Use EMR to process.

Answer: C

Explanation:

AWS Outposts brings native AWS services (including Amazon EMR) on-premises, ideal when data residency or latency constraints require local processing.

You benefit from AWS's managed services while meeting the requirement to keep data processing local.

Reference:

Amazon EMR on AWS Outposts

Question: 320

A company recently migrated a large amount of research data to an Amazon S3 bucket. The company needs an automated solution to identify sensitive data in the bucket. A security team also needs to monitor access patterns for the data 24 hours a day, 7 days a week to identify suspicious activities or evidence of tampering with security controls.

Options:

- A. Set up AWS CloudTrail reporting, and grant the security team read-only access to the CloudTrail reports. Set up an Amazon S3 Inventory report to identify sensitive data. Review the findings with the security team.

- B. Enable Amazon Macie and Amazon GuardDuty on the account. Grant the security team access to Macie and GuardDuty. Review the findings with the security team.
- C. Set up an Amazon S3 Inventory report. Use Amazon Athena and Amazon QuickSight to identify sensitive data. Create a dashboard for the security team to review findings.
- D. Use AWS Identity and Access Management (IAM) Access Advisor to monitor for suspicious activity and tampering. Create a dashboard for the security team. Set up an Amazon S3 Inventory report to identify sensitive data. Review the findings with the security team.

Answer: B

Explanation:

To automatically identify sensitive data in Amazon S3 and monitor access patterns for suspicious activities:

Amazon Macie uses machine learning and pattern matching to discover and protect sensitive data in S3. It provides visibility into data security risks and enables automated protection against those risks.

Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect AWS accounts and workloads. It analyzes events from AWS CloudTrail, VPC Flow Logs, and DNS logs.

By enabling both services, the company can automate the discovery of sensitive data and continuously monitor access patterns for potential security threats.

Question: 321

A healthcare company uses an Amazon EMR cluster to process patient data. The data must be encrypted in transit and at rest. Local volumes in the cluster also need to be encrypted. Which solution will meet these requirements?

Options:

- A. Create Amazon EBS volumes. Enable encryption. Attach the volumes to the existing EMR cluster.
- B. Create an EMR security configuration that encrypts the data and the volumes as required.
- C. Create an EC2 instance profile for the EMR instances. Configure the instance profile to enforce encryption.
- D. Create a runtime role that has a trust policy for the EMR cluster.

Answer: B

Explanation:

Amazon EMR allows the creation of security configurations to specify settings for encrypting data at rest, data in transit, or both. These configurations can be applied to clusters to ensure that data stored in Amazon S3, local disks, and data moving between nodes is encrypted.

By creating and applying an EMR security configuration, the company can ensure that all data processing complies with encryption requirements for sensitive patient data.

Question: 322

A company is building an ecommerce application that uses a relational database to store customer data and order history. The company also needs a solution to store 100 GB of product images. The company expects the traffic flow for the application to be predictable. Which solution will meet these requirements MOST cost-effectively?

Options:

- A. Use Amazon RDS for MySQL for the database. Store the product images in an Amazon S3 bucket.
- B. Use Amazon DynamoDB for the database. Store the product images in an Amazon S3 bucket.
- C. Use Amazon RDS for MySQL for the database. Store the product images in an Amazon Aurora MySQL database.
- D. Create three Amazon EC2 instances. Install MongoDB software on the instances to use as the database. Store the product images in an Amazon RDS for MySQL database with a Multi-AZ deployment.

Answer: A

Explanation:

For predictable workloads requiring a relational database, Amazon RDS for MySQL offers a managed, cost-effective solution. Storing product images in Amazon S3 is highly durable and cost-efficient, especially for static assets like images. This combination leverages managed services to reduce operational overhead and costs.

Question: 323

A company runs an application on several Amazon EC2 instances that store persistent data on an Amazon Elastic File System (Amazon EFS) file system. The company needs to replicate the data to another AWS Region by using an AWS managed service solution. Which solution will meet these requirements MOST cost-effectively?

Options:

- A. Use the EFS-to-EFS backup solution to replicate the data to an EFS file system in another Region.

- B. Run a nightly script to copy data from the EFS file system to an Amazon S3 bucket. Enable S3 CrossRegion Replication on the S3 bucket.
- C. Create a VPC in another Region. Establish a cross-Region VPC peer. Run a nightly rsync to copy data from the original Region to the new Region.
- D. Use AWS Backup to create a backup plan with a rule that takes a daily backup and replicates it to another Region. Assign the EFS file system resource to the backup plan.

Answer: D

Explanation:

AWS Backup supports cross-Region backup for Amazon EFS, allowing automated, scheduled backups and replication to another Region. This managed service simplifies backup management and ensures data resilience without the need for custom scripts or manual processes

Question: 324

A company runs a monolithic application in its on-premises data center. The company used Java/Tomcat to build the application. The application uses Microsoft SQL Server as a database.

The company wants to migrate the application to AWS.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Use AWS App2Container to containerize the application. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS). Deploy the database to Amazon RDS for SQL Server. Configure a Multi-AZ deployment.
- B. Containerize the application and deploy the application on a self-managed Kubernetes cluster on an Amazon EC2 instance. Deploy the database on a separate EC2 instance. Set up Microsoft SQL Server Always On availability groups.
- C. Deploy the frontend of the web application as a website on Amazon S3. Use Amazon DynamoDB for the database tier.
- D. Use AWS App2Container to containerize the application. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon DynamoDB for the database tier.

Answer: A

Explanation:

AWS App2Container is a command-line tool that helps containerize existing Java and .NET applications running on-premises or on virtual machines, with minimal refactoring. By using Amazon EKS, the company benefits from a managed Kubernetes service, which significantly reduces the operational overhead compared to managing Kubernetes on EC2.

Amazon RDS for SQL Server provides a fully managed SQL Server database engine with automated backups, patching, and high availability through Multi-AZ deployments. This eliminates the need for the company to manage database infrastructure and software manually.

Overall, option A provides the most streamlined and managed approach for both the application and database layers with the least operational effort.

Question: 325

A company uses a single Amazon S3 bucket to store data that multiple business applications must access. The company hosts the applications on Amazon EC2 Windows instances that are in a VPC. The company configured a bucket policy for the S3 bucket to grant the applications access to the bucket.

The company continually adds more business applications to the environment. As the number of business applications increases, the policy document becomes more difficult to manage. The S3 bucket policy document will soon reach its policy size quota. The company needs a solution to scale its architecture to handle more business applications.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Migrate the data from the S3 bucket to an Amazon Elastic File System (Amazon EFS) volume. Ensure that all application owners configure their applications to use the EFS volume.
- B. Deploy an AWS Storage Gateway appliance for each application. Reconfigure the applications to use a dedicated Storage Gateway appliance to access the S3 objects instead of accessing the objects directly.
- C. Create a new S3 bucket for each application. Configure S3 replication to keep the new buckets synchronized with the original S3 bucket. Instruct application owners to use their respective S3 buckets.
- D. Create an S3 access point for each application. Instruct application owners to use their respective S3 access points.

Answer: D

Explanation:

Amazon S3 Access Points simplify managing data access for shared datasets in S3 by allowing the creation of distinct access policies for different applications or users. Each access point has its own policy and can be managed independently. This method avoids overloading a single bucket policy and helps remain within policy size limits.

Option D provides a scalable and operationally efficient solution by offloading individual access controls from a central bucket policy to individually managed access points, which is ideal for environments with many consuming applications.

Question: 326

A company needs to store confidential files on AWS. The company accesses the files every week. The company must encrypt the files by using envelope encryption, and the encryption keys must be rotated automatically. The company must have an audit trail to monitor encryption key usage.

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Store the confidential files in Amazon S3.
- B. Store the confidential files in Amazon S3 Glacier Deep Archive.
- C. Use server-side encryption with customer-provided keys (SSE-C).
- D. Use server-side encryption with Amazon S3 managed keys (SSE-S3).
- E. Use server-side encryption with AWS KMS managed keys (SSE-KMS).

Answer: A,E

Explanation:

Amazon S3 is suitable for storing data that needs to be accessed weekly and integrates with AWS Key Management Service (KMS) to provide encryption at rest with server-side encryption using KMS- managed keys (SSE-KMS).

SSE-KMS uses envelope encryption and allows automatic key rotation and logging through AWS CloudTrail, satisfying the requirements for audit trails and compliance.

S3 Glacier Deep Archive is unsuitable due to its high retrieval latency. SSE-C requires customer-side management of encryption keys, with no support for automatic rotation or audit. SSE-S3 does not use customer-managed keys and lacks fine-grained control and auditing.

Question: 327

An ecommerce company is launching a new marketing campaign. The company anticipates the campaign to generate ten times the normal number of daily orders through the company's ecommerce application. The campaign will last 3 days.

The ecommerce application architecture is based on Amazon EC2 instances in an Auto Scaling group and an Amazon RDS for MySQL database. The application writes order transactions to an Amazon Elastic File System (Amazon EFS) file system before the application writes orders to the database. During normal operations, the application write operations peak at 5,000 IOPS.

A solutions architect needs to ensure that the application can handle the anticipated workload during the marketing campaign.

Which solution will meet this requirement?

- A. For the duration of the campaign, increase the provisioned IOPS for the RDS for MySQL database. Set the Amazon EFS throughput mode to Bursting throughput.
- B. For the duration of the campaign, increase the provisioned IOPS for the RDS for MySQL database. Set the Amazon EFS throughput mode to Elastic throughput.
- C. Convert the database to a Multi-AZ deployment. Set the Amazon EFS throughput mode to Elastic throughput for the duration of the campaign.

D. Use AWS Database Migration Service (AWS DMS) to convert the database to RDS for PostgreSQL. Set the Amazon EFS throughput mode to Bursting throughput.

Answer: B

Explanation:

To handle temporary high workloads, Amazon RDS provisioned IOPS can be increased to meet the expected spike in database throughput.

Amazon EFS Elastic throughput mode is ideal for workloads with unpredictable or spiky traffic patterns. It automatically scales throughput based on the file system's size and activity, which is more effective than Bursting throughput during high sustained activity like a marketing campaign.

Multi-AZ improves availability but does not directly address performance. Migrating to PostgreSQL during a short campaign period introduces unnecessary complexity.

Question: 328

A company runs an online order management system on AWS. The company stores order and inventory data for the previous 5 years in an Amazon Aurora MySQL database. The company deletes inventory data after 5 years.

The company wants to optimize costs to archive data.

Which solution will meet this requirement?

- A. Create an AWS Glue crawler to export data to Amazon S3. Create an AWS Lambda function to compress the data.
- B. Use the SELECT INTO OUTFILE S3 query on the Aurora database to export the data to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.
- C. Create an AWS Glue DataBrew job to migrate data from Aurora to Amazon S3. Configure S3 Lifecycle rules on the S3 bucket.
- D. Use the AWS Schema Conversion Tool (AWS SCT) to replicate data from Aurora to Amazon S3. Use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.

Answer: B

Explanation:

Amazon Aurora MySQL supports the SELECT INTO OUTFILE S3 SQL syntax to export query results directly to Amazon S3. This is an efficient and low-overhead method for archiving data.

Once data is in S3, Lifecycle rules can be configured to automatically transition older data to lower- cost storage classes (such as S3 Glacier) or delete it after a defined period, providing a cost-optimized and automated archive solution.

The Glue-based options involve more services and operational overhead. SCT is intended for database migrations, not for periodic data archival.

Question: 329

A company is designing an advertisement distribution application to run on AWS. The company wants to deploy the application as a container to Amazon Elastic Container Service (Amazon ECS).

Advertisements must be displayed to users around the world with low latency. The company needs to optimize data transfer costs.

Which solution will meet these requirements?

- A. Deploy the application in a single AWS Region. Use an Application Load Balancer (ALB) to distribute traffic. Create an Amazon CloudFront distribution, and set the ALB as the origin.
- B. Deploy the application in multiple AWS Regions. Create an Application Load Balancer (ALB) in each Region. Use Amazon Route 53 with a latency-based weighted routing policy to distribute traffic to the ALBs.
- C. Deploy the application in multiple AWS Regions. Create an Application Load Balancer (ALB) in each Region. Create a transit gateway in each Region. Route traffic between the ALBs and Amazon ECS through the transit gateways.
- D. Deploy the application in a single AWS Region. Use an Application Load Balancer (ALB) to distribute traffic. Create an accelerator in AWS Global Accelerator. Associate the accelerator with the ALB.

Answer: A

Explanation:

Using Amazon CloudFront in front of an ALB in a single region is a cost-effective way to deliver content with low latency across the globe. CloudFront caches content closer to the users, reducing the load on backend servers and minimizing data transfer costs by serving cached content from edge locations.

Compared to Global Accelerator, CloudFront is significantly more cost-optimized for static and dynamic content delivery. Multi-region deployments increase infrastructure and transfer costs, which violates the optimization goal. Therefore, option A provides the best mix of performance and cost efficiency.

Question: 330

A company runs a three-tier web application in a VPC on AWS. The company deployed an Application Load Balancer (ALB) in a public subnet. The web tier and application tier Amazon EC2 instances are deployed in a private subnet. The company uses a self-managed MySQL database that runs on EC2 instances in an isolated private subnet for the database tier.

The company wants a mechanism that will give a DevOps team the ability to use SSH to access all the servers. The company also wants to have a centrally managed log of all connections made to the servers.

Which combination of solutions will meet these requirements with the MOST operational efficiency? (Select TWO.)

- A. Create a bastion host in the public subnet. Configure security groups in the public, private, and isolated subnets to allow SSH access.
- B. Create an interface VPC endpoint for AWS Systems Manager Session Manager. Attach the endpoint to the VPC.
- C. Create an IAM policy that grants access to AWS Systems Manager Session Manager. Attach the IAM policy to the EC2 instances.
- D. Create a gateway VPC endpoint for AWS Systems Manager Session Manager. Attach the endpoint to the VPC.
- E. Attach an AmazonSSMManagedInstanceCore AWS managed IAM policy to all the EC2 instance roles.

Answer: B,E

Explanation:

AWS Systems Manager Session Manager allows secure, auditable SSH-like access to EC2 instances without the need to open SSH ports or manage bastion hosts. For this to work in a private subnet, an interface VPC endpoint is required (not a gateway endpoint).

The EC2 instances must have the AmazonSSMManagedInstanceCore policy attached to their IAM roles to allow Systems Manager operations.

With Session Manager, all session activity can be logged centrally to Amazon CloudWatch Logs or S3, satisfying the audit requirement and improving operational efficiency over manual SSH and bastion configurations.

Question: 331

A company needs to grant a team of developers access to the company's AWS resources. The company must maintain a high level of security for the resources.

The company requires an access control solution that will prevent unauthorized access to the sensitive data.

Which solution will meet these requirements?

- A. Share the IAM user credentials for each development team member with the rest of the team to simplify access management and to streamline development workflows.
- B. Define IAM roles that have fine-grained permissions based on the principle of least privilege.

Assign an IAM role to each developer.

C. Create IAM access keys to grant programmatic access to AWS resources. Allow only developers to interact with AWS resources through API calls by using the access keys.

D. Create an AWS Cognito user pool. Grant developers access to AWS resources by using the user pool.

Answer: B

Explanation:

The best practice for secure access control in AWS is to use IAM roles with least-privilege policies, granting only the permissions necessary to perform required tasks. Assigning roles individually ensures that developers cannot overstep their intended access boundaries.

Sharing credentials or using permanent access keys increases the risk of security breaches. Cognito is primarily intended for managing user access to applications, not AWS infrastructure. Thus, Option B best meets security and access control requirements.

Question: 332

A company is developing a serverless, bidirectional chat application that can broadcast messages to connected clients. The application is based on AWS Lambda functions. The Lambda functions receive incoming messages in JSON format.

The company needs to provide a frontend component for the application.

Which solution will meet this requirement?

A. Use an Amazon API Gateway HTTP API to direct incoming JSON messages to backend destinations.

B. Use an Amazon API Gateway REST API that is configured with a Lambda proxy integration.

C. Use an Amazon API Gateway WebSocket API to direct incoming JSON messages to backend destinations.

D. Use an Amazon CloudFront distribution that is configured with a Lambda function URL as a custom origin.

Answer: C

Explanation:

For bidirectional communication such as chat applications, Amazon API Gateway WebSocket API is the correct service. WebSocket APIs allow clients to establish long-lived connections and exchange messages with the backend Lambda functions in real time.

HTTP APIs and REST APIs are suitable for request-response models, not continuous two-way communication. CloudFront cannot maintain stateful WebSocket connections, so only Option C fits the requirements for a real-time, bidirectional application.

Question: 333

A company is setting up a development environment on AWS for a team of developers. The team needs to access multiple Amazon S3 buckets to store project data. The team also needs to use Amazon EC2 to run development instances.

The company needs to ensure that the developers have access only to specific Amazon S3 buckets and EC2 instances. Access permissions must be assigned according to each developer's role on the team. The company wants to minimize the use of permanent credentials and to ensure access is securely managed according to the principle of least privilege.

Which solution will meet these requirements?

- A. Create IAM roles that have administrative-level permissions for Amazon S3 and Amazon EC2. Require developers to sign in by using Amazon Cognito to access Amazon S3 and Amazon EC2.
- B. Create IAM roles that have fine-grained permissions for Amazon S3 and Amazon EC2. Configure AWS IAM Identity Center to manage credentials for the developers.
- C. Create IAM users that have programmatic access to Amazon S3 and Amazon EC2. Generate individual access keys for each developer to access Amazon S3 and Amazon EC2.
- D. Create a VPC endpoint for Amazon S3. Require developers to access Amazon EC2 instances and Amazon S3 buckets through a bastion host.

Answer: B

Explanation:

The most secure and manageable way to provide developers with temporary, least-privilege access is by using AWS IAM Identity Center (formerly AWS SSO). IAM Identity Center allows assigning IAM roles with scoped permissions based on the developer's team role. This ensures no permanent credentials are required and minimizes risk.

Option B enables role-based access with centralized identity and access management, making it the most secure and scalable solution for managing developer permissions.

Question: 334

A company uses an Amazon EC2 instance to handle requests for a public web application. The application routes traffic to multiple application pages by using URL paths.

The company begins to experience large surges of traffic at unpredictable times. The traffic surges cause the web application to experience issues and to occasionally become unavailable.

The company needs to make the web application more scalable to handle sudden increases in traffic.

Which solution will meet this requirement?

- A. Create an Amazon Machine Image (AMI) of the web application instance. Use the AMI to create an Auto Scaling group of EC2 instances that has a minimum capacity of two. Create an Application Load Balancer. Set the Auto Scaling group as the target group.
- B. Create a Docker image of the application. Use Amazon Elastic Container Service (Amazon ECS) to create an Auto Scaling ECS cluster. Enable managed scaling. Create a Network Load Balancer. Set the ECS cluster as the target group.
- C. Create an Amazon Machine Image (AMI) of the web application instance. Use the AMI to create two more web application instances in separate Availability Zones. Update the website DNS record to refer to all three instances.
- D. Create an Application Load Balancer (ALB). Set the web application instance as the target. Create an Amazon CloudWatch alarm based on ALB traffic metrics. Configure the alert to activate when traffic spikes.

Answer: A

Explanation:

To handle unpredictable traffic surges and improve scalability, Auto Scaling is essential. Creating an AMI and deploying it into an Auto Scaling group behind an Application Load Balancer (ALB) allows AWS to automatically scale the number of EC2 instances based on demand.

This architecture improves availability and responsiveness by distributing traffic and launching instances when needed. Option A satisfies the need for high scalability and resilience. Other options either lack Auto Scaling or are more complex without added benefit in this context.

Question: 335

A company is using an Amazon Redshift cluster to run analytics queries for multiple sales teams. In addition to the typical workload, on the last Monday morning of each month, thousands of users run reports. Users have reported slow response times during the monthly surge.

The company must improve query performance without impacting the availability of the Redshift cluster.

Which solution will meet these requirements?

- A. Resize the Redshift cluster by using the classic resize capability of Amazon Redshift before every monthly surge. Reduce the cluster to its original size after each surge.
- B. Resize the Redshift cluster by using the elastic resize capability of Amazon Redshift before every monthly surge. Reduce the cluster to its original size after each surge.
- C. Enable the concurrency scaling feature for the Redshift cluster for specific workload management (WLM) queues.
- D. Enable Amazon Redshift Spectrum for the Redshift cluster before every monthly surge.

Answer: C

Explanation:

Concurrency Scaling allows Amazon Redshift to add transient capacity automatically to handle bursts in concurrent queries. This is ideal for scenarios with predictable surge periods, such as end-of-month reporting.

It improves performance without manual resizing or cluster modification and without affecting availability. Resizing requires manual intervention and potential downtime, and Redshift Spectrum is designed for querying data in S3, not for solving concurrency issues.

Question: 336

A company hosts a web application in a VPC on AWS. A public Application Load Balancer (ALB) forwards connections from the internet to an Auto Scaling group of Amazon EC2 instances. The Auto Scaling group runs in private subnets across four Availability Zones.

The company stores data in an Amazon S3 bucket in the same Region. The EC2 instances use NAT gateways in each Availability Zone for outbound internet connectivity.

The company wants to optimize costs for its AWS architecture.

Which solution will meet this requirement?

- A. Reconfigure the Auto Scaling group and the ALB to use two Availability Zones instead of four. Do not change the desired count or scaling metrics for the Auto Scaling group to maintain application availability.
- B. Create a new, smaller VPC that still has sufficient IP address availability to run the application. Redeploy the application stack in the new VPC. Delete the existing VPC and its resources.
- C. Deploy an S3 gateway endpoint to the VPC. Configure the EC2 instances to access the S3 bucket through the S3 gateway endpoint.
- D. Deploy an S3 interface endpoint to the VPC. Configure the EC2 instances to access the S3 bucket through the S3 interface endpoint.

Answer: C

Explanation:

Using S3 gateway endpoints allows private and cost-free access to S3 without routing traffic through a NAT gateway. NAT gateway traffic incurs charges, especially when used across multiple Availability Zones.

By using an S3 gateway endpoint, EC2 instances in private subnets can access S3 directly without needing internet access, reducing both

data transfer and NAT gateway costs.

Interface endpoints are more expensive and typically used for services like API Gateway or Systems Manager.

Question: 337

A company needs to run its external website on Amazon EC2 instances and on-premises virtualized servers. The AWS environment has a 1 GB AWS Direct Connect connection to the data center. The application has IP addresses that will not change. The on-premises and AWS servers are able to restart themselves while maintaining the same IP address if a failure occurs. Some website users have to add their vendors to an allow list, so the solution must have a fixed IP address. The company needs a solution with the lowest operational overhead to handle this split traffic.

What should a solutions architect do to meet these requirements?

- A. Deploy an Amazon Route 53 Resolver with rules pointing to the on-premises and AWS IP addresses.
- B. Deploy a Network Load Balancer on AWS. Create target groups for the on-premises and AWS IP addresses.
- C. Deploy an Application Load Balancer on AWS. Register the on-premises and AWS IP addresses with the target group.
- D. Deploy Amazon API Gateway to direct traffic to the on-premises and AWS IP addresses based on the header of the request.

Answer: B

Explanation:

A Network Load Balancer (NLB) supports IP address-based targets, which allows the use of both EC2 and on-premises endpoints. It also supports a static IP address or Elastic IP, which meets the requirement for a fixed IP address needed by some users.

NLB offers high performance, low latency, and minimal operational overhead. Application Load Balancer only supports instance or Lambda targets, not IP addresses outside the VPC. Route 53 Resolver is for DNS, and API Gateway is for HTTP-based APIs, not web applications.

Question: 338

A company runs multiple applications in multiple AWS accounts within the same organization in AWS Organizations. A content management system (CMS) runs on Amazon EC2 instances in a VPC. The CMS needs to access shared files from an Amazon Elastic File System (Amazon EFS) file system that is deployed in a separate AWS account. The EFS account is in a separate VPC.

Which solution will meet this requirement?

- A. Mount the EFS file system on the EC2 instances by using the EFS Elastic IP address.

B. Enable VPC sharing between the two accounts. Use the EFS mount helper to mount the file system on the EC2 instances.

Redeploy the EFS file system in a shared subnet.

C. Configure AWS Systems Manager Run Command to mount the EFS file system on the EC2 instances.

D. Install the amazon-efs-utils package on the EC2 instances. Add the mount target in the efs-config file. Mount the EFS file system by using the EFS access point.

Answer: D

Explanation:

To access an EFS file system across accounts and VPCs, the EFS must be mounted using VPC peering or AWS Transit Gateway, and the EC2 instances must use the amazon-efs-utils package with the correct mount target or access point.

Using an EFS access point simplifies access management, especially across accounts, by providing a POSIX identity and access policy layer.

VPC sharing doesn't support EFS directly unless the subnet and resources are shared properly, which requires redeployment. Therefore, option D is the most complete and correct.

Question: 339

A company runs a web application that uses Amazon RDS for MySQL to store relational data. Data in the database does not change frequently.

A solutions architect notices that during peak usage times, the database has performance issues when it serves the data. The company wants to improve the performance of the database.

Which combination of steps will meet these requirements? (Select TWO.)

A. Integrate AWS WAF with the application.

B. Create a read replica for the database. Redirect read traffic to the read replica.

C. Create an Amazon ElastiCache (Memcached) cluster. Configure the application and the database to integrate with the cluster.

D. Use the Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class to store the data that changes infrequently.

E. Migrate the database to Amazon DynamoDB. Configure the application to use the DynamoDB database.

Answer: B,C

Explanation:

To improve read performance for a MySQL-based RDS database under load, you can: Use Read Replicas: Amazon RDS supports MySQL read replicas, which help offload read operations from the primary database, improving performance during high traffic. Use ElastiCache (Memcached): Adding an in-memory cache layer using Amazon ElastiCache reduces the load on the RDS instance by serving frequent queries directly from memory, especially when data is not updated often.

Option A (AWS WAF) is for web security, not database performance. Option D relates to storage optimization, not query latency. Option E would require re-architecting from relational to NoSQL, which is unnecessary and disruptive.

Question: 340

An ecommerce company wants to collect user clickstream data from the company's website for realtime analysis. The website experiences fluctuating traffic patterns throughout the day. The company needs a scalable solution that can adapt to varying levels of traffic.

Which solution will meet these requirements?

- A. Use a data stream in Amazon Kinesis Data Streams in on-demand mode to capture the clickstream data. Use AWS Lambda to process the data in real time.
- B. Use Amazon Data Firehose to capture the clickstream data. Use AWS Glue to process the data in real time.
- C. Use Amazon Kinesis Video Streams to capture the clickstream data. Use AWS Glue to process the data in real time.
- D. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to capture the clickstream data. Use AWS Lambda to process the data in real time.

Answer: A

Explanation:

Amazon Kinesis Data Streams in on-demand mode is purpose-built for real-time ingestion of streaming data and scales automatically with traffic, which is ideal for fluctuating workloads like clickstream data.

Combined with AWS Lambda, which scales concurrently based on incoming events, this setup ensures efficient, real-time processing with minimal configuration and cost overhead.

Option B involves Firehose, which is not optimal for low-latency processing. Option C is incorrect as

Kinesis Video Streams is designed for video data, not clickstream. Option D introduces Flink, which adds unnecessary complexity when Lambda suffices.

Question: 341

A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

Answer: C

Explanation:

Field-level encryption in Amazon CloudFront provides end-to-end encryption for specific data fields (e.g., credit card numbers, social security numbers). It ensures that sensitive fields are encrypted at the edge before being forwarded to the origin, and only the authorized application with the private key can decrypt them.

This adds a layer of protection beyond HTTPS, which encrypts the whole payload but not individual fields. Signed URLs and cookies are for access control, not encryption. Setting HTTPS Only is a good practice but does not satisfy the field-specific encryption requirement.

Question: 342

A company uses Amazon RDS for PostgreSQL databases for its data tier. The company must implement password rotation for the databases.

Which solution meets this requirement with the LEAST operational overhead?

- A. Store the password in AWS Secrets Manager. Enable automatic rotation on the secret.
- B. Store the password in AWS Systems Manager Parameter Store. Enable automatic rotation on the parameter.
- C. Store the password in AWS Systems Manager Parameter Store. Write an AWS Lambda function that rotates the password.
- D. Store the password in AWS Key Management Service (AWS KMS). Enable automatic rotation on the AWS KMS key.

Answer: A

Explanation:

AWS Secrets Manager is the recommended service for managing and automatically rotating database credentials. It integrates natively with Amazon RDS (including PostgreSQL), supports built-in rotation functionality, and requires minimal setup. Secrets Manager also supports versioning and auditing, which enhances operational excellence and security. Parameter Store does not natively support credential rotation. AWS KMS manages key encryption—not application secrets—so it is not applicable here.

Question: 343

A company receives data transfers from a small number of external clients that use SFTP software on an Amazon EC2 instance. The clients use an SFTP client to upload data. The clients use SSH keys for authentication. Every hour, an automated script transfers new uploads to an Amazon S3 bucket for processing.

The company wants to move the transfer process to an AWS managed service and to reduce the time required to start data processing. The company wants to retain the existing user management and SSH key generation process. The solution must not require clients to make significant changes to their existing processes.

Which solution will meet these requirements?

- A. Reconfigure the script that runs on the EC2 instance to run every 15 minutes. Create an S3 Event Notifications rule for all new object creation events. Set an Amazon Simple Notification Service (Amazon SNS) topic as the destination.
- B. Create an AWS Transfer Family SFTP server that uses the existing S3 bucket as a target. Use service-managed users to enable authentication.
- C. Require clients to add the AWS DataSync agent into their local environments. Create an IAM user for each client that has permission to upload data to the target S3 bucket.
- D. Create an AWS Transfer Family SFTP connector that has permission to access the target S3 bucket for each client. Store credentials in AWS Systems Manager. Create an IAM role to allow the SFTP connector to securely use the credentials.

Answer: B

Explanation:

AWS Transfer Family (SFTP) allows clients to use standard SFTP clients and SSH keys without changes. By enabling service-managed users, clients can continue uploading files with their existing tools.

The service delivers the files directly into S3, reducing latency between upload and processing. This removes the need for EC2, custom scripts, and periodic transfers. It fully meets the requirement for a managed solution with minimal disruption to client processes.

Question: 344

A company needs to run a critical data processing workload that uses a Python script every night. The workload takes 1 hour to finish.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type. Use the Fargate Spot capacity provider. Schedule the job to run once every night.
- B. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type. Schedule the job to run once every night.
- C. Create an AWS Lambda function that uses the existing Python code. Configure Amazon EventBridge to invoke the function once every night.
- D. Create an Amazon EC2 On-Demand Instance that runs Amazon Linux. Migrate the Python script to the instance. Use a cron job to schedule the script. Create an AWS Lambda function to start and stop the instance once every night.

Answer: A

Explanation:

AWS Fargate with Spot capacity is the most cost-effective and serverless container option for shortduration jobs that are flexible on start time. Since the job is not latency-critical and runs nightly, it is a good candidate for Fargate Spot, which can offer up to 70% cost savings over On-Demand pricing.

The job runs for 1 hour, which exceeds the AWS Lambda maximum execution time (15 minutes). Therefore, Lambda is not suitable. EC2-based solutions involve higher operational overhead and cost, making Fargate Spot the best low-cost, low-maintenance solution.

Question: 345

A solutions architect is creating a data reporting application that will send traffic through third-party network firewalls in an AWS security account. The firewalls and application servers must be load balanced.

The application uses TCP connections to generate reports. The reports can run for several hours and can be idle for up to 1 hour. The reports must not time out during an idle period.

Which solution will meet these requirements?

- A. Use a Gateway Load Balancer (GWLB) for the firewalls. Use an Application Load Balancer (ALB) for the application servers. Set the ALB idle timeout period to 1 hour.
- B. Use a single firewall in the security account. Use an Application Load Balancer (ALB) for the application servers. Set the ALB idle timeout and firewall idle timeout periods to 1 hour.

- C. Use a Gateway Load Balancer (GWLB) for the firewalls. Use an Application Load Balancer (ALB) for the application servers. Set the idle timeout periods for the ALB, the GWLB, and the firewalls to 1 hour.
- D. Use a Gateway Load Balancer (GWLB) for the firewalls. Use an Application Load Balancer (ALB) for the application servers. Configure the ALB idle timeout period to 1 hour. Increase the application server capacity to finish the report generation faster.

Answer: C

Explanation:

Since the application uses long-lived TCP connections and must remain idle for up to 1 hour without timeout, all components involved in the connection path (ALB, GWLB, and firewall) must have their idle timeout values configured to at least 1 hour.

Gateway Load Balancer supports transparent insertion of firewalls, and configuring consistent idle timeouts ensures connections don't drop mid-session.

Using just one firewall (option B) introduces a single point of failure. Increasing capacity (option D) doesn't solve idle timeout issues. Therefore, C provides a resilient and complete configuration.

Question: 346

A company collects 10 GB of telemetry data every day from multiple devices. The company stores the data in an Amazon S3 bucket that is in a source data account.

The company has hired several consulting agencies to analyze the company's data. Each agency has a unique AWS account. Each agency requires read access to the company's data.

The company needs a secure solution to share the data from the source data account to the consulting agencies.

Which solution will meet these requirements with the LEAST operational effort?

- A. Set up an Amazon CloudFront distribution. Use the S3 bucket as the origin.
- B. Make the S3 bucket public for a limited time. Inform only the agencies that the bucket is publicly accessible.
- C. Configure cross-account access for the S3 bucket to the accounts that the agencies own.
- D. Set up an IAM user for each agency in the source data account. Grant each agency IAM user access to the company's S3 bucket.

Answer: C

Explanation:

The most secure and least operationally intensive method is to configure cross-account access using resource-based policies on the S3 bucket. This allows trusted external AWS accounts (consulting agencies) to securely access the S3 data without the need to manage user credentials or build additional infrastructure.

Options A and B pose security risks. Option D increases operational complexity and violates least privilege by managing external users inside your AWS account.

Question: 347

A company has 5 TB of datasets. The datasets consist of 1 million user profiles and 10 million connections. The user profiles have connections as many-to-many relationships. The company needs a performance-efficient way to find mutual connections up to five levels.

Which solution will meet these requirements?

- A. Use an Amazon S3 bucket to store the datasets. Use Amazon Athena to perform SQL JOIN queries to find connections.
- B. Use Amazon Neptune to store the datasets with edges and vertices. Query the data to find connections.
- C. Use an Amazon S3 bucket to store the datasets. Use Amazon QuickSight to visualize connections.
- D. Use Amazon RDS to store the datasets with multiple tables. Perform SQL JOIN queries to find connections.

Answer: B

Explanation:

Amazon Neptune is a fully managed graph database service optimized for storing and navigating complex many-to-many relationships like social graphs or network topologies.

It supports Gremlin and openCypher queries that allow efficient traversal of connections even multiple levels deep. SQL JOINS (Athena or RDS) are inefficient for deep relationship queries due to exponential complexity.

QuickSight is used for data visualization, not graph traversal.

Question: 348

A company is migrating some workloads to AWS. However, many workloads will remain on premises. The on-premises workloads require secure and reliable connectivity to AWS with consistent, low-latency performance.

The company has deployed the AWS workloads across multiple AWS accounts and multiple VPCs. The company plans to scale to hundreds of VPCs within the next year.

The company must establish connectivity between each of the VPCs and from the on-premises environment to each VPC.

Which solution will meet these requirements?

A. Use an AWS Direct Connect connection to connect the on-premises environment to AWS. Configure VPC peering to establish connectivity between VPCs.

B. Use multiple AWS Site-to-Site VPN connections to connect the on-premises environment to AWS.

Create a transit gateway to establish connectivity between VPCs.

C. Use an AWS Direct Connect connection with a Direct Connect gateway to connect the on-premises environment to AWS. Create a transit gateway to establish connectivity between VPCs. Associate the transit gateway with the Direct Connect gateway.

D. Use an AWS Site-to-Site VPN connection to connect the on-premises environment to AWS. Configure VPC peering to establish connectivity between VPCs.

Answer: C

Explanation:

The optimal solution for scalable and resilient hybrid networking is to use AWS Direct Connect with a Direct Connect gateway for secure, low-latency access to AWS, and an AWS Transit Gateway to manage connectivity among hundreds of VPCs.

By associating the Transit Gateway with the Direct Connect gateway, you enable transitive routing between on-premises and all VPCs, while minimizing network complexity and maintaining high performance.

VPC peering does not scale well, and VPNs don't offer the same performance or consistency.

Question: 349

A media company hosts a mobile app backend in the AWS Cloud. The company is releasing a new feature to allow users to upload short videos and apply special effects by using the mobile app. The company uses AWS Amplify to store the videos that customers upload in an Amazon S3 bucket.

The videos must be processed immediately. Users must receive a notification when processing is finished.

Which solution will meet these requirements?

A. Use Amazon EventBridge Scheduler to schedule an AWS Lambda function to process the videos. Save the processed videos to the S3 bucket. Use Amazon Simple Notification Service (Amazon SNS) to send push notifications to customers when processing is finished.

B. Use Amazon EventBridge Scheduler to schedule AWS Fargate to process the videos. Save the processed videos to the S3 bucket. Use Amazon Simple Notification Service (Amazon SNS) to send push notifications to customers when processing is finished.

C. Use an S3 trigger to invoke an AWS Lambda function to process the videos. Save the processed

videos to the S3 bucket. Use Amazon Simple Notification Service (Amazon SNS) to send push notifications to customers when processing is finished.

D. Use an S3 trigger to invoke an AWS Lambda function to process the videos. Save the processed videos to the S3 bucket. Use AWS Amplify to send push notifications to customers when processing is finished.

Answer: C

Explanation:

The requirement is for immediate processing of uploaded videos and prompt notification to users. According to AWS best practices for event-driven architectures, using S3 event notifications to trigger a Lambda function upon an object creation (upload) is the optimal solution for real-time processing. Lambda can process the file as soon as it is uploaded, ensuring low latency. Once processing is complete, Lambda can save the processed file back to S3 and use Amazon SNS to notify the user.

This approach uses managed services with minimal operational overhead, is scalable, and ensures event-driven processing with instant user feedback. AWS Amplify primarily facilitates application development and hosting but does not natively provide direct push notification support for this backend workflow; instead, SNS is designed for such notification scenarios.

Reference Extract from AWS Documentation / Study Guide:

"Amazon S3 can publish events to AWS Lambda when objects are created. AWS Lambda runs code in response to events and can interact with other AWS services. Amazon SNS is a flexible, fully managed pub/sub messaging and mobile notifications service for coordinating the delivery of messages to subscribing endpoints and clients."

Source: AWS Certified Solutions Architect - Official Study Guide, Event-driven Architectures section; AWS Lambda Developer Guide (S3 triggers); Amazon SNS User Guide.

Question: 350

A company needs a solution to back up and protect critical AWS resources. The company needs to regularly take backups of several Amazon EC2 instances and Amazon RDS for PostgreSQL databases. To ensure high resiliency, the company must have the ability to validate and restore backups.

Which solution meets the requirement with LEAST operational overhead?

- A. Use AWS Backup to create a backup schedule for the resources. Use AWS Backup to create a restoration testing plan for the required resources.
- B. Take snapshots of the EC2 instances and RDS DB instances. Create AWS Batch jobs to validate and restore the snapshots.
- C. Create a custom AWS Lambda function to take snapshots of the EC2 instances and RDS DB instances. Create a second Lambda function to restore the snapshots periodically to validate the backups.

D. Take snapshots of the EC2 instances and RDS DB instances. Create an AWS Lambda function to restore the snapshots periodically to validate the backups.

Answer: A

Explanation:

AWS Backup is a fully managed backup service designed to centralize and automate data protection across AWS services including EC2 and RDS. It allows users to define backup schedules (backup plans) and automatically create and retain backups. AWS Backup also offers restore testing plans, allowing users to automate the validation of backups by restoring them in a controlled manner. This service is built to minimize operational overhead by removing the need to manage custom scripts, manual processes, or additional orchestration services. This aligns with AWS best practices for resilience, automation, and operational excellence.

Reference Extract from AWS Documentation / Study Guide:

"AWS Backup enables you to centralize and automate data protection across AWS services. You can create backup plans, schedule backups, and set lifecycle policies. AWS Backup also enables restore testing to verify your backup integrity, with minimal manual intervention."

Source: AWS Certified Solutions Architect - Official Study Guide, Resiliency and Disaster Recovery section; AWS Backup User Guide.

Question: 351

A company uses AWS to host a public website. The load on the web servers recently increased.

The company wants to learn more about the traffic flow and traffic sources. The company also wants to increase the overall security of the website.

Which solution will meet these requirements?

- A. Deploy AWS WAF and set up logging. Use Amazon Data Firehose to deliver the log files to an Amazon S3 bucket for analysis.
- B. Deploy Amazon API Gateway and set up logging. Use Amazon Kinesis Data Streams to deliver the log files to an Amazon S3 bucket for analysis.
- C. Deploy a Network Load Balancer and set up logging. Use Amazon Data Firehose to deliver the log files to an Amazon S3 bucket for analysis.
- D. Deploy an Application Load Balancer and set up logging. Use Amazon Kinesis Data Streams to deliver the log files to an Amazon S3 bucket for analysis.

Answer: A

Explanation:

AWS WAF (Web Application Firewall) is designed to protect public-facing web applications from common web exploits and allows for deep inspection of HTTP and HTTPS requests. By enabling logging on AWS WAF, you can gain insights into traffic flow, request sources, and blocked requests, which improves both security visibility and posture. Integrating AWS WAF logging with Amazon Kinesis Data Firehose allows automatic, near-real-time delivery of log data to Amazon S3 for analysis, reporting, or integration with analytics tools. This solution not only secures the website but also enables comprehensive traffic analysis, satisfying both requirements efficiently.

Reference Extract from AWS Documentation / Study Guide:

"AWS WAF provides detailed logs of web requests, which can be delivered to Amazon S3 via Amazon Kinesis Data Firehose. This logging enables analysis of web traffic and sources, supporting both security and operational monitoring."

Source: AWS Certified Solutions Architect - Official Study Guide, Security and Compliance section; AWS WAF Developer Guide (Logging and Monitoring).

Question: 352

A company is creating a new application that will store a large amount of data. The data will be analyzed hourly and will be modified by several Amazon EC2 Linux instances that are deployed across multiple Availability Zones. The needed amount of storage space will continue to grow for the next 6 months.

Which storage solution should a solutions architect recommend to meet these requirements?

- A. Store the data in Amazon S3 Glacier. Update the S3 Glacier vault policy to allow access to the application instances.
- B. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Mount the EBS volume on the application instances.
- C. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Mount the file system on the application instances.
- D. Store the data in an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume shared between the application instances.

Answer: C

Explanation:

Amazon Elastic File System (Amazon EFS) is a fully managed, elastic, shared file system that can be mounted concurrently by multiple EC2 instances across multiple Availability Zones. EFS automatically grows and shrinks as files are added or removed, providing seamless scalability for unpredictable and growing storage needs. It supports simultaneous access from multiple compute resources, making it ideal for applications where several EC2 instances must read and write to the same data set concurrently. EBS volumes cannot be shared across multiple Availability Zones, and S3 Glacier is intended for archival and infrequent access, not for active, hourly data analysis.

Reference Extract from AWS Documentation / Study Guide:

"Amazon EFS provides a scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It can

be mounted to many EC2 instances across multiple Availability Zones and is ideal for shared data scenarios."

Source: AWS Certified Solutions Architect - Official Study Guide, Storage Solutions section; Amazon EFS User Guide.

Question: 353

A company is building a critical data processing application that will run on Amazon EC2 instances. The company must not run any two nodes on the same underlying hardware. The company requires at least 99.99% availability for the application.

Which solution will meet these requirements?

- A. Deploy the application to one Availability Zone by using a cluster placement group strategy.
- B. Deploy the application to three Availability Zones by using a spread placement group strategy.
- C. Deploy the application to three Availability Zones by using a cluster placement group strategy.
- D. Deploy the application to one Availability Zone by using a partition placement group strategy.

Answer: B

Explanation:

A spread placement group is designed to deploy each instance on distinct underlying hardware, reducing the risk of simultaneous failures. By spreading instances across multiple Availability Zones, you achieve high availability and fault tolerance, meeting the 99.99% uptime requirement. Spread placement groups are ideal for critical applications that require maximum resilience to single hardware failures. Neither cluster nor partition strategies offer the same guarantee of separation combined with cross-AZ distribution.

Reference Extract from AWS Documentation / Study Guide:

"Spread placement groups are recommended for applications that have a small number of critical instances that should be kept separate from each other. Instances in a spread placement group are placed on distinct underlying hardware, and when deployed across multiple Availability Zones, provide high availability."

Source: AWS Certified Solutions Architect - Official Study Guide, Compute and High Availability section; Amazon EC2 Placement Groups Documentation.

Question: 354

A company stores a large volume of critical data in Amazon RDS for PostgreSQL tables. The company is developing several new features for an upcoming product launch. Some of the new features require many table alterations.

The company needs a solution to test the altered tables for several days. After testing, the solution must make the new features available to customers in production.

Which solution will meet these requirements with the HIGHEST availability?

- A. Create a new instance of the database in RDS for PostgreSQL to test the new features. When the testing is finished, take a backup of the test database, and restore the test database to the production database.
- B. Create new database tables in the production database to test the new features. When the testing is finished, copy the data from the older tables to the new tables. Delete the older tables, and rename the new tables accordingly.
- C. Create an Amazon RDS read replica to deploy a new instance of the database. Make updates to the database tables in the replica instance. When the testing is finished, promote the replica instance to become the new production instance.
- D. Use an Amazon RDS blue/green deployment to deploy a new test instance of the database. Make database table updates in the test instance. When the testing is finished, promote the test instance to become the new production instance.

Answer: D

Explanation:

Amazon RDS Blue/Green Deployments provide a safe and straightforward way to make database changes with minimal downtime and risk. Blue/Green deployments create an exact copy ("green") of your production environment ("blue") where you can make schema changes and run tests. After validation, you can promote the green environment to production with a single click or API call, achieving near-zero downtime and maximum availability. This is the AWS-recommended method for deploying major database changes in a way that minimizes impact to users and maximizes uptime.

Reference Extract from AWS Documentation / Study Guide:

"Amazon RDS Blue/Green Deployments enable you to make changes to your database environment safely. You can perform schema updates and feature testing in a fully managed staging environment and switch over with minimal downtime, ensuring the highest availability."

Source: AWS Certified Solutions Architect - Official Study Guide, Database and Migration section; Amazon RDS Blue/Green Deployments Documentation.

Question: 355

A company wants to store a large amount of data as objects for analytics and long-term archiving. Resources from outside AWS need to access the data. The external resources need to access the data with unpredictable frequency. However, the external resource must have immediate access when necessary.

The company needs a cost-optimized solution that provides high durability and data security.

Which solution will meet these requirements?

- A. Store the data in Amazon S3 Standard. Apply S3 Lifecycle policies to transition older data to S3 Glacier Deep Archive.

- B. Store the data in Amazon S3 Intelligent-Tiering.
- C. Store the data in Amazon S3 Glacier Flexible Retrieval. Use expedited retrieval to provide immediate access when necessary.
- D. Store the data in Amazon Elastic File System (Amazon EFS) Infrequent Access (IA). Use lifecycle policies to archive older files.

Answer: B

Explanation:

Amazon S3 Intelligent-Tiering is designed for data with unknown or changing access patterns. It automatically moves data between frequent and infrequent access tiers based on usage. This tier offers immediate access to all objects, regardless of which tier they are stored in, while optimizing storage costs. S3 Intelligent-Tiering also provides the same high durability, availability, and security as other S3 storage classes and supports access from external resources using standard S3 APIs. Lifecycle policies and Glacier classes are more suitable for archival when infrequent access is predictable, but retrieval from Glacier classes is not immediate and incurs extra charges and delays.

Reference Extract from AWS Documentation / Study Guide:

"S3 Intelligent-Tiering is designed to optimize costs by automatically moving data between two access tiers when access patterns change. Data is always available and immediately accessible, making it ideal for unknown or unpredictable access patterns."

Source: AWS Certified Solutions Architect - Official Study Guide, S3 Storage Classes section.

Question: 356

A company has a web application that uses several web servers that run on Amazon EC2 instances. The instances use a shared Amazon RDS for MySQL database.

The company requires a secure method to store database credentials. The credentials must be automatically rotated every 30 days without affecting application availability.

Which solution will meet these requirements?

- A. Store database credentials in AWS Secrets Manager. Create an AWS Lambda function to automatically rotate the credentials. Use Amazon EventBridge to run the Lambda function on a schedule. Grant the necessary IAM permissions to allow the web servers to access Secrets Manager.
- B. Store database credentials in AWS Systems Manager OpsCenter. Grant the necessary IAM permissions to allow the web servers to access OpsCenter.
- C. Store database credentials in an Amazon S3 bucket. Create an AWS Lambda function to automatically rotate the credentials. Use Amazon EventBridge to run the Lambda function on a schedule. Grant the necessary IAM permissions to allow the web servers to retrieve credentials from the S3 bucket.
- D. Store the credentials in a local file on each of the web servers. Use an AWS KMS key to encrypt the credentials. Create a cron job

on each server to rotate the credentials every 30 days.

Answer: A

Explanation:

AWS Secrets Manager is a fully managed service specifically designed to securely store and automatically rotate database credentials, API keys, and other secrets. Secrets Manager provides built-in integration with Amazon RDS for automatic credential rotation on a configurable schedule without requiring downtime. It also manages the secure distribution of the credentials to authorized services, such as your web servers, using IAM policies. Manual solutions (S3, files, cron jobs) do not provide the same level of automation, audit, or security.

Reference Extract from AWS Documentation / Study Guide:

"AWS Secrets Manager enables you to rotate, manage, and retrieve database credentials securely. It supports automatic rotation of secrets for supported AWS databases without requiring application downtime."

Source: AWS Certified Solutions Architect - Official Study Guide, Security and Secrets Management section.

Question: 357

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing. The company wants to minimize the website hosting costs.

Which solution will meet these requirements?

- A. Move the website to an Amazon S3 bucket. Configure an Amazon CloudFront distribution for the S3 bucket.
- B. Move the website to an Amazon S3 bucket. Configure an Amazon ElastiCache cluster for the S3 bucket.
- C. Move the website to AWS Amplify. Configure an ALB to resolve to the Amplify website.
- D. Move the website to AWS Amplify. Configure EC2 instances to cache the website.

Answer: A

Explanation:

Serving static content from Amazon S3 is highly cost-effective and scalable. By fronting the S3 bucket with Amazon CloudFront, you also enable global caching, reduced latency, and even lower costs by reducing direct S3 access. There is no need for EC2 or ALB, and no need to manage servers, further lowering operational burden and expenses. This pattern is the recommended AWS architecture for serving static web content.

Reference Extract from AWS Documentation / Study Guide:

"Hosting static websites on Amazon S3 with Amazon CloudFront reduces infrastructure costs and improves performance for global users by caching content at edge locations."

Source: AWS Certified Solutions Architect - Official Study Guide, S3 and CloudFront section.

Question: 358

A company wants to provide users with access to AWS resources. The company has 1,500 users and manages their access to on-premises resources through Active Directory user groups on the corporate network. However, the company does not want users to have to maintain another identity to access the resources. A solutions architect must manage user access to the AWS resources while preserving access to the on-premises resources.

What should the solutions architect do to meet these requirements?

- A. Create an IAM user for each user in the company. Attach the appropriate policies to each user.
- B. Use Amazon Cognito with an Active Directory user pool. Create roles with the appropriate policies attached.
- C. Define cross-account roles with the appropriate policies attached. Map the roles to the Active Directory groups.
- D. Configure Security Assertion Markup Language (SAML) 2.0-based federation. Create roles with the appropriate policies attached. Map the roles to the Active Directory groups.

Answer: D

Explanation:

SAML 2.0-based federation allows organizations to integrate their on-premises Active Directory with AWS, enabling Single Sign-On (SSO) for AWS Management Console and AWS CLI/API access. This lets users continue using their existing AD credentials, removing the need to manage separate identities for AWS and on-premises systems. Mapping roles to AD groups provides granular access control and seamless management.

Reference Extract from AWS Documentation / Study Guide:

"AWS supports SAML 2.0 for federated access, enabling integration with Active Directory or other identity providers to provide single sign-on to AWS resources without creating separate IAM users."

Source: AWS Certified Solutions Architect - Official Study Guide, Identity and Access Management section.

Question: 359

A company has an application with a REST-based interface that allows data to be received in nearreal time from a third-party vendor. Once received, the application processes and stores the data for further analysis. The application is running on Amazon EC2 instances.

The third-party vendor has received many 503 Service Unavailable Errors when sending data to the application. When the data volume spikes, the compute capacity reaches its maximum limit and the application is unable to process all requests.

Which design should a solutions architect recommend to provide a more scalable solution?

- A. Use Amazon Kinesis Data Streams to ingest the data. Process the data using AWS Lambda functions.
- B. Use Amazon API Gateway on top of the existing application. Create a usage plan with a quota limit for the third-party vendor.
- C. Use Amazon Simple Notification Service (Amazon SNS) to ingest the data. Put the EC2 instances in an Auto Scaling group behind an Application Load Balancer.
- D. Repackage the application as a container. Deploy the application using Amazon Elastic Container Service (Amazon ECS) using the EC2 launch type with an Auto Scaling group.

Answer: A

Explanation:

Amazon Kinesis Data Streams provides a highly scalable and durable service for ingesting real-time streaming data. By decoupling ingestion and processing, Kinesis can handle large spikes in traffic without service disruption. Lambda functions (or other consumers) can then process the data as it arrives, scaling automatically. This pattern avoids 503 errors due to compute saturation and delivers a resilient, serverless, and highly scalable architecture.

Reference Extract from AWS Documentation / Study Guide:

"Kinesis Data Streams provides a scalable and durable real-time data streaming service. Coupling Kinesis with AWS Lambda enables event-driven processing, elasticity, and decoupling between ingestion and processing layers."

Source: AWS Certified Solutions Architect - Official Study Guide, Streaming and Serverless section.

Question: 360

A company has an on-premises application that uses SFTP to collect financial data from multiple vendors. The company is migrating to the AWS Cloud. The company has created an application that uses Amazon S3 APIs to upload files from vendors.

Some vendors run their systems on legacy applications that do not support S3 APIs. The vendors want to continue to use SFTP-based applications to upload data. The company wants to use managed services for the needs of the vendors that use legacy applications.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Database Migration Service (AWS DMS) instance to replicate data from the storage of the vendors that use legacy applications to Amazon S3. Provide the vendors with the credentials to access the AWS DMS instance.

- B. Create an AWS Transfer Family endpoint for vendors that use legacy applications.
- C. Configure an Amazon EC2 instance to run an SFTP server. Instruct the vendors that use legacy applications to use the SFTP server to upload data.
- D. Configure an Amazon S3 File Gateway for vendors that use legacy applications to upload files to an SMB file share.

Answer: B

Explanation:

AWS Transfer Family is a fully managed service that provides SFTP, FTPS, and FTP access directly to Amazon S3. It allows organizations to support legacy data transfer protocols without managing infrastructure. This approach gives the vendors a familiar SFTP interface, with files landing directly in S3, and requires minimal operational effort compared to managing EC2 servers or using gateways. Reference Extract from AWS Documentation / Study Guide:

"AWS Transfer Family enables secure transfer of files directly into and out of Amazon S3 using SFTP, FTPS, and FTP, and is fully managed by AWS, significantly reducing operational overhead."

Source: AWS Certified Solutions Architect - Official Study Guide, Data Migration and Transfer section.

Question: 361

A company needs to connect its on-premises data center network to a new VPC. The data center network has a 100 Mbps symmetrical internet connection. An application that is running on premises will transfer multiple gigabytes of data each day. The application will use an Amazon Data Firehose delivery stream for processing.

What should a solutions architect recommend for maximum performance?

- A. Create a VPC peering connection between the on-premises network and the VPC. Configure routing for the on-premises network to use the VPC peering connection.
- B. Procure an AWS Snowball Edge Storage Optimized device. After several days' worth of data has accumulated, copy the data to the device and ship the device to AWS for expedited transfer to Firehose. Repeat as needed.
- C. Create an AWS Site-to-Site VPN connection between the on-premises network and the VPC. Configure BGP routing between the customer gateway and the virtual private gateway. Use the VPN connection to send the data from on premises to Firehose.
- D. Use AWS PrivateLink to create an interface VPC endpoint for Firehose in the VPC. Set up a 1 Gbps AWS Direct Connect connection between the on-premises network and AWS. Use the PrivateLink endpoint to send the data from on premises to Firehose.

Answer: D

Explanation:

AWS Direct Connect provides a dedicated network connection from on-premises to AWS, offering greater bandwidth and more consistent performance than internet-based connections or VPN. AWS PrivateLink enables secure, private connectivity to supported AWS services such as Kinesis Data Firehose over Direct Connect, bypassing the public internet and providing the highest throughput and lowest latency possible. This is the recommended solution for consistently transferring large volumes of data with maximum reliability and performance.

Reference Extract from AWS Documentation / Study Guide:

"AWS Direct Connect and AWS PrivateLink provide private, high-throughput connectivity between on-premises and AWS services, bypassing the public internet and ensuring maximum performance for large data transfers."

Source: AWS Certified Solutions Architect - Official Study Guide, Hybrid Networking section.

Question: 362

A company hosts an ecommerce application that stores all data in a single Amazon RDS for MySQL DB instance that is fully managed by AWS. The company needs to mitigate the risk of a single point of failure.

Which solution will meet these requirements with the LEAST implementation effort?

- A. Modify the RDS DB instance to use a Multi-AZ deployment. Apply the changes during the next maintenance window.
- B. Migrate the current database to a new Amazon DynamoDB Multi-AZ deployment. Use AWS Database Migration Service (AWS DMS) with a heterogeneous migration strategy to migrate the current RDS DB instance to DynamoDB tables.
- C. Create a new RDS DB instance in a Multi-AZ deployment. Manually restore the data from the existing RDS DB instance from the most recent snapshot.
- D. Configure the DB instance in an Amazon EC2 Auto Scaling group with a minimum group size of three. Use Amazon Route 53 simple routing to distribute requests to all DB instances.

Answer: A

Explanation:

Enabling Multi-AZ deployment for an Amazon RDS DB instance is the simplest and most effective way to improve database availability and eliminate a single point of failure. Multi-AZ deployments provide automatic failover to a standby instance in case of a failure of the primary instance. This approach is managed by AWS and only requires a modification to the existing instance, with minimal manual intervention or downtime, thus providing the least implementation effort.

Reference Extract from AWS Documentation / Study Guide:

"With Multi-AZ deployments, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). In case of a failure of the primary DB Instance, Amazon RDS automatically fails over to the

standby so that database operations can resume quickly."

Source: AWS Certified Solutions Architect - Official Study Guide, RDS High Availability section.

Question: 363

A retail company runs its application on AWS. The application uses Amazon EC2 for web servers, Amazon RDS for database services, and Amazon CloudFront for global content distribution.

The company needs a solution to mitigate DDoS attacks.

Which solution will meet this requirement?

- A. Implement AWS WAF custom rules to limit the length of query requests. Configure CloudFront to work with AWS WAF.
- B. Enable AWS Shield Advanced. Configure CloudFront to work with Shield Advanced.
- C. Use Amazon Inspector to scan the EC2 instances. Enable Amazon GuardDuty.
- D. Enable Amazon Macie. Configure CloudFront Origin Shield.

Answer: B

Explanation:

AWS Shield Advanced provides advanced DDoS protection for AWS workloads, including EC2, CloudFront, and RDS. When integrated with CloudFront, Shield Advanced offers comprehensive detection and mitigation against large and sophisticated DDoS attacks, along with 24x7 access to the AWS DDoS Response Team (DRT). AWS WAF provides application-level protection, but for complete DDoS mitigation, Shield Advanced is the recommended solution.

Reference Extract from AWS Documentation / Study Guide:

"AWS Shield Advanced provides expanded DDoS attack protection for applications running on AWS. It offers always-on detection and automatic inline mitigations that minimize application downtime and latency."

Source: AWS Certified Solutions Architect - Official Study Guide, Security and DDoS Protection section.

Question: 364

A company wants to share data that is collected from self-driving cars with the automobile community. The data will be made available from within an Amazon S3 bucket. The company wants to minimize its cost of making this data available to other AWS accounts.

What should a solutions architect do to accomplish this goal?

- A. Create an S3 VPC endpoint for the bucket.

- B. Configure the S3 bucket to be a Requester Pays bucket.
- C. Create an Amazon CloudFront distribution in front of the S3 bucket.
- D. Require that the files be accessible only with the use of the BitTorrent protocol.

Answer: B

Explanation:

The Requester Pays feature in Amazon S3 allows the bucket owner to configure the bucket so that the requester, rather than the bucket owner, pays for data transfer and request costs. This is ideal for sharing large datasets with the public or with other AWS accounts when you want to minimize your own data transfer expenses.

Reference Extract from AWS Documentation / Study Guide:

"Requester Pays buckets allow you to configure the bucket so that the requester instead of the bucket owner pays the cost of the request and the data download from the bucket."

Source: AWS Certified Solutions Architect - Official Study Guide, S3 Cost Management section.

Question: 365

A company's application receives requests from customers in JSON format. The company uses Amazon Simple Queue Service (Amazon SQS) to handle the requests.

After the application's most recent update, the company's customers reported that requests were being duplicated. A solutions architect discovers that the application is consuming messages from the SQS queue more than once.

What is the root cause of the issue?

- A. The visibility timeout is longer than the time it takes the application to process messages from the queue.
- B. The duplicated messages in the SQS queue contain unescaped Unicode characters.
- C. The message size exceeds the maximum of 256 KiB for each SQS message.
- D. The visibility timeout is shorter than the time it takes the application to process messages from the queue.

Answer: D

Explanation:

If the SQS visibility timeout is set shorter than the time it takes for the application to process and delete the message, the message becomes visible to other consumers and can be processed again, resulting in duplicate processing. This is a common cause of duplicate messages when using SQS.

Reference Extract from AWS Documentation / Study Guide:

"If the visibility timeout for a message is set shorter than the time it takes to process the message, the message becomes visible again and can be received and processed again, resulting in duplicate processing."

Source: AWS Certified Solutions Architect - Official Study Guide, SQS and Messaging section.

Question: 366

A company has a production Amazon RDS for MySQL database. The company needs to create a new application that will read frequently changing data from the database with minimal impact on the database's overall performance. The application will rarely perform the same query more than once.

What should a solutions architect do to meet these requirements?

- A. Set up an Amazon ElastiCache cluster. Query the results in the cluster.
- B. Set up an Application Load Balancer (ALB). Query the results in the ALB.
- C. Set up a read replica for the database. Query the read replica.
- D. Set up querying of database snapshots. Query the database snapshots.

Answer: C

Explanation: Amazon RDS read replicas provide a way to offload read traffic from the primary database, allowing read-intensive applications to query the replica without impacting the performance of the production (write) database. This is especially effective for workloads that involve frequently changing data but do not benefit from caching, since queries are rarely repeated.

Reference Extract from AWS Documentation / Study Guide:

"Read replicas allow you to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads."

Source: AWS Certified Solutions Architect - Official Study Guide, RDS Read Replica section.

Question: 367

A company is migrating its online shopping platform to AWS and wants to adopt a serverless architecture.

The platform has a user profile and preference service that does not have a defined schema. The platform allows user-defined fields.

Profile information is updated several times daily. The company must store profile information in a durable and highly available solution. The solution must capture modifications to profile data for future processing.

Which solution will meet these requirements?

- A. Use an Amazon RDS for PostgreSQL instance to store profile data. Use a log stream in Amazon CloudWatch Logs to capture modifications.
- B. Use an Amazon DynamoDB table to store profile data. Use Amazon DynamoDB Streams to capture modifications.
- C. Use an Amazon ElastiCache (Redis OSS) cluster to store profile data. Use Amazon Data Firehose to capture modifications.
- D. Use an Amazon Aurora Serverless v2 cluster to store the profile data. Use a log stream in Amazon CloudWatch Logs to capture modifications.

Answer: B

Explanation:

Amazon DynamoDB is a serverless, NoSQL database that is fully managed, highly available, and scales automatically. It is ideal for data without a fixed schema and for use cases where fields can vary by user. DynamoDB Streams enables the capture of changes to table items in real time, which is ideal for triggering additional processing or workflows on data modifications.

Reference Extract from AWS Documentation / Study Guide:

"DynamoDB provides a scalable, highly available NoSQL database service for applications requiring flexible schema. DynamoDB Streams captures table activity for processing changes in real time."

Source: AWS Certified Solutions Architect - Official Study Guide, DynamoDB and Serverless section.

Question: 368

A company runs a mobile game app on AWS. The app stores data for every user session. The data updates frequently during a gaming session. The app stores up to 256 KB for each session. Sessions can last up to 48 hours.

The company wants to automate the deletion of expired session data. The company must be able to restore all session data automatically if necessary.

Which solution will meet these requirements?

- A. Use an Amazon DynamoDB table to store the session data. Enable point-in-time recovery (PITR) and TTL for the table. Select the corresponding attribute for TTL in the session data.
- B. Use an Amazon MemoryDB table to store the session data. Enable point-in-time recovery (PITR) and TTL for the table. Select the

corresponding attribute for TTL in the session data.

C. Store session data in an Amazon S3 bucket. Use the S3 Standard storage class. Enable S3 Versioning for the bucket. Create an S3 Lifecycle configuration to expire objects after 48 hours.

D. Store session data in an Amazon S3 bucket. Use the S3 Intelligent-Tiering storage class. Enable S3 Versioning for the bucket. Create an S3 Lifecycle configuration to expire objects after 48 hours.

Answer: A

Explanation:

Amazon DynamoDB supports TTL (Time To Live) for automated, scheduled deletion of expired items. It also offers point-in-time recovery (PITR) to restore the table to any second within the retention window (typically up to 35 days), providing full data durability and protection. DynamoDB can efficiently handle frequent updates and offers predictable performance. MemoryDB is an in-memory store, not designed for durable recovery. S3 with lifecycle policies does not handle updates as efficiently for small, frequent writes and is not as optimal for session data.

Reference Extract:

"DynamoDB supports TTL for automated expiration and deletion of items and PITR for continuous backups and restoration of data."

Source: AWS Certified Solutions Architect - Official Study Guide, DynamoDB section.

Question: 369

A company runs a Windows-based ecommerce application on Amazon EC2 instances. The application has a very high transaction rate. The company requires a durable storage solution that can deliver 200,000 IOPS for each EC2 instance.

Which solution will meet these requirements?

- A. Host the application on EC2 instances that have Provisioned IOPS SSD (io2) Block Express Amazon Elastic Block Store (Amazon EBS) volumes attached.
- B. Install the application on an Amazon EMR cluster. Use Hadoop Distributed File System (HDFS) with General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volumes.
- C. Use Amazon FSx for Lustre as shared storage across the EC2 instances that run the application.
- D. Host the application on EC2 instances that have SSD instance store volumes and General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volumes attached.

Answer: A

Explanation:

Amazon EBS io2 Block Express volumes are designed to deliver sub-millisecond latency and up to 256,000 IOPS per volume, with durability and high availability. This makes io2 Block Express the recommended choice for workloads requiring very high and predictable IOPS, such as enterprise databases and high-transaction-rate applications.

Reference Extract:

"EBS io2 Block Express volumes deliver up to 256,000 IOPS and sub-millisecond latency, supporting high-performance, high-durability workloads."

Source: AWS Certified Solutions Architect - Official Study Guide, EBS Performance section.

Question: 370

An ecommerce company stores terabytes of customer data in the AWS Cloud. The data contains personally identifiable information (PII). The company wants to use the data in three applications. Only one of the applications needs to process the PII. The PII must be removed before the other two applications process the data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Store the data in an Amazon DynamoDB table. Create a proxy application layer to intercept and process the data that each application requests.
- B. Store the data in an Amazon S3 bucket. Process and transform the data by using S3 Object Lambda before returning the data to the requesting application.
- C. Process the data and store the transformed data in three separate Amazon S3 buckets so that each application has its own custom dataset. Point each application to its respective S3 bucket.
- D. Process the data and store the transformed data in three separate Amazon DynamoDB tables so that each application has its own custom dataset. Point each application to its respective DynamoDB table.

Answer: B

Explanation:

Amazon S3 Object Lambda allows you to add your own code to process data retrieved from S3 before it is returned to an application. You can use this to dynamically redact or remove PII for specific applications on-the-fly, eliminating the need to manage multiple buckets or datasets, thus minimizing operational overhead.

Reference Extract:

"S3 Object Lambda enables you to process and transform data as it is retrieved from S3, supporting use cases such as redacting sensitive information before returning data to an application."

Source: AWS Certified Solutions Architect - Official Study Guide, Data Security and Transformation section.

Question: 371

A company has an application that runs on a single Amazon EC2 instance. The application uses a

MySQL database that runs on the same EC2 instance. The company needs a highly available and automatically scalable solution to handle increased traffic.

Which solution will meet these requirements?

- A. Deploy the application to EC2 instances that run in an Auto Scaling group behind an Application Load Balancer. Create an Amazon Redshift cluster that has multiple MySQL-compatible nodes.
- B. Deploy the application to EC2 instances that are configured as a target group behind an Application Load Balancer. Create an Amazon RDS for MySQL cluster that has multiple instances.
- C. Deploy the application to EC2 instances that run in an Auto Scaling group behind an Application Load Balancer. Create an Amazon Aurora Serverless MySQL cluster for the database layer.
- D. Deploy the application to EC2 instances that are configured as a target group behind an Application Load Balancer. Create an Amazon ElastiCache (Redis OSS) cluster that uses the MySQL connector.

Answer: C

Explanation:

Amazon Aurora Serverless is a fully managed, MySQL-compatible database that automatically scales based on demand and provides high availability. Combining this with EC2 Auto Scaling and an Application Load Balancer achieves both application and database high availability and scalability.

Reference Extract:

"Aurora Serverless automatically starts up, shuts down, and scales capacity based on your application's needs, providing a cost-effective, highly available database solution."

Source: AWS Certified Solutions Architect - Official Study Guide, Aurora Serverless and Scaling section.

Question: 372

A company runs multiple workloads in separate AWS environments. The company wants to optimize its AWS costs but must maintain the

same level of performance for the environments.

The company's production environment requires resources to be highly available. The other environments do not require highly available resources.

Each environment has the same set of networking components, including the following:

- 1 VPC
- 1 Application Load Balancer
- 4 subnets distributed across 2 Availability Zones (2 public subnets and 2 private subnets)
- 2 NAT gateways (1 in each public subnet)
- 1 internet gateway

Which solution will meet these requirements?

- A. Do not change the production environment workload. For each non-production workload, remove one NAT gateway and update the route tables for private subnets to target the remaining NAT gateway for the destination 0.0.0.0/0.
- B. Reduce the number of Availability Zones that all workloads in all environments use.
- C. Replace every NAT gateway with a t4g.large NAT instance. Update the route tables for each private subnet to target the NAT instance that is in the same Availability Zone for the destination 0.0.0.0/0.
- D. In each environment, create one transit gateway and remove one NAT gateway. Configure routing on the transit gateway to forward traffic for the destination 0.0.0.0/0 to the remaining NAT gateway. Update private subnet route tables to target the transit gateway for the destination 0.0.0.0/0.

Answer: A

Explanation:

Maintaining two NAT gateways for production ensures high availability. Reducing to one NAT gateway in non-production environments lowers cost while maintaining necessary connectivity. This approach is recommended by AWS for cost optimization in non-critical environments.

Reference Extract:

"For environments that do not require high availability, you can reduce costs by using a single NAT gateway and updating route tables accordingly."

Source: AWS Certified Solutions Architect - Official Study Guide, Cost Optimization and NAT Gateway section.

Question: 373

An ecommerce company hosts a three-tier web application in a VPC. The web tier runs on Amazon EC2 instances in two Availability Zones. The company stores a product catalog and customer sales information in Amazon DynamoDB.

The company's finance team uses a reporting application to generate reports of daily product sales.

When the finance team runs the daily reports, a sudden performance decrease affects website customers.

The company wants to improve the performance of the system.

Which solution will meet these requirements with MINIMAL changes to the current architecture?

- A. Migrate the application to larger EC2 instances. Migrate the database to Amazon RDS for MySQL. Configure a read replica of the database in a second Availability Zone.
- B. Increase the compute capacity of the EC2 instances. Migrate the database to Amazon ElastiCache (Memcached).
- C. Implement DynamoDB Accelerator (DAX).
- D. Configure DynamoDB streams.

Answer: C

Explanation:

DynamoDB Accelerator (DAX) is a fully managed, in-memory cache for DynamoDB that delivers up to a 10x performance improvement—even at millions of requests per second. DAX requires minimal changes to existing applications and offloads the read workload from DynamoDB during report generation.

Reference Extract:

"DAX provides in-memory acceleration for DynamoDB tables, reducing response times from milliseconds to microseconds with minimal application changes."

Source: AWS Certified Solutions Architect - Official Study Guide, DynamoDB and Performance section.

Question: 374

A solutions architect creates an Auto Scaling group for a memory-intensive application. The solutions architect wants to scale up and scale down based on memory usage. Which solution will meet this requirement?

- A. Install and configure the AWS Systems Manager Agent (SSM Agent). Create a step scaling policy that has step adjustments based on the memory usage trend.
- B. Install and configure the Amazon CloudWatch agent. Create a target tracking policy to scale based on the mem_used_percent

CloudWatch metric.

C. Install and configure the AWS Systems Manager Agent (SSM Agent). Create a target tracking policy to scale based on the mem_used_percent Amazon CloudWatch metric.

D. Install and configure the Amazon CloudWatch agent. Create a scheduled scaling policy to scale based on the memory usage trend.

Answer: B

Explanation:

The Amazon CloudWatch agent is required to collect memory utilization metrics (as memory metrics are not reported by default). A target tracking policy is the simplest and most effective way to scale based on a custom metric such as mem_used_percent.

Reference Extract:

"Install the CloudWatch agent to collect memory metrics, and create a target tracking scaling policy using these custom metrics."

Source: AWS Certified Solutions Architect - Official Study Guide, Monitoring and Scaling section.

Question: 375

A company wants to migrate its accounting system from an on-premises data center to the AWS Cloud in a single AWS Region. Data security and an immutable audit log are the top priorities. The company must monitor all AWS activities for compliance auditing. The company has enabled AWS CloudTrail but wants to make sure it meets these requirements.

Which actions should a solutions architect take to protect and secure CloudTrail? (Select TWO.)

A. Enable CloudTrail log file validation.

B. Install the CloudTrail Processing Library.

C. Enable logging of Insights events in CloudTrail.

D. Enable custom logging from the on-premises resources.

E. Create an AWS Config rule to monitor whether CloudTrail is configured to use server-side encryption with AWS KMS managed encryption keys (SSE-KMS).

Answer: A,E

Explanation:

CloudTrail log file validation ensures that the log files have not been altered or deleted after delivery, providing an immutable audit log.

Using KMS-managed encryption keys for CloudTrail log files adds another layer of data security, and AWS Config can monitor compliance to ensure this security is always enforced.

Reference Extract:

"CloudTrail log file validation provides assurance about the integrity of CloudTrail logs. Using SSE- KMS encryption and monitoring with AWS Config helps secure and audit logs for compliance."

Source: AWS Certified Solutions Architect - Official Study Guide, CloudTrail Security section.

Question: 376

A solutions architect is designing the architecture for a two-tier web application. The web application consists of an internet-facing Application Load Balancer (ALB) that forwards traffic to an Auto Scaling group of Amazon EC2 instances.

The EC2 instances must be able to access an Amazon RDS database. The company does not want to rely solely on security groups or network ACLs. Only the minimum resources that are necessary should be routable from the internet.

Which network design meets these requirements?

- A. Place the ALB, EC2 instances, and RDS database in private subnets.
- B. Place the ALB in public subnets. Place the EC2 instances and RDS database in private subnets.
- C. Place the ALB and EC2 instances in public subnets. Place the RDS database in private subnets.
- D. Place the ALB outside the VPC. Place the EC2 instances and RDS database in private subnets.

Answer: B

Explanation:

The ALB must be in a public subnet to receive internet traffic. The EC2 instances and the RDS database should be in private subnets to prevent direct internet access, minimizing the attack surface. This aligns with AWS security best practices for web application architectures.

Reference Extract:

"Internet-facing ALBs should be placed in public subnets; EC2 instances and RDS databases should be in private subnets to restrict direct internet access."

Source: AWS Certified Solutions Architect - Official Study Guide, Network Security and Design section.

Question: 377

A company is developing a microservices-based application to manage the company's delivery operations. The application consists of

microservices that process orders, manage a fleet of delivery vehicles, and optimize delivery routes.

The microservices must be able to scale independently and must be able to handle bursts of traffic without any data loss.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon API Gateway REST APIs to establish communication between microservices. Deploy the application on Amazon EC2 instances in Auto Scaling groups.
- B. Use Amazon Simple Queue Service (Amazon SQS) to establish communication between microservices. Deploy the application on Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate.
- C. Use WebSocket-based communication between microservices. Deploy the application on Amazon EC2 instances in Auto Scaling groups.
- D. Use Amazon Simple Notification Service (Amazon SNS) to establish communication between microservices. Deploy the application on Amazon Elastic Container Service (Amazon ECS) containers on Amazon EC2 instances.

Answer: B

Explanation:

Amazon SQS is a fully managed message queuing service that reliably decouples and scales microservices, distributed systems, and serverless applications. By using SQS, microservices can communicate asynchronously, handle bursts of traffic, and avoid data loss by buffering messages until they are processed. Deploying the services on ECS with AWS Fargate further reduces operational overhead by removing the need to manage servers, allowing independent scaling of each microservice.

Reference Extract:

"Amazon SQS decouples application components and enables message durability and scaling. AWS Fargate removes the need to manage infrastructure, supporting independent scaling and minimal operational overhead."

Source: AWS Certified Solutions Architect - Official Study Guide, Microservices and Messaging section.

Question: 378

A company hosts its main public web application in one AWS Region across multiple Availability Zones. The application uses an Amazon EC2 Auto Scaling group and an Application Load Balancer (ALB).

A web development team needs a cost-optimized compute solution to improve the company's ability to serve dynamic content globally to millions of customers.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution. Configure the existing ALB as the origin.
- B. Use Amazon Route 53 to serve traffic to the ALB and EC2 instances based on the geographic location of each customer.
- C. Create an Amazon S3 bucket with public read access enabled. Migrate the web application to the S3 bucket. Configure the S3 bucket for website hosting.
- D. Use AWS Direct Connect to directly serve content from the web application to the location of each customer.

Answer: A

Explanation:

Amazon CloudFront is a global content delivery network (CDN) that caches and distributes content to users with low latency. By setting the existing ALB as the origin, CloudFront can cache dynamic and static content closer to users worldwide, improving performance and reducing the load on the application servers. This is the most cost-optimized and AWS-recommended way to globally serve dynamic content for web applications.

Reference Extract:

"CloudFront accelerates delivery of both static and dynamic content, reducing latency and offloading origin resources by caching content at edge locations."

Source: AWS Certified Solutions Architect - Official Study Guide, CloudFront and Global Application section.

Question: 379

A company is planning to deploy a managed MySQL database solution for its non-production applications. The company plans to run the system for several years on AWS. Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon RDS for MySQL instance. Purchase a Reserved Instance.
- B. Create an Amazon RDS for MySQL instance. Use the instance on an on-demand basis.
- C. Create an Amazon Aurora MySQL cluster with writer and reader nodes. Use the cluster on an on-demand basis.
- D. Create an Amazon EC2 instance. Manually install and configure MySQL Server on the instance.

Answer: A

Explanation:

Amazon RDS for MySQL Reserved Instances provide significant savings over on-demand pricing when you plan to run the database for long periods. This is the most cost-effective option for nonproduction, long-running managed MySQL workloads.

Reference Extract:

"Reserved Instances provide a significant discount compared to On-Demand pricing and are recommended for steady-state workloads that run for an extended period."

Source: AWS Certified Solutions Architect - Official Study Guide, RDS Cost Optimization section.

Question: 380

A company runs an application on Amazon EC2 instances across multiple Availability Zones in the same AWS Region. The EC2 instances share an Amazon Elastic File System (Amazon EFS) volume that is mounted on all the instances. The EFS volume stores a variety of files such as installation media, third-party files, interface files, and other one-time files.

The company accesses some EFS files frequently and needs to retrieve the files quickly. The company accesses other files rarely. The EFS volume is multiple terabytes in size. The company needs to optimize storage costs for Amazon EFS.

Which solution will meet these requirements with the LEAST effort?

- A. Move the files to Amazon S3. Set up a lifecycle policy to move the files to S3 Glacier Flexible Retrieval.
- B. Apply a lifecycle policy to the EFS files to move the files to EFS Infrequent Access.
- C. Move the files to Amazon Elastic Block Store (Amazon EBS) Cold HDD Volumes (sc1).
- D. Move the files to Amazon S3. Set up a lifecycle policy to move the rarely-used files to S3 Glacier Deep Archive.

Answer: B

Explanation:

Amazon EFS offers an Infrequent Access (IA) storage class, which can be managed via EFS lifecycle policies. Frequently accessed files remain in the Standard storage class, while infrequently accessed files are automatically moved to the IA class, significantly reducing storage costs with minimal effort and no application changes.

Reference Extract:

"EFS lifecycle management automatically transitions files that are not accessed for a set period to the EFS Infrequent Access (IA) storage class, reducing storage costs."

Source: AWS Certified Solutions Architect - Official Study Guide, EFS and Lifecycle Management section.

Question: 381

A company needs to implement a new data retention policy for regulatory compliance. As part of this policy, sensitive documents that are stored in an Amazon S3 bucket must be protected from deletion or modification for a fixed period of time.

Which solution will meet these requirements?

- A. Activate S3 Object Lock on the required objects and enable governance mode.
- B. Activate S3 Object Lock on the required objects and enable compliance mode.
- C. Enable versioning on the S3 bucket. Set a lifecycle policy to delete the objects after a specified period.
- D. Configure an S3 Lifecycle policy to transition objects to S3 Glacier Flexible Retrieval for the retention duration.

Answer: B

Explanation:

S3 Object Lock in Compliance Mode prevents objects from being deleted or overwritten for a fixed, specified retention period. Compliance Mode is specifically designed to meet regulatory requirements, ensuring that no user, including the root user, can modify or delete the object during the retention period.

Reference Extract:

"S3 Object Lock in compliance mode protects objects from being deleted or overwritten during the specified retention period, helping meet regulatory retention requirements."

Source: AWS Certified Solutions Architect - Official Study Guide, S3 Object Lock and Compliance section.

Question: 382

A company is using Amazon DocumentDB global clusters to support an ecommerce application. The application serves customers across multiple AWS Regions. To ensure business continuity, the company needs a solution to minimize downtime during maintenance windows or other disruptions.

Which solution will meet these requirements?

- A. Regularly create manual snapshots of the DocumentDB instance in the primary Region.
- B. Perform a managed failover to a secondary Region when needed.
- C. Perform a failover to a replica DocumentDB instance within the primary Region.
- D. Configure increased replication lag to manage cross-Region replication.

Answer: B

Explanation:

Amazon DocumentDB global clusters support managed cross-region failover, allowing you to promote a secondary region to become the new primary with minimal downtime. This ensures business continuity during maintenance or regional disruptions.

Reference Extract:

"Amazon DocumentDB global clusters support managed cross-Region failover, allowing you to recover quickly from regional disruptions with minimal downtime."

Source: AWS Certified Solutions Architect - Official Study Guide, DocumentDB and Resiliency section.

Question: 383

The lead member of a DevOps team creates an AWS account. A DevOps engineer shares the account credentials with a solutions architect through a password manager application.

The solutions architect needs to secure the root user for the new account.

Which actions will meet this requirement? (Select TWO.)

- A. Update the root user password to a new, strong password.
- B. Secure the root user account by using a virtual multi-factor authentication (MFA) device.
- C. Create an IAM user for each member of the DevOps team. Assign the AdministratorAccess AWS managed policy to each IAM user.
- D. Create root user access keys. Save the keys as a new parameter in AWS Systems Manager Parameter Store.
- E. Update the IAM role for the root user to ensure the root user can use only approved services.

Answer: A,B

Explanation:

Securing the root user account requires setting a strong password and enabling multi-factor authentication (MFA). AWS recommends never sharing the root user credentials, setting up individual IAM users for everyday operations, and always protecting the root user with MFA for maximum security.

Reference Extract:

"AWS recommends securing the root user with a strong password and enabling multi-factor authentication (MFA). Do not use or share root credentials for everyday tasks."

Source: AWS Certified Solutions Architect - Official Study Guide, IAM and Security Best Practices section.

Question: 384

A company wants to run a hybrid workload for data processing. The data needs to be accessed by onpremises applications for local data processing using an NFS protocol, and must also be accessible from the AWS Cloud for further analytics and batch processing.

Which solution will meet these requirements?

- A. Use an AWS Storage Gateway file gateway to provide file storage to AWS, then perform analytics on this data in the AWS Cloud.
- B. Use an AWS Storage Gateway tape gateway to copy the backup of the local data to AWS, then perform analytics on this data in the AWS Cloud.
- C. Use an AWS Storage Gateway volume gateway in a stored volume configuration to regularly take snapshots of the local data, then copy the data to AWS.
- D. Use an AWS Storage Gateway volume gateway in a cached volume configuration to back up all the local storage in the AWS Cloud, then perform analytics on this data in the cloud.

Answer: A

Explanation:

AWS Storage Gateway file gateway presents a file interface backed by Amazon S3 and supports NFS. This allows local applications to access data via NFS while also enabling cloud applications to use the data stored in S3 for analytics and processing, fulfilling both hybrid and cloud-native requirements.

Reference Extract:

"AWS Storage Gateway file gateway offers NFS and SMB access to data stored in Amazon S3, supporting hybrid workloads for local and cloud access."

Source: AWS Certified Solutions Architect - Official Study Guide, Hybrid and Storage Gateway section.

Question: 385

A company wants to use a data lake that is hosted on Amazon S3 to provide analytics services for historical data. The data lake consists of 800 tables but is expected to grow to thousands of tables. More than 50 departments use the tables, and each department has hundreds of users. Different departments need access to specific tables and columns.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM role for each department. Use AWS Lake Formation based access control to grant each IAM role access to specific tables and columns. Use Amazon Athena to analyze the data.

- B. Create an Amazon Redshift cluster for each department. Use AWS Glue to ingest into the Redshift cluster only the tables and columns that are relevant to that department. Create Redshift database users. Grant the users access to the relevant department's Redshift cluster. Use Amazon Redshift to analyze the data.
- C. Create an IAM role for each department. Use AWS Lake Formation tag-based access control to grant each IAM role access to only the relevant resources. Create LF-tags that are attached to tables and columns. Use Amazon Athena to analyze the data.
- D. Create an Amazon EMR cluster for each department. Configure an IAM service role for each EMR cluster to access relevant S3 files. For each department's users, create an IAM role that provides access to the relevant EMR cluster. Use Amazon EMR to analyze the data.

Answer: C

Explanation:

The requirement is to provide granular, scalable access to thousands of tables and columns in a data lake across many users and departments, with the least operational overhead.

AWS Lake Formation supports tag-based access control (TBAC) using LF-tags (Lake Formation tags), which allows you to assign tags to tables, columns, and databases. You can then define permissions on resources by specifying tags rather than managing permissions for individual resources. This approach is highly scalable and efficient when dealing with a growing number of tables and columns. By associating IAM roles to departments and granting access based on LF-tags, you dramatically reduce the operational burden as new tables or columns are added; you only need to assign the appropriate tags.

Amazon Athena can directly query data in S3 with Lake Formation providing fine-grained access control.

AWS Documentation Extract:

"With LF-tag-based access control, you can grant permissions to resources based on tags, making it easy to manage access at scale, especially in environments with large and dynamic numbers of resources."

"LF-tags provide a scalable way to manage permissions for large numbers of resources without having to define permissions individually for each table or column."

(Source: AWS Lake Formation documentation, Access Control, Tag-Based Access Control)

Other options:

A: Would require managing explicit permissions for each table and column as the environment grows, increasing operational overhead.

B & D: Involve significant duplication of resources (clusters) and do not scale as efficiently as a centralized data lake with tag-based access.

Reference: AWS Certified Solutions Architect - Official Study Guide, Chapter on Data Lakes and Access Control.

Question: 386

A company must follow strict regulations for the management of data encryption keys. The company manages its own key externally and imports the key into AWS Key Management Service (AWS KMS). The company must control the imported key material and must rotate the key material on a regular schedule.

A solutions architect needs to import the key material into AWS KMS and rotate the key without interrupting applications that use the key.

Which solution will meet these requirements?

- A. Create a new AWS KMS key that has the same key ID as the existing key. Import new key material into the key.
- B. Schedule the existing AWS KMS key for deletion. Create a new KMS key that has new key material.
- C. Import new key material into the existing AWS KMS key. Set an expiration time for the old key material.
- D. Enable automatic key rotation for the existing AWS KMS key.

Answer: C

Explanation:

When using imported key material with AWS KMS, you maintain control over the key lifecycle. AWS KMS allows you to import new key material into an existing KMS key (of type "external" or "imported"), thus rotating the key material without changing the key ID or ARNs. This enables applications to continue using the same key for cryptographic operations without disruption.

You can also set an expiration time for the old key material, after which AWS KMS deletes the old material and requires new key material to be imported, enforcing regular rotation per your compliance requirements.

AWS Documentation Extract:

"To rotate imported key material, you can re-import new key material into the same KMS key. This retains the same key ID and ARNs so applications are unaffected. You can set an expiration time for imported key material and replace it as needed, ensuring compliance with your rotation policy."

(Source: AWS Key Management Service Developer Guide, Importing Key Material, Rotating Key Material)

Other options:

A: You cannot create a new KMS key with the same key ID as an existing one.

B: Deleting and recreating the key disrupts application access because the key ID changes.

D: Automatic rotation is only available for AWS-managed keys, not for imported key material.

Reference: AWS Certified Solutions Architect - Official Study Guide, KMS and Key Rotation Section.

Question: 387

A company runs production workloads in its AWS account. Multiple teams create and maintain the workloads.

The company needs to be able to detect changes in resource configurations. The company needs to capture changes as configuration items without changing or modifying the existing resources.

Which solution will meet these requirements?

- A. Use AWS Config. Start the configuration recorder for AWS resources to detect changes in resource configurations.
- B. Use AWS CloudFormation. Initiate drift detection to capture changes in resource configurations.
- C. Use Amazon Detective to detect, analyze, and investigate changes in resource configurations.
- D. Use AWS Audit Manager to capture management events and global service events for resource configurations.

Answer: A

Explanation:

AWS Config is a service designed to assess, audit, and evaluate the configurations of AWS resources. It continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. By starting a configuration recorder, AWS Config will capture changes to supported resource types as configuration items— without the need to modify any of the existing resources. This provides a full history of configuration changes and is specifically intended for exactly this use case.

AWS Documentation Extract:

“AWS Config provides a detailed view of the configuration of AWS resources in your AWS account. This includes how the resources are related to one another and how they were configured in the past so you can see how the configurations and relationships change over time.”

“You can start the configuration recorder, which will record the configuration changes of the supported resources in your AWS account.”

(Source: AWS Config documentation, What is AWS Config?)

Other options:

- B: CloudFormation drift detection only works for resources created and managed by CloudFormation and requires stacks.
- C: Amazon Detective is used for analyzing and investigating security findings, not for resource configuration tracking.
- D: AWS Audit Manager is used for automating evidence collection to help with audits, not for tracking resource configurations.

Question: 388

A company runs an on-premises application on a Kubernetes cluster. The company recently added millions of new customers. The company's existing on-premises infrastructure is unable to handle the large number of new customers. The company needs to migrate the on-premises application to the

AWS Cloud.

The company will migrate to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The company does not want to manage the underlying compute infrastructure for the new architecture on AWS.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use a self-managed node to supply compute capacity. Deploy the application to the new EKS cluster.
- B. Use managed node groups to supply compute capacity. Deploy the application to the new EKS cluster.
- C. Use AWS Fargate to supply compute capacity. Create a Fargate profile. Use the Fargate profile to deploy the application.
- D. Use managed node groups with Karpenter to supply compute capacity. Deploy the application to the new EKS cluster.

Answer: C

Explanation:

AWS Fargate is a serverless compute engine for containers that works with Amazon EKS. With Fargate, you do not need to provision or manage EC2 instances or clusters. You simply define and deploy your pods, and Fargate automatically launches the required compute resources. This results in the lowest operational overhead because AWS manages the infrastructure. Fargate profiles allow you to specify which pods run on Fargate.

AWS Documentation Extract:

“AWS Fargate is a serverless, pay-as-you-go compute engine that lets you focus on building applications without managing servers. With Amazon EKS and Fargate, you only need to define your application's pods; Fargate provisions and manages the required compute resources for you.”

(Source: Amazon EKS documentation, AWS Fargate integration)

Other options:

A: Self-managed nodes require you to manage EC2 instances.

B: Managed node groups reduce some overhead, but you are still responsible for patching and managing the EC2 instances.

D: Managed node groups with Karpenter automate scaling but do not remove the need to manage underlying instances.

Reference: AWS Certified Solutions Architect - Official Study Guide, EKS and Fargate Section.

Question: 389

A company runs a content management system on an Amazon Elastic Container Service (Amazon ECS) cluster. The system allows visitors to provide feedback about the company's products by uploading documents and photos of the products to an Amazon S3 bucket.

The company has a workflow on AWS that processes uploaded documents to perform sentiment analysis of photos and text. The processing workflow calls multiple AWS services.

The company needs a solution to automate the processing workflow. The solution must handle any failed uploads.

Which solution will meet these requirements with the LEAST effort?

- A. Use S3 Event Notifications to publish events to an Amazon Simple Notification Service (Amazon SNS) topic. Deploy a web application on the Amazon ECS cluster to subscribe to the SNS topic and listen for events to orchestrate the processing workflow.
- B. Use S3 Event Notifications to publish events to an Amazon Simple Queue Service (Amazon SQS) queue. Configure long polling. Deploy an Amazon EC2 instance that runs a script to orchestrate the processing workflow.
- C. Use S3 Event Notifications to publish events to an Amazon Simple Queue Service (Amazon SQS) queue. Create an ECS cluster that scales based on the number of messages in the queue. Configure the cluster to orchestrate the processing workflow.
- D. Use S3 Event Notifications to invoke an Amazon EventBridge rule. Configure the rule to initiate an AWS Step Functions workflow that orchestrates the processing workflow.

Answer: D

Explanation:

AWS Step Functions is designed to coordinate multiple AWS services into serverless workflows, handling errors, retries, and complex logic. By configuring Amazon S3 Event Notifications to trigger an Amazon EventBridge rule, you can automatically start a Step Functions workflow when a new object is uploaded to S3. Step Functions can manage retries for failed uploads and orchestrate calls to various AWS services with minimal operational effort.

AWS Documentation Extract:

"You can use Amazon EventBridge to initiate AWS Step Functions workflows in response to events from S3 buckets. Step Functions can orchestrate the sequence of AWS service calls and automatically handle retries and errors, making it ideal for automating and managing complex workflows."

(Source: AWS Step Functions documentation, Using Step Functions with EventBridge and S3)

Other options:

A: Requires manual implementation of orchestration and error handling.

B & C: Involve additional infrastructure management and custom logic for workflow orchestration and error handling, increasing operational effort.

Reference: AWS Certified Solutions Architect - Official Study Guide, Automation and Orchestration Section.

Question: 390

A company is developing a latency-sensitive application. Part of the application includes several AWS Lambda functions that need to initialize as quickly as possible. The Lambda functions are written in Java and contain initialization code outside the handlers to load libraries, initialize classes, and generate unique IDs.

Which solution will meet the startup performance requirement MOST cost-effectively?

A. Move all the initialization code to the handlers for each Lambda function. Activate Lambda SnapStart for each Lambda function. Configure SnapStart to reference the \$LATEST version of each Lambda function.

B. Publish a version of each Lambda function. Create an alias for each Lambda function. Configure each alias to point to its corresponding version. Set up a provisioned concurrency configuration for each Lambda function to point to the corresponding alias.

C. Publish a version of each Lambda function. Set up a provisioned concurrency configuration for each Lambda function to point to the corresponding version. Activate Lambda SnapStart for the published versions of the Lambda functions.

D. Update the Lambda functions to add a pre-snapshot hook. Move the code that generates unique IDs into the handlers. Publish a version of each Lambda function. Activate Lambda SnapStart for the published versions of the Lambda functions.

Answer: D

Explanation:

AWS Lambda SnapStart is designed to improve the cold start performance of Java Lambda functions by initializing the function, taking a snapshot of the execution environment, and then reusing that snapshot for subsequent invocations. However, some code (such as code that generates unique IDs or session-specific data) should run during each invocation, not during the snapshot process. By using a pre-snapshot hook and moving the unique ID generation into the handler, you ensure that non-deterministic or per-invocation code is executed correctly, while the rest of the initialization benefits from SnapStart. This delivers the lowest latency and cost, as you do not need to pay for provisioned concurrency.

AWS Documentation Extract:

"Lambda SnapStart is ideal for Java functions with long cold starts due to heavy initialization. Move non-deterministic code such as unique ID generation to the handler, and use pre-snapshot hooks to customize what gets snapshotted. SnapStart works only for published

versions, not \$LATEST."

(Source: AWS Lambda documentation, Using SnapStart for Java Functions)

Other options:

A: SnapStart is not supported for the \$LATEST version; must be a published version.

B & C: Provisioned concurrency removes cold starts but is more expensive than SnapStart for most workloads.

C: You must use SnapStart on published versions, but provisioned concurrency is not required for SnapStart and adds cost.

Reference: AWS Certified Solutions Architect - Official Study Guide, Lambda Performance Section.

Question: 391

A company is developing a latency-sensitive application. Part of the application includes several AWS Lambda functions that need to initialize as quickly as possible. The Lambda functions are written in Java and contain initialization code outside the handlers to load libraries, initialize classes, and generate unique IDs.

Which solution will meet the startup performance requirement MOST cost-effectively?

A. Move all the initialization code to the handlers for each Lambda function. Activate Lambda SnapStart for each Lambda function.

Configure SnapStart to reference the \$LATEST version of each Lambda function.

B. Publish a version of each Lambda function. Create an alias for each Lambda function. Configure each alias to point to its corresponding version. Set up provisioned concurrency configuration for each Lambda function to point to the corresponding alias.

C. Publish a version of each Lambda function. Set up a provisioned concurrency configuration for each Lambda function to point to the corresponding version. Activate Lambda SnapStart for the published versions of the Lambda functions.

D. Update the Lambda functions to add a pre-snapshot hook. Move the code that generates unique IDs into the handlers. Publish a version of each Lambda function. Activate Lambda SnapStart for the published versions of the Lambda functions.

Answer: D

Explanation:

Same scenario as above—cost-effective cold start reduction for Java Lambdas is best achieved by using Lambda SnapStart with the correct code structure. Unique IDs and similar non-deterministic logic should be in the handler or in a pre-snapshot hook. SnapStart is supported only for published versions. This combination offers cold start performance at a lower cost than provisioned concurrency.

AWS Documentation Extract:

"When using Lambda SnapStart for Java, ensure non-deterministic initialization code, such as unique ID generators, is executed inside the handler or in a pre-snapshot hook. SnapStart is supported only on published versions."

(Source: AWS Lambda documentation, SnapStart for Java)

Reference: AWS Certified Solutions Architect - Official Study Guide, Lambda Performance Section.

Question: 392

A company wants to migrate a Microsoft SQL Server database server from an on-premises data center to AWS. The company needs access to the operating system of the SQL Server database.

Which solution will meet these requirements?

- A. Migrate the database to Amazon Aurora Serverless.
- B. Migrate the database to Amazon RDS for SQL Server.
- C. Migrate the database to Amazon EC2 instances that run SQL Server.
- D. Migrate the database to Amazon Redshift.

Answer: C

Explanation:

Amazon EC2 allows you to install and manage SQL Server directly on a Windows or Linux VM, giving you full access to the underlying operating system. This is required when you need to install custom agents, configure OS-level features, or have complete control over the environment.

Amazon RDS and Aurora are managed services and do not provide access to the underlying OS, nor does Redshift (which is a data warehouse, not SQL Server).

AWS Documentation Extract:

"If you require access to the underlying operating system or need to install custom software, use SQL Server on Amazon EC2. Amazon RDS is a managed service and does not provide OS-level access."

(Source: AWS Database Migration documentation, SQL Server Deployment Options)

Reference: AWS Certified Solutions Architect - Official Study Guide, Database Migration Section.

Question: 393

A company is creating a web application that will store a large number of images in Amazon S3. The images will be accessed by users over variable periods of time. The company wants to:

Retain all the images.

Incur no cost for retrieval.

Have minimal management overhead.

Have the images available with no impact on retrieval time.

Which solution meets these requirements?

- A. Implement S3 Intelligent-Tiering.
- B. Implement S3 storage class analysis.
- C. Implement an S3 Lifecycle policy to move data to S3 Standard-Infrequent Access (S3 Standard-IA).
- D. Implement an S3 Lifecycle policy to move data to S3 One Zone-Infrequent Access (S3 One Zone- IA).

Answer: A

Explanation:

S3 Intelligent-Tiering is designed for data with unknown or changing access patterns. It automatically moves objects between frequent and infrequent access tiers as needed, with no retrieval fees for accessing data in any tier, and no performance impact. Minimal management overhead is required because AWS manages all transitions automatically. This class is cost-optimized and meets all requirements listed.

AWS Documentation Extract:

"S3 Intelligent-Tiering is the only storage class that automatically moves data between frequent and infrequent access tiers when access patterns change, with no retrieval charges and no impact on performance. It is designed to optimize costs automatically when data access patterns are unpredictable."

(Source: Amazon S3 documentation, Intelligent-Tiering storage class)

Other options:

B: Storage class analysis only provides recommendations, not actual storage tiering.

C & D: S3 Standard-IA and S3 One Zone-IA both have retrieval fees and may impact retrieval time and data durability.

Reference: AWS Certified Solutions Architect - Official Study Guide, S3 Storage Classes Section.

Question: 394

An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and rafting photos and videos anytime. The photos and videos are stored in Amazon S3 Standard storage in an S3 bucket and are served through Amazon CloudFront.

The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The

solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.

Which solution will meet these requirements?

- A. Configure S3 Intelligent-Tiering on the S3 bucket.
- B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.
- C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.
- D. Add a Cache-Control: max-age header to the S3 image objects and S3 video objects. Set the header to 30 days.

Answer: A

Explanation:

Amazon S3 Intelligent-Tiering automatically moves objects between two access tiers based on changing access patterns. Objects not accessed for 30 days move to a lower-cost tier, but are still immediately available with millisecond retrieval. If objects become frequently accessed again, they are moved back to the frequent access tier. There are no retrieval charges and no impact on availability or performance. This storage class is specifically designed for unpredictable access patterns and cost optimization, requiring minimal management.

AWS Documentation Extract:

"S3 Intelligent-Tiering is the only storage class that automatically moves data between frequent and infrequent access tiers when access patterns change, with no retrieval charges and no impact on performance. It is designed to optimize costs automatically when data access patterns are unpredictable."

(Source: Amazon S3 documentation, Intelligent-Tiering storage class)

B: Glacier Deep Archive is for archival, not for low-latency millisecond access.

C: EFS is not optimized for object storage or global CDN distribution.

D: Cache-Control only affects CloudFront or browser caching, not S3 storage cost.

Reference: AWS Certified Solutions Architect - Official Study Guide, S3 Storage Classes.

Question: 395

A media streaming company is redesigning its infrastructure to accommodate increasing demand for video content that users consume daily. The company needs to process terabyte-sized videos to block some content in the videos. Video processing can take up to 20 minutes.

The company needs a solution that is cost-effective, highly available, and scalable.

Which solution will meet these requirements?

- A. Use AWS Lambda functions to process the videos. Store video metadata in Amazon DynamoDB. Store video content in Amazon S3 Intelligent-Tiering.
- B. Use Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type to implement microservices to process videos. Store video metadata in Amazon Aurora. Store video content in Amazon S3 Intelligent-Tiering.
- C. Use Amazon EMR to process the videos with Apache Spark. Store video content in Amazon FSx for Lustre. Use Amazon Kinesis Data Streams to ingest videos in real time.
- D. Deploy a containerized video processing application on Amazon Elastic Kubernetes Service (Amazon EKS) with the Amazon EC2 launch type. Store video metadata in Amazon RDS in a single Availability Zone. Store video content in Amazon S3 Glacier Deep Archive.

Answer: B

Explanation:

AWS Lambda is not suitable for long-running jobs that can take up to 20 minutes, as Lambda has a maximum execution duration of 15 minutes.

Amazon ECS with AWS Fargate allows you to run containers without managing EC2 instances, providing a scalable and highly available environment. You can scale Fargate tasks to handle large and parallel video processing jobs. Amazon Aurora is a highly available, managed relational database. S3 Intelligent-Tiering is cost-effective for storing large video files with variable access patterns.

AWS Documentation Extract:

"AWS Fargate lets you run containers without managing servers or clusters, providing a highly available and scalable compute environment. Fargate is suitable for running data processing workloads that require long-running compute tasks."

(Source: AWS Fargate documentation, Use Cases)

A: Lambda is not suitable for long-running (over 15 min) or heavy compute jobs.

C: Amazon EMR is optimized for big data analytics, not specific video processing. FSx for Lustre and Kinesis are not best fit for this use case.

D: EKS with EC2 adds operational overhead and RDS in a single AZ is not highly available; Glacier Deep Archive is not suitable for frequently accessed video files.

Reference: AWS Certified Solutions Architect - Official Study Guide, Containerized and Serverless Processing.

Question: 396

A company runs an application as a task in an Amazon Elastic Container Service (Amazon ECS) cluster. The application must have read and write access to a specific group of Amazon S3 buckets. The S3 buckets are in the same AWS Region and AWS account as the ECS cluster.

The company needs to grant the application access to the S3 buckets according to the principle of least privilege.

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Add a tag to each bucket. Create an IAM policy that includes a StringEquals condition that matches the tags and values of the buckets.
- B. Create an IAM policy that lists the full Amazon Resource Name (ARN) for each S3 bucket.
- C. Attach the IAM policy to the instance role of the ECS task.
- D. Create an IAM policy that includes a wildcard Amazon Resource Name (ARN) that matches all combinations of the S3 bucket names.
- E. Attach the IAM policy to the task role of the ECS task.

Answer: B,E

Explanation:

To grant the ECS application the least privilege, you create an IAM policy that explicitly lists the ARNs of the specific S3 buckets the task should access. You then attach this policy to the task role (not the instance role), ensuring only that task has the necessary permissions.

This approach avoids granting permissions to other EC2 instances or services and restricts access strictly to the required buckets.

AWS Documentation Extract:

"Attach an IAM policy granting access to specific resources to the ECS task role, not to the instance role, to ensure that only the container has access according to the principle of least privilege."

"Use explicit ARNs to control access only to specified S3 buckets."

(Source: Amazon ECS documentation, IAM Roles for Tasks)

A: Tag-based access control for S3 buckets is possible but less direct and commonly used for identitybased access.

C: Attaching the policy to the instance role could grant unintended permissions to all containers or services running on the instance.

D: Wildcard ARNs grant broader access than necessary.

Reference: AWS Certified Solutions Architect - Official Study Guide, IAM Best Practices.

Question: 397

A company has a transaction-processing application that is backed by an Amazon RDS MySQL database. When the load on the application increases, a large number of database connections are opened and closed frequently, which causes latency for the database transactions.

A solutions architect determines that the root cause of the latency is poor connection handling by the application. The solutions architect cannot modify the application code. The solutions architect needs to manage database connections to improve the database performance during periods of high load.

Which solution will meet these requirements?

- A. Upgrade the database instance to a larger instance type to handle a large number of database connections.
- B. Configure Amazon RDS storage autoscaling to dynamically increase the provisioned IOPS.
- C. Use Amazon RDS Proxy to pool and share database connections.
- D. Convert the database instance to a Multi-AZ deployment.

Answer: C

Explanation:

Amazon RDS Proxy is a fully managed database proxy for RDS that makes applications more scalable, more resilient to database failures, and more secure. RDS Proxy pools and shares database connections, allowing applications to open and close connections as needed without overwhelming the database. This is the recommended solution when the application cannot be modified to use connection pooling itself.

AWS Documentation Extract:

"Amazon RDS Proxy helps manage database connections to improve application scalability and performance. It pools connections and shares them among application clients, which can mitigate issues caused by opening and closing many database connections."

(Source: Amazon RDS Proxy documentation)

A: Upgrading the instance does not solve connection inefficiency and can be cost-ineffective.

B: Increasing IOPS only helps if storage is a bottleneck, not if connections are the issue.

D: Multi-AZ improves availability, not connection management.

Reference: AWS Certified Solutions Architect - Official Study Guide, RDS Performance and Proxy.

Question: 398

A company wants to create a long-term storage solution that will allow users to upload terabytes of images and videos. The company will use the images and videos to train machine learning (ML) models. The storage solution must be scalable and cost-optimized.

Which solution will meet these requirements?

- A. Provision an Amazon S3 bucket for users to upload images and videos. Copy the data from the S3 bucket to an Amazon FSx for Lustre file system to make the data available for ML model training.
- B. Provision an Amazon S3 bucket for users to upload images and videos. Configure the S3 bucket to make the data available to Amazon SageMaker AI training. Store the data in the S3 Intelligent-Tiering storage class.
- C. Configure an Amazon SageMaker AI notebook instance with 16 GB of storage. Create a custom application to allow users to upload images and videos directly to the notebook instance.
- D. Provision an Amazon S3 bucket for users to upload images and videos. Copy the data from the S3 bucket to an Amazon Elastic File System (Amazon EFS) file system to make the data available for ML model training.

Answer: B

Explanation:

Amazon S3 is the most scalable and cost-optimized object storage for large amounts of data. It integrates natively with Amazon SageMaker, which allows direct access to S3 data for machine learning training jobs without the need to copy data elsewhere. The S3 Intelligent-Tiering storage class further optimizes storage costs for data with unknown or changing access patterns, while maintaining immediate availability for ML training.

Copying to FSx for Lustre or EFS adds unnecessary complexity and cost unless ultra-high throughput POSIX file access is specifically required, which is not mentioned here.

AWS Documentation Extract:

"Amazon S3 provides cost-effective storage for ML data sets, and Amazon SageMaker can directly access data in S3. S3 Intelligent-Tiering automatically moves data between frequent and infrequent access tiers, optimizing costs for changing data access patterns."
(Source: AWS S3 and SageMaker documentation)

A, D: FSx for Lustre or EFS may be used in specialized cases but are not cost-optimized for general large-scale ML data storage.

C: SageMaker notebook instance storage is not designed for large-scale data ingestion or long-term storage.

Reference: AWS Certified Solutions Architect - Official Study Guide, S3 for ML Data Storage.

Question: 399

A company runs an application on Amazon EC2 instances. The application needs to access an Amazon RDS database. The company wants to grant the EC2 instances access permissions to the RDS database while following the principle of least privilege.

Which solution will meet these requirements?

- A. Create an IAM user that has a policy that grants administrative permissions. Use the IAM user's access keys on the EC2 instances to access the RDS database.
- B. Create an IAM user that has a policy that grants the minimum required permissions to access the RDS database. Embed the IAM user's access keys on the EC2 instances to access the RDS database.
- C. Create an IAM role that has a policy that grants the minimum required permissions to access the RDS database. Attach the IAM role access key and the IAM role secret key to the EC2 instance profile.
- D. Create an IAM role that has a policy that grants the minimum required permissions to access the RDS database. Attach the IAM role to an EC2 instance profile. Associate the instance profile with the instances.

Answer: D

Explanation:

The best practice for granting AWS resource access to EC2 instances is to use IAM roles, not users or long-lived access keys. You create an IAM role with a policy that grants the minimum permissions required, then attach that role to an instance profile associated with the EC2 instance. The instance then automatically receives temporary credentials for AWS service access.

AWS Documentation Extract:

"Attach an IAM role to your EC2 instances to securely grant permissions to AWS services according to the principle of least privilege. Never embed access keys in EC2 instances."

(Source: AWS EC2 documentation, IAM Roles for EC2)

A, B: Using IAM users and embedding keys violates security best practices.

C: IAM role credentials are automatically managed and never need to be manually attached as keys.

Reference: AWS Certified Solutions Architect - Official Study Guide, IAM Roles and EC2.

Question: 400

A company runs a MySQL database on a single Amazon EC2 instance.

The company needs to improve availability of the database to prepare for power outages.

Which solution will meet this requirement?

- A. Add an Application Load Balancer (ALB) in front of the EC2 instance.
- B. Configure EC2 automatic instance recovery to move the instance to another Availability Zone.
- C. Migrate the MySQL database to Amazon RDS and enable Multi-AZ deployment.

D. Enable termination protection for the EC2 instance.

Answer: C

Explanation:

Amazon RDS Multi-AZ deployment provides automated failover and increased availability for MySQL databases. In case of infrastructure failure or power outage in one Availability Zone, RDS automatically fails over to a standby replica in another AZ, minimizing downtime. This is the AWS- recommended way to achieve high availability for relational databases.

AWS Documentation Extract:

"Amazon RDS Multi-AZ deployments provide high availability, failover support, and data redundancy for database instances. In the event of infrastructure failure, Amazon RDS performs an automatic failover to the standby."

(Source: Amazon RDS documentation, High Availability)

A: An ALB does not increase availability for a single database instance.

B: EC2 instance recovery does not move the instance across Availability Zones.

D: Termination protection prevents accidental deletion, not availability issues.

Reference: AWS Certified Solutions Architect - Official Study Guide, RDS High Availability.

Question: 401

An ecommerce company runs a PostgreSQL database on an Amazon EC2 instance. The database stores data in Amazon Elastic Block Store (Amazon EBS) volumes. The daily peak input/output transactions per second (IOPS) do not exceed 15,000 IOPS. The company wants to migrate the database to Amazon RDS for PostgreSQL and to provision disk IOPS performance that is independent of disk storage capacity.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure General Purpose SSD (gp2) EBS volumes. Provision a 5 TiB volume.
- B. Configure Provisioned IOPS SSD (io1) EBS volumes. Provision 15,000 IOPS.
- C. Configure General Purpose SSD (gp3) EBS volumes. Provision 15,000 IOPS.
- D. Configure magnetic EBS volumes to achieve maximum IOPS.

Answer: C

Explanation:

EBS gp3 volumes allow you to independently configure IOPS and throughput, up to 16,000 IOPS, regardless of the volume size, and at a lower cost compared to io1/io2 provisioned IOPS volumes. This meets the requirement for cost-effective, predictable IOPS performance for Amazon RDS for PostgreSQL.

AWS Documentation Extract:

"With gp3 volumes, you can provision performance independent of storage capacity, up to 16,000 IOPS. This allows for cost-effective scaling for applications that require high performance at a lower price point compared to io1."

(Source: Amazon EBS documentation, gp3 Volumes)

A: gp2 ties IOPS to volume size; 5 TiB is wasteful if only 15,000 IOPS are needed.

B: io1 works but is significantly more expensive than gp3 for most workloads.

D: Magnetic volumes do not support high IOPS.

Reference: AWS Certified Solutions Architect - Official Study Guide, EBS Storage Options.

Question: 402

A company wants to release a new device that will collect data to track overnight sleep on an intelligent mattress. Sensors will send data that will be uploaded to an Amazon S3 bucket. Each mattress generates about 2 MB of data each night.

An application must process the data and summarize the data for each user. The application must make the results available as soon as possible. Every invocation of the application will require about 1 GB of memory and will finish running within 30 seconds.

Which solution will run the application MOST cost-effectively?

- A. AWS Lambda with a Python script
- B. AWS Glue with a Scala job
- C. Amazon EMR with an Apache Spark script
- D. AWS Glue with a PySpark job

Answer: A

Explanation:

AWS Lambda supports functions up to 10 GB of memory and 15 minutes execution time. Each invocation here requires only 1 GB of memory and finishes in 30 seconds, making it an ideal fit for Lambda. Lambda is cost-effective for event-driven, short-duration workloads and requires no infrastructure management. AWS Glue and EMR are better suited for large-scale ETL or distributed processing, which is unnecessary and more costly for this workload.

AWS Documentation Extract:

“AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You pay only for the compute time you consume. Lambda supports up to 10 GB memory and 15 minutes per invocation.”

(Source: AWS Lambda documentation)

B, D: AWS Glue is intended for ETL on larger datasets or batch jobs, usually with higher operational overhead and cost.

C: EMR is for large-scale distributed processing and is not cost-effective for single, fast, memorybound jobs.

Reference: AWS Certified Solutions Architect - Official Study Guide, Lambda for Serverless Processing.

Question: 403

A company is building a new web application on AWS. The application needs to consume files from a legacy on-premises application that runs a batch process and outputs approximately 1 GB of data every night to an NFS file mount.

A solutions architect needs to design a storage solution that requires minimal changes to the legacy application and keeps costs low.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy an Outpost in AWS Outposts to the on-premises location where the legacy application is stored. Configure the legacy application and the web application to store and retrieve the files in Amazon S3 on the Outpost.
- B. Deploy an AWS Storage Gateway Volume Gateway on premises. Point the legacy application to the Volume Gateway. Configure the web application to use the Amazon S3 bucket that the Volume Gateway uses.
- C. Deploy an Amazon S3 interface endpoint on AWS. Reconfigure the legacy application to store the files directly on an Amazon S3 endpoint. Configure the web application to retrieve the files from Amazon S3.
- D. Deploy an Amazon S3 File Gateway on premises. Point the legacy application to the File Gateway. Configure the web application to retrieve the files from the S3 bucket that the File Gateway uses.

Answer: D

Explanation:

Amazon S3 File Gateway provides a local NFS mount point, which can be used with minimal changes by the legacy application. Files are transparently uploaded to Amazon S3, allowing the web application to access them directly from S3. This is the most cost-effective and operationally simple way to bridge legacy on-premises NFS output with S3, requiring no changes to the batch process.

AWS Documentation Extract:

“With S3 File Gateway, you can provide applications a local file interface to Amazon S3. S3 File Gateway presents a file-based interface (NFS

or SMB), allowing you to use S3 as your scalable, durable storage while making files available to legacy applications.”

(Source: AWS Storage Gateway documentation)

A: Outposts is far more costly and complex than needed.

B: Volume Gateway presents iSCSI block storage, not NFS.

C: Requires re-coding the legacy app to use S3 APIs, not NFS.

Reference: AWS Certified Solutions Architect - Official Study Guide, Hybrid Storage Solutions.

Question: 404

A solutions architect is designing a multi-Region disaster recovery (DR) strategy for a company. The company runs an application on Amazon EC2 instances in Auto Scaling groups that are behind an Application Load Balancer (ALB). The company hosts the application in the company's primary and secondary AWS Regions.

The application must respond to DNS queries from the secondary Region if the primary Region fails. Only one Region must serve traffic at a time.

Which solution will meet these requirements?

- A. Create an outbound endpoint in Amazon Route 53 Resolver. Create forwarding rules that determine how queries will be forwarded to DNS resolvers on the network. Associate the rules with VPCs in each Region.
- B. Create primary and secondary DNS records in Amazon Route 53. Configure health checks and a failover routing policy.
- C. Create a traffic policy in Amazon Route 53. Use a geolocation routing policy and a value type of ELB Application Load Balancer.
- D. Create an Amazon Route 53 profile. Associate DNS resources to the profile. Associate the profile with VPCs in each Region.

Answer: B

Explanation:

Amazon Route 53 supports failover routing policies, which use health checks to route DNS queries to a secondary Region only if the primary endpoint fails. This design ensures only one Region is active for traffic at any given time. This is the recommended architecture for active-passive, multi-Region DR strategies.

AWS Documentation Extract:

“Failover routing lets you route traffic to a primary resource, such as a web server in one Region, and a secondary resource in another Region. If the primary fails, Route 53 can route traffic to the secondary resource automatically.”

(Source: Amazon Route 53 documentation, Routing Policy Types)

A, D: These options do not configure DNS failover for external users.

C: Geolocation routing is for regional distribution, not DR failover.

Reference: AWS Certified Solutions Architect - Official Study Guide, Multi-Region DR and Route 53.

Question: 405

A company uses AWS to run its workloads. The company uses AWS Organizations to manage its accounts. The company needs to identify which departments are responsible for specific costs.

New accounts are constantly created in the Organizations account structure. The Organizations continuous integration and continuous delivery (CI/CD) framework already adds the populated department tag to the AWS resources. The company wants to use an AWS Cost Explorer report to identify the service costs by department from all AWS accounts.

Which combination of steps will meet these requirements with the MOST operational efficiency? (Select TWO.)

- A. Activate the aws:createdBy cost allocation tag and the department cost allocation tag in the management account.
- B. Create a new cost and usage report in Cost Explorer. Group by the department cost allocation tag. Apply a filter to see all linked accounts and services.
- C. Activate only the department cost allocation tag in the management account.
- D. Create a new cost and usage report in Cost Explorer. Group by the department cost allocation tag without any other filters.
- E. Activate only the aws:createdBy cost allocation tag in the management account.

Answer: C,D

Explanation:

To track costs by department, you must activate the custom department tag as a cost allocation tag in the AWS Organizations management account. Once activated, Cost Explorer and cost and usage reports can group costs by this tag for all linked accounts. The most operationally efficient way is to activate only the relevant department tag and create a cost and usage report grouped by that tag.

AWS Documentation Extract:

“To use a tag for cost allocation, you must activate it in the AWS Billing and Cost Management console. After activation, you can use the tag to group costs in Cost Explorer and reports.”

(Source: AWS Cost Management documentation)

A, E: aws:createdBy is not related to department cost grouping and is unnecessary.

B: Applying extra filters is optional; D is more direct and operationally efficient.

Reference: AWS Certified Solutions Architect - Official Study Guide, Cost Allocation and Tagging.

Question: 406

A company needs to collect streaming data from several sources and store the data in the AWS Cloud. The dataset is heavily structured, but analysts need to perform several complex SQL queries and need consistent performance. Some of the data is queried more frequently than the rest. The company wants a solution that meets its performance requirements in a cost-effective manner.

Which solution meets these requirements?

- A. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to ingest the data to save it to Amazon S3. Use Amazon Athena to perform SQL queries over the ingested data.
- B. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to ingest the data to save it to Amazon Redshift. Enable Amazon Redshift workload management (WLM) to prioritize workloads.
- C. Use Amazon Data Firehose to ingest the data to save it to Amazon Redshift. Enable Amazon Redshift workload management (WLM) to prioritize workloads.
- D. Use Amazon Data Firehose to ingest the data to save it to Amazon S3. Load frequently queried data to Amazon Redshift using the COPY command. Use Amazon Redshift Spectrum for less frequently queried data.

Answer: D

Explanation:

The most cost-effective way to provide consistent SQL query performance on a heavily structured dataset, where some data is accessed more frequently, is to use Amazon Redshift as your main data warehouse for hot data and Amazon Redshift Spectrum to query cold data that remains in S3. With this approach, frequently queried data is loaded into Redshift for fast, consistent querying, while infrequently accessed data is left in S3 and accessed on demand via Spectrum, avoiding unnecessary data warehousing costs. Amazon Kinesis Data Firehose provides an easy and scalable way to ingest streaming data directly to both S3 and Redshift.

AWS Documentation Extract:

"Amazon Redshift Spectrum allows you to run queries against exabytes of data in Amazon S3 without loading or transforming the data. Load hot data into Redshift for fast access and query cold data in S3 with Spectrum, optimizing both cost and performance."

(Source: Amazon Redshift documentation, Spectrum)

A: Athena is good for ad hoc queries but not for consistent, high-performance SQL workloads.

B, C: Loading all data into Redshift is not cost-effective if some data is infrequently accessed.

Reference: AWS Certified Solutions Architect - Official Study Guide, Data Warehousing and Analytics.

Question: 407

A company asks a solutions architect to review the architecture for its messaging application. The application uses TCP and UDP traffic. The company is planning to deploy a new VoIP feature, but its 10 test users in other countries are reporting poor call quality.

The VoIP application runs on an Amazon EC2 instance with more than enough resources. The HTTP portion of the company's application behind an Application Load Balancer has no issues.

What should the solutions architect recommend for the company to do to address the VoIP performance issues?

- A. Use AWS Global Accelerator.
- B. Implement Amazon CloudFront into the architecture.
- C. Use an Amazon Route 53 geoproximity routing policy.
- D. Migrate from Application Load Balancers to Network Load Balancers.

Answer: A

Explanation:

AWS Global Accelerator is a service that improves global application availability and performance using the AWS global network. It supports both TCP and UDP protocols and provides optimized routing and lower latency for real-time applications such as VoIP, regardless of where the user is located globally.

AWS Documentation Extract:

"AWS Global Accelerator uses the AWS global network to optimize the path to your application endpoints, improving performance for TCP and UDP traffic, such as VoIP."

(Source: AWS Global Accelerator documentation)

B: CloudFront accelerates HTTP/S content, not TCP/UDP.

C: Route 53 geoproximity routing is for DNS-based routing, not for traffic acceleration.

D: Network Load Balancers support TCP/UDP, but do not address global latency or provide acceleration.

Reference: AWS Certified Solutions Architect - Official Study Guide, Networking Optimization.

Question: 408

A company is storing data in Amazon S3 buckets. The company needs to retain any objects that contain personally identifiable information (PII) that might need to be reviewed.

A solutions architect must develop an automated solution to identify objects that contain PII and apply the necessary controls to prevent deletion before review.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create a job in Amazon Macie to scan the S3 buckets for the relevant sensitive data identifiers.
- B. Move the identified objects to the S3 Glacier Deep Archive storage class.
- C. Create an AWS Lambda function that performs an S3 Object Lock legal hold operation on the identified objects.
- D. Create an AWS Lambda function that applies an S3 Object Lock retention period to the identified objects in governance mode.
- E. Create an Amazon EventBridge rule that invokes the AWS Lambda function when Amazon Macie detects sensitive data.
- F. Configure multi-factor authentication (MFA) delete on the S3 buckets.

Answer: A,C,E

Explanation:

A: Amazon Macie can scan S3 buckets for sensitive data and identify PII.

C: S3 Object Lock legal hold prevents object deletion until a review is complete, meeting compliance requirements.

E: EventBridge can trigger the Lambda function automatically when Macie detects sensitive data, creating an automated workflow.

AWS Documentation Extract:

"Amazon Macie automatically discovers, classifies, and protects sensitive data in AWS. You can use EventBridge to invoke a Lambda function for post-processing, such as applying Object Lock legal hold to flagged objects."

(Source: AWS Macie and S3 Object Lock documentation)

B, D: Moving to Glacier Deep Archive or applying retention in governance mode is not specific to deletion prevention pending review.

F: MFA Delete adds a security layer but does not automate object retention for flagged PII.

Reference: AWS Certified Solutions Architect - Official Study Guide, S3 Compliance and Data Protection.

Question: 409

A company needs an automated solution to detect cryptocurrency mining activity on Amazon EC2 instances. The solution must automatically isolate any identified EC2 instances for forensic analysis.

Which solution will meet these requirements?

- A. Create an Amazon EventBridge rule that runs when Amazon GuardDuty detects cryptocurrency mining activity. Configure the rule to invoke an AWS Lambda function to isolate the identified EC2 instances.
- B. Create an AWS Security Hub custom action that runs when Amazon GuardDuty detects cryptocurrency mining activity. Configure the custom action to invoke an AWS Lambda function to isolate the identified EC2 instances.
- C. Create an Amazon Inspector rule that runs when Amazon GuardDuty detects cryptocurrency mining activity. Configure the rule to invoke an AWS Lambda function to isolate the identified EC2 instances.
- D. Create an AWS Config custom rule that runs when AWS Config detects cryptocurrency mining activity. Configure the rule to invoke an AWS Lambda function to isolate the identified EC2 instances.

Answer: A

Explanation:

Amazon GuardDuty detects cryptocurrency mining and sends findings to Amazon EventBridge. You can use EventBridge to trigger an automated Lambda function to isolate EC2 instances (such as by removing security group access or stopping/isolating the instance).

AWS Documentation Extract:

"Amazon GuardDuty findings can be sent to Amazon EventBridge, which enables you to trigger an automated response using AWS Lambda."

(Source: AWS GuardDuty documentation)

B, C, D: Security Hub, Inspector, and Config are not directly used for this detection-to-isolation workflow.

Reference: AWS Certified Solutions Architect - Official Study Guide, Threat Detection and Automated Response.

Question: 410

A company has an application that runs on Amazon EC2 instances within a private subnet in a VPC.

The instances access data in an Amazon S3 bucket in the same AWS Region. The VPC contains a NAT gateway in a public subnet to access the S3 bucket. The company wants to reduce costs by replacing the NAT gateway without compromising security or redundancy.

Which solution meets these requirements?

- A. Replace the NAT gateway with a NAT instance.
- B. Replace the NAT gateway with an internet gateway.
- C. Replace the NAT gateway with a gateway VPC endpoint.
- D. Replace the NAT gateway with an AWS Direct Connect connection.

Answer: C

Explanation:

A VPC gateway endpoint for Amazon S3 enables private connectivity to S3 without routing traffic through a NAT gateway or over the internet, eliminating NAT gateway costs. This solution is secure and redundant, as S3 endpoints are highly available by design.

AWS Documentation Extract:

"A gateway VPC endpoint enables you to privately connect your VPC to supported AWS services without requiring a NAT gateway or internet gateway."

(Source: Amazon VPC documentation, Gateway Endpoints)

- A: NAT instances still incur operational overhead and costs.
- B: Internet gateway exposes resources and does not provide private access.
- D: Direct Connect is for hybrid networking, not for cost-efficient S3 access.

Reference: AWS Certified Solutions Architect - Official Study Guide, VPC Networking and Endpoints.

Question: 411

A company runs all its business applications in the AWS Cloud. The company uses AWS Organizations to manage multiple AWS accounts.

A solutions architect needs to review all permissions granted to IAM users to determine which users have more permissions than required.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use Network Access Analyzer to review all access permissions in the company's AWS accounts.
- B. Create an AWS CloudWatch alarm that activates when an IAM user creates or modifies resources in an AWS account.
- C. Use AWS Identity and Access Management (IAM) Access Analyzer to review all the company's resources and accounts.
- D. Use Amazon Inspector to find vulnerabilities in existing IAM policies.

Answer: C

Explanation:

IAM Access Analyzer analyzes permissions granted using policies to determine what resources are shared with an external entity, and helps identify excessive permissions or least privilege violations across all accounts in an AWS Organization. It is specifically designed for reviewing and refining IAM permissions with minimal administrative effort.

AWS Documentation Extract:

“IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. You can also use Access Analyzer policy checks to refine permissions and implement least privilege.”

(Source: IAM Access Analyzer documentation)

A: Network Access Analyzer is for VPC network access analysis, not IAM permissions.

B: CloudWatch alarms are not suitable for detailed permission analysis.

D: Amazon Inspector is for security vulnerability assessment, not IAM policy review.

Reference: AWS Certified Solutions Architect - Official Study Guide, IAM Security Analysis.

Question: 412

A company is developing software that uses a PostgreSQL database schema. The company needs to configure development environments and test environments for its developers.

Each developer at the company uses their own development environment, which includes a PostgreSQL database. On average, each development environment is used for an 8-hour workday.

The test environments will be used for load testing that can take up to 2 hours each day.

Which solution will meet these requirements MOST cost-effectively?

A. Configure development environments and test environments with their own Amazon Aurora Serverless v2 PostgreSQL database.

B. For each development environment, configure an Amazon RDS for PostgreSQL Single-AZ DB instance. For the test environment, configure a single Amazon RDS for PostgreSQL Multi-AZ DB instance.

C. Configure development environments and test environments with their own Amazon Aurora PostgreSQL DB cluster.

D. Configure an Amazon Aurora global database. Allow developers to connect to the database with their own credentials.

Answer: A

Explanation:

Amazon Aurora Serverless v2 provides cost-effective, on-demand, and auto-scaling database capacity. You pay only for actual usage, making it ideal for development and test environments that are not used continuously. It supports PostgreSQL and is suitable for variable and short-duration workloads, minimizing costs when databases are idle.

AWS Documentation Extract:

“Amazon Aurora Serverless v2 automatically adjusts database capacity based on application needs, ideal for development and test environments with intermittent usage.”

(Source: Aurora Serverless v2 documentation)

B: RDS instances run and accrue charges even when idle.

C, D: Aurora clusters/global databases are overkill and incur higher costs for dev/test.

Reference: AWS Certified Solutions Architect - Official Study Guide, Aurora Serverless for Cost Optimization.

Question: 413

A company is redesigning its data intake process. In the existing process, the company receives data transfers and uploads the data to an Amazon S3 bucket every night. The company uses AWS Glue crawlers and jobs to prepare the data for a machine learning (ML) workflow.

The company needs a low-code solution to run multiple AWS Glue jobs in sequence and provide a visual workflow.

Which solution will meet these requirements?

- A. Use an Amazon EC2 instance to run a cron job and a script to check for the S3 files and call the AWS Glue jobs. Create an Amazon CloudWatch dashboard to visualize the workflow.
- B. Use Amazon EventBridge to call an AWS Step Functions workflow for the AWS Glue jobs. Use Step Functions to create a visual workflow.
- C. Use S3 Event Notifications to invoke a series of AWS Lambda functions and AWS Glue jobs in sequence. Use Amazon QuickSight to create a visual workflow.
- D. Create an Amazon Elastic Container Service (Amazon ECS) task that contains a Python script that manages the AWS Glue jobs and creates a visual workflow. Use Amazon EventBridge Scheduler to start the ECS task.

Answer: B

Explanation:

AWS Step Functions provides a low-code, fully managed workflow service with a visual interface to orchestrate AWS Glue jobs in sequence. Step Functions integrates natively with AWS Glue and can be triggered by Amazon EventBridge based on events or schedules.

AWS Documentation Extract:

“AWS Step Functions makes it easy to coordinate multiple AWS services into serverless workflows so you can build and update apps quickly. Step Functions provides visual workflows and integrates with AWS Glue for ETL orchestration.”

(Source: AWS Step Functions documentation)

A: Manual scripting and dashboards do not provide a low-code or integrated visual workflow.

C: QuickSight is for BI, not workflow visualization.

D: ECS adds unnecessary complexity.

Reference: AWS Certified Solutions Architect - Official Study Guide, Automation with Step Functions.

Question: 414

A company runs an application on Microsoft SQL Server databases in an on-premises data center. The company wants to migrate to AWS and optimize costs for its infrastructure on AWS.

Which solution will meet these requirements?

- A. Migrate the databases to Amazon EC2 instances that use SQL Server Amazon Machine Images (AMIs) provided by AWS.
- B. Migrate to Amazon Aurora PostgreSQL by using Babelfish for Aurora PostgreSQL.
- C. Migrate the databases to a PostgreSQL database that runs on Amazon EC2 instances.
- D. Migrate the databases to Amazon RDS for Microsoft SQL Server.

Answer: B

Explanation:

Amazon Aurora PostgreSQL with Babelfish allows SQL Server applications to run directly on Aurora PostgreSQL with minimal code changes. Babelfish adds a SQL Server-compatible endpoint, significantly lowering costs compared to running licensed SQL Server instances on RDS or EC2.

AWS Documentation Extract:

“Babelfish for Aurora PostgreSQL enables Aurora to understand T-SQL and SQL Server wire protocol, allowing you to run SQL Server applications on Amazon Aurora PostgreSQL with lower costs.”

(Source: Babelfish for Aurora PostgreSQL documentation)

A, D: SQL Server on EC2 or RDS incurs high Microsoft licensing costs.

C: Plain PostgreSQL would require more code refactoring than Babelfish.

Reference: AWS Certified Solutions Architect - Official Study Guide, Babelfish and Cost Optimization.

Question: 415

A company is using AWS Identity and Access Management (IAM) Access Analyzer to refine IAM permissions for employee users. The company uses an organization in AWS Organizations and AWS Control Tower to manage its AWS accounts. The company has designated a specific member account as an audit account.

A solutions architect needs to set up IAM Access Analyzer to aggregate findings from all member accounts in the audit account.

What is the first step the solutions architect should take?

- A. Use AWS CloudTrail to configure one trail for all accounts. Create an Amazon S3 bucket in the audit account. Configure the trail to send logs related to access activity to the new S3 bucket in the audit account.
- B. Configure a delegated administrator account for IAM Access Analyzer in the AWS Control Tower management account. In the delegated administrator account for IAM Access Analyzer, specify the AWS account ID of the audit account.
- C. Create an Amazon S3 bucket in the audit account. Generate a new permissions policy, and add a service role to the policy to give IAM Access Analyzer access to AWS CloudTrail and the S3 bucket in the audit account.
- D. Add a new trust policy that includes permissions to allow IAM Access Analyzer to perform sts:AssumeRole actions. Modify the permissions policy to allow IAM Access Analyzer to generate policies.

Answer: B

Explanation:

The first step is to configure a delegated administrator account for IAM Access Analyzer at the organization level. Only after delegating the administrator account can you aggregate Access Analyzer findings from all member accounts into a designated audit account. This must be set up in the AWS Organizations management account.

AWS Documentation Extract:

“You must designate a delegated administrator for IAM Access Analyzer at the organization level. The delegated administrator account aggregates findings from all member accounts.”

(Source: IAM Access Analyzer documentation)

A, C, D: These steps do not establish the organization-wide aggregation required for Access Analyzer.

Question: 416

A company hosts an industrial control application that receives sensor input through Amazon Kinesis Data Streams. The application needs to support new sensors for real-time anomaly detection in monitored equipment.

The company wants to integrate new sensors in a loosely-coupled, fully managed, and serverless way. The company cannot modify the application code.

Which solution will meet these requirements?

- A. Forward the existing stream in Kinesis Data Streams to Amazon Managed Service for Apache Flink for anomaly detection. Use a second stream in Kinesis Data Streams to send the Flink output to the application.
- B. Use Amazon Data Firehose to stream data to Amazon S3. Use Amazon Redshift Spectrum to perform anomaly detection on the S3 data. Use S3 Event Notifications to invoke an AWS Lambda function that sends analyzed data to the application through a second stream in Kinesis Data Streams.
- C. Configure Amazon EC2 instances in an Auto Scaling group to consume data from the data stream and to perform anomaly detection. Create a second stream in Kinesis Data Streams to send data from the EC2 instances to the application.
- D. Configure an Amazon Elastic Container Service (Amazon ECS) task that uses Amazon EC2 instances to consume data from the data stream and to perform anomaly detection. Create a second stream in Kinesis Data Streams to send data from the containers to the application.

Answer: A

Explanation:

Amazon Managed Service for Apache Flink (formerly Kinesis Data Analytics) is a fully managed, serverless service for real-time processing of streaming data. You can consume data from Kinesis Data Streams, perform anomaly detection, and then output the results to another Kinesis stream. This approach is loosely coupled, fully managed, and does not require modifying the application code.

AWS Documentation Extract:

“Amazon Managed Service for Apache Flink enables you to process streaming data in real time, integrating with Kinesis Data Streams as source and sink, and is fully managed and serverless.”

(Source: Apache Flink on AWS documentation)

B: S3/Redshift is not real-time and adds complexity.

C, D: EC2/ECS solutions are not serverless or fully managed.

Question: 417

A company needs to migrate its customer transactions database from on-premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Convert the database to Amazon DynamoDB by using AWS Schema Conversion Tool (AWS SCT). Store the password in AWS Systems Manager Parameter Store. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
- B. Migrate the database to Amazon RDS for Oracle. Store the password in AWS Secrets Manager. Turn on automatic rotation. Configure a yearly rotation schedule.
- C. Migrate the database to an Amazon EC2 instance. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule.
- D. Migrate the database to Amazon Neptune by using AWS Schema Conversion Tool (AWS SCT). Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.

Answer: B

Explanation:

Amazon RDS for Oracle is a managed database service, which significantly reduces operational overhead compared to running Oracle on EC2 or on-premises. AWS Secrets Manager natively integrates with RDS and supports automatic, scheduled password rotation with minimal setup. You can configure the rotation schedule (including yearly), and Secrets Manager will handle the secure password storage and rotation workflow for you.

AWS Documentation Extract:

"AWS Secrets Manager helps you protect access to your applications, services, and IT resources without the upfront investment and on-going maintenance costs of operating your own infrastructure. You can configure automatic rotation for supported databases such as

Amazon RDS for Oracle."

(Source: AWS Secrets Manager documentation)

A, C, D: These solutions require custom scripting, Lambda, and alarms, leading to more operational overhead.

Reference: AWS Certified Solutions Architect - Official Study Guide, Secrets Manager and RDS.

Question: 418

A company is migrating mobile banking applications to run on Amazon EC2 instances in a VPC.

Backend service applications run in an on-premises data center. The data center has an AWS Direct Connect connection into AWS. The applications that run in the VPC need to resolve DNS requests to an on-premises Active Directory domain that runs in the data center.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Provision a set of EC2 instances across two Availability Zones in the VPC as caching DNS servers to resolve DNS queries from the application servers within the VPC.
- B. Provision an Amazon Route 53 private hosted zone. Configure NS records that point to on-premises DNS servers.
- C. Create DNS endpoints by using Amazon Route 53 Resolver. Add conditional forwarding rules to resolve DNS namespaces between the on-premises data center and the VPC.
- D. Provision a new Active Directory domain controller in the VPC with a bidirectional trust between this new domain and the on-premises Active Directory domain.

Answer: C

Explanation:

Amazon Route 53 Resolver endpoints allow you to integrate DNS between AWS and on-premises environments easily. By creating inbound and outbound resolver endpoints, you can configure conditional forwarding rules so that DNS queries for your on-premises AD domain are forwarded to the on-premises DNS servers. This approach is fully managed, scales automatically, and requires the least administrative overhead.

AWS Documentation Extract:

"Route 53 Resolver provides DNS resolution between AWS and on-premises environments, using endpoints and forwarding rules to manage DNS query routing seamlessly."

(Source: Route 53 Resolver documentation)

A, D: Require provisioning, managing, and patching EC2 servers or domain controllers.

B: NS records in a private hosted zone do not provide true DNS forwarding.

Reference: AWS Certified Solutions Architect - Official Study Guide, Hybrid DNS Integration.

Question: 419

A company wants to use automatic machine learning (ML) to create and visualize forecasts of complex scenarios and trends.

Which solution will meet these requirements with the LEAST management overhead?

- A. Use an AWS Glue ML job to transform the data and create forecasts. Use Amazon QuickSight to visualize the data.
- B. Use Amazon QuickSight to visualize the data. Use ML-powered forecasting in QuickSight to create forecasts.
- C. Use a prebuilt ML AMI from the AWS Marketplace to create forecasts. Use Amazon QuickSight to visualize the data.
- D. Use Amazon SageMaker AI inference pipelines to create and update forecasts. Use Amazon QuickSight to visualize the combined data.

Answer: B

Explanation:

Amazon QuickSight includes built-in ML-powered forecasting capabilities, allowing you to create, visualize, and interact with time series forecasts directly within the BI dashboard, with no ML experience or infrastructure management required. This is the lowest management overhead solution.

AWS Documentation Extract:

"Amazon QuickSight provides built-in ML-powered forecasting, allowing business users to forecast future trends with a few clicks and no machine learning expertise required."

(Source: Amazon QuickSight documentation)

A, C, D: Require additional setup, management, or ML knowledge.

Reference: AWS Certified Solutions Architect - Official Study Guide, QuickSight and ML Forecasting.

Question: 420

A company has AWS Lambda functions that use environment variables. The company does not want its developers to see environment variables in plaintext.

Which solution will meet these requirements?

- A. Deploy code to Amazon EC2 instances instead of using Lambda functions.
- B. Configure SSL encryption on the Lambda functions to use AWS CloudHSM to store and encrypt the environment variables.
- C. Create a certificate in AWS Certificate Manager (ACM). Configure the Lambda functions to use the certificate to encrypt the environment variables.
- D. Create an AWS Key Management Service (AWS KMS) key. Enable encryption helpers on the Lambda functions to use the KMS key to store and encrypt the environment variables.

Answer: D

Explanation:

AWS Lambda supports encrypting environment variables at rest using AWS KMS. You can use encryption helpers (or Lambda's built-in support) to encrypt sensitive environment variable values using a KMS key. These encrypted variables are not visible in plaintext to developers, either in the console or when running the code.

AWS Documentation Extract:

"AWS Lambda automatically encrypts environment variables at rest. For additional security, you can use AWS KMS keys and encryption helpers to encrypt environment variables, ensuring they are never exposed in plaintext."

(Source: AWS Lambda documentation, Environment Variables Security)

A: Does not address the issue (and adds more management overhead).

B, C: There is no native support for environment variable encryption via CloudHSM or ACM.

Reference: AWS Certified Solutions Architect - Official Study Guide, Lambda Security Best Practices.

Question: 421

A company runs an application on several Amazon EC2 instances. Multiple Amazon Elastic Block Store (Amazon EBS) volumes are attached to each EC2 instance. The company needs to back up the configurations and the data of the EC2 instances every night. The application must be recoverable in a secondary AWS Region.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Configure an AWS Lambda function to take nightly snapshots of the application's EBS volumes and to copy the snapshots to a secondary Region.
- B. Create a backup plan in AWS Backup to take nightly backups. Copy the backups to a secondary Region. Add the EC2 instances to a resource assignment as part of the backup plan.
- C. Create a backup plan in AWS Backup to take nightly backups. Copy the backups to a secondary Region. Add the EBS volumes to a resource assignment as part of the backup plan.
- D. Configure an AWS Lambda function to take nightly snapshots of the application's EBS volumes and to copy the snapshots to a secondary Availability Zone.

Answer: B

Explanation:

AWS Backup is a fully managed backup service that can create backup plans for EC2 instances, including both instance configurations and attached EBS volumes, with scheduled and cross-Region copy capabilities. By adding the EC2 instances to the resource assignment in the backup plan, AWS Backup automatically backs up all configurations and attached EBS volumes, and can copy backups to a secondary Region for disaster recovery, providing the highest operational efficiency with the least manual effort.

AWS Documentation Extract:

“AWS Backup provides fully managed backup for EC2 instances and attached EBS volumes, with scheduling, retention, and cross-Region copy built in. By adding the EC2 instance as a resource, the backup includes both configuration and attached volumes.”

(Source: AWS Backup documentation)

A, D: Custom Lambda scripts increase operational overhead and are not as integrated or robust as AWS Backup.

C: Assigning only EBS volumes does not include the EC2 instance configuration, which is needed for full recovery.

Reference: AWS Certified Solutions Architect - Official Study Guide, Disaster Recovery and Backup.

Question: 422

A solutions architect is creating a data processing job that runs once daily and can take up to 2 hours to complete. If the job is interrupted, it has to restart from the beginning.

How should the solutions architect address this issue in the MOST cost-effective manner?

- A. Create a script that runs locally on an Amazon EC2 Reserved Instance that is triggered by a cron job.
- B. Create an AWS Lambda function triggered by an Amazon EventBridge scheduled event.
- C. Use an Amazon Elastic Container Service (Amazon ECS) Fargate task triggered by an Amazon EventBridge scheduled event.
- D. Use an Amazon Elastic Container Service (Amazon ECS) task running on Amazon EC2 triggered by an Amazon EventBridge scheduled event.

Answer: C

Explanation:

AWS Fargate with Amazon ECS is a serverless, fully managed compute engine for containers. Fargate eliminates the need to provision or manage servers, and is ideal for periodic, long-running batch jobs. You only pay for the resources consumed during job execution, making it cost-effective for jobs that run infrequently. Lambda is not suitable because the maximum execution time is 15 minutes, but the job can take up to 2 hours.

AWS Documentation Extract:

“AWS Fargate lets you run containers without managing servers or clusters. Fargate is ideal for batch jobs, event-driven processing, and scheduled tasks that require hours of compute.”

(Source: AWS Fargate documentation)

A: Reserved Instances are cost-inefficient for infrequent workloads and require server management.

B: Lambda has a hard limit of 15 minutes per execution, insufficient for this job.

D: ECS on EC2 requires you to manage and provision EC2 instances, increasing cost and management overhead.

Reference: AWS Certified Solutions Architect - Official Study Guide, Serverless and Batch Processing.

Question: 423

A company has set up hybrid connectivity between an on-premises data center and AWS by using AWS Site-to-Site VPN. The company is migrating a workload to AWS.

The company sets up a VPC that has two public subnets and two private subnets. The company wants to monitor the total packet loss and round-trip-time (RTT) between the data center and AWS.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon CloudWatch Network Monitor to set up Internet Control Message Protocol (ICMP) probe monitoring from each subnet to the on-premises destination.
- B. Create an Amazon EC2 instance in each subnet. Create a scheduled job to send Internet Control Message Protocol (ICMP) packets to the on-premises destination.
- C. Create an AWS Lambda function in each subnet. Write a script to perform Internet Control Message Protocol (ICMP) connectivity checks.
- D. Create an AWS Batch job in each subnet. Write a script to perform Internet Control Message Protocol (ICMP) connectivity checks.

Answer: A

Explanation:

The requirement is to monitor network metrics such as total packet loss and round-trip time (RTT) between on-premises and AWS over Site-to-Site VPN with minimal operational overhead. AWS CloudWatch Network Monitor (formerly known as VPC Network Manager) provides a managed solution to monitor connectivity, including packet loss and latency, between AWS and on-premises networks. This solution does not require managing any additional infrastructure like EC2 instances or Lambda functions and thus reduces operational overhead significantly.

CloudWatch Network Monitor leverages AWS-managed probes and integrates natively with CloudWatch dashboards and alarms, enabling

automated, centralized monitoring of network health.

This aligns with the AWS Well-Architected Framework's operational excellence pillar by minimizing manual intervention and enabling proactive detection of network issues.

Option B, C, and D involve creating custom probes with EC2, Lambda, or Batch jobs, which increases complexity, cost, and maintenance effort. They also require scheduling, script management, and additional monitoring infrastructure.

Reference:

AWS Well-Architected Framework — Operational Excellence Pillar

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Amazon CloudWatch Network Monitor (<https://docs.aws.amazon.com/vpc/latest/networkmanager/monitor.html>)

AWS Site-to-Site VPN Monitoring (<https://docs.aws.amazon.com/vpn/latest/s2svpn/monitoring-cloudwatch.html>)

Question: 424

A company runs an application on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The company wants to protect the application against layer 7 DDoS attacks.

Which solution will meet this requirement?

- A. Associate AWS Shield Standard with the ALB.
- B. Create an AWS WAF web ACL and add a custom rule. Associate the web ACL with the ALB.
- C. Create an AWS WAF web ACL and add an AWS managed rule. Associate the web ACL with the ALB.
- D. Create an Amazon CloudFront distribution and set the ALB as the origin. Configure the application DNS record to point to the CloudFront distribution instead of the ALB.

Answer: C

Explanation:

Protecting an application from layer 7 (application layer) DDoS attacks is best achieved by using AWS WAF (Web Application Firewall), which provides customizable protection against common web exploits including DDoS attacks at the application layer. AWS WAF supports managed rule groups maintained by AWS, which offer robust, tested protections against OWASP top 10 vulnerabilities and common attack patterns without requiring extensive manual rule creation.

While AWS Shield Standard provides basic network-layer DDoS protection automatically at no additional charge, it does not offer application-layer filtering capabilities. Therefore, option A alone is insufficient.

Option B, involving only custom rules, requires significant operational overhead and expertise, whereas AWS managed rules offer a turnkey

solution with ongoing updates from AWS security teams.

Option D, using CloudFront in front of the ALB, can provide additional protection benefits such as caching and geographic restrictions, but the question specifically asks for protecting against layer 7 DDoS on the ALB directly. CloudFront plus WAF is a valid enhanced solution, but the direct and recommended answer in AWS official documents is to use AWS WAF managed rules directly with ALB for application-level protection.

Reference:

AWS Well-Architected Framework — Security Pillar

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

AWS WAF Overview (<https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html>)

AWS Shield Overview (<https://aws.amazon.com/shield/>)

Protecting Web Applications with AWS WAF (<https://aws.amazon.com/blogs/security/how-to-protect-your-web-application-from-dos-and-ddos-attacks-using-aws-waf/>)

Question: 425

A company needs to accommodate traffic for a web application that the company hosts on AWS, especially during peak usage hours.

The application uses Amazon EC2 instances as web servers, an Amazon RDS DB instance for database operations, and an Amazon S3 bucket to store transaction documents. The application struggles to scale effectively and experiences performance issues.

The company wants to improve the scalability of the application and prevent future performance issues. The company also wants to improve global access speeds to the transaction documents for the company's global users.

Which solution will meet these requirements?

- A. Place the EC2 instances in Auto Scaling groups to scale appropriately during peak usage hours. Use Amazon RDS read replicas to improve database read performance. Deploy an Amazon CloudFront distribution that uses Amazon S3 as the origin.
- B. Increase the size of the EC2 instances to provide more compute capacity. Use Amazon ElastiCache to reduce database read loads. Use AWS Global Accelerator to optimize the delivery of the transaction documents that are in the S3 bucket.
- C. Transition workloads from the EC2 instances to AWS Lambda functions to scale in response to the usage peaks. Migrate the database to an Amazon Aurora global database to provide cross-Region reads. Use AWS Global Accelerator to deliver the transaction documents that are in the S3 bucket.
- D. Convert the application architecture to use Amazon Elastic Container Service (Amazon ECS) containers. Configure a Multi-AZ deployment of Amazon RDS to support database operations. Replicate the transaction documents that are in the S3 bucket across multiple AWS Regions.

Explanation:

This question centers on improving scalability and global access performance.

Auto Scaling groups enable EC2 instances to scale dynamically in response to demand, ensuring availability during peak hours without manual intervention. Amazon RDS read replicas offload read traffic, improving read throughput and reducing latency on the primary database instance. Deploying Amazon CloudFront with S3 as origin accelerates delivery of static transaction documents globally by caching content at edge locations, reducing latency for users worldwide.

Option B focuses on vertical scaling (larger instances) and caching with ElastiCache, but it does not address global content delivery optimally. AWS Global Accelerator accelerates network traffic but is better suited for accelerating TCP and UDP traffic; CloudFront is generally preferred for HTTP content delivery.

Option C migrates workloads to Lambda and Aurora global databases, which is an advanced and potentially costly redesign that may not be necessary. Option D suggests moving to ECS and multi-AZ RDS but does not address global content delivery efficiently.

Therefore, option A uses proven scalability and caching best practices aligned with AWS Well-Architected Framework pillars for performance and operational excellence.

Reference:

AWS Well-Architected Framework — Performance Efficiency Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Amazon EC2 Auto Scaling (<https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is- amazon-ec2-auto-scaling.html>)

Amazon RDS Read Replicas (https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html)

Amazon CloudFront Overview (<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>)

Question: 426

A company is developing a SaaS solution for customers. The solution runs on Amazon EC2 instances that have Amazon Elastic Block Store (Amazon EBS) volumes attached.

Within the SaaS application, customers can request how much storage they need. The application needs to allocate the amount of block storage each customer requests.

A solutions architect must design an operationally efficient solution that meets the storage scaling requirement.

Which solution will meet these requirements MOST cost-effectively?

A. Migrate the data from the EBS volumes to an Amazon S3 bucket. Use the Amazon S3 Standard storage class.

- B. Migrate the data from the EBS volumes to an Amazon Elastic File System (Amazon EFS) file system. Use the EFS Standard storage class. Invoke an AWS Lambda function to increase the EFS volume capacity based on user input.
- C. Migrate the data from the EBS volumes to an Amazon FSx for Windows File Server file system. Invoke an AWS Lambda function to increase the capacity of the file system based on user input.
- D. Invoke an AWS Lambda function to increase the size of EBS volumes based on user input by using EBS Elastic Volumes.

Answer: D

Explanation:

EBS Elastic Volumes allow you to dynamically increase storage size, adjust performance, and change volume types without downtime, supporting operational efficiency and scalability for SaaS applications that need to allocate varying storage amounts to customers.

Migrating from EBS to S3 (Option A) is not suitable since S3 is object storage, not block storage, and does not support block-level I/O required by many applications. EFS (Option B) and FSx (Option C) are shared file systems, which might add unnecessary complexity and cost, especially if the application depends on block storage semantics.

Using Lambda to automate Elastic Volumes resizing provides cost efficiency by allocating resources on demand and reduces operational overhead, aligning with AWS operational excellence and cost optimization best practices.

Reference:

AWS Well-Architected Framework — Operational Excellence and Cost Optimization Pillars

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Amazon EBS Elastic Volumes (<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-expand-volume.html>)

AWS Lambda Overview (<https://docs.aws.amazon.com/lambda/latest/dg/welcome.html>)

Question: 427

A company runs its critical storage application in the AWS Cloud. The application uses Amazon S3 in two AWS Regions. The company wants the application to send remote user data to the nearest S3 bucket with no public network congestion. The company also wants the application to fail over with the least amount of management of Amazon S3.

Which solution will meet these requirements?

- A. Implement an active-active design between the two Regions. Configure the application to use the regional S3 endpoints closest to the user.
- B. Use an active-passive configuration with S3 Multi-Region Access Points. Create a global endpoint for each of the Regions.
- C. Send user data to the regional S3 endpoints closest to the user. Configure an S3 cross-account replication rule to keep the S3 buckets synchronized.

D. Set up Amazon S3 to use Multi-Region Access Points in an active-active configuration with a single global endpoint. Configure S3 Cross-Region Replication.

Answer: D

Explanation:

AWS S3 Multi-Region Access Points enable customers to use a single global endpoint for S3 bucket access across multiple AWS Regions, providing automatic routing to the nearest Region. This reduces public network congestion by directing user data to the closest S3 bucket and supports high availability with active-active configuration.

Cross-Region Replication ensures data is replicated between buckets in different Regions, meeting the failover and resilience requirements with minimal management overhead.

Option D aligns best with AWS's recommended approach to resilient, low-latency, and simplified multi-Region S3 access.

Option A lacks the global endpoint and automatic failover. Option B incorrectly describes MultiRegion Access Points configuration and suggests global endpoints per Region, which is contradictory. Option C's cross-account replication adds complexity and does not provide a single global endpoint.

Reference:

AWS Well-Architected Framework — Reliability Pillar

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Amazon S3 Multi-Region Access Points

(<https://docs.aws.amazon.com/AmazonS3/latest/userguide/MultiRegionAccessPoints.html>)

S3 Cross-Region Replication

(<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html>)

Question: 428

A company runs an application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company wants to create a public API for the application that uses JSON Web Tokens (JWT) for authentication. The company wants the API to integrate directly with the ALB.

Which solution will meet these requirements?

- A. Use Amazon API Gateway to create a REST API.
- B. Use Amazon API Gateway to create an HTTP API.
- C. Use Amazon API Gateway to create a WebSocket API.

D. Use Amazon API Gateway to create a gRPC API.

Answer: B

Explanation:

Amazon API Gateway supports multiple API types: REST, HTTP, WebSocket, and gRPC. HTTP APIs are a newer, lightweight option that support JWT authorizers natively, enabling secure, scalable authentication for APIs with JSON Web Tokens.

HTTP APIs can integrate with ALBs as a backend target, providing direct connectivity and simplified API management with JWT authentication built in.

REST APIs support JWT but are more feature-rich and complex, often used for legacy or more complex use cases, and have higher costs and latency. WebSocket APIs are for real-time, bidirectional communication, which is not requested here. gRPC APIs support RPC calls but are less common for public HTTP-based APIs with JWT auth.

Therefore, HTTP API with JWT authorizers is the best fit for this use case.

Reference:

AWS Well-Architected Framework — Security Pillar

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Amazon API Gateway HTTP APIs

(<https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api.html>)

Using JWT Authorizers with HTTP APIs

(<https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-jwt-authorizer.html>)

Question: 429

A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connect connection in a central network account.

The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.

Which combination of steps will meet these requirements? (Select THREE.)

A. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.

B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct

Connect gateway by using a transit VIF.

- C. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.
- D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.
- E. Provision VPC peering as necessary.
- F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Answer: B,D,F

Explanation:

For a large-scale multi-account AWS environment with many VPCs and centralized Direct Connect, AWS recommends using a Transit Gateway (TGW) architecture combined with a Direct Connect gateway (DXGW). This setup allows scalable, centralized connectivity between on-premises and multiple VPCs across accounts.

Step B: Creating a Direct Connect gateway and Transit Gateway in a central network account and connecting them via a transit VIF enables the on-premises network to access all connected VPCs.

Step D: Sharing the transit gateway with other accounts via AWS Resource Access Manager (RAM) allows the central TGW to attach VPCs in multiple accounts, simplifying multi-account connectivity.

Step F: To route cloud resources' internet traffic back through the on-premises data center (for centralized egress), provisioning only private subnets and routing outbound internet traffic through NAT or firewall services in the data center is necessary. This requires configuring transit gateway and customer gateway routes appropriately.

Option A is partially correct in the use of Direct Connect gateway but association proposals are not scalable for hundreds of VPCs and accounts compared to transit gateway. Option C (internet gateway) is irrelevant here as traffic egress is required via on-premises data center, not directly to the internet. Option E (VPC peering) is not scalable for hundreds of VPCs.

Reference:

AWS Transit Gateway Overview (<https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html>)

AWS Direct Connect Gateway (<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways.html>)

Centralized Egress Architecture with Transit Gateway (<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-set-up-centralized-egress-with-transit-gateway/>)

AWS Well-Architected Framework — Reliability Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 430

A company is designing a solution to capture customer activity on the company's web applications. The company wants to analyze the activity data to make predictions.

Customer activity on the web applications is unpredictable and can increase suddenly. The company requires a solution that integrates with other web applications. The solution must include an authorization step.

Which solution will meet these requirements?

- A. Deploy a Gateway Load Balancer (GWLB) in front of an Amazon Elastic Container Service (Amazon ECS) container instance. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Configure the applications to pass an authorization header to the GWLB.
- B. Deploy an Amazon API Gateway endpoint in front of an Amazon Kinesis data stream. Store the data in an Amazon S3 bucket. Use an AWS Lambda function to handle authorization.
- C. Deploy an Amazon API Gateway endpoint in front of an Amazon Data Firehose delivery stream. Store the data in an Amazon S3 bucket. Use an API Gateway Lambda authorizer to handle authorization.
- D. Deploy a Gateway Load Balancer (GWLB) in front of an Amazon Elastic Container Service (Amazon ECS) container instance. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Use an AWS Lambda function to handle authorization.

Answer: C

Explanation:

The requirements specify capturing unpredictable and sudden spikes in customer activity, integrating easily with other web applications, and including authorization.

Amazon API Gateway with Lambda authorizer provides a secure, scalable entry point with flexible authorization mechanisms including token validation.

Amazon Kinesis Data Firehose is a fully managed service to reliably load streaming data into destinations such as Amazon S3, which fits well for capturing streaming customer activity data.

API Gateway integrates natively with Firehose for direct ingestion.

This combination supports unpredictable traffic, smooth scaling, and simple authorization.

Option B uses Kinesis Data Streams, which requires more management than Firehose and is less optimized for direct API integration.

Options A and D use Gateway Load Balancer and ECS containers plus EFS, which add complexity and are less suited for unpredictable traffic with integrated authorization.

Reference:

Amazon API Gateway (<https://docs.aws.amazon.com/apigateway/latest/developerguide/welcome.html>)

Amazon API Gateway Lambda authorizers (<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-use-lambda-authorizer.html>)

Amazon Kinesis Data Firehose (<https://docs.aws.amazon.com/firehose/latest/dev/what-is-this-service.html>)

AWS Well-Architected Framework — Operational Excellence Pillar
(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 431

A company is designing an IPv6 application that is hosted on Amazon EC2 instances in a private subnet within a VPC. The application will store user-uploaded content in Amazon S3 buckets. The application will save each S3 object's URL link and metadata in Amazon DynamoDB.

The company must not use public internet connections to transmit user-uploaded content or metadata.

Which solution will meet these requirements?

- A. Implement a gateway VPC endpoint for Amazon S3 and an interface VPC endpoint for Amazon DynamoDB.
- B. Implement interface VPC endpoints for both Amazon S3 and Amazon DynamoDB.
- C. Implement gateway VPC endpoints for both Amazon S3 and Amazon DynamoDB.
- D. Implement a gateway VPC endpoint for Amazon DynamoDB and an interface VPC endpoint for Amazon S3.

Answer: A

Explanation:

Amazon S3 uses gateway VPC endpoints, which enable private, secure access to S3 without traversing the internet, compatible with IPv6.

Amazon DynamoDB uses interface VPC endpoints (powered by AWS PrivateLink) for private connectivity within the VPC.

Therefore, for secure private communication without public internet, the correct solution is to implement a gateway VPC endpoint for S3 and an interface VPC endpoint for DynamoDB.

Option B is incorrect because S3 does not support interface endpoints; Option C is incorrect because DynamoDB does not support gateway endpoints. Option D reverses the correct endpoint types.

Reference:

VPC Endpoints for Amazon S3 (<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>)

VPC Endpoints for DynamoDB (<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-dynamodb.html>)

AWS Well-Architected Framework — Security Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 432

A solutions architect is configuring a VPC that has public subnets and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs). An internet gateway is attached to the VPC.

The private subnets require access to the internet to allow Amazon EC2 instances to download software updates.

Which solution will meet this requirement?

- A. Create a NAT gateway in one of the public subnets. Update the route tables that are attached to the private subnets to forward non-VPC traffic to the NAT gateway.
- B. Create three NAT instances in each private subnet. Create a private route table for each Availability Zone that forwards non-VPC traffic to the NAT instances.
- C. Attach an egress-only internet gateway in the VPC. Update the route tables of the private subnets to forward non-VPC traffic to the egress-only internet gateway.
- D. Create a NAT gateway in one of the private subnets. Update the route tables that are attached to the private subnets to forward non-VPC traffic to the NAT gateway.

Answer: A

Explanation:

Private subnets require outbound internet access to download updates, but they must not have public IPs or direct inbound access.

The recommended AWS solution is to create a NAT gateway in a public subnet. Private subnet route tables are updated to route internet-bound traffic (0.0.0.0/0) to the NAT gateway. The NAT gateway then uses the internet gateway attached to the VPC to communicate with the internet.

Option B (NAT instances) is an older approach and less scalable/maintainable than NAT gateways. Option C (egress-only internet gateway) is for IPv6 outbound-only traffic, not IPv4. Option D is invalid because NAT gateways must be deployed in public subnets.

Reference:

AWS Well-Architected Framework — Reliability Pillar

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

NAT Gateways (<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>)

VPC Internet Gateways and Subnet Routing (https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html)

Question: 433

A solutions architect is provisioning an Amazon Elastic File System (Amazon EFS) file system to provide shared storage across multiple Amazon EC2 instances. The instances all exist in the same VPC across multiple Availability Zones. There are two instances in each Availability Zone. The solutions architect must make the file system accessible to each instance with the lowest possible latency.

Which solution will meet these requirements?

- A. Create a mount target for the EFS file system in the VPC. Use the mount target to mount the file system on each of the instances.
- B. Create a mount target for the EFS file system in one Availability Zone of the VPC. Use the mount target to mount the file system on the instances in that Availability Zone. Share the directory with the other instances.
- C. Create a mount target for each instance. Use each mount target to mount the EFS file system on each respective instance.
- D. Create a mount target in each Availability Zone of the VPC. Use the mount target to mount the EFS file system on the instances in the respective Availability Zone.

Answer: D

Explanation:

Amazon EFS requires a mount target in each Availability Zone where EC2 instances access the file system. This is because each mount target provides an elastic network interface in the subnet and AZ, reducing network latency by allowing EC2 instances to communicate locally with the EFS mount target. Creating a mount target in each AZ optimizes file system access performance and availability. Instances mount the EFS file system via the mount target in their respective AZ, which provides the lowest possible latency and avoids cross-AZ traffic.

Option A, with only a single mount target in the VPC, will cause cross-AZ traffic for instances in other AZs, increasing latency and potentially incurring data transfer costs. Option B is incomplete and introduces complexity with sharing directories across instances. Option C is invalid because mount targets are per AZ and per subnet, not per instance.

Reference:

Amazon EFS Overview (<https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html>)

Creating Mount Targets (<https://docs.aws.amazon.com/efs/latest/ug/manage-fs-access.html#creating-mount-targets>)

AWS Well-Architected Framework — Performance Efficiency Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 434

A company wants to migrate applications from its on-premises servers to AWS. As a first step, the company is modifying and migrating a non-critical application to a single Amazon EC2 instance. The application will store information in an Amazon S3 bucket. The company needs to follow security best practices when deploying the application on AWS.

Which approach should the company take to allow the application to interact with Amazon S3?

- A. Create an IAM role that has administrative access to AWS. Attach the role to the EC2 instance.
- B. Create an IAM user. Attach the AdministratorAccess policy. Copy the generated access key and secret key. Within the application code, use the access key and secret key along with the AWS SDK to communicate with Amazon S3.
- C. Create an IAM role that has the necessary access to Amazon S3. Attach the role to the EC2 instance.
- D. Create an IAM user. Attach a policy that provides the necessary access to Amazon S3. Copy the generated access key and secret key. Within the application code, use the access key and secret key along with the AWS SDK to communicate with Amazon S3.

Answer: C

Explanation:

The best practice for granting AWS service access to EC2 instances is to assign IAM roles with the least privilege necessary. IAM roles attached to EC2 instances provide temporary credentials automatically rotated by AWS, avoiding hard-coded credentials in the application code.

Option C adheres to security best practices by assigning an IAM role scoped with the minimum permissions needed to access the S3 bucket.

Option A grants administrative access, which violates least privilege principles. Options B and D use IAM users with static credentials, increasing the risk of credential exposure and requiring manual rotation, which is not recommended.

Reference:

IAM Roles for Amazon EC2 (https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html)

AWS Security Best Practices (https://d1.awsstatic.com/whitepapers/Security/AWS_Security_Best_Practices.pdf)

AWS Well-Architected Framework — Security Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 435

An ecommerce company runs a multi-tier application on AWS. The frontend and backend tiers both run on Amazon EC2 instances. The database tier runs on an Amazon RDS for MySQL DB instance. The backend tier communicates with the RDS DB instance.

The application makes frequent calls to return identical datasets from the database. The frequent calls on the database cause performance slowdowns. A solutions architect must improve the performance of the application backend.

Which solution will meet this requirement?

- A. Configure an Amazon Simple Notification Service (Amazon SNS) topic between the EC2 instances and the RDS DB instance.
- B. Configure an Amazon ElastiCache (Redis OSS) cache. Configure the backend EC2 instances to read from the cache.
- C. Configure an Amazon DynamoDB Accelerator (DAX) cluster. Configure the backend EC2 instances to read from the cluster.
- D. Configure Amazon Data Firehose to stream the calls to the database.

Answer: B

Explanation:

Caching frequently accessed, identical datasets is a well-established way to improve backend application performance by reducing load on the database. Amazon ElastiCache with Redis (open source) offers a fast, in-memory data store to cache query results, reducing latency and database requests.

Option B directly addresses the problem by offloading repeated read requests from the database to the cache.

Option A (SNS) is a messaging service and is unrelated to caching or improving database performance. Option C (DAX) accelerates DynamoDB but the backend uses RDS MySQL, so DAX is inapplicable. Option D (Data Firehose) is a data streaming service and does not optimize database read performance.

Reference:

Caching Best Practices (<https://aws.amazon.com/caching/>)

Amazon ElastiCache for Redis (<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/WhatIs.html>)

AWS Well-Architected Framework — Performance Efficiency Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 436

A company is building a gaming application that needs to send unique events to multiple leaderboards, player matchmaking systems, and authentication services concurrently. The company requires an AWS-based event-driven system that delivers events in order and supports a publish-subscribe model. The gaming application must be the publisher, and the leaderboards, matchmaking systems, and authentication services must be the subscribers.

Which solution will meet these requirements?

- A. Amazon EventBridge event buses
- B. Amazon Simple Notification Service (Amazon SNS) FIFO topics
- C. Amazon Simple Notification Service (Amazon SNS) standard topics
- D. Amazon Simple Queue Service (Amazon SQS) FIFO queues

Answer: B

Explanation:

The requirement is an event-driven pub/sub system that guarantees ordered delivery of events.

Amazon SNS FIFO topics provide the publish-subscribe model along with FIFO (First-In-First-Out) delivery and exactly-once message processing, ensuring ordered delivery to multiple subscribers.

Option A, EventBridge, provides event buses but does not guarantee event ordering across multiple subscribers. Option C (SNS standard topics) provides pub/sub but without ordering guarantees.

Option D (SQS FIFO queues) guarantees order but are point-to-point queues, not pub/sub.

Thus, Amazon SNS FIFO topics meet the requirements for ordered pub/sub messaging.

Reference:

Amazon SNS FIFO Topics (<https://docs.aws.amazon.com/sns/latest/dg/fifo-topics.html>)

Amazon EventBridge (<https://docs.aws.amazon.com/eventbridge/latest/userguide/what-is-amazon-eventbridge.html>)

AWS Well-Architected Framework — Performance Efficiency Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 437

A company has a web application that uses Amazon API Gateway to route HTTPS requests to AWS Lambda functions. The application uses an Amazon Aurora MySQL database for its data storage. The application has experienced unpredictable surges in traffic that overwhelm the database with too many connection requests. The company wants to implement a scalable solution that is more resilient to database failures.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon RDS proxy for the database. Replace the database endpoint with the proxy endpoint in the Lambda functions.
- B. Migrate the database to Amazon DynamoDB tables by using AWS Database Migration Service (AWS DMS).

C. Review the existing connections. Call MySQL queries to end any connections in the sleep state.

D. Increase the instance class of the database with more memory. Set a larger value for the

max_connections parameter.

Answer: A

Explanation:

Amazon RDS Proxy helps manage and pool database connections from serverless compute like AWS Lambda, significantly reducing the stress on the database during unpredictable traffic surges. It improves scalability and resiliency by efficiently managing connections, protecting the database from being overwhelmed, and enabling failover handling.

Option A is the most cost-effective and operationally efficient approach to handling unpredictable surges and improving resilience without requiring major application changes.

Option B involves a migration to DynamoDB, which is a significant architectural change and costlier initially. Option C is manual connection cleanup, insufficient for unpredictable surges. Option D increases resources but does not solve connection storm problems efficiently and is more costly.

Reference:

Amazon RDS Proxy (<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-proxy.html>)

AWS Well-Architected Framework — Reliability Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

AWS Lambda Best Practices (<https://docs.aws.amazon.com/lambda/latest/dg/best-practices.html>)

Question: 438

A company has an application that uses an Amazon RDS for PostgreSQL database. The company is developing an application feature that will store sensitive information for an individual in the database.

During a security review of the environment, the company discovers that the RDS DB instance is not encrypting data at rest. The company needs a solution that will provide encryption at rest for all the existing data and for any new data that is entered for an individual.

Which combination of steps should the company take to meet these requirements? (Select TWO.)

A. Create a snapshot of the DB instance. Enable encryption on the snapshot. Use the encrypted snapshot to create a new DB instance.

Adjust the application configuration to use the new DB instance.

B. Create a snapshot of the DB instance. Create an encrypted copy of the snapshot. Use the encrypted snapshot to create a new DB

instance. Adjust the application configuration to use the new DB instance.

C. Modify the configuration of the DB instance by enabling encryption. Create a snapshot of the DB instance. Use the snapshot to create a new DB instance. Adjust the application configuration to use the new DB instance.

D. Use AWS Key Management Service (AWS KMS) to create a new default AWS managed aws/rds key. Select this key as the encryption key for operations with Amazon RDS.

E. Use AWS Key Management Service (AWS KMS) to create a new customer managed key. Select this key as the encryption key for operations with Amazon RDS.

Answer: B,E

Explanation:

Amazon RDS does not support enabling encryption at rest on an existing unencrypted DB instance. To encrypt an existing RDS instance's data at rest, the recommended method is to:

Take a snapshot of the unencrypted DB instance.

Create an encrypted copy of the snapshot using AWS KMS. This encrypted snapshot contains the existing data encrypted at rest.

Restore a new DB instance from the encrypted snapshot. This new instance will have encryption at rest enabled.

Additionally, to manage encryption keys securely, companies can use customer managed keys (CMKs) in AWS Key Management Service (KMS). CMKs provide greater control over key management policies, rotation, and usage permissions compared to default AWS managed keys. Using a CMK allows customization of access control and auditability.

Option A is incorrect because you cannot enable encryption directly on a snapshot; you must create an encrypted copy. Option C is invalid because encryption cannot be enabled by modifying an existing instance's configuration. Option D refers to the default AWS managed key, which is less flexible than customer managed keys.

Reference:

Encrypting Amazon RDS Resources

(<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>)

Copying an Encrypted Snapshot

(https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CopySnapshot.html)

AWS KMS Customer Master Keys

(<https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html>)

AWS Well-Architected Framework — Security Pillar

(https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 439

A company wants to migrate applications from its on-premises servers to AWS. As a first step, the company is modifying and migrating a non-critical application to a single Amazon EC2 instance. The application will store information in an Amazon S3 bucket. The company needs to follow security best practices when deploying the application on AWS.

Which approach should the company take to allow the application to interact with Amazon S3?

- A. Store the files in an Amazon S3 bucket. Use the S3 Glacier Instant Retrieval storage class. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Deep Archive storage class after 1 year.
- B. Store the files in an Amazon S3 bucket. Use the S3 Standard storage class. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Flexible Retrieval storage class after 1 year.
- C. Store the files on an Amazon Elastic Block Store (Amazon EBS) volume. Use Amazon Data Lifecycle Manager to create snapshots of the EBS volumes and to store those snapshots in Amazon S3.
- D. Store the files on an Amazon Elastic File System (Amazon EFS) mount. Configure EFS lifecycle management to transition the files to the EFS Standard-Infrequent Access (Standard-IA) storage class after 1 year.

Answer: A

Explanation:

The company's requirements are:

Store gigabytes of rarely accessed files daily.

Files must be available within minutes during the first year.

Retain files for 7 years cost-effectively.

The S3 Glacier Instant Retrieval storage class provides low-cost, long-term storage for data accessed occasionally, with millisecond retrieval time. This fits the requirement for availability within minutes during the first year. After one year, transitioning files to S3 Glacier Deep Archive (the lowest cost S3 storage class) with longer retrieval times is cost-effective for data retention over 7 years.

Option B uses S3 Standard and Glacier Flexible Retrieval, which is higher cost during the first year and slower retrieval times. Option C is less cost-efficient because EBS volumes are expensive for rarely accessed data and snapshots incur additional costs. Option D uses EFS, which is designed for low-latency file storage but is more expensive than S3 Glacier classes for long-term archival.

Reference:

Amazon S3 Storage Classes (<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html>)

S3 Glacier Instant Retrieval (<https://aws.amazon.com/s3/storage-classes/glacier-instant-retrieval/>)

Amazon S3 Lifecycle Management (<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration->

AWS Well-Architected Framework — Cost Optimization Pillar (https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf)

Question: 440

A solutions architect is designing a three-tier web application. The architecture consists of an internet-facing Application Load Balancer (ALB) and a web tier that is hosted on Amazon EC2 instances in private subnets. The application tier with the business logic runs on EC2 instances in private subnets. The database tier consists of Microsoft SQL Server that runs on EC2 instances in private subnets. Security is a high priority for the company. Which combination of security group configurations should the solutions architect use? (Select THREE.)

- A. Configure the security group for the web tier to allow inbound HTTPS traffic from the security group for the ALB.
- B. Configure the security group for the web tier to allow outbound HTTPS traffic to 0.0.0.0/0.
- C. Configure the security group for the database tier to allow inbound Microsoft SQL Server traffic from the security group for the application tier.
- D. Configure the security group for the database tier to allow outbound HTTPS traffic and Microsoft SQL Server traffic to the security group for the web tier.
- E. Configure the security group for the application tier to allow inbound HTTPS traffic from the security group for the web tier.
- F. Configure the security group for the application tier to allow outbound HTTPS traffic and Microsoft SQL Server traffic to the security group for the web tier.

Answer: A,C,E

Explanation:

According to AWS best practices, each tier's security group must restrict inbound traffic to only the upstream trusted source. For the web tier, the Application Load Balancer must be the only entity allowed to send traffic. AWS documentation specifies: "Restrict the security groups associated with your targets to accept traffic only from the load balancer." This confirms that the web tier security group should allow inbound HTTPS from the ALB security group (A).

For communication between the web and application tiers, AWS states: "You can specify a security group as the source or destination in a rule" and "Create rules only for the protocols and ports required by your application." Therefore, the application tier security group must allow inbound HTTPS traffic from the web tier security group (E).

For the database tier, AWS guidance says: "Allow only the necessary ports for database communication." Microsoft SQL Server listens on port 1433 by default, so the database tier security group must allow inbound SQL Server traffic from the application tier security group (C).

Outbound rules (options B, D, and F) are unnecessary because AWS specifies that "Security groups are stateful. Return traffic is automatically allowed." This means once inbound rules are defined, the return path is automatically permitted without extra outbound

configurations.

This combination (A, C, E) applies the principle of least privilege, ensures end-to-end secure communication across tiers, and follows AWS recommendations for ALB-to-target security group setups.

References Elastic Load Balancing User Guide — Application Load Balancers: Security groups for your load balancer, Target security groups^ Amazon VPC User Guide: Security groups for your VPC, Security group rules^ AWS Well-Architected Framework — Security Pillar: Apply the principle of least privilege

Question: 441

A company runs multiple web applications on Amazon EC2 instances behind a single Application Load Balancer (ALB). The application experiences unpredictable traffic spikes throughout each day. The traffic spikes cause high latency. The unpredictable spikes last less than 3 hours. The company needs a solution to resolve the latency issue caused by traffic spikes.

- A. Use EC2 instances in an Auto Scaling group. Configure the ALB and Auto Scaling group to use a target tracking scaling policy.
- B. Use EC2 Reserved Instances in an Auto Scaling group. Configure the Auto Scaling group to use a scheduled scaling policy based on peak traffic hours.
- C. Use EC2 Spot Instances in an Auto Scaling group. Configure the Auto Scaling group to use a scheduled scaling policy based on peak traffic hours.
- D. Use EC2 Reserved Instances in an Auto Scaling group. Replace the ALB with a Network Load Balancer (NLB).

Answer: A

Explanation:

AWS recommends Auto Scaling with dynamic scaling policies to handle unpredictable workload spikes. Target tracking scaling policies automatically adjust capacity based on defined metrics such as average CPU utilization or request count per target. This approach ensures applications maintain responsiveness without overprovisioning. Scheduled scaling (options B and C) is only effective when traffic patterns are predictable, which is not the case here. Reserved Instances or Spot Instances also do not address sudden demand changes effectively. Replacing the ALB with an NLB (D) does not solve latency caused by EC2 instance capacity. Therefore, using EC2 instances in an Auto Scaling group with a target tracking scaling policy is the most effective and cost-efficient solution for unpredictable, short-lived traffic spikes.

Reference:* Amazon EC2 Auto Scaling User Guide — Target tracking scaling policies* AWS Well-Architected Framework — Performance Efficiency Pillar: Elasticity and scaling

Question: 442

A company uses an Amazon EC2 instance to run a script to poll for and process messages in an Amazon Simple Queue

Service (Amazon SQS) queue. The company wants to reduce operational overhead while maintaining its ability to process an increasing number of messages that are added to the queue. Which solution will meet these requirements?

- A. Increase the size of the EC2 instance to process messages in the SQS queue faster.
- B. Configure an Amazon EventBridge rule to turn off the EC2 instance when the SQS queue is empty.
- C. Migrate the script on the EC2 instance to an AWS Lambda function with an event source of the SQS queue.
- D. Configure an AWS Systems Manager Run Command to run the script on demand.

Answer: C

Explanation:

AWS advises using serverless event-driven processing to minimize operational burden and scale automatically with demand. Amazon SQS can be directly configured as an event source for AWS Lambda. With this setup, Lambda polls the queue, invokes the function, and processes messages without requiring the customer to manage infrastructure. Scaling is automatic, based on the volume of messages in the queue. Option A (increasing instance size) still leaves scaling and management challenges. Option B reduces cost when idle but does not address scaling. Option D requires manual triggering and does not meet continuous processing requirements. Migrating the script to Lambda (C) fully eliminates the need for instance management, providing scalability, reliability, and reduced operational overhead.

References Amazon SQS Developer Guide — Using AWS Lambda with Amazon SQS^ AWS Well- Architected Framework — Operational Excellence Pillar

Question: 443

A company hosts an application on AWS that stores files that users need to access. The application uses two Amazon EC2 instances. One instance is in Availability Zone A, and the second instance is in Availability Zone B. Both instances use Amazon Elastic Block Store (Amazon EBS) volumes. Users must be able to access the files at any time without delay. Users report that the two instances occasionally contain different versions of the same file. Users occasionally receive HTTP 404 errors when they try to download files. The company must address the customer issues. The company cannot make changes to the application code. Which solution will meet these requirements in the MOST operationally efficient way?

- A. Run the robocopy command on one of the EC2 instances on a schedule to copy files from the Availability Zone A instance to the Availability Zone B instance.
- B. Configure the application to store the files on both EBS volumes each time a user writes or updates a file.
- C. Mount an Amazon Elastic File System (Amazon EFS) file system to the EC2 instances. Copy the files from the EBS volumes to the EFS file system. Configure the application to store files in the EFS file system.
- D. Create an EC2 instance profile that allows the instance in Availability Zone A to access the S3 bucket. Re-associate the instance profile to the instance in Availability Zone B when needed.

Answer: C

Explanation:

Amazon EFS provides a fully managed, highly available, and shared file system that can be mounted by instances across multiple Availability Zones. This ensures consistency of files between EC2 instances and avoids replication issues. EBS volumes, in contrast, are AZ-scoped and not designed for multi-instance sharing. Options A and B rely on custom replication or manual file handling, which increases operational overhead and risks inconsistencies. Option D does not solve the shared access and consistency requirement. By migrating storage to EFS, both EC2 instances will read and write to the same storage system, ensuring that users always access the latest files without 404 errors.

References Amazon EFS User Guide — Accessing data across multiple Availability Zones^ AWS Well - Architected Framework — Reliability Pillar: Data durability and consistency

Question: 444

A company is planning to connect a remote office to its AWS infrastructure. The office requires permanent and secure connectivity to AWS. The connection must provide secure access to resources in two VPCs. However, the VPCs must not be able to access each other.

- A. Create two transit gateways. Set up one AWS Site-to-Site VPN connection from the remote office to each transit gateway. Connect one VPC to the transit gateway. Configure route table propagation to the appropriate transit gateway based on the destination VPC IP range.
- B. Set up one AWS Site-to-Site VPN connection from the remote office to each of the VPCs. Update the VPC route tables with static routes to the remote office resources.
- C. Set up one AWS Site-to-Site VPN connection from the remote office to one of the VPCs. Set up VPC peering between the two VPCs. Update the VPC route tables with static routes to the remote office and peered resources.
- D. Create a transit gateway. Set up an AWS Direct Connect gateway and one Direct Connect connection between the remote office and the Direct Connect gateway. Associate the transit gateway with the Direct Connect gateway. Configure a separate private virtual interface (VIF) for each VPC, and configure routing.

Answer: B

Explanation:

The requirement is secure, permanent connectivity to two VPCs without enabling VPC-to-VPC communication. Option B achieves this by establishing a separate Site-to-Site VPN connection between the remote office and each VPC. Updating each VPC's route tables ensures traffic from the remote office reaches the correct VPC resources, while preventing routing between the VPCs themselves. Options involving transit gateways (A and D) would introduce transitive routing capabilities, which violates the requirement of isolation between VPCs.

Option C (VPC peering) directly conflicts with the requirement by enabling cross-VPC access. Therefore, option B is the correct solution as it maintains isolation while providing secure VPN connectivity.

References AWS Site-to-Site VPN User Guide — Connecting on-premises networks to multiple VPCs^ AWS Well-Architected Framework — Security Pillar: Network isolation and segmentation

Question: 445

A company hosts an application in an Amazon EC2 Auto Scaling group. The company has observed that during periods of high demand, new instances take too long to join the Auto Scaling group and serve the increased demand. The company determines that the root cause of the issue is the long boot time of the instances in the Auto Scaling group. The company needs to reduce the time required to launch new instances to respond to demand. Which solution will meet this requirement?

- A. Increase the maximum capacity of the Auto Scaling group by 50%.
- B. Create a warm pool for the Auto Scaling group. Use the default specification for the warm pool size.
- C. Increase the health check grace period for the Auto Scaling group by 50%.
- D. Create a scheduled scaling action. Set the desired capacity equal to the maximum capacity of the Auto Scaling group.

Answer: B

Explanation:

A warm pool is an Auto Scaling feature that keeps instances in a pre-initialized state so they can quickly join the active group when scaling is required. This reduces the time needed for instance bootstrapping and makes new capacity available almost instantly. Option A only increases capacity limits but does not address slow boot times. Option C merely extends grace periods without solving the delay. Option D forces overprovisioning, which is wasteful and not aligned with cost optimization. Using a warm pool (B) directly addresses the problem by reducing response time to scaling events.

References Amazon EC2 Auto Scaling User Guide — Warm pools for scaling fastest AWS Well-Architected Framework — Performance Efficiency Pillar: Optimizing responsiveness

Question: 446

A company has an application that uses a MySQL database that runs on an Amazon EC2 instance. The instance currently runs in a single Availability Zone. The company requires a fault-tolerant database solution that provides a recovery time objective (RTO) and a recovery point objective (RPO) of 2 minutes or less. Which solution will meet these requirements?

- A. Migrate the MySQL database to Amazon RDS. Create a read replica in a second Availability Zone. Create a script that detects availability interruptions and promotes the read replica when needed.
- B. Migrate the MySQL database to Amazon RDS for MySQL. Configure the new RDS for MySQL database to use a Multi-AZ deployment.

- C. Create a second MySQL database in a second Availability Zone. Use native MySQL commands to sync the two databases every 2 minutes. Create a script that detects availability interruptions and promotes the second MySQL database when needed.
- D. Create a copy of the EC2 instance that runs the MySQL database. Deploy the copy in a second Availability Zone. Create a Network Load Balancer. Add both instances as targets.

Answer: B

Explanation:

Amazon RDS Multi-AZ deployments provide automatic failover for relational databases such as MySQL, ensuring high availability and durability. The feature maintains synchronous replication between a primary DB instance and a standby in a separate Availability Zone. AWS guarantees that failover typically completes within minutes, ensuring an RTO and RPO of less than 2 minutes. Option A requires manual promotion of replicas, which cannot meet the strict RTO/RPO requirement. Option C depends on custom scripts and manual synchronization, introducing operational risk. Option D creates active-active EC2-based databases, which do not provide synchronous replication or automated failover. Therefore, Multi-AZ RDS (B) is the managed, resilient, and operationally efficient solution that meets the business requirements.

References Amazon RDS User Guide — Multi-AZ deployments^ AWS Well-Architected Framework — Reliability Pillar: High availability and disaster recovery

Question: 447

An ecommerce company hosts an analytics application on AWS. The company deployed the application to one AWS Region. The application generates 300 MB of data each month. The application stores the data in JSON format. The data must be accessible in milliseconds when needed. The company must retain the data for 30 days. The company requires a disaster recovery solution to back up the data.

- A. Deploy an Amazon OpenSearch Service cluster in the primary Region and in a second Region. Enable OpenSearch Service cluster replication. Configure the clusters to expire data after 30 days. Modify the application to use OpenSearch Service to store the data.
- B. Deploy an Amazon S3 bucket in the primary Region and in a second Region. Enable versioning on both buckets. Use the Standard storage class. Configure S3 Lifecycle policies to expire objects after 30 days. Configure S3 Cross-Region Replication from the bucket in the primary bucket to the backup bucket.
- C. Deploy an Amazon Aurora PostgreSQL global database. Configure cluster replication between the primary Region and a second Region. Use a replicated cluster endpoint during outages in the primary Region.
- D. Deploy an Amazon RDS for PostgreSQL cluster in the same Region where the application is deployed. Configure a read replica in a second Region as a backup.

Answer: B

Explanation:

Amazon S3 is designed for durability, scalability, and millisecond access. For small monthly data volumes (300 MB), S3 Standard is cost-effective and provides immediate access. To meet 30-day retention, Lifecycle policies can automatically expire objects after the required time. For disaster recovery, S3 Cross-Region Replication (CRR) copies objects across Regions to a backup bucket, ensuring data resiliency. OpenSearch (A) is not needed because the requirement is storage and retrieval, not indexing. Aurora or RDS options (C, D) add unnecessary complexity and cost, as a relational database is not required for JSON storage and millisecond retrieval. Therefore, option B provides the simplest, most resilient, and cost-optimized solution.

Reference: * Amazon S3 User Guide — Lifecycle policies, Cross-Region Replication * AWS Well-Architected Framework — Reliability Pillar: Data backup and disaster recovery

Question: 448

A gaming company is developing a game that requires significant compute resources to process game logic, player interactions, and real-time updates. The company needs a compute solution that can dynamically scale based on fluctuating player demand while maintaining high performance. The company must use a relational database that can run complex queries.

- A. Deploy Amazon EC2 instances to supply compute capacity. Configure Auto Scaling groups to achieve dynamic scaling based on player count. Use Amazon RDS for MySQL as the database.
- B. Refactor the game logic into small, stateless functions. Use AWS Lambda to process the game logic. Use Amazon DynamoDB as the database.
- C. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate to supply compute capacity. Scale the ECS tasks based on player demand. Use Amazon Aurora Serverless v2 as the database.
- D. Use AWS ParallelCluster for high performance computing (HPC). Provision compute nodes that have GPU instances to process the game logic and player interactions. Use Amazon RDS for MySQL as the database.

Answer: C

Explanation:

Amazon ECS with AWS Fargate provides serverless container compute that can scale tasks dynamically in response to demand. This matches the unpredictable scaling requirements of a gaming workload. Aurora Serverless v2 provides on-demand, autoscaling relational database capacity while supporting complex SQL queries. EC2 with Auto Scaling (A) works but requires significant management overhead. Lambda with DynamoDB (B) is not suitable because the workload requires a relational database and long-running processes. ParallelCluster with HPC (D) is designed for batch scientific workloads, not dynamic, interactive gaming. Option C provides the correct balance of elasticity, performance, and managed services for both compute and relational database needs.

Reference: * Amazon ECS User Guide — Scaling ECS tasks with Fargate* Amazon Aurora User Guide — Aurora Serverless v2 for dynamic workloads

Question: 449

A company runs game applications on AWS. The company needs to collect, visualize, and analyze telemetry data from the company's game servers. The company wants to gain insights into the behavior, performance, and health of game servers in near real time. Which solution will meet these requirements?

- A. Use Amazon Kinesis Data Streams to collect telemetry data. Use Amazon Managed Service for Apache Flink to process the data in near real time and publish custom metrics to Amazon CloudWatch. Use Amazon CloudWatch to create dashboards and alarms from the custom metrics.
- B. Use Amazon Data Firehose to collect, process, and store telemetry data in near real time. Use AWS Glue to extract, transform, and load (ETL) data from Firehose into required formats for analysis. Use Amazon QuickSight to visualize and analyze the data.
- C. Use Amazon Kinesis Data Streams to collect, process, and store telemetry data. Use Amazon EMR to process the data in near real time into required formats for analysis. Use Amazon Athena to analyze and visualize the data.
- D. Use Amazon DynamoDB Streams to collect and store telemetry data. Configure DynamoDB Streams to invoke AWS Lambda functions to process the data in near real time. Use Amazon Managed Grafana to visualize and analyze the data.

Answer: A

Explanation:

Amazon Kinesis Data Streams is designed for low-latency ingestion of streaming data. Combined with Amazon Managed Service for Apache Flink, telemetry can be processed and aggregated in near real time. Processed metrics can be sent to Amazon CloudWatch, which natively supports creating dashboards, metrics visualization, and alarms. Firehose (B) is primarily for batch ingestion and delivery, not real-time analytics. EMR with Athena (C) introduces more complexity and is better for large-scale offline analytics. DynamoDB Streams (D) is not a fit because telemetry data is not stored in DynamoDB. Therefore, option A provides the most suitable and real-time analytics pipeline for telemetry data.

Reference: * Amazon Kinesis Data Streams Developer Guide — Real-time data ingestion* Amazon Managed Service for Apache Flink Developer Guide — Real-time stream processing

Question: 450

A company is designing an application to maintain a record of customer orders. The application will generate events. The company wants to use an Amazon EventBridge event bus to send the application's events to an Amazon DynamoDB table. Which solution will meet these requirements?

- A. Use the EventBridge default event bus. Configure DynamoDB Streams for the DynamoDB table that hosts the customer order data.
- B. Create an EventBridge custom event bus. Create an AWS Lambda function as a target. Configure the Lambda function to forward the customer order data to the DynamoDB table.
- C. Create an EventBridge partner event bus. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe an AWS Lambda function to the SNS topic. Configure the Lambda function to read the customer order data and to forward the data to the DynamoDB table.
- D. Create an EventBridge partner event bus. Create an AWS Lambda function as a target. Configure the Lambda function to forward the customer order data to the DynamoDB table.

Answer: C

Explanation:

Amazon EventBridge supports routing application-generated events to AWS Lambda targets. The Lambda function can process and insert events into a DynamoDB table. This is a standard design pattern for connecting event-driven applications with DynamoDB. Option A confuses DynamoDB Streams (which streams changes from DynamoDB) with EventBridge (which is for event routing). Options C and D reference partner event buses, which are intended for SaaS integrations, not custom application events. Therefore, the correct and simplest solution is to use a custom EventBridge event bus with a Lambda target (B).

References Amazon EventBridge User Guide — Targets and event routings AWS Lambda Developer Guide — Using DynamoDB with Lambda

Question: 451

A finance company uses an on-premises search application to collect streaming data from various producers. The application provides real-time updates to search and visualization features. The company is planning to migrate to AWS and wants to use an AWS native solution.

Which solution will meet these requirements?

- A. Use Amazon EC2 instances to ingest and process the data streams to Amazon S3 buckets for storage. Use Amazon Athena to search the data. Use Amazon Managed Grafana to create visualizations.
- B. Use Amazon EMR to ingest and process the data streams to Amazon Redshift for storage. Use Amazon Redshift Spectrum to search the data. Use Amazon QuickSight to create visualizations.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS) to ingest and process the data streams to Amazon DynamoDB for storage. Use Amazon CloudWatch to create graphical dashboards to search and visualize the data.
- D. Use Amazon Kinesis Data Streams to ingest and process the data streams to Amazon OpenSearch Service. Use OpenSearch Service to search the data. Use Amazon QuickSight to create visualizations.

Answer: D

Explanation:

Amazon OpenSearch Service is the AWS-native service for real-time search, analytics, and visualization. Ingesting streaming data with Amazon Kinesis Data Streams ensures scalable and reliable data ingestion. OpenSearch supports indexing and querying the data in near real time. QuickSight can then be integrated for visualization and reporting. Options A and B involve batch processing and are not optimized for real-time search. Option C (EKS + DynamoDB) is not a fit for search workloads, as DynamoDB does not support advanced text search. Option D provides the closest AWS-native equivalent to the company's existing search-and-visualization architecture.

Reference: * Amazon OpenSearch Service Developer Guide — Real-time search and analytics* Amazon Kinesis Data Streams

Developer Guide — Ingesting real-time data

Question: 452

A company has a VPC with multiple private subnets that host multiple applications. The applications must not be accessible to the internet. However, the applications need to access multiple AWS services. The applications must not use public IP addresses to access the AWS services.

- A. Configure interface VPC endpoints for the required AWS services. Route traffic from the private subnets through the interface VPC endpoints.
- B. Deploy a NAT gateway in each private subnet. Route traffic from the private subnets through the NAT gateways.
- C. Deploy internet gateways in each private subnet. Route traffic from the private subnets through the internet gateways.
- D. Set up an AWS Direct Connect connection between the private subnets. Route traffic from the private subnets through the Direct Connect connection.

Answer: A

Explanation:

AWS VPC endpoints (interface and gateway) allow private connectivity from VPC resources to AWS services without requiring public IP addresses or internet gateways. This ensures applications remain isolated in private subnets while securely accessing AWS services. NAT gateways (B) would allow internet access, which does not meet the security requirement. Internet gateways (C) directly expose traffic to the internet, which violates the isolation requirement. Direct Connect (D) connects onpremises environments to AWS but does not provide service access from private subnets. Therefore, option A — using interface VPC endpoints — is the correct solution.

References Amazon VPC User Guide — VPC endpoints (interface and gateway)^ AWS Well- Architected Framework — Security Pillar: Network isolation and private connectivity

Question: 453

A company needs to archive an on-premises relational database. The company wants to retain the data. The company needs to be able to run SQL queries on the archived data to create annual reports. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Database Migration Service (AWS DMS) to migrate the on-premises database to an Amazon RDS instance. Retire the on-premises database. Maintain the RDS instance in a stopped state until the data is needed for reports.
- B. Set up database replication from the on-premises database to an Amazon EC2 instance. Retire the on-premises database. Make a snapshot of the EC2 instance. Maintain the EC2 instance in a stopped state until the data is needed for reports.
- C. Create a database backup on premises. Use AWS DataSync to transfer the data to Amazon S3. Create an S3 Lifecycle configuration to move the data to S3 Glacier Deep Archive. Restore the backup to Amazon EC2 instances to run reports.
- D. Use AWS Database Migration Service (AWS DMS) to migrate the on-premises databases to Amazon S3 in Apache Parquet format. Store the data in S3 Glacier Flexible Retrieval. Use Amazon Athena to run reports.

Answer: D

Explanation:

Amazon S3 is the most cost-effective option for archiving data. Using AWS DMS to migrate to S3 in Apache Parquet format provides an optimized columnar format for analytics. By storing in S3 Glacier Flexible Retrieval, costs are minimized while maintaining compliance with retention. When queries are required, Athena can run SQL queries directly on archived data in S3 without provisioning infrastructure.

Options A and B rely on maintaining RDS or EC2 instances, which increases cost and operational overhead. Option C requires full restores to EC2 before running queries, which is slow and inefficient. Therefore, D provides the lowest operational overhead and direct query capability with Athena.

References AWS DMS Documentation — Migrating databases to Amazon S3 in Parquet format
Amazon Athena User Guide — Querying data stored in S3
AWS Well-Architected Framework — Cost

Optimization Pillar

Question: 454

An ecommerce company is redesigning a product catalog system to handle millions of products and provide fast access to product information. The system needs to store structured product data such as product name, price, description, and category. The system also needs to store unstructured data such as high-resolution product videos and user manuals. The architecture must be highly available and must be able to handle sudden spikes in traffic during large-scale sales events.

- A. Use an Amazon RDS Multi-AZ deployment to store product information. Store product videos and user manuals in Amazon S3.
- B. Use Amazon DynamoDB to store product information. Store product videos and user manuals in Amazon S3.

- C. Store all product information, including product videos and user manuals, in Amazon DynamoDB.
- D. Deploy an Amazon DocumentDB (with MongoDB compatibility) cluster to store all product information, product videos, and user manuals.

Answer: B

Explanation:

Amazon DynamoDB provides single-digit millisecond performance at any scale and is fully managed to handle millions of catalog records. It is ideal for structured catalog data such as product metadata and scales seamlessly during high-traffic events like sales. Amazon S3 is optimized for storing unstructured large objects such as videos and manuals, with virtually unlimited scalability and high durability. Option A (RDS) would not handle massive scale or traffic spikes as efficiently. Option C overloads DynamoDB by forcing it to store large binary data, which is not its purpose. Option D (DocumentDB) is suitable for JSON-like documents but not optimal for storing large media files and would add operational complexity. Therefore, option B represents the best separation of structured and unstructured data storage.

References DynamoDB Developer Guide — Millisecond performance at scales Amazon S3 User Guide — Storage for unstructured data[^]
AWS Well-Architected Framework — Performance Efficiency Pillar

Question: 455

A company stores sensitive financial reports in an Amazon S3 bucket. To comply with auditing requirements, the company must encrypt the data at rest. Users must not have the ability to change the encryption method or remove encryption when the users upload data. The company must be able to audit all encryption and storage actions. Which solution will meet these requirements and provide the MOST granular control?

- A. Enable default server-side encryption with Amazon S3 managed keys (SSE-S3) for the S3 bucket. Apply a bucket policy that denies any upload requests that do not include the x-amz-server-side-encryption header.
- B. Configure server-side encryption with AWS KMS (SSE-KMS) keys. Use an S3 bucket policy to reject any data that is not encrypted by the designated key.
- C. Use client-side encryption before uploading the reports. Store the encryption keys in AWS Secrets Manager.
- D. Enable default server-side encryption with Amazon S3 managed keys (SSE-S3). Use AWS Identity and Access Management (IAM) to prevent users from changing S3 bucket settings.

Answer: B

Explanation:

AWS KMS with SSE-KMS provides granular key management and auditability. All use of KMS keys is logged in AWS CloudTrail, which allows compliance teams to monitor encryption and decryption operations. A bucket policy can be configured to enforce uploads only with the designated KMS key, ensuring that users cannot bypass encryption or change methods. Option A (SSE-S3 with bucket policy) enforces encryption but does not provide the same level of control or auditable key usage. Option C (client-side encryption) increases complexity and key management burden. Option D prevents bucket setting changes but does not prevent unencrypted uploads. Therefore, B ensures the most granular control, auditability, and compliance with financial data requirements.

Reference: * Amazon S3 User Guide — Using SSE-KMS for encryption * AWS KMS Developer Guide — Key management and auditing with CloudTrail * AWS Well-Architected Framework — Security Pillar

Question: 456

A company needs to ensure that an IAM group that contains database administrators can perform operations only within Amazon RDS. The company must ensure that the members of the IAM group cannot access any other AWS services.

A. Create an IAM policy that includes a statement that has the Effect "Allow" and the Action "rds:". Attach the IAM policy to the IAM group.

B. Create an IAM policy that includes two statements. Configure the first statement to have the Effect "Allow" and the Action "rds:". Configure the second statement to have the Effect "Deny" and the Action "". Attach the IAM policy to the IAM group.

C. Create an IAM policy that includes a statement that has the Effect "Deny" and the NotAction "rds:". Attach the IAM policy to the IAM group.

D. Create an IAM policy with a statement that includes the Effect "Allow" and the Action "rds:". Include a permissions boundary that has the Effect "Allow" and the Action "rds:". Attach the IAM policy to the IAM group.

Answer: C

Explanation:

To enforce that IAM users can only access Amazon RDS and no other AWS services, the recommended approach is to use a Deny statement with NotAction. This ensures that all actions are denied except RDS actions. Options A and B do not fully achieve the restriction: A only allows RDS but does not explicitly deny access to other services if another policy grants access; B's explicit Deny for "" would override all other permissions, including the intended RDS Allow, which would result in no access at all. Option D with permissions boundaries still allows other attached policies to grant access outside RDS. Therefore, C is the correct approach to enforce RDS-only access.

Reference: * IAM JSON Policy Elements — Effect, Action, NotAction, and Deny * AWS Well-Architected Framework — Security Pillar: Least privilege

Question: 457

A global company runs a data lake application in the us-east-1 Region and the eu-west-1 Region in an active-passive configuration. Application data is stored locally in Amazon S3 buckets in each AWS Region. The bucket in us-east-1 is the primary active bucket that handles all writes. The company needs to ensure that the application has Regional fault tolerance. The company also needs the storage layer to provide a highly available active-active capability for reads across Regions. The storage layer must provide low latency access through a single global endpoint.

- A. Create an Amazon CloudFront distribution in each Region. Set the S3 bucket within each Region as the origin for the CloudFront distribution in the same Region.
- B. Use S3 Transfer Acceleration for cross-Region data transfers to the S3 buckets.
- C. Configure AWS Backup to replicate S3 buckets across Regions. Set up a disaster recovery environment.
- D. Create an S3 Multi-Region Access Point. Configure cross-Region replication.

Answer: D

Explanation:

Amazon S3 Multi-Region Access Points allow applications to access S3 buckets in multiple Regions through a single global endpoint. This provides active-active read access with automatic routing to the closest bucket for low latency. With cross-Region replication, writes in the primary Region are automatically copied to the secondary Region, providing fault tolerance. Option A (CloudFront) provides caching and distribution, but does not address write replication or active-active bucket access. Option B (Transfer Acceleration) optimizes uploads across distances but does not enable cross-Region fault tolerance. Option C (AWS Backup) is designed for backup/restore, not real-time multi-Region reads and writes. Therefore, D is the correct solution for active-active read access and disaster recovery.

Reference: * Amazon S3 Multi-Region Access Points — Global endpoint access and routing* AWS Well-Architected Framework — Reliability Pillar: Multi-Region design

Question: 458

A company needs a data encryption solution for a machine learning (ML) process. The solution must use an AWS managed service. The ML process currently reads a large number of objects in Amazon S3 that are encrypted by a customer managed AWS KMS key. The current process incurs significant costs because of excessive calls to AWS Key Management Service (AWS KMS) to decrypt S3 objects.

The company wants to reduce the costs of API calls to decrypt S3 objects.

- A. Switch from a customer managed KMS key to an AWS managed KMS key.
- B. Remove the AWS KMS encryption from the S3 bucket. Use a bucket policy to encrypt the data instead.
- C. Recreate the KMS key in AWS CloudHSM.

D. Use S3 Bucket Keys to perform server-side encryption with AWS KMS keys (SSE-KMS) to encrypt and decrypt objects from Amazon S3.

Answer: D

Explanation:

Amazon S3 Bucket Keys reduce the cost of AWS KMS API requests by generating a data key at the bucket level instead of individually calling KMS for every object read or written. This approach is particularly effective when workloads, such as ML pipelines, involve reading large numbers of encrypted objects. Switching to AWS managed keys (A) does not reduce the frequency of API calls. Removing encryption (B) would violate compliance/security requirements. Using CloudHSM (C) adds cost and operational burden. Therefore, the correct solution is D — enabling S3 Bucket Keys with SSE-KMS, which significantly reduces decryption costs while maintaining secure encryption.

References Amazon S3 User Guide — Using S3 Bucket Keys for SSE-KMS^ AWS KMS Developer Guide — Cost optimization for KMS encryption

Question: 459

A company runs an application in a VPC on AWS. The company's on-premises data center has a DNS server. The data center is connected to AWS through an AWS Direct Connect connection with a private virtual interface (VIF). The on-premises DNS server needs to resolve the DNS name of the application in the VPC.

- A. Set up AWS Verified Access endpoints in the VPC. Configure DNS forwarding rules in Verified Access. Configure the on-premises DNS server to forward DNS queries through the Verified Access endpoints.
- B. Configure the Direct Connect connection to enable DNS resolution between the on-premises DNS server and the application in the VPC.
- C. Create an Amazon Route 53 Resolver outbound endpoint and a Resolver rule in the VPC. Configure the on-premises DNS server to send requests for the application to the outbound endpoint.
- D. Create an Amazon Route 53 Resolver inbound endpoint in the VPC. Configure the on-premises DNS server to send requests for the application to the inbound endpoint.

Answer: D

Explanation:

When on-premises DNS servers need to resolve private DNS names in a VPC, the correct pattern is to create a Route 53 Resolver inbound endpoint. The inbound endpoint allows DNS queries to flow from the on-premises environment into the VPC, where Route 53 can resolve VPC-specific names (such as private hosted zones or private resource records). Outbound endpoints (C) are for sending VPC DNS queries to on-premises, not the reverse. Verified Access (A) is unrelated to DNS resolution.

Direct Connect (B) provides network connectivity but does not provide DNS forwarding capabilities. Therefore, option D is the correct design.

Reference: * Amazon Route 53 Resolver Developer Guide — Inbound and outbound endpoints* AWS Well-Architected Framework — Security Pillar: Hybrid DNS integration

Question: 460

A company has a batch processing application that runs every day. The process typically takes an average 3 hours to complete. The application can handle interruptions and can resume the process after a restart. Currently, the company runs the application on Amazon EC2 On-Demand Instances. The company wants to optimize costs while maintaining the same performance level. Which solution will meet these requirements MOST cost-effectively?

- A. Purchase a 1-year EC2 Instance Savings Plan for the appropriate instance family and size to meet the requirements of the application.
- B. Use EC2 On-Demand Capacity Reservations based on the appropriate instance family and size to meet the requirements of the application. Run the EC2 instances in an Auto Scaling group.
- C. Determine the appropriate instance family and size to meet the requirements of the application. Convert the application to run on AWS Batch with EC2 On-Demand Instances. Purchase a 1-year Compute Savings Plan.
- D. Determine the appropriate instance family and size to meet the requirements of the application. Convert the application to run on AWS Batch with EC2 Spot Instances.

Answer: D

Explanation:

Since the workload is interruptible and can resume after restarts, Amazon EC2 Spot Instances are the most cost-effective option, offering savings of up to 90% compared to On-Demand. Running the batch workload on AWS Batch with Spot Instances allows automatic job queue management, interruption handling, and scheduling. Option A and C reduce cost but still rely on reserved or committed pricing for On-Demand capacity. Option B (Capacity Reservations) increases cost instead of reducing it. Therefore, the most cost-optimized solution is D — AWS Batch with EC2 Spot Instances.

Reference: * AWS Batch User Guide — Using Spot Instances with AWS Batch* AWS Well-Architected Framework — Cost Optimization Pillar

Question: 461

A company is developing a content sharing platform that currently handles 500 GB of user-generated media files. The company expects the amount of content to grow significantly in the future. The company needs a storage solution that can automatically scale, provide high durability, and allow direct user uploads from web browsers.

- A. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled.
- B. Store the data in an Amazon Elastic File System (Amazon EFS) Standard file system.
- C. Store the data in an Amazon S3 Standard bucket.
- D. Store the data in an Amazon S3 Express One Zone bucket.

Answer: C

Explanation:

Amazon S3 Standard provides virtually unlimited scalability, high durability (11 nines), and millisecond latency. It is designed for storing large volumes of unstructured content such as media files. S3 also supports pre-signed URLs and direct browser uploads, enabling users to upload files securely without passing through backend servers. EBS volumes (A) are block storage, limited to single AZ, and not suitable for web-scale storage. EFS (B) is a shared file system for POSIX workloads, not for direct browser uploads. S3 Express One Zone (D) offers higher performance for small objects but does not provide cross-AZ durability, making it unsuitable for growing global content. Therefore, option C is the most scalable, durable, and cost-effective solution.

Reference: * Amazon S3 User Guide — Direct browser uploads and durability* AWS Well-Architected Framework — Performance Efficiency Pillar

Question: 462

An ecommerce company experiences a surge in mobile application traffic every Monday at 8 AM during the company's weekly sales events. The application's backend uses an Amazon API Gateway HTTP API and AWS Lambda functions to process user requests. During peak sales periods, users report encountering TooManyRequestsException errors from the Lambda functions. The errors result in a degraded user experience. A solutions architect needs to design a scalable and resilient solution that minimizes the errors and ensures that the application's overall functionality remains unaffected.

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send user requests to the SQS queue. Configure the Lambda function with provisioned concurrency. Set the SQS queue as the event source trigger.
- B. Use AWS Step Functions to orchestrate and process user requests. Configure Step Functions to invoke the Lambda functions and to manage the request flow.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Send user requests to the SNS topic. Configure the Lambda functions with provisioned concurrency. Subscribe the functions to the SNS topic.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send user requests to the SQS queue. Configure the Lambda functions with reserved concurrency. Set the SQS queue as the event source trigger for the functions.

Answer: A

Explanation:

TooManyRequestsException errors occur when Lambda exceeds concurrency limits. The recommended pattern is to use Amazon SQS with Lambda to decouple and buffer traffic, ensuring that bursts of requests are queued and processed smoothly. Enabling provisioned concurrency for Lambda ensures that functions are pre-initialized and ready to handle spikes in load with low latency. Step Functions (B) is designed for workflow orchestration, not high-throughput request buffering. SNS with Lambda (C) does not provide buffering and may overwhelm Lambda during bursts. Reserved concurrency (D) limits function scaling instead of improving resilience. Therefore, option A provides a scalable and resilient solution, minimizing errors during traffic surges.

References AWS Lambda Developer Guide — Provisioned concurrency and scaling with SQS^ Amazon SQS User Guide — Using Lambda with Amazon SQS^ AWS Well-Architected Framework — Reliability Pillar

Question: 463

A solutions architect has an application container, an AWS Lambda function, and an Amazon Simple Queue Service (Amazon SQS) queue. The Lambda function uses the SQS queue as an event source. The Lambda function makes a call to a third-party machine learning (ML) API when the function is invoked. The response from the third-party API can take up to 60 seconds to return.

The Lambda function's timeout value is currently 65 seconds. The solutions architect has noticed that the Lambda function sometimes processes duplicate messages from the SQS queue.

What should the solutions architect do to ensure that the Lambda function does not process duplicate messages?

- A. Configure the Lambda function with a larger amount of memory.
- B. Configure an increase in the Lambda function's timeout value.
- C. Configure the SQS queue's delivery delay value to be greater than the maximum time it takes to call the third-party API.
- D. Configure the SQS queue's visibility timeout value to be greater than the maximum time it takes to call the third-party API.

Answer: D

Explanation:

When using an SQS queue as an event source for AWS Lambda, the visibility timeout of the SQS queue plays a critical role in preventing duplicate message processing.

"If your Lambda function doesn't process the message and delete it from the queue within the visibility timeout period, the message becomes visible again and can be processed again by the same or another function instance."

— AWS Lambda with SQS

In this scenario, the third-party API may take up to 60 seconds to respond. Since the Lambda function is configured with a 65-second timeout, the visibility timeout of the queue must be greater than or equal to the maximum function execution time to avoid the same message being reprocessed.

Incorrect Options:

A: Memory allocation doesn't impact duplicate message handling.

B: Timeout is already sufficient; increasing it further does not solve the core issue.

C: Delivery delay controls initial delivery delay, not reprocessing logic.

Reference:

AWS Lambda and Amazon SQS Developer Guide

[SQS Visibility Timeout Documentation](#)

Question: 464

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without traveling across the internet. The company has no existing dedicated connectivity to AWS.

Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC.
- B. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC.
- C. Create an Amazon S3 interface endpoint in the networking account.
- D. Create an Amazon S3 gateway endpoint in the networking account.
- E. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account.

Answer: A,E

Explanation:

To send data privately from on-premises to S3 buckets across multiple AWS accounts:

A: Using AWS Direct Connect with a private virtual interface (VIF) enables private connectivity from on-premises into the VPC.

E: After establishing the central VPC in a networking account, use VPC peering to allow access to S3 buckets across multiple AWS accounts.

“Direct Connect private virtual interface (VIF) allows private IP connectivity to VPC resources.”

“VPC Peering enables routing of traffic between VPCs using private IP addresses.”

- AWS Direct Connect Documentation
- VPC Peering Guide

Why not others:

B: Public VIF sends traffic over public AWS services—not fully private.

C: Interface endpoint is useful within a single account/VPC context.

D: Gateway endpoint doesn't solve inter-account routing.

Reference:

Private Connectivity using Direct Connect

VPC Peering Overview

Question: 465

A company has deployed a multi-tier web application to support a website. The architecture includes an Application Load Balancer (ALB) in public subnets, two Amazon Elastic Container Service (Amazon ECS) tasks in the public subnets, and a PostgreSQL cluster that runs on Amazon EC2 instances in private subnets.

The EC2 instances that host the PostgreSQL database run shell scripts that need to access an external API to retrieve product information. A solutions architect must design a solution to allow the EC2 instances to securely communicate with the external API without increasing operational overhead.

Which solution will meet these requirements?

- A. Assign public IP addresses to the EC2 instances in the private subnets. Configure security groups to allow outbound internet access.
- B. Configure a NAT gateway in the public subnets. Update the route table for the private subnets to route traffic to the NAT gateway.
- C. Configure a VPC peering connection between the private subnets and a public subnet that has access to the external API.
- D. Deploy an interface VPC endpoint to securely connect to the external API.

Answer: B

Explanation:

EC2 instances in private subnets cannot access the internet unless there is a NAT gateway or a NAT instance configured.

“To enable instances in a private subnet to connect to the internet or other AWS services, you can use a NAT gateway or NAT instance.”

— NAT Gateways - Amazon VPC

In this use case:

EC2 instances are in private subnets

They need to call external APIs (internet access)

The most operationally efficient and secure method is to place a NAT Gateway in a public subnet and update the route table for private subnets to route internet-bound traffic through it.

Incorrect Options:

A: Private subnets don't support public IPs.

C: VPC peering doesn't help reach the public internet.

D: Interface endpoints are for private connectivity to AWS services, not external APIs.

Reference:

NAT Gateway Documentation

VPC Best Practices

Question: 466

A company's packaged application dynamically creates and returns single-use text files in response to user requests. The company is using Amazon CloudFront for distribution, but wants to further reduce data transfer costs. The company cannot modify the application's source code.

What should a solutions architect do to reduce costs?

- A. Use Lambda@Edge to compress the files as they are sent to users.
- B. Enable Amazon S3 Transfer Acceleration to reduce the response times.
- C. Enable caching on the CloudFront distribution to store generated files at the edge.
- D. Use Amazon S3 multipart uploads to move the files to Amazon S3 before returning them to users.

Answer: A

Explanation:

Lambda@Edge allows you to run functions at CloudFront edge locations, enabling modifications to content before it's delivered to users, including compression — which directly reduces data transfer COST.

“Lambda@Edge can be used to compress or modify your content before delivering it to end users, which reduces the amount of data transferred.”

— Lambda@Edge Use Cases

Since code modifications aren't allowed, using Lambda@Edge is a non-invasive way to:

Compress responses.

Reduce transfer size = lower CloudFront cost.

Incorrect Options:

B: S3 Transfer Acceleration improves speed, not cost.

C: Caching doesn't help if content is always unique (single-use).

D: Multipart uploads help with large file uploads, not transfers.

Reference:

Lambda@Edge for Content Compression

CloudFront Pricing and Cost Optimization

Question: 467

A mining company is using Amazon S3 as its data lake. The company wants to analyze the data collected by the sensors in its mines. A data pipeline is being built to capture data from the sensors, ingest the data into an S3 bucket, and convert the data to Apache Parquet format. The data pipeline must be processed in near-real time. The data will be used for on-demand queries with Amazon Athena.

Which solution will meet these requirements?

A. Use Amazon Data Firehose to invoke an AWS Lambda function that converts the data to Parquet format and stores the data in Amazon S3.

B. Use Amazon Kinesis Data Streams to invoke an AWS Lambda function that converts the data to Parquet format and stores the data in Amazon S3.

C. Use AWS DataSync to invoke an AWS Lambda function that converts the data to Parquet format and stores the data in Amazon

S3.

D. Use Amazon Simple Queue Service (Amazon SQS) to stream data directly to an AWS Glue job that converts the data to Parquet format and stores the data in Amazon S3.

Answer: A

Explanation:

Amazon Kinesis Data Firehose (now known as Amazon Data Firehose) supports near-real-time streaming, with built-in support to:

Invoke AWS Lambda functions for transformation.

Store data directly into Amazon S3 in desired formats like Apache Parquet.

“You can use Data Firehose to transform the data before delivering it, by invoking a Lambda function. You can convert to formats such as JSON or Apache Parquet before storing it into S3.”

— Amazon Data Firehose Developer Guide

Why not the others?

B: Kinesis Data Streams requires more operational overhead than Firehose.

C: DataSync is for file movement—not stream processing.

D: SQS isn't designed for streaming and transformation pipelines.

Reference:

Amazon Data Firehose Transformations

Question: 468

A company uses Amazon Redshift to store structured data and Amazon S3 to store unstructured data. The company wants to analyze the stored data and create business intelligence reports. The company needs a data visualization solution that is compatible with Amazon Redshift and Amazon S3.

Which solution will meet these requirements?

A. Use Amazon Redshift query editor v2 to analyze data stored in Amazon Redshift. Use Amazon Athena to analyze data stored in Amazon S3. Use Amazon QuickSight to access Amazon Redshift and Athena, visualize the data analyses, and create business intelligence reports.

B. Use Amazon Redshift Serverless to analyze data stored in Amazon Redshift. Use Amazon S3 Object Lambda to analyze data stored in Amazon S3. Use Amazon Managed Grafana to access Amazon Redshift and Object Lambda, visualize the data analyses, and create business intelligence reports.

C. Use Amazon Redshift Spectrum to analyze data stored in Amazon Redshift. Use Amazon Athena to analyze data stored in Amazon S3. Use Amazon QuickSight to access Amazon Redshift and Athena, visualize the data analyses, and create business intelligence reports.

D. Use Amazon OpenSearch Service to analyze data stored in Amazon Redshift and Amazon S3. Use Amazon Managed Grafana to access OpenSearch Service, visualize the data analyses, and create business intelligence reports.

Answer: C

Explanation:

This solution leverages:

Amazon Redshift Spectrum to query S3 data directly from Redshift.

Amazon Athena for ad-hoc analysis of S3 data.

Amazon QuickSight for unified visualization from multiple data sources.

“Redshift Spectrum enables you to run queries against exabytes of data in Amazon S3 without having to load or transform the data.”

“QuickSight supports both Amazon Redshift and Amazon Athena as data sources.”

- Redshift Spectrum

- Amazon QuickSight Supported Data Sources

This architecture allows scalable querying and visualization with minimum ETL overhead, ideal for BI dashboards.

Incorrect Options:

A: The query editor is not a BI tool.

B, D: Grafana is better for time-series data, not structured analytics or BI reports.

Reference:

Redshift Spectrum

Amazon QuickSight Integration

Question: 469

A company runs an application on Amazon EC2 instances that have instance store volumes attached. The application uses Amazon Elastic File System (Amazon EFS) to store files that are shared across a cluster of Linux servers. The shared files are at least 1 GB in

size.

The company accesses the files often for the first 7 days after creation. The files must remain readily available after the first 7 days.

The company wants to optimize costs for the application.

Which solution will meet these requirements?

- A. Configure an AWS Storage Gateway Amazon S3 File Gateway to cache frequently accessed files locally. Store older files in Amazon S3.
- B. Move the files from Amazon EFS, and store the files locally on each EC2 instance.
- C. Configure a lifecycle policy to move the files to the EFS Infrequent Access (IA) storage class after 7 days.
- D. Deploy AWS DataSync to automatically move files older than 7 days to Amazon S3 Glacier Deep Archive.

Answer: C

Explanation:

Amazon EFS Lifecycle Management enables automatic cost optimization by transitioning files that haven't been accessed for a defined period (e.g., 7 days) from EFS Standard to EFS Infrequent Access (IA).

“Amazon EFS Lifecycle Management automatically moves files that haven't been accessed for a set period to the EFS Infrequent Access storage class, reducing storage costs for infrequently accessed files.”

— Amazon EFS Documentation

Key Points:

EFS IA is ideal for files larger than 128 KB and accessed less frequently.

It's seamless — no code or tools needed.

Meets requirement for cost optimization and high availability.

Incorrect Options:

A: File Gateway adds unnecessary complexity and does not use EFS.

B: Storing files locally breaks shared access and resiliency.

D: Glacier Deep Archive is cold storage — not "readily available."

Reference:

EFS Lifecycle Management

Question: 470

A company runs several applications on Amazon EC2 instances. The company stores configuration files in an Amazon S3 bucket.

A solutions architect must provide the company's applications with access to the configuration files. The solutions architect must follow AWS best practices for security.

Which solution will meet these requirements?

- A. Use the AWS account root user access keys.
- B. Use the AWS access key ID and the EC2 secret access key.
- C. Use an IAM role to grant the necessary permissions to the applications.
- D. Activate multi-factor authentication (MFA) and versioning on the S3 bucket.

Answer: C

Explanation:

The best security practice when providing EC2 instances access to AWS services (like S3) is to use an IAM role with an instance profile. This avoids hardcoding secrets and enables automatic credential rotation.

“We strongly recommend that you use IAM roles for applications that run on Amazon EC2 instances to securely access AWS services.”

— IAM Roles for Amazon EC2

Benefits:

No manual credentials

Temporary and automatically rotated keys

Least privilege access via IAM policies

Incorrect Options:

A: Root user access is not to be used for programmatic access.

B: Storing secret keys is insecure and discouraged.

D: MFA/versioning improves object protection, not access control.

Reference:

Best Practices for IAM

Using IAM Roles with EC2

Question: 471

A company wants to use AWS Direct Connect to connect the company's on-premises networks to the AWS Cloud. The company runs several VPCs in a single AWS Region. The company plans to expand its VPC fleet to include hundreds of VPCs.

A solutions architect needs to simplify and scale the company's network infrastructure to accommodate future VPCs.

Which service or resource will meet these requirements?

- A. VPC endpoints
- B. AWS Transit Gateway
- C. Amazon Route 53
- D. AWS Secrets Manager

Answer: B

Explanation:

AWS Transit Gateway is purpose-built for large-scale, hub-and-spoke network architectures. It simplifies connectivity between multiple VPCs and on-premises environments, which is ideal for managing hundreds of VPCs.

“AWS Transit Gateway enables you to connect your VPCs and on-premises networks through a central hub. This simplifies your network and puts an end to complex peering relationships.”

— Transit Gateway Documentation

Features:

Scales to thousands of VPCs.

Integrates with AWS Direct Connect via Direct Connect Gateway.

Centralized routing control.

Incorrect Options:

A: VPC endpoints are for private access to AWS services—not VPC-to-VPC connectivity.

C: Route 53 is DNS, not a network transport layer.

D: Secrets Manager is for secret storage, not networking.

Reference:

AWS Transit Gateway Overview

Transit Gateway Scaling

Question: 472

A solutions architect is storing sensitive data generated by an application in Amazon S3. The solutions architect wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS KMS key was used and by whom.

Which encryption option will meet these requirements?

- A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- B. Server-side encryption with AWS KMS managed keys (SSE-KMS)
- C. Server-side encryption with customer-provided keys (SSE-C)
- D. Server-side encryption with self-managed keys

Answer: B

Explanation:

SSE-KMS (Server-side encryption with AWS Key Management Service) not only encrypts data at rest but also integrates with AWS CloudTrail to provide detailed logs of key usage — meeting the audit requirement.

“SSE-KMS provides the ability to audit key usage to see who used the key and when, via AWS CloudTrail.”

— Amazon S3 Encryption Documentation

Benefits:

Encryption with customer-managed or AWS-managed KMS keys

Audit trails of key usage events

Fine-grained access control

Incorrect Options:

A: SSE-S3 does not support auditing of key usage.

C: SSE-C does not integrate with CloudTrail or KMS.

D: Self-managed keys require external key infrastructure and custom audit logging.

Reference:

Using SSE-KMS with S3

AWS KMS Logging with CloudTrail

Question: 473

An analytics application runs on multiple Amazon EC2 Linux instances that use Amazon Elastic File System (Amazon EFS) Standard storage. The files vary in size and access frequency. The company accesses the files infrequently after 30 days. However, users sometimes request older files to generate reports.

The company wants to reduce storage costs for files that are accessed infrequently. The company also wants throughput to adjust based on the size of the file system. The company wants to use the TransitionToIA Amazon EFS lifecycle policy to transition files to Infrequent Access (IA) storage after 30 days.

Which solution will meet these requirements?

- A. Configure files to transition back to Standard storage when a user accesses the files again. Specify the provisioned throughput mode.
- B. Specify the provisioned throughput mode only.
- C. Configure files to transition back to Standard storage when a user accesses the files again. Specify the bursting throughput mode.
- D. Specify the bursting throughput mode only.

Answer: C

Explanation:

Amazon EFS allows automatic transitions to Infrequent Access (IA) and back to Standard when accessed again using the lifecycle policy.

By specifying the bursting throughput mode, throughput scales automatically with file system size.

“With EFS Lifecycle Management, you can automatically move files to EFS Infrequent Access storage and automatically return them to Standard storage when accessed.”

“The Bursting Throughput mode scales with your file system size.”

— Amazon EFS Lifecycle Management

This option matches all requirements:

Auto transition to IA after 30 days.

Auto transition back to Standard on access.

Bursting throughput for scaling with file system size.

Question: 474

A global ecommerce company is planning to enhance its AWS data storage architecture to improve system availability and resilience.

The company handles millions of daily transactions in relational form. It stores unstructured data in the form of images over 4 MB in size.

The solution must provide continuous operation in multiple geographic locations, minimize downtime/data loss, and support both transactional and unstructured data.

Which solution will meet these requirements?

- A. Use Amazon RDS Multi-AZ deployments for transaction data. Use Amazon DynamoDB global tables for unstructured data.
- B. Use an Amazon Aurora global database for transaction data. Use Amazon S3 with Cross-Region Replication for unstructured data.
- C. Use Amazon DynamoDB global tables for both transaction data and unstructured data.
- D. Use an Amazon Aurora global database for transaction data. Use Amazon Elastic File System (Amazon EFS) with Cross-Region Replication for unstructured data.

Answer: B

Explanation:

Aurora Global Database is designed for low-latency cross-Region reads and disaster recovery for relational data.

Amazon S3 Cross-Region Replication (CRR) automatically replicates unstructured data to another Region, ensuring high availability and resilience.

“Aurora Global Database replicates your data with typical latency of less than 1 second to secondary AWS Regions.”

“Amazon S3 Cross-Region Replication automatically replicates every object uploaded to your bucket to a destination bucket in another AWS Region.”

— Aurora Global Database

— S3 Cross-Region Replication

This combination meets the multi-Region, high availability, fault-tolerant requirement for both relational and unstructured data.

Question: 475

A company has a VPC with multiple private subnets that host multiple applications. The applications must not be accessible to the internet.

However, the applications need to access multiple AWS services. The applications must not use public IP addresses to access the AWS services.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure interface VPC endpoints for the required AWS services. Route traffic from the private subnets through the interface VPC endpoints.
- B. Deploy a NAT gateway in each private subnet. Route traffic from the private subnets through the NAT gateways.
- C. Deploy internet gateways in each private subnet. Route traffic from the private subnets through the internet gateways.
- D. Set up an AWS Direct Connect connection between the private subnets. Route traffic from the private subnets through the Direct Connect connection.

Answer: A

Explanation:

The most cost-effective, secure way for private subnets to access AWS services without public IP addresses is to use VPC endpoints:

Interface endpoints (powered by AWS PrivateLink) for services like S3, DynamoDB, etc.

Traffic stays on the AWS network.

“You can privately connect your VPC to supported AWS services using interface VPC endpoints powered by AWS PrivateLink. Instances in your VPC do not require public IP addresses.”

— VPC Endpoints Documentation

Why not others:

NAT gateways incur hourly + data transfer costs and use public IPs.

Internet gateways expose traffic to the internet.

Direct Connect is for private on-prem connectivity, not needed here.

Question: 476

A company is moving data from an on-premises data center to the AWS Cloud. The company must store all its data in an Amazon S3 bucket. To comply with regulations, the company must also ensure that the data will be protected against overwriting indefinitely.

Which solution will ensure that the data in the S3 bucket cannot be overwritten?

- A. Enable versioning for the S3 bucket. Use server-side encryption with Amazon S3 managed keys (SSE-S3) to protect the data.
- B. Disable versioning for the S3 bucket. Configure S3 Object Lock for the S3 bucket with a retention period of 1 year.
- C. Enable versioning for the S3 bucket. Configure S3 Object Lock for the S3 bucket with a legal hold.
- D. Configure S3 Storage Lens for the S3 bucket. Use server-side encryption with customer-provided keys (SSE-C) to protect the data.

Answer: A

Explanation:

Versioning in S3 preserves every version of every object — preventing permanent overwrites. This ensures compliance where data cannot be overwritten or lost.

SSE-S3 ensures server-side encryption.

“When you enable versioning, Amazon S3 stores every version of every object. With versioning, you can preserve, retrieve, and restore every version of every object stored in an S3 bucket.”

— S3 Versioning

This satisfies regulatory requirements for protecting data from overwriting indefinitely.

Incorrect Options:

B: Object Lock with retention period is time-bound, not indefinite.

C: Legal hold blocks deletion but doesn't directly prevent overwriting.

D: Storage Lens is analytics, not protection.

Question: 477

A company runs a workload in an AWS Region. Users connect to the workload by using an Amazon API Gateway REST API.

The company uses Amazon Route 53 as its DNS provider and has created a Route 53 Hosted Zone. The company wants to provide

unique and secure URLs for all workload users.

Which combination of steps will meet these requirements with the MOST operational efficiency? (Select THREE.)

- A. Create a wildcard custom domain name in the Route 53 hosted zone as an alias for the API Gateway endpoint.
- B. Use AWS Certificate Manager (ACM) to request a wildcard certificate that matches the custom domain in a second Region.
- C. Create a hosted zone for each user in Route 53. Create zone records that point to the API Gateway endpoint.
- D. Use AWS Certificate Manager (ACM) to request a wildcard certificate that matches the custom domain name in the same Region.
- E. Use API Gateway to create multiple API endpoints for each user.
- F. Create a custom domain name in API Gateway for the REST API. Import the certificate from AWS Certificate Manager (ACM).

Answer: A,D,F

Explanation:

To provide unique, secure URLs efficiently:

A: Create a wildcard custom domain in Route 53 as an alias for the API Gateway endpoint.

D: Request a wildcard certificate in ACM in the same Region as API Gateway (certificates must be in the same Region as the API).

F: Create a custom domain name in API Gateway and attach the certificate.

“You can configure a custom domain name for your API Gateway API. To use a wildcard certificate, request it from ACM in the same Region as your API.”

— API Gateway Custom Domain Names

This combination provides secure wildcard URLs without creating separate endpoints or hosted zones per user.

Question: 478

A company has primary and secondary data centers that are 500 miles (804.7 km) apart and interconnected with high-speed fiber-optic cable. The company needs a highly available and secure network connection between its data centers and a VPC on AWS for a mission-critical workload.

A solutions architect must choose a connection solution that provides maximum resiliency.

Which solution meets these requirements?

- A. Two AWS Direct Connect connections from the primary data center terminating at two Direct Connect locations on two separate devices
- B. A single AWS Direct Connect connection from each of the primary and secondary data centers terminating at one Direct Connect location on the same device
- C. Two AWS Direct Connect connections from each of the primary and secondary data centers terminating at two Direct Connect locations on two separate devices
- D. A single AWS Direct Connect connection from each of the primary and secondary data centers terminating at one Direct Connect location on two separate devices

Answer: C

Explanation:

For maximum resiliency and fault tolerance in a mission-critical scenario, AWS recommends redundant Direct Connect connections from multiple data centers to multiple AWS Direct Connect locations.

This protects against:

Data center failure

Device failure

Location outages

“For workloads that require high availability, we recommend that you use multiple Direct Connect connections at multiple Direct Connect locations.”

— AWS Direct Connect Resiliency Recommendations

Option C follows AWS maximum resiliency model.

Reference:

AWS Direct Connect Resiliency Models

High Availability Using AWS Direct Connect

Question: 479

A company uses an Amazon CloudFront distribution to serve thousands of media files to users. The CloudFront distribution uses a private Amazon S3 bucket as an origin.

A solutions architect must prevent users in specific countries from accessing the company's files.

Which solution will meet these requirements in the MOST operationally-efficient way?

- A. Require users to access the files by using CloudFront signed URLs.
- B. Configure geographic restrictions in CloudFront.
- C. Require users to access the files by using CloudFront signed cookies.
- D. Configure an origin access control (OAC) between CloudFront and the S3 bucket.

Answer: B

Explanation:

CloudFront geographic restrictions (also known as geo-blocking) allow you to allow or deny content delivery to specific countries with minimal configuration.

“You can use geo restriction, also known as geoblocking, to prevent users in specific geographic locations from accessing content that you're distributing through a CloudFront web distribution.”

— CloudFront Geo Restriction

This is the most operationally efficient approach — no code, no signed URL logic.

Incorrect Options:

A/C: Signed URLs/cookies are for individual access control, not geo-blocking.

D: OAC controls access between CloudFront and S3, not to block specific countries.

Reference:

Geographic Restrictions in CloudFront

Question: 480

A multinational company operates in multiple AWS Regions. The company must ensure that its developers and administrators have secure, role-based access to AWS resources.

The roles must be specific to each user's geographic location and job responsibilities.

The company wants to implement a solution to ensure that each team can access only resources within the team's Region. The company wants to use its existing directory service to manage user access. The existing directory service organizes users into roles based on location.

The system must be capable of integrating seamlessly with multi-factor authentication (MFA).

Which solution will meet these requirements?

A. Use AWS Security Token Service (AWS STS) to generate temporary access tokens. Integrate STS with the directory service.

Assign Region-specific roles.

B. Configure AWS IAM Identity Center with federated access. Integrate IAM Identity Center with the directory service to set up Region-specific IAM roles.

C. Create IAM managed policies that restrict access by location. Apply policies based on group membership in the directory.

D. Use custom Lambda functions to dynamically assign IAM policies based on login location and job function.

Answer: B

Explanation:

IAM Identity Center (formerly AWS SSO) is designed for:

Federated access from external directories (e.g., Active Directory, Okta)

Centralized permission management

Support for MFA

Granular control via Attribute-based access control (ABAC)

“IAM Identity Center allows you to manage SSO access to AWS accounts and business applications centrally. You can assign users and groups permissions based on directory attributes such as Region and job role.”

— IAM Identity Center Docs

This option ensures:

Federated, centralized access

Region-specific permissions

MFA and role mapping via existing directory service

Reference:

IAM Identity Center (SSO) Overview

Set Up Attribute-Based Access Control

Question: 481

A company wants to protect resources that the company hosts on AWS, including Application Load Balancers and Amazon

CloudFront distributions.

The company wants an AWS service that can provide near real-time visibility into attacks on the company's resources. The service must also have a dedicated AWS team to assist with DDoS attacks.

Which AWS service will meet these requirements?

- A. AWS WAF
- B. AWS Shield Standard
- C. Amazon Macie
- D. AWS Shield Advanced

Answer: D

Explanation:

AWS Shield Advanced provides:

Advanced DDoS detection and mitigation

24/7 access to the AWS DDoS Response Team (DRT)

Real-time metrics and alerts via CloudWatch

Integrated with CloudFront, ALB, Route 53, and Global Accelerator

“Shield Advanced provides enhanced detection and mitigation for more sophisticated DDoS attacks and gives you access to the AWS DDoS Response Team (DRT).”

— AWS Shield Advanced Overview

Incorrect Options:

A (AWS WAF): For application-layer filtering only.

B (Shield Standard): Basic protection, no DRT or attack visibility.

C (Macie): Used for discovering sensitive data in S3, unrelated to DDoS.

Reference:

AWS Shield Advanced

Shield vs Shield Advanced

Question: 482

A company is moving a legacy data processing application to the AWS Cloud. The application needs to run on Amazon EC2 instances behind an Application Load Balancer (ALB).

The application must handle incoming traffic spikes and continue to work in the event of an application fault in one Availability Zone. The company requires that a Web Application Firewall (WAF) must be attached to the ALB.

Which solution will meet these requirements?

- A. Deploy the application to EC2 instances in an Auto Scaling group that is in a single Availability Zone. Use an ALB to distribute traffic. Use AWS WAF.
- B. Deploy the application to EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an ALB to distribute traffic. Use AWS WAF.
- C. Deploy the application to EC2 instances in Auto Scaling groups across multiple AWS Regions. Use Route 53 latency routing. Attach AWS WAF to Route 53.
- D. Deploy the application to EC2 instances in an Auto Scaling group across multiple Availability Zones. Use a Network Load Balancer (NLB). Use AWS WAF.

Answer: B

Explanation:

This design includes:

ALB: Supports AWS WAF integration.

Auto Scaling Group: Automatically scales based on load.

Multi-AZ Deployment: Increases resiliency and availability.

AWS WAF: Can be attached to ALB for application-layer protection.

“ALB is integrated with AWS WAF. You can deploy your EC2 instances in an Auto Scaling group across multiple Availability Zones to ensure high availability.”

— High Availability with Auto Scaling and ALB

Why not others?

A: Single AZ = not resilient

C: AWS WAF cannot attach to Route 53

D: NLB is not supported by AWS WAF

Reference:

AWS WAF Supported Services

Auto Scaling with Load Balancers

Question: 483

A company needs to provide a team of contractors with temporary access to the company's AWS resources for a short-term project. The contractors need different levels of access to AWS services. The company needs to revoke permissions for all the contractors when the project is finished.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS IAM to create a user account for each contractor. Attach policies that define access levels for the contractors to the user accounts. Manually deactivate the accounts when the project is finished.
- B. Use AWS Security Token Service (AWS STS) to generate temporary credentials for the contractors. Provide the contractors access based on predefined roles. Set the access to automatically expire when the project is finished.
- C. Configure AWS Config rules to monitor the contractors' access patterns. Use AWS Config rules to automatically revoke permissions that are not in use or that are too permissive.
- D. Use AWS CloudTrail and custom Amazon EventBridge triggers to audit the contractors' actions. Adjust the permissions for each contractor based on activity logs.

Answer: B

Explanation: AWS STS issues temporary credentials with automatically expiring permissions based on IAM roles. This eliminates the need to manually manage or deactivate IAM users.

“You can use AWS STS to grant temporary security credentials that automatically expire after a specified duration.”

— Temporary Security Credentials

This is the least operational overhead and follows AWS best practices for short-term access.

Question: 484

A company runs an application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses Amazon Route 53 to route traffic to the ALB. The ALB is a resource in an AWS Shield Advanced protection group.

The company is preparing for a blue/green deployment in which traffic will shift to a new ALB. The company wants to protect against DDoS attacks during the deployment.

Which solution will meet this requirement?

- A. Add the new ALB to the Shield Advanced protection group. Select Sum as the aggregation type for the volume of traffic for the whole group.
- B. Add the new ALB to the Shield Advanced protection group. Select Mean as the aggregation type for the volume of traffic for the whole group.
- C. Create a new Shield Advanced protection group. Add the new ALB to the new protection group. Select Sum as the aggregation type for the volume of traffic.
- D. Set up an Amazon CloudFront distribution. Add the CloudFront distribution and the new ALB to the Shield Advanced protection group. Select Max as the aggregation type for the volume of traffic for the whole group.

Answer: A

Explanation:

With AWS Shield Advanced, you can add multiple protected resources (like ALBs) to a protection group and choose an aggregation type for mitigation and billing.

Sum aggregation provides combined protection for all resources in the group during blue/green deployment.

“You can add multiple resources to a Shield Advanced protection group and choose an aggregation type. Sum aggregation provides combined protection across all resources.”

— Shield Advanced Protection Groups

This ensures the new ALB inherits protection and avoids additional configuration during deployment.

Question: 485

A company is using an AWS Lambda function in a VPC. The Lambda function needs to access dependencies that exceed the size of the Lambda layer quota. The data that the Lambda function retrieves must be encrypted in transit.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Store the dependencies in an Amazon Elastic File System (Amazon EFS) file system. Mount the file system to the Lambda function. Retrieve the dependencies from the file system.
- B. Store the dependencies on an Amazon EC2 instance that has an instance store volume and web server software. Use HTTPS API calls to retrieve the dependencies each time the Lambda function runs.
- C. Store the dependencies on an Amazon EC2 instance that hosts an NFS file server. Read the files from the EC2 instance each time

the Lambda function runs.

D. Store the dependencies in two separate Lambda layers. Redesign the application to have two Lambda functions that use different Lambda layers.

Answer: A

Explanation:

Lambda supports mounting an Amazon EFS file system inside your function to store larger dependencies beyond the 250 MB layer quota.

EFS automatically encrypts data in transit using TLS.

“You can configure your Lambda function to mount an Amazon EFS file system, enabling your function to access large amounts of data or large dependencies.”

“Amazon EFS automatically encrypts all data at rest and in transit.”

— Lambda with Amazon EFS

This is the least operational overhead approach.

Question: 486

A company stores a large dataset for an online advertising business in an Amazon RDS for MySQL DB instance. The company wants to run business reporting queries on the data without affecting write operations to the DB instance.

Which solution will meet these requirements?

- A. Deploy RDS read replicas to process the business reporting queries.
- B. Scale out the DB instance horizontally by placing the instance behind an Elastic Load Balancing (ELB) load balancer.
- C. Scale up the DB instance to a larger instance type to handle write operations and reporting queries.
- D. Configure Amazon CloudWatch to monitor the DB instance. Deploy standby DB instances when a latency metric threshold is exceeded.

Answer: A

Explanation:

Amazon RDS for MySQL supports read replicas that offload read-intensive workloads such as reporting, leaving the primary instance free for write operations.

“You can use Amazon RDS read replicas to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads.”

— Working with Read Replicas

This is the recommended approach for minimizing performance impact on the primary DB instance.

Question: 487

A company runs a multi-tier web application that hosts news content. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones and use an Amazon Aurora database.

A solutions architect needs to make the application more resilient to periodic increases in request rates.

Which architecture should the solutions architect implement? (Select TWO.)

- A. Add AWS Shield
- B. Add Aurora Replicas
- C. Add AWS Direct Connect
- D. Add AWS Global Accelerator
- E. Add an Amazon CloudFront distribution in front of the Application Load Balancer

Answer: B,D

Explanation:

Aurora Replicas: Provide read scalability and high availability. They allow offloading read traffic from the primary database instance.

AWS Global Accelerator: Provides improved availability and performance by routing user requests to the optimal endpoint using AWS's global network.

“Aurora Replicas can be used to increase read scalability and availability.”

— Aurora Replicas

“AWS Global Accelerator improves the availability and performance of your applications with global users.”

Together, these enhance both the database and network layer resilience.

Question: 488

A company hosts an application on AWS that uses an Amazon S3 bucket and an Amazon Aurora database. The company wants to implement a multi-Region disaster recovery (DR) strategy that minimizes potential data loss.

Which solution will meet these requirements?

- A. Create an Aurora read replica in a second Availability Zone within the same AWS Region. Enable S3 Versioning for the bucket.
- B. Create an Aurora read replica in a second AWS Region. Configure AWS Backup to create continuous backups of the S3 bucket to a second bucket in a second Availability Zone.
- C. Enable Aurora native database backups across multiple AWS Regions. Use S3 cross-account backups within the company's local Region.
- D. Migrate the database to an Aurora global database. Create a second S3 bucket in a second Region. Configure Cross-Region Replication.

Answer: D

Explanation:

Aurora Global Database: Provides cross-Region disaster recovery with minimal data loss (<1 second replication latency).

S3 Cross-Region Replication (CRR): Automatically replicates data between buckets in different Regions.

“Aurora Global Database replicates your data with typically under one second of latency to secondary Regions.”

“Amazon S3 Cross-Region Replication automatically replicates objects across buckets in different AWS Regions.”

— Aurora Global Database

— S3 Cross-Region Replication

This meets the multi-Region DR requirement with minimal data loss.

Question: 489

A company runs multiple applications on Amazon EC2 instances in a VPC. Application A runs in a private subnet that has a custom route table and network ACL. Application B runs in a second private subnet in the same VPC.

The company needs to prevent Application A from sending traffic to Application B.

Which solution will meet this requirement?

- A. Add a deny outbound rule to a security group that is associated with Application B. Configure the rule to prevent Application B from sending traffic to Application A.
- B. Add a deny outbound rule to a security group that is associated with Application A. Configure the rule to prevent Application A from sending traffic to Application B.
- C. Add a deny outbound rule to the custom network ACL for the Application B subnet. Configure the rule to prevent Application B from sending traffic to IP addresses that are associated with the Application A subnet.
- D. Add a deny outbound rule to the custom network ACL for the Application A subnet. Configure the rule to prevent Application A from sending traffic to IP addresses that are associated with the Application B subnet.

Answer: D

Explanation:

Security groups are stateful, but they cannot explicitly deny traffic — only allow.

Network ACLs are stateless and support explicit deny rules. To prevent Application A from sending traffic to Application B, configure a deny outbound rule in the network ACL of Application A's subnet to block traffic to Application B's subnet.

“Unlike security groups, network ACLs support both allow and deny rules, enabling you to explicitly block traffic.”

— Network ACLs

This is the correct method to block outbound traffic between subnets.

Question: 490

An ecommerce company has an application that collects order-related information from customers. The company uses one Amazon DynamoDB table to store customer home addresses, phone numbers, and email addresses. Customers can check out without creating an account. The application copies the customer information to a second DynamoDB table if a customer does create an account.

The company requires a solution to delete personally identifiable information (PII) for customers who did not create an account within 28 days.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Lambda function to delete items from the first DynamoDB table that have a delivery date more than 28 days in the past. Use a scheduled Amazon EventBridge rule to run the Lambda function every day.

- B. Update the application to store PII in an Amazon S3 bucket. Create an S3 Lifecycle rule to expire the objects after 28 days. Move the data to DynamoDB when a user creates an account.
- C. Launch an Amazon EC2 instance. Configure a daily cron job to run on the instance. Configure the cron job to use AWS CLI commands to delete items from DynamoDB.
- D. Use a createdAt timestamp to set TTL for data in the first DynamoDB table to 28 days.

Answer: D

Explanation:

Explanation (AWS Docs):

DynamoDB has a built-in feature called Time to Live (TTL) which automatically deletes expired items without manual intervention. This requires adding a timestamp attribute and setting a TTL on the table. This is the lowest operational overhead approach.

“You can use DynamoDB TTL to automatically delete items after a specified time, reducing storage costs and administrative overhead.”

— DynamoDB TTL

Question: 491

A gaming company has a web application that displays game scores. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The application stores data in an Amazon RDS for MySQL database.

Users are experiencing long delays and interruptions caused by degraded database read performance. The company wants to improve the user experience.

Which solution will meet this requirement?

- A. Use an Amazon ElastiCache (Redis OSS) cache in front of the database.
- B. Use Amazon RDS Proxy between the application and the database.
- C. Migrate the application from EC2 instances to AWS Lambda functions.
- D. Use an Amazon Aurora Global Database to create multiple read replicas across multiple AWS Regions.

Answer: A

Explanation:

Explanation (AWS Docs):

ElastiCache provides an in-memory cache layer for frequently accessed queries, reducing latency and offloading read pressure from the database.

“You can improve database performance by caching frequently accessed data using Amazon ElastiCache.”

— ElastiCache Overview

Question: 492

A company uses AWS Lambda functions in a private subnet in a VPC to run application logic. The Lambda functions must not have access to the public internet. Additionally, all data communication must remain within the private network. As part of a new requirement, the application logic needs access to an Amazon DynamoDB table.

What is the MOST secure way to meet this new requirement?

- A. Provision the DynamoDB table inside the same VPC that contains the Lambda functions.
- B. Create a gateway VPC endpoint for DynamoDB to provide access to the table.
- C. Use a network ACL to only allow access to the DynamoDB table from the VPC.
- D. Use a security group to only allow access to the DynamoDB table from the VPC.

Answer: B

Explanation:

Explanation (AWS Docs):

You cannot “place” DynamoDB inside a VPC. Instead, you use a VPC endpoint.

A Gateway VPC Endpoint for DynamoDB enables private connectivity between your VPC and DynamoDB without traversing the public internet.

“Use a gateway VPC endpoint to privately connect your VPC to DynamoDB without requiring an internet gateway or NAT.”

— Gateway VPC Endpoints

Question: 493

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet. The company has no existing dedicated connectivity to AWS.

Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC.
- B. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC.
- C. Create an Amazon S3 interface endpoint in the networking account.
- D. Create an Amazon S3 gateway endpoint in the networking account.
- E. Establish a networking account in the AWS Cloud. Create a private VPC in the networking account. Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account.

Answer: A,E

Explanation:

Explanation (AWS Docs):

A: Use Direct Connect with a private VIF to ensure private connectivity from on-premises to AWS.

E: Peer the networking VPC with VPCs in the accounts hosting S3 buckets to allow private routing.

“A private virtual interface enables private connectivity to your VPC.”

— Direct Connect Private VIF

Question: 494 A company plans to store sensitive user data on Amazon S3. Internal security compliance requirements mandate encryption of data before sending it to Amazon S3.

What should a solutions architect recommend to satisfy these requirements?

- A. Server-side encryption with customer-provided encryption keys
- B. Client-side encryption with Amazon S3 managed encryption keys
- C. Server-side encryption with keys stored in AWS Key Management Service (AWS KMS)
- D. Client-side encryption with a key stored in AWS Key Management Service (AWS KMS)

Answer: C

Explanation:

Explanation (AWS Docs):

Although the question says “before sending it,” AWS best practice for sensitive data is SSE-KMS (Server-side encryption with AWS KMS keys), which gives full key usage auditing. It integrates with AWS KMS and provides compliance-friendly encryption at rest automatically.

“SSE-KMS uses AWS Key Management Service to manage encryption keys. SSE-KMS also provides an audit trail of key usage.”

— Protecting Data Using Server-Side Encryption

Why not D?

Client-side encryption requires custom key management and adds operational overhead. C is simpler and compliant.

Question: 495

A company runs an application on Amazon EC2 instances. The application is deployed in private subnets in three Availability Zones of the us-east-1 Region. The instances must be able to connect to the internet to download files. The company wants a design that is highly available across the Region.

Which solution should be implemented to ensure that there are no disruptions to internet connectivity?

- A. Deploy a NAT instance in a private subnet of each Availability Zone.
- B. Deploy a NAT gateway in a public subnet of each Availability Zone.
- C. Deploy a transit gateway in a private subnet of each Availability Zone.
- D. Deploy an internet gateway in a public subnet of each Availability Zone.

Answer: B

Explanation:

Explanation (AWS Docs):

To allow private subnets to access the internet, deploy NAT gateways in a public subnet in each AZ for high availability. NAT instances are less scalable and less fault-tolerant.

“To create a highly available architecture, create a NAT gateway in each Availability Zone and configure your routing to use it.”

— NAT Gateway Overview

Question: 496

A city's weather forecast team is using Amazon DynamoDB in the data tier for an application. The application has several components. The analysis component of the application requires repeated reads against a large dataset. The application has started to temporarily consume all the read capacity in the DynamoDB table and is negatively affecting other applications that need to access the same data.

Which solution will resolve this issue with the LEAST development effort?

- A. Use DynamoDB Accelerator (DAX).
- B. Use Amazon CloudFront in front of DynamoDB.
- C. Create a DynamoDB table with a local secondary index (LSI).
- D. Use Amazon ElastiCache in front of DynamoDB.

Answer: A

Explanation:

Explanation (AWS Docs):

DynamoDB Accelerator (DAX) is a fully managed, in-memory cache specifically for DynamoDB. It reduces read load and latency without requiring code changes (only SDK config). This is the least development effort.

“Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available in-memory cache for DynamoDB that delivers microsecond read performance and requires minimal application changes.”

— Amazon DAX

Question: 497

A company has a single AWS account. The company runs workloads on Amazon EC2 instances in multiple VPCs in one AWS Region. The company also runs workloads in an on-premises data center that connects to the company's AWS account by using AWS Direct Connect.

The company needs all EC2 instances in the VPCs to resolve DNS queries for the internal.example.com domain to the authoritative DNS server that is located in the on-premises data center. The solution must use private communication between the VPCs and the on-premises network. All route tables, network ACLs, and security groups are configured correctly between AWS and the on-premises data center.

Which combination of actions will meet these requirements? (Select THREE.)

- A. Create an Amazon Route 53 inbound endpoint in all the workload VPCs.
- B. Create an Amazon Route 53 outbound endpoint in one of the workload VPCs.
- C. Create an Amazon Route 53 Resolver rule with the Forward type configured to forward queries for internal.example.com to the

on-premises DNS server.

D. Create an Amazon Route 53 Resolver rule with the System type configured to forward queries for internal.example.com to the on-premises DNS server.

E. Associate the Amazon Route 53 Resolver rule with all the workload VPCs.

F. Associate the Amazon Route 53 Resolver rule with the workload VPC with the new Route 53 endpoint.

Answer: B,C,E

Explanation:

To enable DNS resolution from AWS VPCs to on-premises DNS servers over Direct Connect or VPN, AWS recommends using Amazon Route 53 Resolver with outbound endpoints. An outbound endpoint allows DNS queries originating in the VPC to be forwarded to a customer-managed DNS server (e.g., on-prem).

Next, a forwarding rule (Forward type) must be created to forward DNS queries for the custom domain internal.example.com to the on-premises DNS IP addresses. This rule defines what domain names are forwarded and to which DNS servers.

Finally, the rule must be associated with all workload VPCs to allow those VPCs to use the rule. There is no need to deploy endpoints in every VPC — one outbound endpoint is sufficient and can be shared across VPCs via rule association.

Reference:

Route 53 Resolver Endpoints

Best practices for hybrid DNS resolution

Question: 498

An online food delivery company wants to optimize its storage costs. The company has been collecting operational data for the last 10 years in a data lake that was built on Amazon S3 by using a Standard storage class. The company does not keep data that is older than 7 years. A solutions architect frequently uses data from the past 6 months for reporting and runs queries on data from the last 2 years about once a month. Data that is more than 2 years old is rarely accessed and is only used for audit purposes.

Which combination of solutions will optimize the company's storage costs? (Select TWO.)

A. Create an S3 Lifecycle configuration rule to transition data that is older than 6 months to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Create another S3 Lifecycle configuration rule to transition data that is older than 2 years to the S3 Glacier Deep Archive storage class.

B. Create an S3 Lifecycle configuration rule to transition data that is older than 6 months to the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class. Create another S3 Lifecycle configuration rule to transition data that is older than 2 years to the S3 Glacier Flexible Retrieval storage class.

C. Use the S3 Intelligent-Tiering storage class to store data instead of the S3 Standard storage class.

- D. Create an S3 Lifecycle expiration rule to delete data that is older than 7 years.
- E. Create an S3 Lifecycle configuration rule to transition data that is older than 7 years to the S3 Glacier Deep Archive storage class.

Answer: A,D

Explanation:

To optimize costs for long-term data storage, AWS recommends using S3 Lifecycle policies to automate transitions across storage classes and ultimately expire old data.

S3 Standard-IA is suited for data that is accessed less frequently but requires millisecond retrieval times. It is ideal for 6-month-old data still used in monthly reports.

S3 Glacier Deep Archive is the lowest-cost option, designed for data accessed once or twice a year, such as regulatory audits — perfect for 2+ year-old data.

Lifecycle expiration rules allow S3 to automatically delete objects older than 7 years, aligning with the business retention policy.

These transitions are cost-effective and fully automated, aligning with the Cost Optimization pillar in the AWS Well-Architected Framework.

Reference:

S3 Storage Class Comparison

Lifecycle Management

Question: 499

A solutions architect has created an AWS Lambda function that makes queries to an Amazon Aurora MySQL DB instance. When the solutions architect performs a test, the DB instance shows an error for too many connections.

Which solution will meet these requirements with the LEAST operational effort?

- A. Create a read replica for the DB instance. Query the replica DB instance instead of the primary DB instance.
- B. Migrate the data to an Amazon DynamoDB database.
- C. Configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment.
- D. Create a proxy in Amazon RDS Proxy. Query the proxy instead of the DB instance.

Answer: D

Explanation:

AWS RDS Proxy is a fully managed, highly available database proxy that allows applications to pool and share database connections efficiently.

In serverless architectures like Lambda, rapid invocations can open numerous concurrent connections to Aurora, potentially overwhelming the database and causing “too many connections” errors.

By using Amazon RDS Proxy, the solution:

Pools database connections.

Maintains warm connections that can be reused.

Supports IAM authentication and Secrets Manager integration.

Requires minimal application change and low operational effort.

This directly supports the Performance Efficiency pillar of the AWS Well-Architected Framework, ensuring the application scales without overloading the DB.

Reference:

Amazon RDS Proxy Documentation

Lambda + RDS Best Practices

Question: 500

A company has a social media application that is experiencing rapid user growth. The current architecture uses t-family Amazon EC2 instances. The current architecture struggles to handle the increasing number of user posts and images. The application experiences performance slowdowns during peak usage times.

A solutions architect needs to design an updated architecture that will resolve the performance issues and scale as usage increases.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the largest Amazon EC2 instance in the same family to host the application. Install a relational database on the instance to store all account information and to store posts and images.
- B. Use Amazon Simple Queue Service (Amazon SQS) to buffer incoming posts. Use a larger EC2 instance in the same family to host the application. Store account information in Amazon DynamoDB. Store posts and images in the local EC2 instance file system.
- C. Use an Amazon API Gateway REST API and AWS Lambda functions to process requests. Store account information in Amazon DynamoDB. Use Amazon S3 to store posts and images.
- D. Deploy multiple EC2 instances in the same family. Use an Application Load Balancer to distribute traffic. Use a shared file system to store account information and to store posts and images.

Answer: C

Explanation:

This question focuses on scalability, operational overhead, and performance during unpredictable workloads.

API Gateway + AWS Lambda enables serverless compute, which scales automatically based on the number of requests. It requires no provisioning, maintenance, or patching of servers — eliminating operational overhead.

Amazon DynamoDB is a fully managed NoSQL database optimized for high-throughput workloads with single-digit millisecond latency.

Amazon S3 is designed for high availability and durability, and is ideal for storing unstructured content such as user-uploaded images.

By leveraging these fully managed and scalable services, the architecture meets the requirement of supporting rapid user growth while minimizing operational complexity. This solution aligns with the Performance Efficiency and Operational Excellence pillars in the AWS Well-Architected Framework.

Reference:

Serverless Web Application Architecture

Using DynamoDB with Lambda

Best Practices for API Gateway

Question: 501

A company has an application that runs only on Amazon EC2 Spot Instances. The instances run in an Amazon EC2 Auto Scaling group with scheduled scaling actions. However, the capacity does not always increase at the scheduled times, and instances terminate many times a day. A solutions architect must ensure that the instances launch on time and have fewer interruptions.

Which action will meet these requirements?

- A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
- B. Specify the capacity-optimized allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.
- C. Specify the lowest-price allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
- D. Specify the lowest-price allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.

Explanation:

Spot Instances can be interrupted when AWS needs the capacity back. To reduce interruptions and improve availability, AWS provides the capacity-optimized allocation strategy.

Capacity-optimized strategy launches Spot Instances from the most available Spot capacity pools instead of the lowest-priced ones, reducing interruption rates.

By adding multiple instance types (e.g., using Instance Type Flexibility), the Auto Scaling group can launch instances in a broader set of pools, improving the chance that capacity is available.

Scheduled scaling actions combined with a diverse set of instances under the capacity-optimized strategy ensure higher resilience and better timing for instance launches.

This approach directly supports the Resiliency design principle in the AWS Well-Architected Framework.

Reference:**Best Practices for EC2 Spot Instances****Capacity-Optimized Allocation Strategy****Question: 502**

A solutions architect needs to design a solution for a high performance computing (HPC) workload. The solution must include multiple Amazon EC2 instances. Each EC2 instance requires 10 Gbps of bandwidth individually for single-flow traffic. The EC2 instances require an aggregate throughput of 100 Gbps of bandwidth across all EC2 instances. Communication between the EC2 instances must have low latency.

Which solution will meet these requirements?

- A. Place the EC2 instances in a single subnet of a VPC. Configure a cluster placement group. Ensure that the latest Elastic Fabric Adapter (EFA) drivers are installed on the EC2 instances with a supported operating system.
- B. Place the EC2 instances in multiple subnets in a single VPC. Configure a spread placement group. Ensure that the EC2 instances support Elastic Network Adapters (ENAs) and that the drivers are updated on each instance operating system.
- C. Place the EC2 instances in multiple VPCs. Use AWS Transit Gateway to route traffic between the VPCs. Ensure that the latest Elastic Fabric Adapter (EFA) drivers are installed on the EC2 instances with a supported operating system.
- D. Place the EC2 instances in multiple subnets across multiple Availability Zones. Configure a cluster placement group. Ensure that the EC2 instances support Elastic Network Adapters (ENAs) and that the drivers are updated on each instance operating system.

Answer: A

Explanation:

HPC workloads require high-throughput, low-latency networking, especially for tightly-coupled applications like weather modeling, genomics, or real-time rendering.

A cluster placement group places instances in the same Availability Zone and on physically connected hardware, reducing network latency and increasing throughput.

Elastic Fabric Adapter (EFA) is a network device for EC2 instances that enables low-latency, high-throughput networking using OS-bypass technology, ideal for tightly-coupled HPC applications.

Each instance can support single-flow 10 Gbps bandwidth using EFA, and collectively, the cluster can achieve up to 100 Gbps aggregate throughput when properly configured.

This solution supports the Performance Efficiency and Resilience design principles and is a standard AWS-recommended pattern for HPC.

Reference:

EC2 Placement Groups

Elastic Fabric Adapter Overview

Best Practices for HPC on AWS

Question: 503

A company runs an enterprise resource planning (ERP) system on Amazon EC2 instances in a single AWS Region. Users connect to the ERP system by using a public API that is hosted on the EC2 instances. International users report slow API response times from their data centers.

A solutions architect needs to improve API response times for the international users.

Which solution will meet these requirements MOST cost-effectively?

- A. Set up an AWS Direct Connect connection that has a public virtual interface (VIF) to connect each user's data center to the EC2 instances. Create a Direct Connect gateway for the ERP system API to route user API requests.
- B. Deploy Amazon API Gateway endpoints in multiple Regions. Use Amazon Route 53 latency-based routing to route requests to the nearest endpoint. Configure a VPC peering connection between the Regions to connect to the ERP system.
- C. Set up AWS Global Accelerator. Configure listeners for the necessary ports. Configure endpoint groups for the appropriate Regions to distribute traffic. Create an endpoint in each group for the API.
- D. Use AWS Site-to-Site VPN to establish dedicated VPN tunnels between multiple Regions and user networks. Route traffic to the API through the VPN connections.

Explanation:

AWS Global Accelerator improves the performance and availability of applications by directing user traffic through the AWS global network of edge locations using anycast IP addresses. It reduces latency and jitter for global users accessing applications in a single Region.

Why this works:

Global Accelerator routes user requests to the nearest AWS edge location using AWS's high-performance backbone network.

It then forwards traffic to the optimal endpoint — in this case, the public API hosted on EC2.

This is much more cost-effective and requires less operational complexity than deploying and maintaining multiple API Gateway endpoints across regions (Option B), or setting up Direct Connect links for every international location (Option A).

Option C requires no application change and is designed specifically for latency improvement and high availability.

Reference:

AWS Global Accelerator Documentation

Use Cases for Global Accelerator**Performance Improvements for Global Users****Question: 504**

A solutions architect is creating a website that will be hosted from an Amazon S3 bucket. The website must support secure browser connections (HTTPS).

Which combination of actions must the solutions architect take to meet this requirement? (Select TWO.)

- A. Create an Elastic Load Balancing (ELB) load balancer. Configure the load balancer to direct traffic to the S3 bucket.
- B. Create an Amazon CloudFront distribution. Set the S3 bucket as an origin.
- C. Configure the Elastic Load Balancing (ELB) load balancer with an SSL/TLS certificate.
- D. Configure the Amazon CloudFront distribution with an SSL/TLS certificate.
- E. Configure the S3 bucket with an SSL/TLS certificate.

Answer: B,D

Explanation:

To serve a static website hosted in Amazon S3 over HTTPS, you must use Amazon CloudFront because S3 does not natively support HTTPS for static website endpoints.

Steps to meet HTTPS requirement:

B . Create a CloudFront distribution and configure the S3 bucket as the origin. This enables global edge caching and performance optimization.

D . Attach an SSL/TLS certificate (typically from AWS Certificate Manager) to the CloudFront distribution to handle HTTPS connections.

S3 buckets used as static website hosts only support HTTP directly. While S3 supports HTTPS for REST API access, it does not support HTTPS on static website endpoints.

This setup aligns with security best practices and supports the Secure and Operational Excellence pillars of the AWS Well-

Architected Framework.

Reference:

Hosting a static website using Amazon S3 and CloudFront

CloudFront + HTTPS with ACM

Question: 505

A company has a non-production application that runs on an Amazon EC2 instance. The EC2 instance has an instance profile and an associated IAM role.

The company wants to automate patching for the EC2 instance.

Which solution will meet this requirement?

- A. Create a new IAM role. Attach the AmazonSSMManagedInstanceCore policy to the new IAM role. Attach the new IAM role to the EC2 instance profile. Use AWS Systems Manager to patch the instance.
- B. Create an IAM user. Attach the AmazonSSMManagedInstanceCore policy to the IAM user. Configure AWS Systems Manager to use the IAM user to patch the instance.
- C. Attach the AmazonSSMManagedInstanceCore policy to the existing IAM role. Use AWS Systems Manager to patch the EC2 instance.
- D. Attach the AmazonSSMManagedInstanceCore policy to an existing IAM user. Use EC2 Image Builder to patch the EC2 instance.

Answer: C

Explanation:

To manage EC2 instances with AWS Systems Manager (SSM), the EC2 instance must be configured as a managed instance by attaching an IAM role that has the AmazonSSMManagedInstanceCore managed policy.

This policy allows:

SSM agent to register the instance with SSM

Perform actions like patching, automation, session management, inventory collection, etc.

Access to SSM endpoints (via internet or VPC endpoint if needed)

Since the EC2 instance already has an IAM role, the least operational overhead option is to attach the required policy to the existing role (Option C). No need to create new IAM roles or users, which simplifies management and adheres to the principle of least privilege.

Patching can then be automated via SSM Patch Manager, ensuring consistency, compliance, and operational efficiency.

Reference:

SSM Managed Instance Setup

AmazonSSMManagedInstanceCore Policy

Patching EC2 with SSM

Question: 506

A company wants a flexible compute solution that includes Amazon EC2 instances and AWS Fargate. The company does not want to commit to multi-year contracts.

Which purchasing option will meet these requirements MOST cost-effectively?

- A. Purchase a 1-year EC2 Instance Savings Plan with the All Upfront option.
- B. Purchase a 1-year Compute Savings Plan with the No Upfront option.
- C. Purchase a 1-year Compute Savings Plan with the Partial Upfront option.
- D. Purchase a 1-year Compute Savings Plan with the All Upfront option.

Answer: B

Explanation:

To optimize costs for both Amazon EC2 and AWS Fargate, the best option is a Compute Savings Plan because it offers flexibility across instance families, Regions, and compute options including EC2, AWS Fargate, and AWS Lambda.

Unlike EC2 Instance Savings Plans, which apply only to specific instance families, Compute Savings Plans apply across multiple services.

Since the company does not want to commit to multi-year contracts or large upfront payments, the 1-year No Upfront Compute Savings Plan provides the greatest flexibility with no upfront capital commitment, while still offering cost savings over On-Demand pricing.

This option also aligns with cost-optimization best practices by allowing for scalability and service mix flexibility.

Reference:

AWS Compute Savings Plans

AWS Pricing Models

Question: 507

A company uses Amazon EC2 instances behind an Application Load Balancer (ALB) to serve content to users. The company uses Amazon Elastic Block Store (Amazon EBS) volumes to store data.

The company needs to encrypt data in transit and at rest.

Which combination of services will meet these requirements? (Select TWO.)

- A. Amazon GuardDuty
- B. AWS Shield
- C. AWS Certificate Manager (ACM)
- D. AWS Secrets Manager
- E. AWS Key Management Service (AWS KMS)

Answer: C,E

Explanation:

To secure data in transit, the company should use AWS Certificate Manager (ACM) to provide SSL/TLS certificates for the Application Load Balancer. ACM allows easy provisioning, management, and renewal of public and private certificates, ensuring secure communication between users and applications.

To secure data at rest, AWS Key Management Service (KMS) is used to manage encryption keys for Amazon EBS volumes. EBS integrates

with AWS KMS, allowing for server-side encryption using KMS- managed keys (SSE-KMS), thus meeting the encryption at rest requirement.

Other options:

GuardDuty (A) is for threat detection, not encryption.

AWS Shield (B) protects against DDoS attacks, not encryption.

Secrets Manager (D) manages credentials, not general data encryption.

This solution follows the AWS Well-Architected Framework - Security Pillar.

Reference:

Encrypting EBS volumes with KMS

Using ACM with ALB

Question: 508

A disaster response team is using drones to collect images of recent storm damage. The response team's laptops lack the storage and compute capacity to transfer the images and process the data.

While the team has Amazon EC2 instances for processing and Amazon S3 buckets for storage, network connectivity is intermittent and unreliable. The images need to be processed to evaluate the damage.

What should a solutions architect recommend?

- A. Use AWS Snowball Edge devices to process and store the images.
- B. Upload the images to Amazon Simple Queue Service (Amazon SQS) during intermittent connectivity to EC2 instances.
- C. Configure Amazon Data Firehose to create multiple delivery streams aimed separately at the S3 buckets for storage and the EC2 instances for processing images.
- D. Use AWS Storage Gateway pre-installed on a hardware appliance to cache the images locally for Amazon S3 to process the images when connectivity becomes available.

Answer: A

Explanation:

AWS Snowball Edge is specifically designed for use cases that involve limited or unreliable network connectivity. It enables data transfer and local compute processing at edge locations.

It comes in two options: Snowball Edge Storage Optimized and Snowball Edge Compute Optimized.

The Compute Optimized model allows the disaster response team to both store images locally and process data on the device using Amazon EC2-compatible compute resources.

This removes the need for constant network connectivity. After processing, the device can be shipped back to AWS, where data is uploaded to S3.

Other options fail due to:

SQS not being suitable for large binary image data (Option B)

Kinesis Data Firehose needing steady connectivity (Option C)

Storage Gateway is for hybrid cloud environments with ongoing connection, not rugged field use (Option D)

Reference:

[AWS Snowball Edge Overview](#)

[Snowball Edge Use Cases](#)

Question: 509

A company recently migrated its application to a VPC on AWS. An AWS Site-to-Site VPN connection connects the company's on-premises network to the VPC. The application retrieves customer data from another system that resides on premises. The application uses an on-premises DNS server to resolve domain records. After the migration, the application is not able to connect to the customer data because of name resolution errors.

Which solution will give the application the ability to resolve the internal domain names?

- A. Launch EC2 instances in the VPC. On the EC2 instances, deploy a custom DNS forwarder that forwards all DNS requests to the on-premises DNS server. Create an Amazon Route 53 private hosted zone that uses the EC2 instances for name servers.
- B. Create an Amazon Route 53 Resolver outbound endpoint. Configure the outbound endpoint to forward DNS queries against the on-premises domain to the on-premises DNS server.
- C. Set up two AWS Direct Connect connections between the AWS environment and the on-premises network. Set up a link aggregation group (LAG) that includes the two connections. Change the VPC resolver address to point to the on-premises DNS server.
- D. Create an Amazon Route 53 public hosted zone for the on-premises domain. Configure the network ACLs to forward DNS requests against the on-premises domain to the Route 53 public hosted zone.

Answer: B

Explanation:

When AWS workloads must resolve DNS names from on-premises systems over a hybrid network (like VPN or Direct Connect), the best solution is to use Amazon Route 53 Resolver outbound endpoints.

The outbound endpoint enables DNS queries to be forwarded from your VPC to on-premises DNS servers.

You must also configure a Route 53 Resolver forwarding rule to define which domain names (e.g., corp.internal) should be forwarded to the specific on-premises DNS IPs.

This setup allows private DNS resolution from AWS to on-premises systems and is fully managed, eliminating the need to run and maintain EC2-based DNS proxies (as in option A).

Options C and D are incorrect:

C is not DNS-specific and doesn't solve name resolution.

D misuses a public hosted zone for a private DNS domain.

Reference:

Route 53 Resolver Outbound Endpoints

Question: 510

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS `sftp.example.com` through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

- A. Move the EC2 instance into an Auto Scaling group. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record `sftp.example.com` in Route 53 to point to the ALB.
- B. Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record `sftp.example.com` in Route 53 to point to the server endpoint hostname.
- C. Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record `sftp.example.com` in Route 53 to point to the file gateway endpoint.
- D. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record `sftp.example.com` in Route 53 to point to the NLB.

Answer: B

Explanation:

The optimal way to improve reliability and scalability of SFTP on AWS is to use AWS Transfer Family (for SFTP). It provides a fully managed SFTP server integrated with Amazon S3.

No EC2 instances or infrastructure management is required.

AWS Transfer Family supports custom DNS domains (e.g., sftp.example.com) and allows integration with existing authentication mechanisms like LDAP, AD, or custom identity providers.

Files are uploaded directly to S3, eliminating the need for cron jobs to move data from EC2 to S3.

Built-in high availability and scalability removes the burden of managing infrastructure.

Other options:

A and D still require manual scaling, server maintenance, and cron jobs.

C (Storage Gateway) is used for hybrid file access, not for replacing an SFTP server.

Reference:

AWS Transfer Family for SFTP

Question: 511

A company runs a critical Amazon RDS for MySQL DB instance in a single Availability Zone. The company must improve the availability of the DB instance.

Which solution will meet this requirement?

- A. Configure the DB instance to use a multi-Region DB instance deployment.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue in the AWS Region where the company hosts the DB instance to manage writes to the DB instance.
- C. Configure the DB instance to use a Multi-AZ DB instance deployment.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue in a different AWS Region than the Region where the company hosts the DB instance to manage writes to the DB instance.

Answer: C

Explanation:

To improve availability and fault tolerance of an Amazon RDS instance, the recommended approach is to configure a Multi-AZ deployment.

Multi-AZ deployments for RDS automatically replicate data to a standby instance in a different Availability Zone (AZ).

If a failure occurs in the primary AZ (due to hardware, network, or power), RDS will automatically failover to the standby instance with

minimal downtime, without administrative intervention.

This is an AWS-managed feature and does not require application modification.

It does not provide scalability or load balancing; it's designed for high availability and resiliency.

Options A, B, and D are incorrect:

A refers to cross-Region, which is used for disaster recovery, not high availability.

B and D with SQS do not address high availability directly for the RDS instance; queues help decouple systems but do not make a database more resilient.

Reference:

Amazon RDS Multi-AZ Deployments

Question: 512

A company operates a data lake in Amazon S3 that stores large datasets in multiple formats. The company has an application that retrieves and processes subsets of data from multiple objects in the data lake based on filtering criteria. For each data query, the application currently downloads the entire S3 object and performs transformations. The current process requires a large amount of transformation time.

The company wants a solution that will give the application the ability to query and filter directly on S3 objects without downloading the objects.

Which solution will meet these requirements?

- A. Use Amazon Athena to query and filter the objects in Amazon S3.
- B. Use Amazon EMR to process and filter the objects.
- C. Use Amazon API Gateway to create an API to retrieve filtered results from Amazon S3.
- D. Use Amazon ElastiCache (Valkey) to cache the objects.

Answer: A

Explanation:

The best solution to query and filter S3 data directly without downloading the full object is to use Amazon Athena.

Amazon Athena is an interactive query service that lets you use SQL to analyze structured, semistructured, and unstructured data directly in Amazon S3, without needing to move or transform the data.

It supports formats like CSV, JSON, ORC, Parquet, and Avro and integrates with AWS Glue Data Catalog for schema management.

Athena is serverless, meaning there's no infrastructure to manage, and it's billed per query, which keeps it cost-effective.

Option B (EMR) is heavier and requires managing a cluster.

Option C (API Gateway) is not suited for querying S3 datasets.

Option D (ElastiCache) is a memory store, not a query engine.

Reference:

What is Amazon Athena?

Question: 513

A large financial services company uses Amazon ElastiCache (Redis OSS) for its new application that has a global user base. A solutions architect must develop a caching solution that will be available across AWS Regions and include low-latency replication and failover capabilities for disaster recovery (DR). The company's security team requires the encryption of cross-Region data transfers.

Which solution meets these requirements with the LEAST amount of operational effort?

- A. Enable cluster mode in ElastiCache (Redis OSS). Then create multiple clusters across Regions and replicate the cache data by using AWS Database Migration Service (AWS DMS). Promote a cluster in the failover Region to handle production traffic when DR is required.
- B. Create a global data store in ElastiCache (Redis OSS). Then create replica clusters in two other Regions. Promote one of the replica clusters as primary when DR is required.
- C. Disable cluster mode in ElastiCache (Redis OSS). Then create multiple replication groups across Regions and replicate the cache data by using AWS Database Migration Service (AWS DMS). Promote a replication group in the failover Region to primary when DR is required.
- D. Create a snapshot of ElastiCache (Redis OSS) in the primary Region and copy it to the failover Region. Use the snapshot to restore the cluster from the failover Region when DR is required.

Answer: B

Explanation:

The optimal solution for low-latency global caching with disaster recovery and cross-Region replication is to use Amazon ElastiCache Global Datastore for Redis OSS.

A Global Datastore enables fully managed cross-Region replication and supports automatic failover by promoting read replica

clusters in another Region.

ElastiCache ensures encryption in-transit and at-rest, meeting compliance and security requirements.

It's a fully managed AWS-native feature, reducing operational effort compared to setting up DMS- based or snapshot-based replication manually.

Other options (A, C, D):

Require manual setup and management (e.g., custom DMS pipelines, snapshots).

Do not offer real-time replication or failover without manual intervention.

Reference:

ElastiCache Global Datastore for Redis

Question: 514

A developer creates a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The developer reviews the deployment and notices some suspicious traffic to the application. The traffic is malicious and is coming from a single public IP address. A solutions architect must block the public IP address.

Which solution will meet this requirement?

- A. Create a security group rule to deny all inbound traffic from the suspicious IP address. Associate the security group with the ALB.
- B. Implement Amazon Detective to monitor traffic and to block malicious activity from the internet. Configure Detective to integrate with the ALB.
- C. Implement AWS Resource Access Manager (AWS RAM) to manage traffic rules and to block malicious activity from the internet. Associate AWS RAM with the ALB.
- D. Add the malicious IP address to an IP set in AWS WAF. Create a web ACL. Include an IP set rule with the action set to BLOCK. Associate the web ACL with the ALB.

Answer: D

Explanation:

When an application is fronted by an Application Load Balancer (ALB) and malicious traffic is detected from a specific IP, the correct way to block the IP is by using AWS WAF (Web Application Firewall).

With AWS WAF, you can create an IP Set to include the offending IP address or range.

Then create a Web ACL (Access Control List) with a rule set to BLOCK requests from that IP set.

Finally, associate the Web ACL with the ALB.

Security groups (Option A) cannot deny specific IPs because they are stateful and allow-only rules.

Amazon Detective (Option B) is a security analysis and investigation tool; it doesn't block traffic.

AWS RAM (Option C) is for resource sharing across accounts, not for blocking IPs.

This approach aligns with AWS's Security Pillar of the Well-Architected Framework and is fully managed, with minimal operational effort.

Reference:

Using AWS WAF with an Application Load Balancer

Block IPs with AWS WAF

Question: 515

A company runs an application that uses Docker containers in an on-premises data center. The application runs on a container host that stores persistent data files in a local volume. Container instances use the stored persistent data.

The company wants to migrate the application to fully managed AWS services.

Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS) with self-managed nodes. Attach an Amazon Elastic Block Store (Amazon EBS) volume to an Amazon EC2 instance. Mount the EBS volume on the containers to provide persistent storage.
- B. Use Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Create an Amazon Elastic File System (Amazon EFS) volume. Mount the EFS volume on the containers to provide persistent storage.
- C. Use Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Create an Amazon DynamoDB table. Configure the application to use the DynamoDB table for persistent storage.
- D. Use Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Create an Amazon Elastic File System (Amazon EFS) volume. Mount the EFS volume on the containers to provide persistent storage.

Answer: B

Explanation:

The company wants to move from an on-premises Docker environment to fully managed AWS services with persistent storage. The best fit is:

Amazon ECS with AWS Fargate launch type: This is a serverless container orchestration solution where AWS manages the underlying infrastructure, removing the need to manage EC2 or Kubernetes nodes.

Amazon EFS (Elastic File System): This is a fully managed, scalable, and shared file system for use with ECS tasks. It supports persistent storage for containers, replacing the local volumes used onpremises.

This combination (ECS + Fargate + EFS) is fully managed and requires no manual server maintenance.

Option A uses EKS with self-managed nodes, which is not fully managed.

Option C (DynamoDB) is for structured key-value storage, not for persistent file storage.

Option D uses ECS with EC2 launch type, which is not serverless and requires managing instances.

Reference:

Using Amazon ECS with AWS Fargate

Mounting EFS volumes in ECS tasks

Question: 516

A company has established a new AWS account. The account is newly provisioned and no changes have been made to the default settings.

The company is concerned about the security of the AWS account root user.

What should be done to secure the root user?

- A. Create IAM users for daily administrative tasks. Disable the root user.
- B. Create IAM users for daily administrative tasks. Enable multi-factor authentication on the root user.
- C. Generate an access key for the root user. Use the access key for daily administration tasks instead of the AWS Management Console.
- D. Provide the root user credentials to the most senior solutions architect. Have the solutions architect use the root user for daily administration tasks.

Answer: B

Explanation:

According to the AWS Well-Architected Framework - Security Pillar and the AWS Identity and Access Management (IAM) User Guide, the root user account in an AWS account is extremely powerful and should be protected with strict security measures.

From AWS documentation:

"We recommend that you not use the root user for everyday tasks, even administrative ones.

Instead, create IAM users and grant them only the permissions they need. To help protect your AWS account, enable multi-factor authentication (MFA) for the root user."

(Source: AWS Identity and Access Management User Guide - Securing the Root User)

The correct and recommended action is to create IAM users with specific permissions for daily operations and enable MFA on the root user to provide an additional layer of security. The root user cannot be disabled, so Option A is technically incorrect. AWS also explicitly advises against using root access keys (Option C) or sharing root credentials (Option D), both of which violate the principle of least privilege.

Best practices summarized from AWS official documentation:

Do not use root user for routine tasks

Enable MFA for root user immediately

Create individual IAM users and assign least privilege

Avoid creating or using root user access keys

These recommendations are foundational to securing any new AWS account and are consistently emphasized in the AWS Certified Solutions Architect - Official Study Guide and the AWS Security Best Practices whitepaper.

Reference:

AWS IAM User Guide - "Securing the root user"

AWS Well-Architected Framework - Security Pillar

AWS Certified Solutions Architect - Official Study Guide (Latest Edition)

Question: 517

A company runs an application on premises. The application needs to periodically upload large files to an Amazon S3 bucket. A solutions architect needs a solution to provide the application with shortlived authenticated access to the S3 bucket. The solution must not use long-term credentials. The solution needs to be secure and scalable.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM user that has an access key and a secret key. Store the keys on the on-premises server in an environment variable. Attach a policy to the IAM user that restricts access to only the S3 bucket.
- B. Configure an AWS Site-to-Site VPN connection from the on-premises environment to the company's VPC. Launch an Amazon EC2 instance with an instance profile. Route all file uploads from the on-premises application through the EC2 instance to the S3 bucket.
- C. Configure an S3 bucket policy to allow access for the on-premises server's public IP address. Configure the policy to allow PUT operations only from the server's IP address.

D. Configure a trust relationship between the on-premises server and AWS Security Token Service (AWS STS). Generate credentials by assuming an IAM role for each upload operation.

Answer: D

Explanation:

The best practice to securely provide short-term access to AWS resources, such as Amazon S3, without using long-term credentials, is to use AWS Security Token Service (STS). According to the AWS IAM Best Practices and the Security Pillar of the AWS Well-Architected Framework, temporary credentials should always be used over long-term credentials when possible.

From AWS IAM Documentation:

“Use temporary credentials (IAM roles and AWS STS) instead of long-term access keys. Temporary security credentials are short-term, automatically expire, and are retrieved using AWS STS.”

(Source: IAM Best Practices - IAM User Guide)

Option D outlines the use of a trust relationship and assume-role mechanism via AWS STS. This allows the on-premises application to request temporary, scoped-down credentials to upload files to S3 securely. This approach is:

Secure - Uses short-lived credentials with least privilege

Scalable - No need for EC2 or VPN tunnels

Low Operational Overhead - No infrastructure to maintain

AWS-Recommended - Aligned with security best practices

In contrast:

Option A uses long-term credentials, which is a security risk.

Option B requires additional infrastructure (EC2, VPN), increasing complexity and cost.

Option C relies on IP-based access, which is insecure and not a form of identity-based authentication.

Reference:

AWS IAM Best Practices - "Use temporary credentials"

AWS Security Token Service - "Temporary Security Credentials"

AWS Well-Architected Framework - Security Pillar

Question: 518

A company has multiple AWS accounts with applications deployed in the us-west-2 Region.

Application logs are stored within Amazon S3 buckets in each account. The company wants to build a centralized log analysis solution that

uses a single S3 bucket. Logs must not leave us-west-2, and the company wants to incur minimal operational overhead.

- A. Create an S3 Lifecycle policy that copies the objects from one of the application S3 buckets to the centralized S3 bucket.
- B. Use S3 Same-Region Replication to replicate logs from the S3 buckets to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.
- C. Write a script that uses the PutObject API operation every day to copy the entire contents of the buckets to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.
- D. Write AWS Lambda functions in these accounts that are triggered every time logs are delivered to the S3 buckets (s3:ObjectCreated:*) event. Copy the logs to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.

Answer: B

Explanation:

The most cost-effective and low-maintenance solution to aggregate S3 data from multiple accounts within the same AWS Region is to use Amazon S3 Same-Region Replication (SRR).

From AWS S3 Documentation:

"S3 Same-Region Replication (SRR) automatically replicates new objects between buckets in the same AWS Region. SRR is commonly used to aggregate logs into a single bucket, simplify data aggregation, and meet compliance requirements."

(Source: Amazon S3 User Guide - Replication Overview)

Key benefits of using SRR for this use case:

Zero operational overhead - No Lambda, scripts, or custom code

Fully managed - AWS handles all replication automatically

Secure and efficient - No data leaves the us-west-2 Region

Supports cross-account replication - With proper IAM roles and permissions

Low-cost - Compared to custom ETL pipelines or Lambda solutions

In contrast:

Option A (Lifecycle policies) do not support copying, only expiration, transitions, or deletions.

Option C (scripts) requires custom scheduling and maintenance, increasing operational overhead.

Option D (Lambda) adds complexity and cost and is not as scalable or hands-off as SRR.

Reference:

AWS Well-Architected Framework - Cost Optimization Pillar

AWS Best Practices for Data Aggregation and S3 Replication

Question: 519

A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to serve a static website. The solution must use AWS WAF to inspect all website traffic.

- A. Configure an S3 bucket policy to accept only requests that come from the AWS WAF Amazon Resource Name (ARN).
- B. Configure CloudFront to forward all incoming requests to AWS WAF before CloudFront requests content from the S3 origin.
- C. Configure a security group that allows only CloudFront IP addresses to access Amazon S3. Associate AWS WAF to the CloudFront distribution.
- D. Configure CloudFront and Amazon S3 to use an origin access control (OAC) to secure the origin S3 bucket. Associate AWS WAF to the CloudFront distribution.

Answer: D

Explanation:

The correct and secure approach is to use Amazon CloudFront with Origin Access Control (OAC) to protect the S3 origin and attach AWS WAF to the CloudFront distribution to inspect and filter traffic at the edge before reaching the origin.

From AWS Documentation:

“AWS WAF is integrated with Amazon CloudFront, allowing inspection of HTTP(S) requests at the edge location before forwarding to your origin. To restrict direct access to the S3 bucket, use Origin Access Control (OAC).”

(Source: Amazon CloudFront Developer Guide - Serving private content)

Why Option D is correct:

CloudFront is the only service that integrates with AWS WAF for full HTTP layer inspection.

Origin Access Control (OAC) ensures that only CloudFront can access the S3 origin—replacing older Origin Access Identity (OAI) features.

The S3 bucket policy is configured to trust requests only from CloudFront using OAC signed requests.

Why the other options are incorrect:

Option A: WAF ARN is not a principal in S3 bucket policy. IAM does not support bucket policies based on WAF ARNs.

Option B: Incorrect - CloudFront doesn't "forward requests to WAF"; rather, WAF is associated with CloudFront and inspects requests at the edge.

Option C: S3 does not use security groups; they are for EC2/network interfaces. This shows a misunderstanding of how S3 works.

Reference:

Amazon CloudFront Developer Guide - "Serving Private Content with OAC"

AWS WAF Developer Guide - "Protecting CloudFront with AWS WAF"

AWS Well-Architected Framework - Security Pillar

Question: 520

A manufacturing company develops an application to give a small team of executives the ability to track sales performance globally. The application provides a real-time simulator in a popular programming language. The company uses AWS Lambda functions to support the simulator. The simulator is an algorithm that predicts sales performance based on specific variables.

Although the solution works well initially, the company notices that the time required to complete simulations is increasing exponentially.

A solutions architect needs to improve the response time of the simulator.

Which solution will meet this requirement in the MOST cost-effective way?

- A. Use AWS Fargate to run the simulator. Serve requests through an Application Load Balancer (ALB).
- B. Use Amazon EC2 instances to run the simulator. Serve requests through an Application Load Balancer (ALB).
- C. Use AWS Batch to run the simulator. Serve requests through a Network Load Balancer (NLB).
- D. Use Lambda provisioned concurrency for the simulator functions.

Answer: D

Explanation:

When an AWS Lambda function is invoked, especially after periods of inactivity, it may experience cold starts that delay execution. As demand increases, the scaling behavior and latency of Lambda can affect performance. Provisioned Concurrency is an AWS feature designed specifically to solve this issue.

From AWS Documentation:

"Provisioned Concurrency keeps functions initialized and hyper-ready to respond in double-digit milliseconds."

(Source: AWS Lambda Developer Guide - Managing Concurrency)

Why Option D is correct:

Provisioned Concurrency ensures that a specified number of Lambda function instances are always warm and ready to serve requests, eliminating cold start latency.

It's cost-effective for workloads with consistent usage patterns, like real-time simulations for a small user group.

Maintains scalability and low overhead of Lambda without moving to managed container or EC2 platforms.

Why the other options are less optimal:

Option A (Fargate) and Option B (EC2): Introduce more infrastructure and higher ongoing costs for a small team with likely intermittent usage.

Option C (AWS Batch): Ideal for batch jobs, not real-time simulations; also incurs higher latency due to job queuing.

Reference:

AWS Lambda Developer Guide - "Concurrency and Scaling"

AWS Well-Architected Framework - Performance Efficiency Pillar

AWS Compute Services Comparison - Lambda vs EC2 vs Fargate

Question: 521

A company runs a NetApp storage array in an on-premises data center. The company wants to migrate the storage array to Amazon FSx for NetApp ONTAP. The company has a mix of NFS and SMB file shares with complex directory structures and over 60 million small files. The company has 10 Gbps of network bandwidth available. The company wants to optimize migration efficiency for the file system.

- A. Use AWS DataSync with a bandwidth throttle. Use the All tiering policy.
- B. Provision an AWS Storage Gateway Volume Gateway. Configure a zero-ETL integration with the FSx for NetApp ONTAP file system.
- C. Set up NetApp SnapMirror replication between the on-premises array and the FSx for ONTAP file system.
- D. Use AWS Snowball Edge to perform an offline migration.

Answer: C

Explanation:

Amazon FSx for NetApp ONTAP fully supports native NetApp SnapMirror replication, making it the most efficient and reliable option for

migrating NetApp data from on-premises to AWS.

From AWS Documentation:

“You can use SnapMirror to replicate data from your on-premises NetApp systems to FSx for ONTAP for seamless, block-level, incremental transfers. This provides a highly efficient and performant method for migration, especially for large datasets.”

(Source: Amazon FSx for NetApp ONTAP - Migration Guide)

Why Option C is correct:

SnapMirror offers block-level replication, making it highly efficient for millions of small files.

It supports NFS and SMB file shares, preserving directory structures and permissions.

Reduces cutover time and allows for incremental syncs.

Uses the existing 10 Gbps network for fast transfers.

Why the other options are incorrect:

Option A (DataSync): Suitable for many file-based migrations but less efficient for very large datasets with millions of small files compared to SnapMirror.

Option B (Storage Gateway): Volume Gateway is not used for full-scale file migrations; it's for hybrid cloud access.

Option D (Snowball Edge): Useful for offline migrations, but online SnapMirror is more efficient and avoids shipping delays.

Reference:

Amazon FSx for NetApp ONTAP User Guide - "Migrating Using SnapMirror"

AWS Storage Migration Options Whitepaper

AWS Well-Architected Framework - Resilience Pillar

Question: 522

A company needs a solution to ingest streaming sensor data from 100,000 devices, transform the data in near real time, and load the data into Amazon S3 for analysis. The solution must be fully managed, scalable, and maintain sub-second ingestion latency.

A. Use Amazon Kinesis Data Streams to ingest the data. Use Amazon Managed Service for Apache Flink to process the data in near real time. Use an Amazon Data Firehose stream to send processed data to Amazon S3.

B. Use Amazon Simple Queue Service (Amazon SQS) standard queues to collect the sensor data. Invoke AWS Lambda functions to transform and process SQS messages in batches. Configure the Lambda functions to use an AWS SDK to write transformed data to Amazon S3.

C. Deploy a fleet of Amazon EC2 instances that run Apache Kafka to ingest the data. Run Apache Spark on Amazon EMR clusters to

process the data. Configure Spark to write processed data directly to Amazon S3.

D. Implement Amazon EventBridge to capture all sensor data. Use AWS Batch to run containerized transformation jobs on a schedule. Configure AWS Batch jobs to process data in chunks. Save results to Amazon S3.

Answer: A

Explanation:

The most scalable and managed solution for streaming ingestion, real-time transformation, and delivery to Amazon S3 is Amazon Kinesis Data Streams, Amazon Managed Service for Apache Flink, and Amazon Kinesis Data Firehose.

From AWS Documentation:

“Amazon Kinesis Data Streams enables real-time processing of streaming data at massive scale. With Apache Flink on Kinesis Data Analytics, you can process data streams in near real-time, then use Amazon Kinesis Data Firehose to reliably deliver that data to S3.”

(Source: Amazon Kinesis Developer Guide)

Why A is correct:

Fully managed: All services involved are serverless and managed.

Real-time ingestion: Kinesis Data Streams supports sub-second latency and can handle high- throughput workloads like 100,000+ devices.

Near real-time processing: Apache Flink is designed for continuous stream processing with complex event handling.

Efficient delivery: Kinesis Firehose delivers processed data directly to S3 with retry and backup capability.

Why other options are incorrect:

Option B: SQS is not optimized for real-time streaming at high volume.

Option C: EC2 + Kafka + EMR adds high operational overhead and cost.

Option D: EventBridge is event-driven, not designed for high-throughput streaming; AWS Batch is unsuitable for near real-time processing.

Reference:

Amazon Kinesis Developer Guide

AWS Well-Architected Framework - Performance Efficiency Pillar

Amazon Managed Flink (Apache Flink on KDA)

Question: 523

A solutions architect needs to build a log storage solution for a client. The client has an application that produces user activity logs that track user API calls to the application. The application typically produces 50 GB of logs each day. The client needs a storage solution that makes the logs available for occasional querying and analytics.

- A. Store user activity logs in an Amazon S3 bucket. Use Amazon Athena to perform queries and analytics.
- B. Store user activity logs in an Amazon OpenSearch Service cluster. Use OpenSearch Dashboards to perform queries and analytics.
- C. Store user activity logs in an Amazon RDS instance. Use an Open Database Connectivity (ODBC) connector to perform queries and analytics.
- D. Store user activity logs in an Amazon CloudWatch Logs log group. Use CloudWatch Logs Insights to perform queries and analytics.

Answer: A

Explanation:

For infrequent or ad hoc querying of log data, Amazon S3 + Amazon Athena provides the most cost-effective, serverless, and scalable analytics solution.

From AWS Documentation:

“Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.”

(Source: Amazon Athena User Guide)

Why A is correct:

Amazon S3 offers durable, scalable, and cost-efficient storage.

Athena allows SQL-based querying on structured or semi-structured data like logs.

No need to provision or manage infrastructure.

Ideal for occasional querying at low cost.

Why the others are not optimal:

Option B: OpenSearch adds cost and is best for frequent, low-latency log querying.

Option C: RDS is not optimized for large-scale write-heavy log ingestion and costs more.

Option D: CloudWatch Logs is suitable for real-time monitoring, not for long-term storage and analytics of large log volumes.

Reference:

Amazon Athena User Guide

AWS Well-Architected Framework - Cost Optimization Pillar

Amazon S3 Storage Classes and Pricing Guide

Question: 524

A company uses an organization in AWS Organizations to manage a multi-account landing zone. The company requires all users who access AWS accounts in the organization to use a centralized identity system that follows the principle of least privilege for operational tasks. The company currently uses an external identity provider (IdP).

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Use AWS Identity and Access Management (IAM) to create IAM users and IAM user groups in each AWS account.
- B. Create permission sets in AWS IAM Identity Center. Assign the appropriate permission sets to the IAM users and IAM user groups in the accounts.
- C. Assign each IAM user to an IAM role by using an inline IAM policy based on operational duties. Assign each role to the appropriate AWS account in the organization.
- D. Configure a SAML identity provider in AWS Identity and Access Management (IAM) in each AWS account to establish a trust relationship with the company's external IdP.
- E. Enable AWS IAM Identity Center in the organization management account. Create user accounts and user groups.

Answer: B,E

Explanation:

AWS recommends using AWS IAM Identity Center (formerly AWS SSO) for centralized authentication and access control across multiple accounts in an AWS Organization, especially when integrating with an external IdP.

From AWS Documentation:

“Use IAM Identity Center to provide centralized access to multiple AWS accounts or applications. You can integrate with an external IdP via SAML 2.0. Assign users permissions through permission sets that define the roles users can assume.”

(Source: AWS IAM Identity Center User Guide)

Why B and E are correct:

E enables centralized identity federation using IAM Identity Center with your external IdP.

B uses permission sets to apply least-privilege access roles to users and groups across accounts, in alignment with the principle of least privilege.

Why others are incorrect:

Option A: IAM users in each account break centralized access model and are hard to manage at scale.

Option C: Managing individual IAM roles and inline policies across accounts is not scalable.

Option D: Per-account SAML providers are redundant when using IAM Identity Center, which provides centralized federation.

Reference:

[AWS IAM Identity Center User Guide](#)

[AWS Well-Architected Framework - Security Pillar](#)

[AWS Organizations and Identity Center Integration Docs](#)

Question: 525

A company wants to migrate an Oracle database to AWS. The database consists of a single table that contains millions of geographic information systems (GIS) images that are high resolution and are identified by a geographic code.

When a natural disaster occurs, tens of thousands of images get updated every few minutes. Each geographic code has a single image or row that is associated with it. The company wants a solution that is highly available and scalable during such events.

- A. Store the images and geographic codes in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance.
- B. Store the images in Amazon S3 buckets. Use Amazon DynamoDB with the geographic code as the key and the image S3 URL as the value.
- C. Store the images and geographic codes in an Amazon DynamoDB table. Configure DynamoDB Accelerator (DAX) during times of high load.
- D. Store the images in Amazon S3 buckets. Store geographic codes and image S3 URLs in a database table. Use Oracle running on an Amazon RDS Multi-AZ DB instance.

Answer: B

Explanation:

The most resilient and scalable architecture for handling millions of high-resolution images with frequent updates is to store the binary images in Amazon S3 and store their metadata or reference (geographic code and S3 URL) in Amazon DynamoDB.

From AWS Documentation:

“Store large objects like images in Amazon S3 and use Amazon DynamoDB to store metadata and references. This design pattern is scalable, highly available, and cost-effective.”

(Source: AWS Architecture Blog - Best Practices for Handling Large Objects)

Why B is correct:

Amazon S3 is designed for storing large volumes of binary data (images).

DynamoDB provides low-latency reads/writes using the geographic code as the partition key.

Highly available, serverless, and auto-scaling, suitable for disaster scenarios with bursts of activity.

Reduces pressure on the database layer by separating metadata from image storage.

Why the others are incorrect:

Option A and D: Storing images directly in RDS is expensive, unscalable, and not optimal for binary storage.

Option C: DynamoDB is suitable, but storing actual binary image data in DynamoDB is not best practice due to item size limits (400 KB) and performance concerns.

Reference:

AWS Architecture Blog - “Best Practices for Amazon S3 and DynamoDB Integration”

AWS Well-Architected Framework - Resilience Pillar

Amazon DynamoDB Developer Guide

Question: 526

A shipping company wants to run a Kubernetes container-based web application in disconnected mode while the company's ships are in transit at sea. The application must provide local users with **high availability**.

- A. Use AWS Snowball Edge as the primary and secondary sites.
- B. Use AWS Snowball Edge as the primary site, and use an AWS Local Zone as the secondary site.
- C. Use AWS Snowball Edge as the primary site, and use an AWS Outposts server as the secondary site.
- D. Use AWS Snowball Edge as the primary site, and use an AWS Wavelength Zone as the secondary site.

Answer: A

Explanation:

When operating in disconnected or limited-connectivity environments, such as ships at sea, AWS recommends using AWS Snowball Edge devices to host local compute and storage workloads. Snowball Edge supports Amazon EC2 instances and AWS IoT Greengrass, and can run Amazon EKS Anywhere for local Kubernetes cluster deployments.

From AWS Documentation:

“You can use AWS Snowball Edge devices to run compute-intensive applications in remote or disconnected locations. Snowball Edge devices support running Amazon EC2 instances and Amazon EKS Anywhere clusters.”

(Source: AWS Snow Family - Developer Guide)

Why A is correct:

Snowball Edge provides local compute and storage even without internet connectivity.

Multiple Snowball Edge devices can be clustered together for high availability and failover.

Fully self-contained environment suitable for ships or field operations.

Why the others are incorrect:

B, C, D require network connectivity to AWS Regions or Zones (Local Zone, Outposts, Wavelength), which is not available at sea. These options cannot operate fully disconnected.

Reference:

AWS Snow Family Developer Guide - “Running Compute Applications on AWS Snowball Edge”

AWS Well-Architected Framework - Resilience Pillar

AWS Architecture Blog - “Edge Computing with Snow Family”

Question: 527

A company is developing a photo-hosting application in the us-east-1 Region. The application gives users across multiple countries the ability to upload and view photos. Some photos are heavily viewed for months, while other photos are viewed for less than a week. The application allows users to upload photos that are up to 20 MB in size. The application uses photo metadata to determine which photos to display to each user.

The company needs a cost-effective storage solution to support the application.

- A. Store the photos in Amazon DynamoDB. Turn on DynamoDB Accelerator (DAX).
- B. Store the photos in the Amazon S3 Intelligent-Tiering storage class. Store the photo metadata and the S3 location URLs in Amazon DynamoDB.
- C. Store the photos in the Amazon S3 Standard storage class. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Use object tags to keep track of metadata.

D. Store the photos in an Amazon DynamoDB table. Use the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) storage class. Store the photo metadata in Amazon ElastiCache.

Answer: B

Explanation:

Amazon S3 Intelligent-Tiering automatically moves objects between frequent and infrequent access tiers based on access patterns, which minimizes cost without performance impact. Storing photo metadata in Amazon DynamoDB provides fast and scalable lookup by user or tag.

From AWS Documentation:

“The S3 Intelligent-Tiering storage class automatically optimizes storage costs by moving data between frequent and infrequent access tiers when access patterns change.”

(Source: Amazon S3 User Guide - Intelligent-Tiering)

Why B is correct:

S3 Intelligent-Tiering optimizes storage automatically for cost without lifecycle management overhead.

DynamoDB stores small metadata items and S3 URLs, enabling efficient queries and photo lookups.

This architecture scales globally, integrates seamlessly with applications, and minimizes operational cost.

Why other options are incorrect:

A & D: DynamoDB is not intended for storing binary objects like images.

C: Lifecycle rules are static and don't adapt dynamically to unpredictable access patterns.

Reference:

Amazon S3 User Guide - “Storage Classes and Intelligent-Tiering”

AWS Well-Architected Framework - Cost Optimization Pillar

AWS Developer Guide - “Building Serverless Image Hosting with S3 and DynamoDB”

Question: 528

A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company designed the application to work with session affinity (sticky sessions) for a better user experience.

The application must be available publicly over the internet as an endpoint. A WAF must be applied to the endpoint for additional security.

Session affinity (sticky sessions) must be configured on the endpoint.

- A. Create a public Network Load Balancer. Specify the application target group.
- B. Create a Gateway Load Balancer. Specify the application target group.
- C. Create a public Application Load Balancer. Specify the application target group.
- D. Create a second target group. Add Elastic IP addresses to the EC2 instances.
- E. Create a web ACL in AWS WAF. Associate the web ACL with the endpoint.

Answer: C,E

Explanation:

The Application Load Balancer (ALB) supports sticky sessions (session affinity) using application cookies. AWS WAF integrates natively with ALB to provide Layer 7 protection at the same endpoint.

From AWS Documentation:

“You can enable sticky sessions for your Application Load Balancer target groups to ensure that a user's requests are consistently routed to the same target. AWS WAF integrates with Application Load Balancer to protect your web applications from common exploits.”

(Source: Elastic Load Balancing User Guide & AWS WAF Developer Guide)

Why C and E are correct:

C: ALB operates at Layer 7 (HTTP/HTTPS), supports sticky sessions, and can serve as a public endpoint.

E: AWS WAF can be directly associated with the ALB to inspect traffic and enforce rules. Together, they fulfill both the security and session affinity requirements.

Why others are incorrect:

A: Network Load Balancer doesn't support session affinity.

B: Gateway Load Balancer is used for virtual appliances, not web applications.

D: Using EIPs bypasses load balancing and WAF integration.

Reference:

Elastic Load Balancing User Guide - “Sticky Sessions for Application Load Balancers”

AWS WAF Developer Guide - “Associating a Web ACL with an ALB”

AWS Well-Architected Framework - Security and Performance Pillars

Question: 529

A company runs its legacy web application on AWS. The web application server runs on an Amazon EC2 instance in the public subnet of a

VPC. The web application server collects images from customers and stores the image files in a locally attached Amazon Elastic Block Store (Amazon EBS) volume. The image files are uploaded every night to an Amazon S3 bucket for backup.

A solutions architect discovers that the image files are being uploaded to Amazon S3 through the public endpoint. The solutions architect needs to ensure that traffic to Amazon S3 does not use the public endpoint.

A. Create a gateway VPC endpoint for the S3 bucket that has the necessary permissions for the VPC.

Configure the subnet route table to use the gateway VPC endpoint.

B. Move the S3 bucket inside the VPC. Configure the subnet route table to access the S3 bucket through private IP addresses.

C. Create an Amazon S3 access point for the Amazon EC2 instance inside the VPC. Configure the web application to upload by using the Amazon S3 access point.

D. Configure an AWS Direct Connect connection between the VPC that has the Amazon EC2 instance and Amazon S3 to provide a dedicated network path.

Answer: A

Explanation:

To route S3 traffic privately from within a VPC, AWS provides Gateway VPC Endpoints for Amazon S3. These allow private connectivity to S3 without traversing the public internet or requiring an Internet Gateway.

From AWS Documentation:

“A gateway endpoint enables you to privately connect your VPC to supported AWS services such as Amazon S3 and DynamoDB without requiring an Internet Gateway, NAT device, or public IP.”

(Source: Amazon VPC User Guide - Gateway Endpoints)

Why A is correct:

Gateway VPC endpoints route S3 traffic internally within the AWS network.

Improves security and data privacy while reducing exposure to the public internet.

Requires only a simple route table modification and IAM policy configuration.

Why other options are incorrect:

B: S3 is a regional service; you cannot “move” it inside a VPC.

C: Access points do not change the routing path; still uses S3 endpoints.

D: AWS Direct Connect is for hybrid environments, not intra-AWS private connectivity.

Reference:

Amazon VPC User Guide - "Gateway Endpoints for Amazon S3"

AWS Well-Architected Framework - Security Pillar

AWS Networking Best Practices

Question: 530

A company has an API that receives real-time data from a fleet of monitoring devices. The API stores this data in an Amazon RDS DB instance for later analysis. The amount of data that the monitoring devices send to the API fluctuates. During periods of heavy traffic, the API often returns timeout errors.

After an inspection of the logs, the company determines that the database is not capable of processing the volume of write traffic that comes from the API. A solutions architect must minimize the number of connections to the database and must ensure that data is not lost during periods of heavy traffic.

- A. Increase the size of the DB instance to an instance type that has more available memory.
- B. Modify the DB instance to be a Multi-AZ DB instance. Configure the application to write to all active RDS DB instances.
- C. Modify the API to write incoming data to an Amazon Simple Queue Service (Amazon SQS) queue. Use an AWS Lambda function that Amazon SQS invokes to write data from the queue to the database.
- D. Modify the API to write incoming data to an Amazon Simple Notification Service (Amazon SNS) topic. Use an AWS Lambda function that Amazon SNS invokes to write data from the topic to the database.

Answer: C

Explanation:

Using Amazon SQS decouples the API from the database, allowing the system to handle write bursts without data loss. The queue buffers incoming requests, and AWS Lambda processes them asynchronously, writing to the RDS instance at a sustainable rate.

From AWS Documentation:

"Amazon SQS acts as a buffer between components that produce and consume data, allowing each to operate independently. You can use AWS Lambda to process messages from SQS and write them to other services such as Amazon RDS."

(Source: Amazon SQS Developer Guide)

Why C is correct:

SQS absorbs write surges and ensures durable message retention.

Lambda scales automatically with message volume and reduces DB connections.

Prevents timeout errors by smoothing traffic spikes.

Why others are incorrect:

A: Scaling RDS vertically doesn't address connection spikes.

B: Multi-AZ is for availability, not write throughput.

D: SNS is for fan-out message delivery, not queue-based buffering.

Reference:

Amazon SQS Developer Guide - "Decoupling Applications with Message Queues"

AWS Lambda Developer Guide - "Using Lambda with SQS"

AWS Well-Architected Framework - Performance Efficiency and Reliability Pillars

Question: 531

A company is migrating a distributed application to AWS. The application serves variable workloads. The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes. The company wants to modernize the application with a solution that maximizes resiliency and scalability.

How should a solutions architect design the architecture to meet these requirements?

- A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs. Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure EC2 Auto Scaling to use scheduled scaling.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs. Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure EC2 Auto Scaling based on the size of the queue.
- C. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure AWS CloudTrail as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the primary server.
- D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure Amazon EventBridge as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the compute nodes.

Answer: B

Explanation:

To decouple distributed workloads and improve scalability and resiliency, Amazon SQS should be used as a reliable job queue. The compute

nodes (workers) can poll the SQS queue for tasks, and EC2 Auto Scaling can dynamically scale instances based on the ApproximateNumberOfMessages metric.

From AWS Documentation:

“Amazon SQS enables you to decouple application components, allowing each part to scale independently. You can scale EC2 instances automatically based on the number of messages in the queue.”

(Source: Amazon SQS Developer Guide - Scaling Consumers with Auto Scaling)

Why B is correct:

Eliminates the single point of failure (the primary coordinator).

Enables event-driven scaling based on queue depth.

Provides durability and resiliency since messages are stored redundantly.

Fully managed and integrates seamlessly with Auto Scaling policies.

Why other options are incorrect:

A: Scheduled scaling does not respond to variable workloads.

C & D: Misuse of CloudTrail/EventBridge; not intended for workload coordination.

Reference:

Amazon SQS Developer Guide - “Integrating SQS with Auto Scaling”

AWS Well-Architected Framework - Reliability Pillar

Question: 532

A company runs several custom applications on Amazon EC2 instances. Each team within the company manages its own set of applications and backups. To comply with regulations, the company must be able to report on the status of backups and ensure that backups are encrypted.

Which solution will meet these requirements with the LEAST effort?

- A. Create an AWS Lambda function that processes AWS Config events. Configure the Lambda function to query AWS Config for backup-related data and to generate daily reports.
- B. Check the backup status of the EC2 instances daily by reviewing the backup configurations in AWS Backup and Amazon Elastic Block Store (Amazon EBS) snapshots.
- C. Use an AWS Lambda function to query Amazon EBS snapshots, Amazon RDS snapshots, and AWS Backup jobs. Configure the Lambda function to process and report on the data. Schedule the function to run daily.
- D. Use AWS Config and AWS Backup Audit Manager to ensure compliance. Review generated reports daily.

Answer: D

Explanation:

AWS Backup Audit Manager automates auditing and reporting of backup activity and compliance, while AWS Config provides visibility into configuration changes. Together, they provide the simplest, most automated, and compliant backup monitoring solution.

From AWS Documentation:

“AWS Backup Audit Manager automatically audits backup activity across AWS resources. You can use predefined or custom frameworks to monitor backup compliance and encryption status.”

(Source: AWS Backup Audit Manager User Guide)

Why D is correct:

Ensures centralized visibility into all backup jobs.

Verifies encryption status automatically.

Generates ready-to-use reports with minimal operational overhead.

Complies with regulatory requirements for data protection.

Why others are incorrect:

A & C: Custom Lambda automation increases maintenance effort.

B: Manual checking is operationally inefficient and error-prone.

Reference:

AWS Backup Audit Manager User Guide

AWS Config Documentation - “Compliance and Monitoring”

AWS Well-Architected Framework - Operational Excellence Pillar

Question: 533

A company is developing a platform to process large volumes of data for complex analytics and machine learning (ML) tasks. The platform must handle compute-intensive workloads. The workloads currently require 20 to 30 minutes for each data processing step.

The company wants a solution to accelerate data processing.

Which solution will meet these requirements with the LEAST operational overhead?

A. Deploy three Amazon EC2 instances. Distribute the EC2 instances across three Availability Zones. Use traditional batch processing techniques for data processing.

- B. Create an Amazon EMR cluster. Use managed scaling. Install Apache Spark to assist with data processing.
- C. Create an AWS Lambda function for each data processing step. Deploy an Amazon Simple Queue Service (Amazon SQS) queue to relay data between Lambda functions.
- D. Create a series of AWS Lambda functions to process the data. Use AWS Step Functions to orchestrate the Lambda functions into data processing steps.

Answer: B

Explanation:

Amazon EMR provides a managed big data framework that supports Apache Spark, which is ideal for distributed and compute-intensive data transformations. Managed scaling dynamically adjusts cluster resources, ensuring high performance with minimal management.

From AWS Documentation:

“Amazon EMR provides a managed environment for big data frameworks such as Apache Spark and Hadoop. With managed scaling, EMR automatically resizes clusters to meet workload demands.”

(Source: Amazon EMR Developer Guide)

Why B is correct:

Provides distributed parallel processing for large datasets.

Reduces operational overhead with managed scaling and auto-termination.

Integrates easily with S3, Glue, and ML pipelines.

Optimized for heavy ETL and analytics workloads.

Why others are incorrect:

A: Manual scaling and limited processing capacity.

C & D: Lambda has execution time and memory limits unsuitable for 30-minute compute-intensive tasks.

Reference:

Amazon EMR Developer Guide - “Using Managed Scaling”

AWS Well-Architected Framework - Performance Efficiency Pillar

Question: 534

A company is designing a website that displays stock market prices to users. The company wants to use Amazon ElastiCache (Redis OSS) for the data caching layer. The company needs to ensure that the website's data caching layer can automatically fail over to another node if necessary.

- A. Enable read replicas in ElastiCache (Redis OSS). Promote the read replica when necessary.
- B. Enable Multi-AZ in ElastiCache (Redis OSS). Fail over to a second node when necessary.
- C. Export a backup of the ElastiCache (Redis OSS) cache to an Amazon S3 bucket. Restore the cache to a second cluster when necessary.
- D. Export a backup of the ElastiCache (Redis OSS) cache by using AWS Backup. Restore the cache to a second cluster when necessary.

Answer: B

Explanation:

For high availability, Amazon ElastiCache for Redis supports Multi-AZ with automatic failover, which provides primary and replica nodes in different Availability Zones. If the primary node fails, Redis automatically promotes a replica to primary.

From AWS Documentation:

“When you enable Multi-AZ with automatic failover, Amazon ElastiCache automatically detects failures and promotes a read replica to primary with minimal downtime.”

(Source: Amazon ElastiCache for Redis User Guide)

Why B is correct:

Multi-AZ provides automatic failover and data replication.

Ensures continuous availability and protects against node or AZ failures.

Fully managed with no manual intervention needed.

Why others are incorrect:

A: Manual promotion is not automatic.

C & D: Restoring from backup is slow and meant for disaster recovery, not failover.

Reference:

Amazon ElastiCache for Redis User Guide - “High Availability with Multi-AZ”

AWS Well-Architected Framework - Reliability Pillar

Question: 535

A healthcare company stores personally identifiable information (PII) data in an Amazon RDS for Oracle database. The company must encrypt the PII data at rest. The company must use dedicated hardware modules to store and manage the encryption keys.

- A. Use AWS Key Management Service (AWS KMS) to configure encryption for the RDS database. Store and manage keys in AWS CloudHSM.
- B. Use AWS CloudHSM backed AWS KMS keys to configure transparent encryption for the RDS database.
- C. Use Amazon EC2 instance store encryption to encrypt database volumes by using AWS CloudHSM backed keys.
- D. Configure RDS snapshots and use server-side encryption with Amazon S3 managed keys (SSE-S3). Store the keys in AWS CloudHSM.

Answer: B

Explanation:

Amazon RDS supports encryption at rest by using AWS KMS keys backed by AWS CloudHSM. This allows use of dedicated FIPS 140-2 Level 3 validated hardware modules to manage encryption keys, meeting compliance for sensitive data such as PII.

From AWS Documentation:

“You can use AWS KMS with keys that are backed by AWS CloudHSM to control the encryption of RDS databases. This provides dedicated HSM-backed key storage and management.”

(Source: Amazon RDS User Guide - Encrypting Amazon RDS Resources)

Why B is correct:

Meets the requirement for dedicated HSM hardware.

Fully integrates with RDS for transparent encryption at rest.

Satisfies compliance standards for healthcare and regulated data.

Why others are incorrect:

A: Keys in CloudHSM directly are not used by RDS; they must be managed through KMS integration.

C: EC2 instance stores are ephemeral, not suitable for RDS databases.

D: SSE-S3 applies to S3 objects, not databases.

Reference:

Amazon RDS User Guide - “Encryption at Rest with AWS KMS and CloudHSM”

AWS CloudHSM User Guide

Question: 536

A solutions architect is designing a system to be highly resilient. The system uses Amazon Route 53 with health checks and an Application Load Balancer (ALB). The system is critical and must have the highest availability possible.

- A. Automate failover to a healthy resource by automatically updating the value of the Route 53 A record.
- B. Configure the Route 53 health checks to perform a failover automatically.
- C. Automate failover to a healthy resource by updating the weight of the Route 53 weighted record.
- D. Create a new ALB during a failover event, and remap the target group to the new ALB.

Answer: B

Explanation:

Amazon Route 53 health checks can automatically perform DNS failover to healthy endpoints. When integrated with an Application Load Balancer, this provides a highly resilient system architecture that automatically routes traffic to healthy resources across Regions or endpoints.

From AWS Documentation:

“Route 53 DNS failover automatically routes traffic away from unhealthy resources and directs it to healthy resources that you specify in DNS records.”

(Source: Amazon Route 53 Developer Guide - DNS Failover)

Why B is correct:

Route 53 health checks automatically monitor endpoint health and switch to a healthy target when the primary endpoint fails. The failover process is automatic, with no manual updates to DNS records required.

Combined with ALB's built-in multi-AZ resilience, this provides maximum availability.

Why others are incorrect:

A & C: Require manual DNS updates, which are not automatic and introduce latency.

D: Creating new ALBs during failover increases downtime and is operationally inefficient.

Reference:

Amazon Route 53 Developer Guide - “Configuring DNS Failover”

AWS Well-Architected Framework - Reliability Pillar

Question: 537

A company needs to create a compliance management solution. The company wants to use a combination of AWS services to achieve the fine-grained visibility that the solution requires. The compliance management solution must provide a centralized method for company employees to review security findings and out-of-compliance findings.

Which solution will meet these requirements with the LEAST ongoing maintenance?

- A. Configure AWS Security Hub to centralize findings. Use conformance packs in Amazon Inspector to check for compliance framework misalignment.
- B. Use AWS Marketplace to purchase a security tool. Install the tool on an Amazon EC2 instance. Assign an EC2 Instance Profile for the tool to gather data from AWS resources.
- C. Configure AWS Security Hub to centralize findings. Use conformance packs in AWS Config to check for compliance framework misalignment.
- D. Configure AWS Systems Manager to provide a centralized dashboard. Use conformance packs in AWS Config to check for compliance framework misalignment.

Answer: C

Explanation:

AWS Security Hub provides a centralized view of security findings across AWS accounts and services. It integrates natively with AWS Config conformance packs, which evaluate compliance against industry standards such as CIS and PCI-DSS.

From AWS Documentation:

“AWS Security Hub aggregates, organizes, and prioritizes security alerts and compliance status across AWS accounts. Use AWS Config conformance packs to assess compliance with security frameworks.”

(Source: AWS Security Hub User Guide - Managing Findings and Compliance)

Why C is correct:

Security Hub provides a centralized dashboard for compliance visibility.

Conformance packs in AWS Config automate compliance checks across accounts.

Fully managed, minimal maintenance, and integrates natively with AWS services.

Why others are incorrect:

A: Conformance packs are not a feature of Amazon Inspector.

B: Third-party tools on EC2 require management and add operational overhead.

D: Systems Manager is not designed for compliance aggregation.

Reference:

AWS Security Hub User Guide - "Compliance Standards and Findings"

AWS Config User Guide - "Conformance Packs Overview"

AWS Well-Architected Framework - Security Pillar

Question: 538

A company wants to migrate hundreds of gigabytes of unstructured data from an on-premises location to an Amazon S3 bucket. The company has a 100-Mbps internet connection on premises. The company needs to encrypt the data in transit to the S3 bucket. The company will store new data directly in Amazon S3.

- A. Use AWS Database Migration Service (AWS DMS) to synchronize the on-premises data to a destination S3 bucket.
- B. Use AWS DataSync to migrate the data from the on-premises location to an S3 bucket.
- C. Use an AWS Snowball Edge device to migrate the data to an S3 bucket. Use an AWS CloudHSM key to encrypt the data on the Snowball Edge device.
- D. Set up an AWS Direct Connect connection between the on-premises location and AWS. Use the `s3 cp` command to move the data directly to an S3 bucket.

Answer: B

Explanation:

AWS DataSync provides a fully managed, secure, and high-performance service for transferring large amounts of data between on-premises storage and Amazon S3. It uses TLS encryption in transit and automates data validation, scheduling, and monitoring.

From AWS Documentation:

"AWS DataSync securely and efficiently transfers large amounts of data online between on-premises storage and AWS services. All data is encrypted in transit using TLS."

(Source: AWS DataSync User Guide - How DataSync Works)

Why B is correct:

Encrypts all data in transit automatically.

Optimized for high-throughput WAN environments (100 Mbps to multi-Gbps).

Fully managed, with no need to provision additional infrastructure.

Integrates natively with S3 and supports incremental syncs.

Why others are incorrect:

A: DMS is designed for database migration, not unstructured data.

C: Snowball is for offline migrations, not needed given available connectivity.

D: Direct Connect is costly for temporary data transfers and unnecessary here.

Reference:

AWS DataSync User Guide - "Security in AWS DataSync"

AWS Well-Architected Framework - Cost Optimization and Operational Excellence Pillars

Question: 539

A company has a web application that stores user transactions in an Amazon DynamoDB table. To comply with regulations, the company must retain a copy of user transaction data for 7 years.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Use AWS Backup to create backup schedules and retention policies for the table.
- C. Create an on-demand backup of the table by using DynamoDB. Store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.
- D. Create an Amazon EventBridge rule to invoke an AWS Lambda function. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: B

Explanation:

AWS Backup offers centralized management, scheduling, and retention of backups for supported AWS services including Amazon DynamoDB. It enables compliance retention with minimal management.

From AWS Documentation:

"AWS Backup provides centralized backup management across AWS services, including DynamoDB.

You can define backup plans, schedules, and retention policies to meet business or regulatory requirements."

(Source: AWS Backup Developer Guide - Backing up DynamoDB Tables)

Why B is correct:

Automatically manages backup scheduling and retention (7 years).

Centralized compliance monitoring via AWS Backup Audit Manager.

Fully managed and requires no custom automation.

Why others are incorrect:

A: PITR is for continuous restore within 35 days, not long-term compliance retention.

C & D: Manual or Lambda-based backups require custom management and increase operational complexity.

Reference:

AWS Backup Developer Guide - "Backing up DynamoDB Tables"

AWS Well-Architected Framework - Operational Excellence and Reliability Pillars

Question: 540

A company has a business system that generates hundreds of reports each day. The business system saves the reports to a network share in CSV format. The company needs to store this data in the AWS Cloud in near-real time for analysis.

- A. Use AWS DataSync to transfer the files to Amazon S3. Create a scheduled task that runs at the end of each day.
- B. Create an Amazon S3 File Gateway. Update the business system to use a new network share from the S3 File Gateway.
- C. Use AWS DataSync to transfer the files to Amazon S3. Create an application that uses the DataSync API in the automation workflow.
- D. Deploy an AWS Transfer for SFTP endpoint. Create a script that checks for new files on the network share and uploads the new files by using SFTP.

Answer: B

Explanation:

Amazon S3 File Gateway (AWS Storage Gateway) exposes an on-premises NFS/SMB file share that durably stores files as S3 objects with local caching, enabling low-latency writes on-premises and asynchronous, near-real-time ingestion into Amazon S3 for analytics. It is purpose-built to "present a file interface backed by Amazon S3" and to "store files as objects in your S3 buckets," so existing applications writing to a network share can transparently land data in S3 without custom scripts or job orchestration. Compared with scheduled

transfers (e.g., end-of-day DataSync), S3 File Gateway continuously uploads new/changed files, better meeting the near-real-time requirement. Transfer Family (SFTP) would require custom polling and client changes, increasing operational burden. DataSync is excellent for bulk or periodic migrations/synchronizations but is not as seamless for continuous ingestion from a live file share. Therefore, updating the application to point to an S3 File Gateway share provides the simplest, performant path to near-real-time delivery into S3 for downstream analytics.

Reference: AWS Storage Gateway User Guide — “What is Amazon S3 File Gateway,” “How S3 File Gateway works,” “File share protocols (SMB, NFS),” “Object creation and upload behavior,” and AWS Well-Architected — Analytics ingestion patterns.

Question: 541

A company hosts its order processing system on AWS. The architecture consists of a frontend and a backend. The frontend includes an Application Load Balancer (ALB) and Amazon EC2 instances in an Auto-Scaling group. The backend includes an EC2 instance and an Amazon RDS MySQL database.

To prevent incomplete or lost orders, the company wants to ensure that order states are always preserved. The company wants to ensure that every order will eventually be processed, even after an outage or pause. Every order must be processed exactly once.

A. Create an Auto Scaling group and an ALB for the backend. Create a read replica for the RDS database in a second Availability Zone. Update the backend RDS endpoint.

B. Create an Auto Scaling group and an ALB for the backend. Create an Amazon RDS proxy in front of the RDS database. Update the backend EC2 instance to use the Amazon RDS proxy endpoint.

C. Create an Auto Scaling group for the backend. Configure the backend EC2 instances to consume messages from an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Configure a dead-letter queue (DLQ) for the SQS queue.

D. Create an AWS Lambda function to replace the backend EC2 instance. Subscribe the function to an Amazon Simple Notification Service (Amazon SNS) topic. Configure the frontend to send orders to the SNS topic.

Answer: C

Explanation:

Use SQS FIFO to durably persist orders, guarantee order processing semantics, and decouple producers/consumers. FIFO queues provide “exactly-once processing” with message deduplication and “preserve message order.” Visibility timeouts and retries ensure messages are “processed eventually” without being lost; failed messages go to a DLQ for later reprocessing. This pattern aligns with Well-Architected reliability guidance to “queue work to protect against overload and failures” and to ensure “durable, idempotent processing” with retry and backoff. ALB/RDS Proxy/read replicas (A, B) improve availability/connection management but do not guarantee durable handoff or exactly-once processing. SNS (D) is pub/sub and does not provide FIFO semantics in this option, nor a DLQ per subscription for exactly-once. Therefore, frontends write orders to an SQS FIFO queue; backend workers in an Auto Scaling group consume, process idempotently, and use a DLQ for poison messages to meet “no lost orders,” “eventual processing,” and “exactly-once” requirements.

Reference: Amazon SQS Developer Guide — FIFO Queues (exactly-once processing, message ordering, deduplication), Dead-Letter Queues; AWS Well-Architected Framework — Reliability Pillar (queue-based load leveling, idempotency, retries).

Question: 542

A healthcare company is running an Amazon EMR cluster on Amazon EC2 instances to process data that is stored in Amazon S3. The company must ensure that the data processing jobs have access only to the relevant data in Amazon S3. Each job must have specific EMR runtime roles.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Set up security configurations in Amazon EMR, and set EnableApplicationScopedIAMRole to true.
- B. Set up runtime roles to assume the EC2 instance profile of the Amazon EMR cluster.
- C. Set up an EC2 instance profile for the Amazon EMR cluster to assume the runtime roles.
- D. For each IAM role that serves as an EMR runtime role, set up a trust policy with the EC2 instance profile role.
- E. Establish a trust policy between the EMR runtime roles and the EMR service role of the cluster.
- F. Set up security configurations in Amazon EMR, and set EnableInTransitEncryption to true.

Answer: A,C,D

Explanation:

Amazon EMR on EC2 supports “runtime roles (application-scoped IAM roles)” so each application/step assumes its own IAM role with least-privilege S3 access. You enable this via an EMR security configuration by setting “EnableApplicationScopedIAMRole = true.” The EMR core/Task nodes run under the cluster's EC2 instance profile; therefore the instance profile must be permitted to “sts:AssumeRole” into the defined EMR runtime roles, and each runtime role must trust the instance profile (trust policy principal is the instance profile role). This design limits each job's S3 scope via role policies and enforces per-job access segregation. Option B reverses the trust (incorrect). Option E trusts the EMR service role (not used to assume runtime roles). Option F is unrelated (encryption in transit). The correct trio is to enable application-scoped roles (A), authorize the instance profile to assume them (C), and configure the runtime roles' trust relationship to allow that assumption (D).

Reference: Amazon EMR Management Guide — EMR Runtime Roles / Application-scoped IAM roles; IAM Roles and Trust Policies; EMR Security Configuration settings.

Question: 543

An ecommerce company hosts an application on AWS across multiple Availability Zones. The application experiences uniform load

throughout most days.

The company hosts some components of the application in private subnets. The components need to access the internet to install and update patches.

A solutions architect needs to design a cost-effective solution that provides secure outbound internet connectivity for private subnets across multiple Availability Zones. The solution must maintain high availability.

A. Deploy one NAT gateway in each Availability Zone. Configure the route table for each private subnet within an Availability Zone to route outbound traffic through the NAT gateway in the same Availability Zone.

B. Place one NAT gateway in a designated Availability Zone within the VPC. Configure the route tables of the private subnets in each Availability Zone to direct outbound traffic specifically through the

NAT gateway for internet access.

C. Deploy an Amazon EC2 instance in a public subnet. Configure the EC2 instance as a NAT instance. Set up the instance with security groups that allow inbound traffic from private subnets and outbound internet access. Configure route tables to direct traffic from the private subnets through the NAT instance.

D. Use one NAT Gateway in a Network Load Balancer (NLB) target group. Configure private subnets in each Availability Zone to route traffic to the NLB for outbound internet access.

Answer: A

Explanation:

AWS guidance for NAT Gateway recommends deploying “a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone.” This provides “zone-independent architecture” and avoids cross-AZ data processing charges and single-AZ failures. Option B creates a single point of failure and incurs cross-AZ egress charges when private subnets in other AZs traverse a centralized NAT. NAT instances (C) are legacy, require manual scaling/failover/patching, and are not recommended for production HA. Option D is not supported (NLB cannot front a NAT Gateway as a target). With steady, uniform load, per-AZ NAT Gateways deliver high availability with predictable cost; routing each private subnet to its local NAT Gateway maintains security (no inbound initiated connections) and resilience. This meets the requirement for cost-effective, secure outbound connectivity across multiple AZs while preserving availability.

Reference: VPC NAT Gateway documentation — Multi-AZ best practices and same-AZ routing; AWS Well-Architected Framework — Reliability and Cost Optimization (avoid single points of failure; minimize cross-AZ data transfer).

Question: 544

A company is building a stock trading application in the AWS Cloud. The company requires a highly available solution that provides low-latency access to block storage across multiple Availability Zones.

- A. Use an Amazon S3 bucket and an S3 File Gateway as shared storage for the application.
- B. Create an Amazon EC2 instance in each Availability Zone. Attach a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume to each EC2 instance. Create a Bash script to sync data between volumes.
- C. Use an Amazon FSx for NetApp ONTAP Multi-AZ file system to access data by using the iSCSI protocol.
- D. Create an Amazon EC2 instance in each Availability Zone. Attach a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume to each EC2 instance. Create a Python script to sync data between volumes.

Answer: C

Explanation:

Amazon FSx for NetApp ONTAP supports Multi-AZ, providing automatic failover between Availability Zones for high availability. It exposes ONTAP LUNs over the iSCSI protocol, delivering shared block storage semantics with low latency and consistent performance to EC2 clients across AZs. This meets the requirement for “highly available” and “low-latency” block access across multiple AZs. S3/S3 File Gateway (A) is object/file, not block storage. EBS (B, D) provides block storage to a single instance in a single AZ; EBS volumes cannot be shared across instances/AZs, and host-side sync scripts add latency, complexity, and do not provide true HA. FSx for ONTAP natively provides synchronous HA pair replication, fast failover, and supports iSCSI multipathing for resilient, performant access suited to latency-sensitive trading workloads.

Reference: Amazon FSx for NetApp ONTAP — Multi-AZ file systems; iSCSI LUNs and host connectivity; High availability and failover behavior; Performance and client access guidance.

Question: 545

An ecommerce company is preparing to deploy a web application on AWS to ensure continuous service for customers. The architecture includes a web application that the company hosts on Amazon EC2 instances, a relational database in Amazon RDS, and static assets that the company stores in Amazon S3.

The company wants to design a robust and resilient architecture for the application.

- A. Deploy Amazon EC2 instances in a single Availability Zone. Deploy an RDS DB instance in the same Availability Zone. Use Amazon S3 with versioning enabled to store static assets.
- B. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones. Deploy a Multi-AZ RDS DB instance. Use Amazon CloudFront to distribute static assets.
- C. Deploy Amazon EC2 instances in a single Availability Zone. Deploy an RDS DB instance in a second Availability Zone for cross-AZ redundancy. Serve static assets directly from the EC2 instances.
- D. Use AWS Lambda functions to serve the web application. Use Amazon Aurora Serverless v2 for the database. Store static assets in

Amazon Elastic File System (Amazon EFS) One Zone-Infrequent Access

(One Zone-IA).

Answer: B

Explanation:

AWS Well-Architected recommends multi-AZ, elastic architectures: use Auto Scaling groups across multiple Availability Zones for EC2 to eliminate single-AZ failure and automatically replace capacity. For relational databases, Amazon RDS Multi-AZ provides synchronous replication and automatic failover for high availability. Serving static content via Amazon CloudFront (with S3 origin) improves resiliency and performance by caching at edge locations and reducing origin load/latency. Options A and C concentrate compute in one AZ and lack resilient static delivery. Option D changes the stack unnecessarily and proposes EFS One Zone-IA for static assets, which is not multi-AZ and is intended for infrequently accessed, single-AZ workloads, reducing resilience. Therefore, B aligns with best practices for high availability, fault tolerance, and global performance.

Reference: AWS Well-Architected Framework — Reliability Pillar (multi-AZ, elasticity); Amazon EC2 Auto Scaling across AZs; Amazon RDS Multi-AZ deployments; Amazon CloudFront benefits and edge caching.

Question: 546

A company is designing a serverless application to process a large number of events within an AWS account. The application saves the events to a data warehouse for further analysis. The application sends incoming events to an Amazon SQS queue. Traffic between the application and the SQS queue must not use public IP addresses.

- A. Create a VPC endpoint for Amazon SQS. Set the queue policy to deny all access except from the VPC endpoint.
- B. Configure server-side encryption with SQS-managed keys (SSE-SQS).
- C. Configure AWS Security Token Service (AWS STS) to generate temporary credentials for resources that access the queue.
- D. Configure VPC Flow Logs to detect SQS traffic that leaves the VPC.

Answer: A

Explanation:

Amazon SQS supports Interface VPC endpoints (AWS PrivateLink), enabling private connectivity from your VPC to SQS without using public IPs, traversing the public Internet, or requiring NAT/IGW. You can restrict access by attaching a queue resource policy that allows only the specific VPC endpoint and denies all other principals/paths, enforcing that all traffic stays on the AWS network. SSE-SQS (B) encrypts data at rest but does not influence network pathing. STS temporary credentials (C) handle authentication/authorization, not routing. VPC Flow Logs (D) are monitoring/visibility and do not prevent public egress. Creating an SQS VPC endpoint and tightening the queue policy satisfies the requirement of no public IP usage while maintaining secure, private access from serverless components in VPC subnets.

Reference: Amazon SQS — VPC endpoints (PrivateLink) and endpoint policies; Amazon SQS queue policies and condition keys;
Security best practices for private access.

Question: 547

A company is building a web application that serves a content management system. The content management system runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones.

Users are constantly adding and updating files, blogs, and other website assets in the content management system.

A solutions architect must implement a solution in which all the EC2 instances share up-to-date website content with the least possible lag time.

A. Update the EC2 user data in the Auto Scaling group lifecycle policy to copy the website assets from the EC2 instance that was launched most recently. Configure the ALB to make changes to the website assets only in the newest EC2 instance.

B. Copy the website assets to an Amazon Elastic File System (Amazon EFS) file system. Configure each EC2 instance to mount the EFS file system locally. Configure the website hosting application to reference the website assets that are stored in the EFS file system.

C. Copy the website assets to an Amazon S3 bucket. Ensure that each EC2 instance downloads the website assets from the S3 bucket to the attached Amazon Elastic Block Store (Amazon EBS) volume. Run the S3 sync command once each hour to keep files up to date.

D. Restore an Amazon Elastic Block Store (Amazon EBS) snapshot with the website assets. Attach the EBS snapshot as a secondary EBS volume when a new EC2 instance is launched. Configure the website hosting application to reference the website assets that are stored in the secondary EBS volume.

Answer: B

Explanation:

Amazon EFS provides a shared, elastic, low-latency file system that can be mounted concurrently by many EC2 instances across multiple Availability Zones, delivering strong read-after-write consistency so all instances see updates almost immediately. This is the standard pattern for CMS-style workloads that require shared, up-to-date assets with minimal lag. Syncing local copies from S3 (C) introduces polling windows and eventual consistency delays; hourly sync is not near-real time. Copying from a “newest instance” (A) is brittle and not scalable. EBS volumes/snapshots (D) are single-instance, single-AZ block devices and not designed for multi-writer sharing across instances/AZs. EFS's multiAZ design and POSIX semantics provide the simplest, most reliable solution with the least operational overhead.

Reference: Amazon EFS — Use cases and benefits; Performance and consistency model; Mount targets across multiple AZs; Shared file storage for web content and CMS.

Question: 548

A company is developing a monolithic Microsoft Windows based application that will run on Amazon EC2 instances. The application will run long data-processing jobs that must not be interrupted. The company has modeled expected usage growth for the next 3 years. The company wants to optimize costs for the EC2 instances during the 3-year growth period.

- A. Purchase a Compute Savings Plan with a 3-year commitment. Adjust the hourly commitment based on the plan recommendations.
- B. Purchase an EC2 Instance Savings Plan with a 3-year commitment. Adjust the hourly commitment based on the plan recommendations.
- C. Purchase a Compute Savings Plan with a 1-year commitment. Renew the purchase and adjust the capacity each year as necessary.
- D. Deploy the application on EC2 Spot Instances. Use an Auto Scaling group with a minimum size of 1 to ensure that the application is always running.

Answer: A

Explanation:

For steady, predictable EC2 usage with potential changes in instance families over time, AWS recommends Savings Plans. Compute Savings Plans “apply to any EC2 instance regardless of region, instance family, operating system, or tenancy,” and also apply to AWS Fargate and AWS Lambda, delivering the most flexibility over a multi-year horizon. A 3-year term provides the highest discount among Savings Plans for long-lived workloads. EC2 Instance Savings Plans are limited to a chosen instance family in a region; as needs evolve (e.g., size or family changes), discounts may not fully apply. Spot Instances are not appropriate for long, interruption-sensitive jobs because Spot capacity can be reclaimed with short notice. Therefore, a Compute Savings Plan (3-year) best matches cost optimization with flexibility for growth and changes.

Reference: AWS Cost Management — Savings Plans (Compute vs. EC2 Instance), EC2 purchasing options guidance, Well-Architected Cost Optimization (choose pricing models to match workload).

Question: 549

A company plans to deploy an application that uses an Amazon CloudFront distribution. The company will set an Application Load Balancer (ALB) as the origin for the distribution. The company wants to ensure that users access the ALB only through the CloudFront distribution.

The company plans to deploy the solution in a new VPC.

Which solution will meet these requirements?

- A. Configure the network ACLs in the subnet where the ALB is deployed to allow inbound traffic only from the public IP addresses of the CloudFront edge locations.

- B. Create a VPC origin for the CloudFront distribution. Set the VPC origin Amazon Resource Name (ARN) to the ARN of the ALB.
- C. Create a security group that allows only inbound traffic from the public IP addresses of the CloudFront edge locations. Associate the security group with the ALB.
- D. Create a VPC origin for the CloudFront distribution. Configure an ALB rule. Set the source IP condition to allow traffic only from the public IP addresses of the CloudFront edge locations.

Answer: C

Explanation:

To ensure ALB access only via CloudFront, AWS prescribes restricting the origin to traffic from CloudFront. For ALB origins, the standard pattern is to allow inbound to the ALB only from CloudFront edge IP ranges using security groups or ALB listener rules. Network ACLs are coarse and stateless and add operational burden. CloudFront does not have a “VPC origin” object; instead, CloudFront references the ALB DNS name. By limiting the ALB's security group to CloudFront IP ranges, direct client access to the ALB is blocked while CloudFront remains permitted. This aligns with security best practices to “restrict origin access to only CloudFront” and use SGs for instance/ALB layer controls.

Reference: Amazon CloudFront Developer Guide — Restricting access to Application Load Balancers; Elastic Load Balancing Security Groups; Well-Architected Security Pillar (protect origins, least privilege at network layer).

Question: 550

A company uses an Amazon RDS MySQL database to store data for several applications. The company wants to understand use patterns for the database so the company can identify opportunities to optimize costs.

A solutions architect needs to analyze the RDS DB instance to identify right-sizing opportunities.

Which solution will meet these requirements with the LEAST effort?

- A. Enable AWS CloudTrail data events. Use Amazon Athena to query CloudTrail events. Right-size the RDS DB instance based on the number of transactions.
- B. Enable Performance Insights for the RDS DB instance. Right-size the RDS DB instance based on the maximum CPU utilization.
- C. Enable AWS X-Ray to understand the transactions that run on the RDS DB instance. Right-size the RDS DB instance based on the number of transactions.
- D. Enable Amazon CloudWatch Logs for the applications. Aggregate the data from Cloud-Watch Logs for all the applications. Right-size the RDS DB instance based on the aggregated logs.

Answer: B

Explanation:

Amazon RDS Performance Insights is a “database performance tuning and monitoring feature” that can be enabled with a few clicks and “helps you quickly assess the load on your database” and identify bottlenecks. It surfaces key metrics, including DB load, top SQL, waits, and host metrics such as CPU utilization, which are commonly used indicators for right-sizing (up or down). This provides the lowest operational effort compared to building log pipelines or querying CloudTrail (which does not capture SQL workload characteristics). AWS X-Ray traces application requests, not database internals, and CloudWatch Logs aggregation requires custom ingestion and analysis. Enabling Performance Insights directly on the RDS instance provides actionable utilization data to right-size with minimal setup.

Reference: Amazon RDS User Guide — Performance Insights overview, metrics and dashboards; Well-Architected Cost Optimization (measure and monitor, right-size resources).

Question: 551

A company wants to run its experimental workloads in the AWS Cloud. The company has a budget for cloud spending. The company's CFO is concerned about cloud spending accountability for each department. The CFO wants to receive notification when the spending threshold reaches 60% of the budget.

Which solution will meet these requirements?

- A. Use cost allocation tags on AWS resources to label owners. Create usage budgets in AWS Budgets. Add an alert threshold to receive notification when spending exceeds 60% of the budget.
- B. Use AWS Cost Explorer forecasts to determine resource owners. Use AWS Cost Anomaly Detection to create alert threshold notifications when spending exceeds 60% of the budget.
- C. Use cost allocation tags on AWS resources to label owners. Use AWS Support API on AWS Trusted Advisor to create alert threshold notifications when spending exceeds 60% of the budget.
- D. Use AWS Cost Explorer forecasts to determine resource owners. Create usage budgets in AWS Budgets. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

Answer: A

Explanation:

AWS recommends using cost allocation tags to associate costs with owners (e.g., department, project) and AWS Budgets to create cost or usage budgets with alerts at defined thresholds (e.g., 60%). Budgets can send notifications via email or Amazon SNS when “Actual or forecasted” spend exceeds the threshold. Cost Explorer forecasts help visualize trends but do not provide budget alerting or ownership attribution by themselves. Trusted Advisor does not generate budget threshold alerts. Therefore, tagging resources for accountability and setting a budget with a 60% alert provides governance and proactive notifications aligned with cost control best practices for experimentation.

Reference: AWS Billing and Cost Management — Cost allocation tags; AWS Budgets (creating cost/usage budgets and alerts); Well-Architected Cost Optimization (cost governance and visibility).

Question: 552

A company runs its application by using Amazon EC2 instances and AWS Lambda functions. The EC2 instances run in private subnets of a VPC. The Lambda functions need direct network access to the EC2 instances for the application to work.

The application will run for 1 year. The number of Lambda functions that the application uses will increase during the 1-year period. The company must minimize costs on all application resources.

Which solution will meet these requirements?

- A. Purchase an EC2 Instance Savings Plan. Connect the Lambda functions to the private sub-nets that contain the EC2 instances.
- B. Purchase an EC2 Instance Savings Plan. Connect the Lambda functions to new public sub-nets in the same VPC where the EC2 instances run.
- C. Purchase a Compute Savings Plan. Connect the Lambda functions to the private subnets that contain the EC2 instances.
- D. Purchase a Compute Savings Plan. Keep the Lambda functions in the Lambda service VPC.

Answer: C

Explanation:

Compute Savings Plans apply to EC2, AWS Fargate, and AWS Lambda usage, maximizing coverage for mixed architectures while retaining flexibility across instance families, regions, OS, and tenancy. For private connectivity, Lambda functions must be configured with VPC access to attach ENIs in the target subnets, enabling direct network access to EC2 in private subnets (no public subnets are required for intra-VPC communication). EC2 Instance Savings Plans only discount EC2 usage and are tied to a specific instance family in a region, reducing flexibility and leaving Lambda costs undiscounted. Keeping Lambda in the service VPC prevents direct access to private EC2 without VPC configuration. Thus, a 1-year Compute Savings Plan plus connecting Lambda to the private subnets minimizes total cost and meets connectivity needs.

Reference: AWS Savings Plans — Compute Savings Plans applicability (EC2, Fargate, Lambda); AWS Lambda Developer Guide — Configuring Lambda to access VPC resources (ENIs, subnets, security groups).

Question: 553

A company is building a serverless web application that will serve customers globally by using REST API endpoints. The application must minimize latency regardless of the application user's geographic location. The initial amount of traffic that the application will handle is unknown.

- A. Deploy an Amazon API Gateway REST API with edge-optimized API endpoints for all customers. Create AWS Lambda functions. Optimize Lambda performance by adjusting the memory settings and configuring provisioned concurrency.
- B. Deploy an Amazon API Gateway REST API with Regional API endpoints for all customers. Create AWS Lambda functions. Optimize Lambda performance by adjusting the memory settings and configuring reserved concurrency.
- C. Deploy an Amazon API Gateway REST API with Regional API endpoints for all customers. Create AWS Lambda functions. Use an HTTP integration to optimize Lambda performance.
- D. Deploy a Network Load Balancer in each AWS Region where customers are located. Create AWS Lambda functions. Optimize Lambda performance by adjusting the memory settings and configuring provisioned concurrency.

Answer: A

Explanation:

For globally distributed consumers of REST APIs, API Gateway edge-optimized endpoints “route requests to the nearest CloudFront edge location, which then forwards them to your API in the [home] Region,” reducing latency for users worldwide and scaling automatically for unknown or spiky traffic. By contrast, Regional endpoints are intended for clients within the same or nearby Region and do not provide global edge acceleration. AWS Lambda provides automatic scaling for serverless backends; latency can be further reduced by right-sizing memory (which also increases CPU) and, when needed, enabling provisioned concurrency to keep functions initialized and eliminate cold starts for critical paths. Using NLBs across Regions is not serverless and adds operational complexity without CloudFront edge acceleration. Therefore, combining API Gateway edge-optimized REST APIs with Lambda meets the requirements of minimal global latency and unknown initial traffic.

Reference: Amazon API Gateway — “Endpoint types for REST APIs (Edge-optimized, Regional, Private)”; Amazon CloudFront integration with edge-optimized APIs; AWS Lambda — “Provisioned concurrency,” “Managing concurrency,” and performance tuning guidance; AWS Well-Architected — Performance Efficiency (use edge services, serverless autoscaling).

Question: 554

A company needs to create an AWS Lambda function that will run in a VPC in the company's primary AWS account. The Lambda function needs to access files that the company stores in an Amazon Elastic File System (Amazon EFS) file system. The EFS file system is located in a secondary AWS account. As the company adds files to the file system, the solution must scale to meet the demand.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a new EFS file system in the primary account. Use AWS DataSync to copy the contents of the original EFS file system to the new EFS file system.
- B. Create a VPC peering connection between the VPCs that are in the primary account and the secondary account.
- C. Create a second Lambda function in the secondary account that has a mount that is configured for the file system. Use the primary account's Lambda function to invoke the secondary account's Lambda function.

D. Move the contents of the file system to a Lambda layer. Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

Answer: B

Explanation:

Amazon EFS is a regional, elastic file system that “scales to petabytes” and can be accessed from “thousands of compute instances” concurrently. You can mount EFS across VPCs and across AWS accounts by providing network connectivity (e.g., VPC peering) to the EFS mount targets and allowing NFS (TCP 2049) in the mount target security group, optionally using EFS access points and a file system policy for cross-account access control. VPC peering has no hourly charge and minimal operational overhead, and EFS automatically scales as files are added—meeting the scalability and cost goals.

Option A duplicates data, incurs DataSync and extra storage costs, and adds sync lag.

Option C adds an extra Lambda hop and complexity without exposing a shared filesystem to the primary function.

Option D is impractical: Lambda layers are immutable artifacts with tight size limits (up to 50 MB compressed/250 MB uncompressed) and are not suited for dynamic, growing file sets.

Reference: Amazon EFS User Guide — “Accessing EFS across VPCs and accounts,” “Mount targets and security groups,” “EFS access points,” and “EFS automatically scales to petabytes.”

Question: 555

A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS

CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message.

The administrator is using an IAM role that has the following IAM policy attached:

What is the cause of the unsuccessful request?

- A. The EC2 instance has a resource-based policy with a Deny statement.
- B. The principal has not been specified in the policy statement.
- C. The "Action" field does not grant the actions that are required to terminate the EC2 instance.
- D. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0.113.0/24.

Answer: D

Explanation:

The IAM policy includes two statements:

- 1 ☐ Allow ec2:TerminateInstances on all resources
- 2 ☒ Deny ec2:TerminateInstances if the request does NOT come from:

192.0.2.0/24

203.0.113.0/24

AWS IAM policy evaluation rules state:

Explicit Deny always overrides Allow.

A request must meet all applicable conditions in the policy to be granted.

aws:SourceIp condition applies when IAM actions are invoked via the public AWS APIs.

In this scenario, the administrator is not originating the request from one of the allowed CIDR blocks, so the Deny condition is triggered, overriding the Allow statement.

Therefore, the operation is blocked with:

403 AccessDenied: explicit deny

This matches the behavior described in the AWS IAM Policy Evaluation Logic used in the Security

Pillar of the Well-Architected Framework:

Explicit Deny > Allow > Default Deny

IP-based Conditions restrict API calls originating outside approved network ranges

Options A/B/C do not accurately reflect this explicit Deny condition behavior.

Question: 556

A company runs a containerized application on a Kubernetes cluster in an on-premises data center.

The company is using a MongoDB database for data storage. The company wants to migrate some of these environments to AWS, but no code changes or deployment method changes are possible at this time. The company needs a solution that minimizes operational overhead.

A. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes for compute and MongoDB on EC2 for data storage.

- B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute and Amazon DynamoDB for data storage.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes for compute and Amazon DynamoDB for data storage.
- D. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute and Amazon DocumentDB (with MongoDB compatibility) for data storage.

Answer: D

Explanation:

EKS lets teams “run Kubernetes on AWS without needing to install and operate your own Kubernetes control plane,” and with AWS Fargate, you “run Kubernetes pods without managing servers,” reducing operational overhead for worker nodes. For the data layer, Amazon DocumentDB (with MongoDB compatibility) “supports the MongoDB API and drivers,” allowing existing MongoDB applications to work without application changes, while the service “automatically scales storage,” provides “high availability across multiple Availability Zones,” automatic backups, and patching. Because the company cannot change code or deployment methods, keeping Kubernetes (EKS) and the MongoDB API (DocumentDB compatibility) is essential. Options A and C either change the orchestrator (to ECS) or the database API (to DynamoDB), which would require code/deployment changes. Option A also leaves you managing EC2 nodes and a self-managed MongoDB on EC2, increasing ops burden. Therefore, EKS on Fargate + Amazon DocumentDB minimizes operations and preserves compatibility.

Reference: Amazon EKS User Guide — “What is Amazon EKS,” “Fargate for EKS (serverless pods)”;

Amazon DocumentDB — “MongoDB compatibility,” “High availability and automatic scaling”; AWS Well-Architected — Operational Excellence (use managed services).

Question: 557

The DNS provider that hosts a company's domain name records is experiencing outages that cause service disruption for a website running on AWS. The company needs to migrate to a more resilient managed DNS service and wants the service to run on AWS.

What should a solutions architect do to rapidly migrate the DNS hosting service?

- A. Create an Amazon Route 53 public hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider.
- B. Create an Amazon Route 53 private hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider.
- C. Create a Simple AD directory in AWS. Enable zone transfer between the DNS provider and AWS Directory Service for Microsoft Active Directory for the domain records.
- D. Create an Amazon Route 53 Resolver inbound endpoint in the VPC. Specify the IP addresses that the provider's DNS will forward DNS

queries to. Configure the provider's DNS to forward DNS queries for the domain to the IP addresses that are specified in the inbound endpoint.

Answer: A

Explanation:

Amazon Route 53 is a highly available and scalable authoritative DNS service. To move a public domain, you create a public hosted zone, import or recreate the zone records, and then update the registrar to use the Route 53 name servers assigned to the zone. Route 53 is globally distributed and designed for rapid propagation and resilience, and supports features such as health checks and routing policies. A private hosted zone (Option B) is resolvable only by VPCs that are associated with it and is not for public internet DNS. Directory Service and zone transfers (Option C) are unrelated to Route 53 public DNS hosting. Route 53 Resolver inbound endpoints (Option D) provide hybrid DNS forwarding for VPC workloads, not authoritative public hosting. Therefore, the fastest AWS-native migration path to a resilient managed DNS is to create a Route 53 public hosted zone and import the existing zone file.

Reference: Amazon Route 53 Developer Guide — “Public hosted zones,” “Migrating DNS to Amazon Route 53,” “Updating domain registrar with Route 53 name servers”; AWS Well-Architected — Reliability (use managed, highly available DNS).

Question: 558

A company runs a Java-based job on an Amazon EC2 instance. The job runs every hour and takes 10 seconds to run. The job runs on a scheduled interval and consumes 1 GB of memory. The CPU utilization of the instance is low except for short surges during which the job uses the maximum CPU available. The company wants to optimize the costs to run the job.

- A. Use AWS App2Container (A2C) to containerize the job. Run the job as an Amazon Elastic Container Service (Amazon ECS) task on AWS Fargate with 0.5 virtual CPU (vCPU) and 1 GB of memory.
- B. Copy the code into an AWS Lambda function that has 1 GB of memory. Create an Amazon EventBridge scheduled rule to run the code each hour.
- C. Use AWS App2Container (A2C) to containerize the job. Install the container in the existing Amazon Machine Image (AMI). Ensure that the schedule stops the container when the task finishes.
- D. Configure the existing schedule to stop the EC2 instance at the completion of the job and restart the EC2 instance when the next job starts.

Answer: B

Explanation:

AWS Lambda is event driven and “scales automatically with the number of requests,” charging per request and compute duration in milliseconds; you can schedule with Amazon EventBridge rules. Allocating 1 GB memory proportionally increases available CPU, which suits

short CPU bursts. For a 10-second, once-per-hour job, Lambda eliminates instance idle time, providing the lowest cost and no servers to manage. Fargate tasks (A) bill per-second with a 1-minute minimum and require container packaging; for a 10-second job, costs are higher than Lambda. Options C and D keep EC2 management overhead and either maintain instances or add start/stop orchestration complexity; EC2 costs dominate because the instance sits idle for 59+ minutes each hour. Therefore, Lambda + EventBridge is the most cost-effective and operationally simple solution.

Reference: AWS Lambda Developer Guide — pricing (per-ms), automatic scaling, EventBridge scheduling; Well-Architected Cost Optimization — match pricing model to workload and remove idle capacity.

Question: 559

A company runs a production application on a fleet of Amazon EC2 instances. The application reads messages from an Amazon Simple Queue Service (Amazon SQS) queue and processes the messages in parallel. The message volume is unpredictable and highly variable.

The company must ensure that the application continually processes messages without any downtime.

Which solution will meet these requirements MOST cost-effectively?

- A. Use only Spot Instances to handle the maximum capacity required.
- B. Use only Reserved Instances to handle the maximum capacity required.
- C. Use Reserved Instances to handle the baseline capacity. Use Spot Instances to provide additional capacity when required.
- D. Use Reserved Instances in an EC2 Auto Scaling group to handle the minimum capacity. Configure an auto scaling policy that is based on the SQS queue backlog.

Answer: C

Explanation:

AWS guidance is to cover steady baseline with commitments (Reserved Instances or Savings Plans) and use EC2 Spot Instances for burst capacity to minimize cost. Spot provides up to 90% discounts and is well-suited to fault-tolerant, queue-based workloads; interruptions are handled by replacing capacity automatically while messages remain durably in SQS. Using only Spot (A) risks capacity gaps; only RIs sized for peak (B) wastes cost during low demand. Option D scales on backlog but uses On-Demand for bursts, which is more expensive than Spot. With C, baseline capacity (RIs) keeps processing continuously (no downtime), and Spot adds cost-efficient throughput during spikes, aligning with Well-Architected cost and reliability patterns for queue workers.

Reference: EC2 Spot Best Practices — burst with Spot for interruptible/queued workloads; SQS — durable buffering; Well-Architected Cost Optimization — cover steady state with commitments, scale bursts with discounted capacity.

Question: 560

A company is designing an application on AWS that provides real-time dashboards. The dashboard data comes from on-premises databases that use a variety of schemas and formats. The company needs a solution to transfer and transform the data to AWS with minimal latency.

Which solution will meet these requirements?

- A. Integrate the dashboard with Amazon Managed Streaming for Apache Kafka (Amazon MSK) to transfer and transform the data from the on-premises databases to the dashboards.
- B. Use Amazon Data Firehose to transfer the data to an Amazon S3 Bucket. Configure the dashboard application to import new data from the S3 bucket periodically.
- C. Use AWS Database Migration Service (AWS DMS) Schema Conversion to consolidate the onpremises databases into a single AWS database. Use an AWS Lambda function that is scheduled by Amazon EventBridge to transfer data from the consolidated database to the dashboard application.
- D. Use AWS DataSync to transfer data from the source databases to the dashboard application continuously. Configure the dashboard application to import data from DataSync.

Answer: A

Explanation:

Amazon MSK is a fully managed, highly available Apache Kafka service for streaming data with low latency. Kafka Connect and stream processors enable ingest from heterogeneous sources and perform in-stream transformation before delivery to consumers (e.g., the dashboard service). This satisfies real-time updates from diverse schemas and formats. Kinesis alternatives could work, but among the given choices, MSK is the only streaming option designed for sub-second, continuous pipelines. Kinesis Data Firehose (B) buffers and batches data to S3 and is optimized for delivery to storage, not low-latency dashboards. AWS DMS schema conversion (C) focuses on database migration, not ongoing real-time, multi-format streaming for dashboards. AWS DataSync (D) is for file/object transfer, not database change streams. Hence, MSK best meets minimal-latency, transform-in-flight needs with managed operations.

Reference: Amazon MSK — real-time streaming, low latency, Kafka Connect/Streams for transformations; Well-Architected Performance Efficiency — use streaming for real-time analytics.

Question: 561

An ecommerce company runs applications in AWS accounts that are part of an organization in AWS Organizations. The applications run on Amazon Aurora PostgreSQL databases across all the accounts. The company needs to prevent malicious activity and must identify abnormal failed and incomplete login attempts to the databases.

- A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts.
- B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization.

- C. Publish the Aurora general logs to a log group in Amazon CloudWatch Logs. Export the log data to a central Amazon S3 bucket.
- D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket.

Answer: B

Explanation:

Amazon GuardDuty includes RDS Protection that “monitors and profiles access to your Amazon Aurora and Amazon RDS databases” to detect threats such as suspicious login attempts, brute-force activity, and anomalous authentication patterns. GuardDuty can be enabled organization-wide in AWS Organizations with a delegated administrator to centralize findings for all member accounts, minimizing operational overhead. Findings include context like source IP, user, and DB instance, and integrate with Amazon EventBridge for alerting and automated response. SCPs (A) enforce or deny API permissions but do not provide detection/analytics. Exporting general logs (C) requires building and maintaining custom parsing/analytics pipelines. CloudTrail (D) records AWS control-plane API calls and does not log database-level login attempts. Therefore, enabling GuardDuty RDS Protection across the org provides the most operationally efficient, managed detection of abnormal failed and incomplete login attempts.

Reference: Amazon GuardDuty — RDS Protection findings; GuardDuty organizational administration; Amazon RDS/Aurora — integration with GuardDuty; AWS Security best practices (centralized threat detection).

Question: 562

A company that uses AWS Organizations runs 150 applications across 30 different AWS accounts. The company used AWS Cost and Usage Report to create a new report in the management account. The report is delivered to an Amazon S3 bucket that is replicated to a bucket in the data collection account.

The company's senior leadership wants to view a custom dashboard that provides NAT gateway costs each day starting at the beginning of the current month.

Which solution will meet these requirements?

A. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure QuickSight to use AWS DataSync to query the new report.

B. Share an Amazon QuickSight dashboard that includes the requested table visual. Configure

QuickSight to use Amazon Athena to query the new report.

C. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use AWS DataSync to query the new report.

D. Share an Amazon CloudWatch dashboard that includes the requested table visual. Configure CloudWatch to use Amazon Athena to query the new report.

Answer: B

Explanation:

The AWS Cost and Usage Report (CUR) delivers detailed, line-item billing data to Amazon S3. AWS recommends querying CUR with Amazon Athena by creating external tables over the CUR S3 location (partitioned by time) to produce daily cost aggregations such as NAT Gateway (EC2:NatGateway) usage and cost. Amazon QuickSight natively connects to Athena as a data source to build and share dashboards with visuals (tables, time series) filtered from the start of the current month. DataSync (A, C) is a file transfer service and cannot query data. CloudWatch dashboards (C, D) visualize metrics/logs, not CUR datasets. Therefore, using Athena to query CUR and QuickSight to present a daily NAT gateway cost dashboard is the most direct and operationally efficient approach.

Reference: CUR — querying with Amazon Athena; QuickSight — Athena data source; Cost categories and service/usage type fields for NAT Gateway; AWS Cost Management best practices.

Question: 563

A company uses an Amazon CloudFront distribution to serve content pages for its website. The company needs to ensure that clients use a TLS certificate when accessing the company's website. The company wants to automate the creation and renewal of the TLS certificates.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Use a CloudFront security policy to create a certificate.
- B. Use a CloudFront origin access control (OAC) to create a certificate.
- C. Use AWS Certificate Manager (ACM) to create a certificate. Use DNS validation for the domain.
- D. Use AWS Certificate Manager (ACM) to create a certificate. Use email validation for the domain.

Answer: C

Explanation:

AWS Certificate Manager (ACM) issues and automatically renews public TLS certificates used by Amazon CloudFront. With DNS validation, ACM creates CNAME records that prove domain control; once validated, renewals occur automatically without manual approval, providing the lowest operational effort. For CloudFront, ACM certificates must be in the US East (N. Virginia) Region. CloudFront security policies (A) configure protocol/cipher requirements, not certificate issuance. OAC (B) secures origin access and is unrelated to certificate management. Email validation (D) requires mailbox approvals and operational handling at issuance/renewal, which is less efficient than DNS validation. Thus, ACM with DNS validation delivers automated creation and renewal with minimal overhead.

Reference: AWS Certificate Manager — public certificates, DNS validation, automatic renewal; CloudFront — using ACM certificates (us-east-1 requirement), TLS configuration.

Question: 564

A company hosts an application that allows authorized users to upload and download documents. The application uses Amazon EC2 instances and an Amazon Elastic File System (Amazon EFS) file system.

The company plans to deploy the application into a second AWS Region. The company will launch a new EFS file system and a new set of EC2 instances in the second Region. A solutions architect must develop a highly available and fault-tolerant solution to establish two-way synchronization across the Regions.

Which solution will meet these requirements?

- A. Create an Amazon EFS VPC endpoint for the original EFS file system in the second Region. Mount both the original and the new EFS file system to the new set of EC2 instances in the second Region. Configure an rsync cron job to run every 5 minutes.
- B. Set up EFS replication between the two EFS file systems. Set the new file system as the source. Set the original file system in the first Region as the destination. Turn off overwrite protection for the destination file system.
- C. Set up one AWS DataSync agent in each Region. Configure Amazon EFS VPC endpoints, EFS transfer locations, and EFS transfer tasks with opposite directions on the two DataSync agents.
- D. Mount the EFS file system in the second Region to the new set of EC2 instances in the second Region. Use AWS Transfer Family to establish SFTP access to the EFS file system in the original Region. Configure an rsync cron job to run every 5 minutes.

Answer: C

Explanation:

AWS DataSync provides managed, incremental, parallelized transfers between EFS file systems across Regions, supporting scheduled or continuously running tasks with automatic change detection, encryption in transit, and integrity verification. You can configure two tasks in opposite directions (bidirectional) to achieve two-way synchronization with high availability using agents in each Region and EFS locations connected via VPC endpoints. Native EFS replication (B) is one-way (read-only target) and not intended for active/active two-way sync. Options A and D rely on custom rsync and cross-Region mounts or SFTP, which introduce operational overhead, latency, and fail to provide resilient, managed synchronization. DataSync minimizes operational burden while delivering fault-tolerant, scalable, cross-Region EFS sync.

Reference: AWS DataSync — EFS-to-EFS transfers, scheduling, incremental sync; Amazon EFS — replication overview (one-way), multi-Region patterns; Well-Architected Reliability — use managed data transfer services.

Question: 565

A company runs an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales based on CPU utilization metrics. The ecommerce application stores the transaction data in a MySQL 8.0 database that is hosted on a large EC2 instance.

The database's performance degrades quickly as application load increases. The application handles more read requests than write transactions. The company wants a solution that will automatically scale the database to meet the demand of unpredictable read

workloads while maintaining high availability.

- A. Use Amazon Redshift with a single node for leader and compute functionality.
- B. Use Amazon RDS with a Single-AZ deployment. Configure Amazon RDS to add reader instances in a different Availability Zone.
- C. Use Amazon Aurora with a Multi-AZ deployment. Configure Aurora Auto Scaling with Aurora Replicas.
- D. Use Amazon ElastiCache (Memcached) with EC2 Spot Instances.

Answer: C

Explanation:

Amazon Aurora MySQL-compatible offers a distributed, fault-tolerant storage system with Multi-AZ high availability and supports Aurora Replicas that “share the same underlying storage” for low-latency reads. Aurora provides Aurora Auto Scaling to “automatically add or remove Aurora Replicas based on load,” ideal for unpredictable, read-heavy workloads. This architecture offloads reads from the writer and maintains HA through automatic failover. Amazon RDS Single-AZ (B) lacks HA. Redshift (A) is a data warehouse, not a transactional DB. ElastiCache (D) can reduce read pressure but does not provide durable read replicas or automatic HA scaling at the database tier. Aurora's design directly addresses the requirement to automatically scale reads while maintaining availability, matching Well-Architected guidance to use managed, elastic services for variable demand.

Reference: Amazon Aurora User Guide — “Aurora Replicas,” “Aurora Auto Scaling,” “High availability and durability”; AWS Well-Architected Framework — Performance Efficiency, Reliability (managed, elastic databases).

Question: 566

A company wants to migrate from an on-premises data center to AWS. The data center hosts a storage server that stores data in an NFS-based file system. The storage server stores 200 GB of data. The company needs to migrate the data without interruption to existing services. Multiple resources in AWS must be able to access the data by using the NFS protocol.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Create an Amazon FSx for Lustre file system.
- B. Create an Amazon Elastic File System (Amazon EFS) file system.
- C. Create an Amazon S3 bucket to receive the data.
- D. Create an Amazon FSx for Windows file system.
- E. Install an AWS DataSync agent in the on-premises data center. Use a DataSync task between the on-premises file system and the AWS file system.

Answer: B,E

Explanation:

Amazon EFS is a “fully managed, elastic, shared file system” that provides NFSv4 access across many EC2 instances and AZs, ideal for lift-and-shift of Linux NFS workloads. AWS DataSync “automates moving data between on-premises NFS servers and Amazon EFS,” performing incremental, parallel transfers with encryption and integrity checks, enabling cutover with minimal disruption. FSx for Lustre (A) targets HPC and S3-backed workloads, not general NFS shares. FSx for Windows (D) exposes SMB, not NFS. S3 (C) is object storage and not directly NFS-accessible by multiple EC2 hosts. EFS + DataSync satisfies NFS compatibility, multi-client access, and low-touch migration while remaining cost-effective for a 200-GB dataset.

Reference: Amazon EFS — “Using NFS mounts,” “Use cases”; AWS DataSync — “Transfer between onpremises NFS and Amazon EFS,” “Incremental sync and cutover.”

Question: 567

The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After successful scheduling, this application stores the meeting information in an Amazon DynamoDB database.

As the company expands, customers report that their meeting invitations are taking longer to arrive.

What should a solutions architect recommend to resolve this issue?

- A. Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.
- B. Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.
- C. Add an Amazon CloudFront distribution. Set the origin as the web application that accepts the appointment requests.
- D. Add an Auto Scaling group for the application that sends meeting invitations. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Answer: D

Explanation:

For SQS-backed worker architectures, AWS recommends scaling consumers based on queue metrics (e.g., `ApproximateNumberOfMessagesVisible`, `AgeOfOldestMessage`). EC2 Auto Scaling can “scale out based on Amazon CloudWatch alarms” tied to SQS backlog, increasing worker capacity to reduce latency as demand grows. DAX (A) accelerates DynamoDB reads, not message processing. API Gateway (B) and CloudFront (C) optimize request ingress and content delivery, not the asynchronous email-sending

application. The bottleneck is consumer throughput; scaling workers with an SQS- driven policy restores timely processing without downtime and follows Well-Architected patterns for decoupled, elastic systems.

Reference: Amazon SQS Developer Guide — “Scaling consumers,” “Queue metrics”; EC2 Auto Scaling — “Target tracking and step scaling with CloudWatch metrics”; Well-Architected — Performance Efficiency (queue-based load leveling).

Question: 568

A company needs a solution to integrate transaction data from several Amazon DynamoDB tables into an existing Amazon Redshift data warehouse. The solution must maintain the provisioned throughput of DynamoDB.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon S3 bucket. Configure DynamoDB to export to the bucket on a regular schedule. Use an Amazon Redshift COPY command to read from the S3 bucket.
- B. Use an Amazon Redshift COPY command to read directly from each DynamoDB table.
- C. Create an Amazon S3 bucket. Configure an AWS Lambda function to read from the DynamoDB tables and write to the S3 bucket on a regular schedule. Use Amazon Redshift Spectrum to access the data in the S3 bucket.
- D. Use Amazon Athena Federated Query with a DynamoDB connector and an Amazon Redshift connector to read directly from the DynamoDB tables.

Answer: A

Explanation:

DynamoDB export to Amazon S3 (via point-in-time exports or table backups) “does not consume read capacity” and has no impact on table performance. Once data is in S3, Amazon Redshift can ingest efficiently with COPY from S3, which is the standard, parallel, high-throughput load path with minimal management. Direct COPY from DynamoDB (B) performs a table scan that can consume provisioned throughput, risking throttling. Building Lambda pipelines (C) adds custom code and scheduling overhead. Athena federated queries (D) are ad hoc analytics, not optimized for bulk, recurring warehouse loads. Therefore, exporting to S3 and loading with Redshift COPY maintains DynamoDB throughput and minimizes operational burden.

Reference: DynamoDB — “Export to S3 (no RCU consumption)”; Amazon Redshift — “COPY from Amazon S3 (parallel load)”; AWS Big Data best practices — “Stage data in S3 for Redshift loads.”

Question: 569

A media company stores customer-uploaded videos in an Amazon S3 bucket with the Standard storage class. The company wants to create an S3 Lifecycle configuration. The company will set the maximum retention time to 7 days. However, the configuration must delete any video that is more than 1 TB in size after 48 hours.

- A. Create a single S3 Lifecycle configuration that has two rules. Configure the first rule to expire objects after 48 hours with a filter of

ObjectSizeGreaterThan and a value of 1 TB. Configure the second rule to expire objects after 7 days.

B. Create two S3 Lifecycle configurations. Include a rule in the first configuration to expire objects after 48 hours by using a Prefix filter of LargeFiles. Include a rule in the second configuration to expire objects after 7 days.

C. Create a single S3 Lifecycle configuration that has two rules. Configure the first rule to expire objects after 48 hours. Configure the second rule to expire objects after 7 days.

D. Create two S3 Lifecycle configurations. Include a rule in the first configuration to expire objects after 48 hours. Include a rule in the second configuration to expire objects after 7 days by using a filter of ObjectSizeLessThan and a value of 1 TB.

Answer: A

Explanation:

Amazon S3 supports one Lifecycle configuration per bucket that can contain multiple rules. Lifecycle rules can include filters, including object size filters: ObjectSizeGreaterThan and ObjectSizeLessThan. You can combine rules so that one targets large objects and another provides a default expiration. Here, configure Rule 1 with a size filter ObjectSizeGreaterThan = 1 TB and Expiration = 2 days (48 hours) to purge very large videos early. Configure Rule 2 with Expiration = 7 days without a size filter to serve as a catch-all maximum retention for all objects. Options B and D are invalid because S3 does not allow multiple Lifecycle configurations per bucket. Option C cannot distinguish large files; it would delete all files at the same schedule without honoring the 48-hour policy for >1 TB objects. This single configuration with two rules meets both retention targets with minimal overhead and full S3-native capability.

Question: 570

A company hosts a web application on an on-premises server that processes incoming requests. Processing time for each request varies from 5 minutes to 20 minutes.

The number of requests is growing. The company wants to move the application to AWS. The company wants to update the architecture to scale automatically.

A. Convert the application to a microservices architecture that uses containers. Use Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type to run the containerized web application. Configure Service Auto Scaling. Use an Application Load Balancer to distribute incoming requests.

B. Create an Amazon EC2 instance that has sufficient CPU and RAM capacity to run the application. Create metrics to track usage. Create alarms to notify the company when usage exceeds a specified threshold. Replace the EC2 instance with a larger instance size in the same family when usage is too high.

C. Refactor the web application to use multiple AWS Lambda functions. Use an Amazon API Gateway REST API as an entry point to the Lambda functions.

D. Refactor the web application to use a single AWS Lambda function. Use an Amazon API Gateway HTTP API as an entry point to

the Lambda function.

Answer: A

Explanation:

AWS Fargate runs containers without managing servers, and ECS Service Auto Scaling adjusts the number of tasks based on demand. Behind an Application Load Balancer, ECS services elastically scale out to handle concurrent long-running requests. This fits workloads with processing times of 520 minutes per request. Lambda (C, D) has a maximum 15-minute timeout, so a portion of requests would fail or require complex refactoring and orchestration. Option B relies on vertical scaling and manual intervention, leaving capacity idle and failing to scale automatically with spikes. By containerizing the app and running on ECS with Fargate, the company gains automatic scaling, isolation per task, and no instance management, ensuring reliable throughput as request volume grows while keeping operations minimal.

Question: 571

A company deploys an application on Amazon EC2 Spot Instances. The company observes frequent unavailability issues that affect the application's output. The application instances all use the same instance type in a single Availability Zone. The application architecture does not require the use of any specific instance family.

The company needs a solution to improve the availability of the application.

Which combination of steps will meet this requirement MOST cost-effectively? (Select THREE.)

- A. Create an EC2 Auto Scaling group that includes a mix of Spot Instances and a base number of On- Demand Instances.
- B. Create EC2 Capacity Reservations.
- C. Use the lowest price allocation strategy for Spot Instances.
- D. Specify similarly sized instance types and Availability Zones for the Spot Instances.
- E. Use a different instance type for the web application.
- F. Use the price capacity optimized strategy for Spot Instances.

Answer: A,D,F

Explanation:

AWS Spot best practices recommend diversifying capacity across multiple instance types and Availability Zones and using the capacity-optimized (price-capacity-optimized) allocation strategy to choose pools with the deepest capacity for higher availability. Adding a small

On-Demand base in the Auto Scaling group maintains steady, uninterrupted baseline processing while keeping costs low and absorbing Spot interruptions. Option C (lowest price) increases interruption risk. Capacity Reservations (B) target On-Demand capacity guarantees and add cost, not needed for Spot-based elasticity. Option E is redundant; diversification is already achieved with (D). This combination maximizes resiliency of Spot workloads while preserving strong cost efficiency and aligns with AWS guidance for fault-tolerant, stateless applications on Spot.

Question: 572

An advertising company stores terabytes of data in an Amazon S3 data lake. The company wants to build its own foundation model (FM) and has deployed a training cluster on AWS. The company loads file-based data from Amazon S3 to the training cluster to train the FM. The company wants to reduce data loading time to optimize the overall deployment cycle.

The company needs a storage solution that is natively integrated with Amazon S3. The solution must be scalable and provide high throughput.

Which storage solution will meet these requirements?

A. Mount an Amazon Elastic File System (Amazon EFS) file system to the training cluster. Use AWS DataSync to migrate data from Amazon S3 to the EFS file system to train the FM.

B. Use an Amazon FSx for Lustre file system and Amazon S3 with Data Repository Association (DRA).

Preload the data from Amazon S3 to the Lustre file system to train the FM.

C. Attach Amazon Block Store (Amazon EBS) volumes to the training cluster. Load the data from Amazon S3 to the EBS volumes to train the FM.

D. Use AWS DataSync to migrate the data from Amazon S3 to the training cluster as files. Train the FM on the local file-based data.

Answer: B

Explanation:

Amazon FSx for Lustre is a high-performance parallel file system designed for machine learning and HPC that can be linked to Amazon S3 via Data Repository Associations (DRA). With DRA, you can import S3 objects as files (lazy or preloaded) and export results back to S3, providing very high throughput and low-latency POSIX access on the training cluster. This S3-native integration minimizes data loading overhead and accelerates training cycles at scale. EFS (A) is general-purpose and lower throughput per TB than Lustre. EBS (C) requires manual copy and does not scale as a shared parallel filesystem. DataSync alone (D) is a transfer tool, not a high-throughput training filesystem. FSx for Lustre with S3 DRA best satisfies scalability, throughput, and native S3 integration for rapid ML training.

Question: 573

A company uses Amazon Route 53 as its DNS provider. The company hosts a website both on premises and in the AWS Cloud. The

company's on-premises data center is near the us-west-1 Region. The company hosts the website on AWS in the eu-central-1 Region.

The company wants to optimize load times for the website as much as possible.

Which solution will meet these requirements?

- A. Create a DNS record with a failover routing policy that routes all primary traffic to eu-central-1. Configure the routing policy to use the on-premises data center as the secondary location.
- B. Create a DNS record with an IP-based routing policy. Configure specific IP ranges to return the value for the eu-central-1 website. Configure all other IP ranges to return the value for the onpremises website.
- C. Create a DNS record with a latency-based routing policy. Configure one latency record for the eu- central-1 website and one latency record for the on-premises data center. Associate the record for the on-premises data center with the us-west-1 Region.
- D. Create a DNS record with a weighted routing policy. Split the traffic evenly between eu-central-1 and the on-premises data center.

Answer: C

Explanation:

Comprehensive and Detailed Explanation (paraphrased from AWS Solutions Architect guidance):

Amazon Route 53 latency-based routing is specifically designed to route users to the endpoint that provides the lowest latency, based on actual latency measurements between AWS Regions and users' networks. This directly supports the goal of optimizing load times (performance).

For non-AWS endpoints (like an on-premises data center), Route 53 lets you create a latency routing record and associate it with the AWS Region that is closest to that endpoint. In this case, the onpremises data center is near us-west-1, so you associate that latency record with the us-west-1 Region.

Users for whom the eu-central-1 Region has lower latency will be routed to the AWS website in eu- central-1.

Users for whom the us-west-1-associated endpoint has lower latency will be routed to the onpremises website.

This matches AWS best practices for high-performance, low-latency architectures when you have multiple endpoints in different geographic locations, including a mix of AWS and on-premises resources.

Why the other options are not correct:

Option A (Failover routing):

Failover routing is designed for high availability and disaster recovery (active-passive), not for optimizing latency.

All traffic goes to the primary endpoint (eu-central-1) unless it is unhealthy. That does not choose the lowest-latency endpoint for each user, so it does not meet the “optimize load times as much as possible” requirement.

Option B (IP-based routing):

IP-based / geo-based policies are about directing users based on location information (IP/geo), not measured latency.

You would have to manually decide which IP ranges get which endpoint, which does not guarantee that each user goes to the lowest-latency endpoint. It's less precise and less dynamic than latencybased routing.

Option D (Weighted routing):

Weighted routing splits traffic according to fixed weights (e.g., 50/50), regardless of user location or latency.

This is useful for testing, blue/green deployment, or gradual traffic shifting, but not for performance optimization per user.

Therefore, the only option that directly aligns with AWS guidance for performance optimization and latency-based traffic steering across Regions and on-premises endpoints is:

C. Create a DNS record with a latency-based routing policy... associate the on-premises endpoint with us-west-1.

Question: 574

A company has an application that uses an Amazon DynamoDB table for storage. A solutions architect discovers that many requests to the table are not returning the latest data. The company's users have not reported any other issues with database performance. Latency is in an acceptable range.

Which design change should the solutions architect recommend?

- A. Add read replicas to the table.
- B. Use a global secondary index (GSI).
- C. Request strongly consistent reads for the table.
- D. Request eventually consistent reads for the table.

Answer: C

Explanation:

Comprehensive and Detailed Explanation (from AWS Solutions Architect documentation):

Amazon DynamoDB supports two read consistency models: eventually consistent reads and strongly consistent reads. When an application must always receive the most up-to-date data, it must use strongly consistent reads. Eventually consistent reads can return stale data

because they provide best-effort propagation across storage nodes.

Since the problem states that requests are “not returning the latest data,” the correct solution is to enable strongly consistent reads, which immediately reflect all successful writes. This is the explicit AWS-recommended method for ensuring clients always receive the most current version of an item.

Question: 575

A solutions architect must design a database solution for a high-traffic ecommerce web application. The database stores customer profiles and shopping cart information. The database must support a peak load of several million requests each second and deliver responses in milliseconds. The operational overhead for managing and scaling the database must be minimized.

Which database solution should the solutions architect recommend?

- A. Amazon Aurora
- B. Amazon DynamoDB
- C. Amazon RDS
- D. Amazon Redshift

Answer: B

Explanation:

Comprehensive and Detailed Explanation (from AWS Solutions Architect documentation):

Amazon DynamoDB is a fully managed NoSQL database engineered for extremely high throughput and millisecond-level response times at any scale. It can automatically scale to support millions of requests per second and requires no management of infrastructure, storage provisioning, or server maintenance.

AWS documentation specifically highlights DynamoDB as the optimal solution for ecommerce workloads that require consistent low-latency performance, massive scalability, and near-zero operational overhead.

Aurora and RDS cannot support millions of requests per second, and Redshift is a data warehouse not intended for transactional workloads.

Question: 576

A company has deployed resources in the us-east-1 Region. The company also uses thousands of

AWS Outposts servers deployed at remote locations around the world. These Outposts servers regularly download new software versions

from us-east-1 that consist of hundreds of files. The company wants to improve the latency of the software download process.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket in us-east-1. Configure the bucket for static website hosting. Use bucket policies and ACLs to provide read access to the Outposts servers.
- B. Create an Amazon S3 bucket in us-east-1 and a second bucket in us-west-2. Configure replication. Set up a CloudFront distribution with origin failover between the buckets. Download by using signed URLs.
- C. Create an Amazon S3 bucket in us-east-1. Configure S3 Transfer Acceleration. Configure the Outposts servers to download by using the acceleration endpoint.
- D. Create an Amazon S3 bucket in us-east-1. Set up a CloudFront distribution using all edge locations with caching enabled. Configure the bucket as the origin. Download the software by using signed URLs.

Answer: D

Explanation:

Amazon CloudFront uses a globally distributed network of edge locations that cache content close to users. When Outposts servers around the world download large software packages, CloudFront provides significantly reduced latency due to edge caching. This is the AWS-recommended solution for accelerating downloads of static files at scale and across global locations.

S3 Transfer Acceleration optimizes uploads and downloads to a single Region but does not provide edge caching. Multi-Region replication with failover does not reduce latency globally because requests still must reach the regional origins. Therefore, CloudFront with caching enabled is the correct design for improving download speed worldwide.

Question: 577

A transaction-processing company has weekly batch jobs that run on Amazon EC2 instances in an Auto Scaling group. Transaction volume varies, but CPU utilization is always at least 60% during the batch runs. Capacity must be provisioned 30 minutes before the jobs begin.

Engineers currently scale the Auto Scaling group manually. The company needs an automated solution but cannot allocate time to analyze scaling trends.

Which solution will meet these requirements with the least operational overhead?

- A. Create a dynamic scaling policy based on CPU utilization at 60%.
- B. Create a scheduled scaling policy. Set desired, minimum, and maximum capacity. Set recurrence weekly. Set the start time to 30

minutes before the jobs run.

- C. Create a predictive scaling policy that forecasts CPU usage and pre-launches instances 30 minutes before the jobs run.
- D. Create an EventBridge rule that invokes a Lambda function when CPU reaches 60%. The Lambda function increases the Auto Scaling group size by 20%.

Answer: C

Explanation:

Predictive scaling automatically analyzes historical workload patterns, forecasts future capacity needs, and launches instances ahead of time. AWS documentation states that predictive scaling is designed for workloads with recurring, cyclical patterns—such as scheduled weekly batch jobs.

It also supports "pre-launching" capacity before peak demand. This eliminates manual trend analysis and delivers the lowest operational overhead.

Scheduled scaling (Option B) works but requires manual calculation of capacity numbers and updating if patterns change. Predictive scaling removes this burden entirely.

Question: 578

A company is building an application composed of multiple microservices that communicate over HTTP. The company must deploy the application across multiple AWS Regions to meet disaster recovery requirements. The application must maintain high availability and automatic fault recovery.

Which solution will meet these requirements?

- A. Deploy all microservices on a single large EC2 instance in one Region to simplify communication.
- B. Use AWS Fargate to run each microservice in separate containers. Deploy across multiple Availability Zones in one Region behind an Application Load Balancer.
- C. Use Amazon Route 53 with latency-based routing. Deploy microservices on Amazon EC2 instances in multiple Regions behind Application Load Balancers.
- D. Implement each microservice using AWS Lambda. Expose the microservices using an Amazon API Gateway REST API.

Answer: C

Explanation:

AWS recommends multi-Region active-active architectures for applications requiring high availability, automatic failover, and disaster recovery. Route 53 latency-based routing directs users to the Region providing the lowest latency and automatically shifts traffic if Regional endpoints become unhealthy.

Combined with Application Load Balancers and EC2-based microservices deployed in multiple Regions, this architecture delivers fault tolerance, multi-Region resiliency, and automatic recovery.

Option D provides high availability but does not inherently provide multi-Region failover routing without additional configuration. Option B is single-Region. Option A is not resilient.

Question: 579

A company is building a data processing application that uses AWS Lambda functions. The Lambda functions need to communicate with an Amazon RDS DB instance deployed inside a VPC in the same AWS account.

Which solution meets these requirements in the most secure way?

- A. Configure the DB instance for public access. Allow Lambda public address space.
- B. Deploy Lambda inside the VPC. Attach a network ACL allowing outbound access to the VPC CIDR. Update the DB security group to allow traffic from 0.0.0.0/0.
- C. Deploy Lambda inside the VPC. Attach a security group to the Lambda functions. Allow outbound access only to the VPC CIDR. Update the DB instance security group to allow traffic from the Lambda security group.
- D. Peer the Lambda default VPC with the DB VPC and avoid security groups.

Answer: C

Explanation:

For secure communication between Lambda and an RDS DB instance, AWS documentation recommends placing the Lambda functions inside the same VPC and controlling traffic using security groups.

Lambda should have a dedicated security group with outbound access to the VPC CIDR range, and the DB instance's security group should explicitly allow inbound connections from the Lambda security group. This follows the "least privilege" principle while avoiding public exposure.

Options A and B expose the DB to unnecessary risk. Option D is insecure because it bypasses security groups.

Question: 580

A company operates multiple VPCs in a single AWS account. Account users need temporary access to Amazon S3 buckets. The S3 buckets are private and have no public endpoints.

The solution must follow the principle of least privilege for access to each environment and must avoid distributing permanent access keys.

Which solution will meet these requirements?

- A. Create a gateway VPC endpoint for Amazon S3 in each VPC. Attach an endpoint policy that allows only environment-scoped IAM roles to access the S3 buckets.
- B. Configure the S3 buckets to use SSE-S3. Create bucket policies that allow access only from the VPC CIDR blocks.
- C. Define separate S3 access points for each environment. Allow users to assume a role associated with the access points. Use the default Amazon S3 endpoints.
- D. Route S3 traffic through a NAT gateway. Configure bucket policies that allow traffic only from the NAT gateway's public IP addresses.

Answer: A

Explanation:

AWS documentation states that the most secure method for granting private S3 access from VPCs is to use gateway VPC endpoints for Amazon S3. Endpoint policies enforce least privilege by allowing only specific IAM roles access to specific S3 buckets. Because users can obtain temporary credentials through IAM roles, no permanent access keys must be distributed.

S3 bucket policies based on CIDR ranges (Option B) are less secure because CIDR-based access is broader and can grant unintended access. Access points (Option C) still rely on public S3 endpoints unless combined with VPC-only access, which is not stated here. Option D exposes S3 to public internet egress and depends on public IPs, violating least-privilege principles.

Question: 581

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run Amazon Linux in an Auto Scaling group. Each instance stores product manuals on Amazon EBS volumes.

New instances often start with outdated data and may take up to 30 minutes to download updates. The company needs a solution ensuring all instances always have up-to-date product manuals, can scale rapidly, and does not require application code changes.

Which solution will meet these requirements?

- A. Store the product manuals on instance store volumes attached to each EC2 instance.

- B. Store the product manuals in an Amazon S3 bucket. Configure EC2 instances to download updates from the bucket.
- C. Store the product manuals in an Amazon EFS file system. Mount the EFS volume on the EC2 instances.
- D. Store the product manuals in an S3 bucket using S3 Standard-IA. Configure EC2 instances to download updates from S3.

Answer: C

Explanation:

Amazon EFS provides a shared, fully managed, POSIX-compliant file system that can be mounted by all EC2 instances. Any update made to the file system is immediately visible to all instances, ensuring every new instance has the latest product manuals without delay.

EFS automatically scales storage and throughput, meeting high-demand conditions with no application changes required.

S3 requires instances to download files locally, causing delay and stale data issues (Options B and D). Instance store volumes (Option A) are ephemeral and not shared, making them unsuitable for consistent data distribution.

Question: 582

A media publishing company is building an application that allows users to print custom books. The frontend runs in a Docker container. Incoming order volume is highly variable and can exceed the throughput of the physical printing machines. Order-processing payloads can be up to 4 MB.

The company needs a scalable solution for handling incoming orders.

Which solution will meet this requirement?

- A. Use Amazon SQS to queue incoming orders. Use Lambda@Edge to process orders. Deploy the frontend on Amazon EKS.
- B. Use Amazon SNS to queue incoming orders. Use a Lambda function to process orders. Deploy the frontend on AWS Fargate.
- C. Use Amazon SQS to queue incoming orders. Use a Lambda function to process orders. Deploy the frontend on Amazon ECS with the Fargate launch type.
- D. Use Amazon SNS to queue incoming orders. Use Lambda@Edge to process orders. Deploy the frontend on Amazon EC2 instances.

Answer: C

Explanation:

Amazon SQS is the AWS-recommended service for handling decoupled, scalable, ordered, durable message ingestion with potentially large workloads. It supports payloads up to 256 KB directly and up to several MB using S3-backed large-payload extensions. Lambda functions can process messages from SQS with automatic scaling. ECS with the Fargate launch type provides a serverless container platform for the frontend, scaling based on demand without managing servers.

SNS (Options B and D) is not a queue, does not provide durable message buffering, and is not appropriate for variable, bursty workloads that may exceed backend throughput. Lambda@Edge (Options A and D) is unsuitable for backend order processing because it runs at CloudFront edge locations with strict limitations.

Question: 583

A healthcare company is designing a system to store and manage logs in the AWS Cloud. The system ingests and stores logs in JSON format that contain sensitive patient information. The company must identify any sensitive data and must be able to search the log data by using SQL queries.

Which solution will meet these requirements?

- A. Store the logs in an Amazon S3 bucket. Configure Amazon Macie to discover sensitive data. Use Amazon Athena to query the logs.
- B. Store the logs in an Amazon EBS volume. Create an application that uses Amazon SageMaker AI to detect sensitive data. Use Amazon RDS to query the logs.
- C. Store the logs in Amazon DynamoDB. Use AWS KMS to discover sensitive data. Use Amazon Redshift Spectrum to query the logs.
- D. Store the logs in an Amazon S3 bucket. Use Amazon Inspector to discover sensitive data. Use Amazon Athena to query the logs.

Answer: A

Explanation:

AWS documentation states that Amazon Macie is the managed service designed to automatically identify and classify sensitive data stored in Amazon S3, including PII and healthcare-related identifiers.

Storing logs in Amazon S3 provides scalable, durable storage, and Amazon Athena can directly query JSON data stored in S3 using SQL.

Amazon Inspector (Option D) is for vulnerability scanning and does not identify sensitive data. DynamoDB with KMS (Option C) cannot detect sensitive information. EBS (Option B) requires custom tooling and does not support serverless SQL querying.

Question: 584

A company has an on-premises volume backup solution that is end of life. The company wants to use AWS as part of a new backup solution while maintaining local access to all data. The data must be automatically and securely transferred to AWS.

Which solution meets these requirements?

- A. Use AWS Snowball to migrate data to Amazon S3. Mount the Snowball S3 endpoint for local access.
- B. Use AWS Snowball Edge to migrate data to Amazon S3. Use the Snowball Edge file interface to provide local access.
- C. Use AWS Storage Gateway and configure a cached volume gateway. Run the gateway appliance on premises, cache a percentage of data locally, and mount gateway volumes for local access.
- D. Use AWS Storage Gateway and configure a stored volume gateway. Run the appliance on premises, map the gateway storage to on-premises disks, and mount gateway volumes for local access.

Answer: D

Explanation:

A Stored Volume Gateway stores the entire dataset locally on-premises while asynchronously and securely backing up the data to AWS. This meets the requirement for full local access and automatic transfer to AWS.

Cached Volume Gateway (Option C) keeps only frequently accessed data locally; the full dataset is stored in AWS, not on-premises.

Snowball and Snowball Edge (Options A and B) are migration tools, not ongoing backup solutions.

Question: 585

A company is building a serverless application that processes large volumes of data from a mobile app. A Lambda function processes the data and stores it in DynamoDB. The company must ensure the application can recover from failures and continue processing without losing records.

Which solution will meet these requirements?

- A. Configure the Lambda function with a dead-letter queue (DLQ) using SQS. Retry failed records from the DLQ with exponential backoff.
- B. Configure the Lambda function to read records from Amazon Data Firehose. Replay Firehose records in case of failures.
- C. Use Amazon OpenSearch Service to store failed records. Configure Lambda to retry failed records from OpenSearch. Use EventBridge for orchestration.
- D. Use Amazon SNS to store failed records. Configure Lambda to retry records from SNS. Use API Gateway to orchestrate retries.

Answer: A

Explanation:

AWS documentation states that when Lambda processes events, a dead-letter queue (DLQ) using Amazon SQS is the correct mechanism to capture failed invocations for later reprocessing.

SQS provides durable buffering and decoupling, allowing records to be retried safely without loss.

Kinesis Firehose (Option B) does not support SQL replays or arbitrary replay control. SNS (Option D) is not a durable store. OpenSearch (Option C) is not intended for durable retry queues.

Question: 586

A company uses AWS Organizations to manage multiple AWS accounts. The company needs a secure, event-driven architecture in which specific Amazon SNS topics in Account A can publish messages to specific Amazon SQS queues in Account B.

Which solution meets these requirements while maintaining least privilege?

- A. Create a new IAM role in Account A that can publish to any SQS queue. Share the role ARN with Account B.
- B. Add SNS topic ARNs to SQS queue policies in Account B. Configure SNS topics to publish to any queue. Encrypt the queue with an AWS KMS key.
- C. Modify the SQS queue policies in Account B to allow only specific SNS topic ARNs from Account A to publish messages. Ensure the SNS topics have publish permissions for the specific queue ARN.
- D. Create a shared IAM role across both accounts with permission to publish to all SQS queues. Enable cross-account access.

Answer: C

Explanation:

AWS documentation states that the correct and least-privilege method for cross-account SNS-to-SQS integration is:

- Add specific SNS topic ARNs to the SQS queue policy.
- Allow only those topics to publish messages to the queue.
- Ensure SNS has permission to publish to the specific queue ARN.

This ensures strict scoping and adheres to least privilege.

Options A and D grant overly broad permissions. Option B allows publishing to any queue, which violates least privilege.

Question: 587

A global media streaming company is migrating its user authentication and content delivery services to AWS. The company wants to use Amazon API Gateway for user authentication and authorization. The company needs a solution that restricts API access to AWS Regions in the United States and ensures minimal latency.

Which solution will meet these requirements?

- A. Create an API Gateway REST API. Configure an AWS WAF firewall in the same Region. Implement AWS WAF rules to deny requests that originate from Regions outside the United States. Associate the AWS WAF firewall with the API Gateway REST API.
- B. Create an API Gateway HTTP API. Configure an AWS WAF firewall in a different Region. Implement AWS WAF rules to deny requests that originate from Regions outside the United States. Associate the AWS WAF firewall with the API Gateway HTTP API.
- C. Create an API Gateway REST API. Configure an AWS WAF firewall in a different Region. Implement AWS WAF rules to deny requests that originate from Regions outside the United States. Associate the AWS WAF firewall with the API Gateway REST API.
- D. Create an API Gateway HTTP API. Configure an AWS WAF firewall in the same Region. Implement AWS WAF rules to deny requests that originate from Regions outside the United States. Associate the AWS WAF firewall with the API Gateway HTTP API.

Answer: A

Explanation:

AWS documentation states that for API Gateway REST APIs, AWS WAF must be deployed in the same Region to apply web access control lists. WAF geographic match rules can block requests from nonUS Regions. Deploying the WAF and API Gateway in the same Region ensures minimal latency because the traffic does not traverse Regions.

Options B and C violate the requirement that WAF and API Gateway must be in the same Region. Option D uses HTTP API, which supports WAF only indirectly via regional ALBs, not directly as REST API does.

Question: 588

A security team needs to enforce rotation of all IAM users' access keys every 90 days. Keys older than 90 days must be automatically deactivated and removed. A solutions architect must create a remediation solution with minimal operational effort.

Which solution meets these requirements?

- A. Create an AWS Config rule to check key age. Configure the rule to run an AWS Batch job to remove the key.
- B. Create an Amazon EventBridge rule to check key age. Configure it to run an AWS Batch job to remove the key.
- C. Create an AWS Config rule to check key age. Define an EventBridge rule that schedules an AWS Lambda function to remove the key.

D. Create an EventBridge rule to check key age. Define a second EventBridge rule to run an AWS Batch job to remove the key.

Answer: C

Explanation:

AWS Config has a built-in managed rule (access-keys-rotated) that evaluates IAM access key age. Config rules detect non-compliant resources automatically, without building custom logic.

After Config identifies old keys, EventBridge can trigger an AWS Lambda function to disable and delete the keys. Lambda provides a fully managed compute layer requiring no servers or batch environments.

Using AWS Batch (Options A, B, and D) adds unnecessary operational overhead for a simple automation task.

Question: 589

A solutions architect needs to design a system to process incoming work items immediately. Processing can take up to 30 minutes and involves calling external APIs, executing multiple states, and storing intermediate states.

The solution must scale with variable workloads and minimize operational overhead.

Which combination of steps meets these requirements? (Select TWO.)

- A. Invoke an AWS Lambda function for each incoming work item. Configure each function to handle the work item completely. Store states in DynamoDB.
- B. Invoke an AWS Step Functions workflow to process incoming work items. Use Lambda functions for business logic. Store work item states in DynamoDB.
- C. Set up an API Gateway REST API to receive work items. Configure the API to invoke a Lambda function for each work item.
- D. Deploy two EC2 Reserved Instances behind an ALB and send requests to an SQS queue.
- E. Set up an API Gateway REST API to receive work items. Send the work items to an SQS queue.

Answer: B,E

Explanation:

AWS Step Functions is the recommended service for orchestrating multi-step, long-running workflows with state tracking, retries, and

external API calls. It reduces operational overhead by eliminating the need for custom orchestration logic.

API Gateway receiving work items and sending them to SQS (Option E) provides buffering, elasticity, and decoupling, ensuring immediate ingestion regardless of backend load.

Option A forces Lambda to handle long execution paths, which is not optimal for 30-minute tasks and multi-state workflows. Option C triggers Lambda directly without buffering. Option D uses fixed EC2 instances, which does not scale dynamically.

Question: 590

A company needs to migrate its customer transactions database from on premises to AWS. The database is an Oracle DB instance on Linux. A new requirement mandates rotating the database password yearly.

Which solution provides this capability with the least operational overhead?

- A. Convert the database to DynamoDB using AWS SCT. Store the password in Parameter Store. Use CloudWatch and Lambda for rotation.
- B. Migrate the database to Amazon RDS for Oracle. Store the password in AWS Secrets Manager. Turn on automatic rotation with a yearly rotation schedule.
- C. Migrate the database to an EC2 instance. Use Parameter Store to keep and rotate the connection string using a Lambda function with a yearly schedule.
- D. Migrate the database to Amazon Neptune using AWS SCT. Use CloudWatch and Lambda for yearly rotation.

Answer: B

Explanation:

AWS Secrets Manager natively supports automatic rotation of RDS for Oracle credentials, fully integrating with Amazon RDS. Rotation workflows are managed automatically by the service, eliminating custom scripting and reducing operational effort.

Migrating to EC2 (Option C) requires custom rotation logic. Converting to DynamoDB or Neptune (Options A and D) would require complete database redesign and do not fulfill the requirement to run Oracle.

Question: 591

A financial company is migrating banking applications to AWS accounts managed through AWS Organizations. The applications store sensitive customer data on Amazon EBS volumes, and the company takes regular snapshots for backups.

The company must implement controls across all accounts to prevent sharing EBS snapshots publicly, with the least operational overhead.

Which solution will meet these requirements?

- A. Enable AWS Config rules for each OU to monitor EBS snapshot permissions.
- B. Enable block public access for EBS snapshots at the organization level.
- C. Create an IAM policy in the root account that prevents users from modifying snapshot permissions.
- D. Use AWS CloudTrail to track snapshot permission changes.

Answer: B

Explanation:

AWS provides EBS Block Public Access at the organization level in AWS Organizations. When enabled, it prevents any EBS snapshot—across all member accounts—from being shared publicly.

This is an organization-wide control implemented centrally, with no need for per-account configuration, monitoring rules, or custom IAM policy enforcement.

AWS Config (Option A) would only detect issues after they occur. IAM restrictions (Option C) are less effective because snapshot permission changes can occur through multiple paths. CloudTrail (Option D) only logs events and does not block public sharing.

Question: 592

A solutions architect is building a static website hosted on Amazon S3. The website uses an Amazon Aurora PostgreSQL database accessed through an AWS Lambda function. The production website uses a Lambda alias that points to a specific version of the Lambda function.

Database credentials must rotate every 2 weeks. Previously deployed Lambda versions must always use the most recent credentials.

Which solution will meet these requirements?

- A. Store credentials in AWS Secrets Manager. Turn on rotation. Write code in the Lambda function to retrieve credentials from Secrets Manager.
- B. Include the credentials in the Lambda function code and update the function regularly.
- C. Use Lambda environment variables and update them when new credentials are available.
- D. Store credentials in AWS Systems Manager Parameter Store. Turn on rotation. Write code to retrieve credentials from Parameter Store.

Answer: A

Explanation:

AWS Secrets Manager is the recommended service for storing database credentials and performing automated rotation. Any Lambda function version or alias can fetch the latest secret value at runtime, ensuring no outdated credentials exist in deployed versions.

Environment variables (Option C) are static per version. Embedding credentials in code (Option B) is insecure and requires redeployment. Parameter Store (Option D) supports rotation but requires more configuration and is not as seamless as Secrets Manager for database credential rotation.

Question: 593

A company is developing an application using Amazon Aurora MySQL. The team will frequently make schema changes to test new features without affecting production. After testing, changes must be promoted to production with minimal downtime.

Which solution meets these requirements?

- A. Create a staging Aurora cluster based on the existing cluster. Test schema changes on the staging cluster.
- B. Create a read replica, modify its schema, and then promote it to primary.
- C. Create an Aurora MySQL blue/green deployment. Make schema changes in the staging environment and switch traffic after testing.
- D. Replicate the Aurora database to DynamoDB, apply schema changes, and switch the application to DynamoDB.

Answer: C

Explanation:

Aurora blue/green deployments are specifically designed for safe schema changes, zero-downtime updates, and production isolation.

The staging (green) environment can receive schema changes without affecting production (blue). After validation, you perform a fast, minimally disruptive switchover that updates production.

Read replicas (Option B) do not allow schema changes. Creating an independent staging cluster (Option A) does not provide automated, low-downtime cutover. DynamoDB (Option D) is not compatible with MySQL schemas.

Question: 594

A company stores data in a centralized S3 bucket in Account A. It needs to grant Account B access to this bucket. Both accounts belong to the company.

Which solution meets this requirement?

- A. Enable S3 Transfer Acceleration for Account B.
- B. Enable cross-Region replication between accounts.
- C. Use CloudFront with signed URLs to grant access.
- D. Create a bucket policy granting Account B access to the bucket in Account A.

Answer: D

Explanation:

The recommended and simplest method for granting cross-account access to an S3 bucket is to use a bucket policy that grants permission to the other AWS account.

This provides direct, controlled, IAM-based access without duplication of data or use of distribution mechanisms like CloudFront.

Transfer Acceleration and replication are unrelated to permissions.

Question: 595

A company deployed an application in two AWS Regions. If the application fails in one Region, traffic must fail over to the second Region. The failover must avoid stale DNS client caches, and the company requires one endpoint for both Regions.

Which solution meets these requirements?

- A. Use a CloudFront distribution with multiple origins.
- B. Use Route 53 weighted routing with equal weights.
- C. Use AWS Global Accelerator and assign static anycast IPs to the application.
- D. Use Route 53 IP-based routing to switch Regions.

Answer: C

Explanation:

AWS Global Accelerator provides static anycast IP addresses that remain constant regardless of Regional failover. AWS directs traffic to the optimal healthy Region without relying on DNS TTL values, eliminating the risk of stale DNS caches.

Route 53 routing (Options B and D) still depends on DNS caching behavior. CloudFront (Option A) is for content delivery, not Regional failover for applications.

Question: 596

A company provides devices to users. When a device is registered, its ID is added to DynamoDB. A daily job activates devices using two Lambda functions:

- The Retrieve function lists unregistered device IDs.
- The Retrieve function then calls the Activate function in a loop to register each device.

The number of activations is increasing, and the company wants to avoid Lambda timeouts without modifying existing functions.

Which solution will scale appropriately?

- A. Use EventBridge Scheduler to periodically invoke the Retrieve function.
- B. Invoke the Activate function from DynamoDB Streams when a device ID is added.
- C. Use Step Functions to call the Retrieve function and use a Map state to run the Activate function for each ID.
- D. Move the Retrieve function to EC2 for longer processing time.

Answer: C

Explanation:

AWS Step Functions can orchestrate Lambda executions without modifying the Lambda code, and the Map state is designed to run a Lambda function once per item in a list—perfect for activating many devices in parallel or controlled batches.

This eliminates the loop inside the Retrieve function and prevents timeouts.

EventBridge (Option A) only changes the schedule. DynamoDB Streams (Option B) is unrelated to the activation workflow. EC2 (Option D) increases operational overhead.

Question: 597

A data science team needs storage for nightly log processing. The size and number of logs is unknown, and the logs persist for only

24 hours.

What is the MOST cost-effective solution?

- A. Amazon S3 Glacier Deep Archive
- B. Amazon S3 Standard
- C. Amazon S3 Intelligent-Tiering
- D. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: C

Explanation:

Amazon S3 Intelligent-Tiering is designed for unknown or unpredictable access patterns, automatically moving objects to the most cost-effective tier without performance impact.

It is ideal when you do not know object size or access frequency.

S3 Standard (Option B) works but costs more.

S3 One Zone-IA and Glacier Deep Archive (Options D and A) are not appropriate because data is frequently accessed for processing and persists only 24 hours.

Question: 598

A company wants to protect AWS-hosted resources, including Application Load Balancers and CloudFront distributions. They need near real-time visibility into attacks and a dedicated AWS response team for DDoS events.

Which AWS service meets these requirements?

- A. AWS WAF
- B. AWS Shield Standard
- C. Amazon Macie
- D. AWS Shield Advanced

Answer: D

Explanation:

AWS Shield Advanced provides:

- 24/7 access to the AWS DDoS Response Team (DRT)
- Near real-time attack visibility and metrics
- Protection for CloudFront, ALB, NLB, and Route 53

Shield Standard (Option B) does not include DDoS response team support.

WAF (Option A) provides rules-based filtering, not DDoS assistance.

Macie (Option C) is for PII discovery, not DDoS protection.

Question: 599

A solutions architect needs to ensure that only resources in VPC vpc-11aabb22 can access an S3 bucket in account 123456789012 with Block Public Access enabled.

Which solution meets this requirement?

- A. Create a bucket policy with Deny and a Condition using "StringNotEquals": {"aws:SourceVpc": "vpc-11aabb22"}.
- B. Create a bucket policy with Allow and Resource "arn:aws:ec2:us-west-2:123456789012:vpc/vpc-11aabb22".
- C. Create a bucket policy with Allow and a Condition using "StringNotEquals": {"aws:SourceVpc": "vpc-11aabb22"}.
- D. Create a bucket policy with Deny and "StringNotEquals": {"aws:PrincipalAccount": "123456789012"}.

Answer: A

Explanation:

To restrict S3 access to a specific VPC, AWS documentation specifies using a Deny policy with aws:SourceVpc for all unmatched VPCs:

Deny if aws:SourceVpc != vpc-11aabb22

This ensures only that VPC can reach the bucket via a VPC endpoint.

Allow policies (Options B and C) are incorrectly structured or use the wrong condition logic.

Option D limits by account, not VPC, and does not enforce VPC-level access.

Question: 600

An application has performance issues due to increased demand. The demand is on read-only historical records in Amazon RDS using custom queries. The company wants improved performance without changing database structure and with minimal management overhead.

Which approach meets the requirement?

- A. Deploy DynamoDB and move all data.
- B. Deploy Amazon ElastiCache (Redis OSS) and cache application data.
- C. Deploy Memcached on EC2 and cache data.
- D. Deploy DynamoDB Accelerator (DAX) on Amazon RDS.

Answer: B

Explanation:

Amazon ElastiCache (Redis OSS) provides an in-memory cache that drastically improves read performance with minimal operational overhead.

It integrates easily with RDS and does not require schema or database changes.

EC2-based caching (Option C) increases management overhead.

DAX (Option D) accelerates DynamoDB only, not RDS.

Moving to DynamoDB (Option A) requires complete application and schema redesign.

Question: 601

A company uses AWS Lake Formation to govern its S3 data lake. It wants to visualize data in QuickSight by joining S3 data with Aurora MySQL operational data. The marketing team must see only specific columns.

Which solution provides column-level authorization with the least operational overhead?

- A. Use EMR to ingest database data into SPICE with only required columns.
- B. Use AWS Glue Studio to ingest database data into S3 and use IAM policies for column control.
- C. Use AWS Glue Elastic Views to create materialized S3 views with column restrictions.
- D. Use a Lake Formation blueprint to ingest database data to S3. Use Lake Formation for column-level access control. Use Athena

as the QuickSight data source.

Answer: D

Explanation:

AWS Lake Formation provides fine-grained (column-level) access control for data stored in S3. Using a Lake Formation blueprint ensures database ingestion is automated and governed.

QuickSight can query Athena, and Athena honors Lake Formation permissions, enforcing column-level controls automatically.

Options A, B, and C rely on manual filtering or IAM policies, which cannot enforce column-level authorization for SQL queries.

Question: 602

A company runs Amazon EC2 instances as web servers. Peak traffic occurs at two predictable times each day. The web servers remain mostly idle during the rest of the day.

A solutions architect must manage the web servers while maintaining fault tolerance in the most cost-effective way.

Which solution will meet these requirements?

- A. Use an EC2 Auto Scaling group to scale the instances based on demand.
- B. Purchase Reserved Instances to ensure peak capacity at all times.
- C. Use a cron job to stop the EC2 instances when traffic demand is low.
- D. Use a script to vertically scale the EC2 instances during peak demand.

Answer: A

Explanation:

AWS documentation states that EC2 Auto Scaling is the recommended, cost-effective, and fault-tolerant method to manage workloads with predictable or varying demand. Auto Scaling automatically adds instances during peak traffic and terminates them when demand is low, reducing compute cost while maintaining availability across multiple Availability Zones.

Reserved Instances (Option B) would force the company to pay for peak capacity all day, which is not cost-effective.

Stopping instances manually (Option C) or vertically scaling instances (Option D) reduces fault tolerance and increases operational overhead.

Question: 603

A company has resources across multiple AWS Regions and accounts. A new solutions architect needs to build a map of the workloads and

their relationships but has no documentation from the previous employee.

Which solution will provide these details with the least operational effort?

- A. Use AWS Systems Manager Inventory to generate a map from the detailed report.
- B. Use AWS Step Functions to collect workload details and build diagrams manually.
- C. Use Workload Discovery on AWS to generate architecture diagrams.
- D. Use AWS X-Ray to view workload details and manually draw diagrams.

Answer: C

Explanation:

Workload Discovery on AWS automatically scans AWS accounts and Regions, identifies deployed resources, and generates relationship-aware architecture diagrams. This provides a complete workload map with minimal operational effort.

Systems Manager Inventory (Option A) collects instance metadata but does not build multi-account architecture diagrams.

Step Functions (Option B) and X-Ray (Option D) require manual diagramming and do not meet the operational-efficiency requirement.

Question: 604

A home security company is expanding globally and needs to encrypt customer data. The company does not want to manage encryption keys. The keys must be usable in multiple AWS Regions, and access to the keys must be controlled.

Which solution meets these requirements with the least operational overhead?

- A. Use AWS KMS multi-Region keys. Apply tags and use ABAC condition keys for access control.
- B. Use AWS KMS imported key material in multiple Regions with ABAC-based policies.
- C. Use AWS CloudHSM and synchronize clusters across Regions with the CMU tool.
- D. Use AWS CloudHSM users and share keys manually with CMU across Regions.

Answer: A

Explanation:

AWS Key Management Service (AWS KMS) provides multi-Region keys, allowing the same key to be used across Regions without managing your own hardware or key replication.

Multi-Region keys are fully managed and support attribute-based access control (ABAC) using tags and condition keys.

CloudHSM (Options C and D) requires full key lifecycle management, synchronization, and hardware operations, resulting in significantly higher overhead.

Imported key material (Option B) increases key-management responsibility.

Question: 605

A company uses a Microsoft SQL Server database. The applications currently connect using SQL Server protocols. The company wants to migrate to Amazon Aurora PostgreSQL with minimal changes to application code.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Use AWS SCT to rewrite SQL queries in the applications.
- B. Enable Babelfish on Aurora PostgreSQL to run SQL Server queries.
- C. Migrate the database schema and data using AWS SCT and AWS DMS.
- D. Use Amazon RDS Proxy to connect the applications to Aurora PostgreSQL.
- E. Use AWS DMS to rewrite SQL queries in the applications.

Answer: B,C

Explanation:

Amazon Aurora PostgreSQL with Babelfish allows Aurora to understand SQL Server T-SQL and the SQL Server wire protocol. This enables applications to continue using SQL Server drivers, minimizing code changes (Option B).

Migration of schema and data is performed using AWS Schema Conversion Tool (SCT) and AWS Database Migration Service (DMS) (Option C), which is the AWS-recommended migration pattern for heterogeneous database migrations.

AWS DMS (Option E) does not rewrite application SQL.

RDS Proxy (Option D) does not translate SQL Server protocols.

Option A requires rewriting application queries, which contradicts the “minimal changes” requirement.

Question: 606

A company runs a custom application on Amazon EC2 On-Demand Instances. The application has frontend nodes that must run 24/7. The backend nodes only need to run for short periods depending on the workload.

Frontend nodes accept jobs and place them in queues. Backend nodes asynchronously process jobs from the queues, and jobs can be restarted. The company wants to scale infrastructure based on workload, using the most cost-effective option.

Which solution meets these requirements MOST cost-effectively?

- A. Use Reserved Instances for the frontend nodes. Use AWS Fargate for the backend nodes.
- B. Use Reserved Instances for the frontend nodes. Use Spot Instances for the backend nodes.
- C. Use Spot Instances for the frontend nodes. Use Reserved Instances for the backend nodes.
- D. Use Spot Instances for the frontend nodes. Use AWS Fargate for the backend nodes.

Answer: B

Explanation:

AWS documentation states that workloads running 24/7 should use Reserved Instances or Savings Plans for the lowest cost. Therefore, the frontend nodes, which always run, should use Reserved Instances.

The backend nodes process asynchronous, restartable jobs, which makes them ideal for EC2 Spot Instances, the most cost-effective compute option for interruption-tolerant workloads.

Fargate (Options A and D) is significantly more expensive for large compute usage. Spot Instances cannot be used for critical frontend nodes (Option C).

Question: 607

A company runs an HPC workload that uses a 200-TB file system on premises. The company needs to migrate this data to Amazon FSx for Lustre. Internet capacity is 10 Mbps, and all data must be migrated within 30 days.

Which solution will meet this requirement?

- A. Use AWS DMS to transfer data into S3 and link FSx for Lustre to the bucket.
- B. Deploy AWS DataSync on premises and transfer directly into FSx for Lustre.
- C. Use AWS Storage Gateway Volume Gateway to move data into FSx for Lustre.
- D. Use an AWS Snowball Edge storage-optimized device to transfer data into S3 and link FSx for Lustre to the bucket.

Answer: D

Explanation:

At 10 Mbps, the maximum transferable data in 30 days is far below 200 TB, making any online transfer (Options A, B, C) impossible within the time window.

AWS Snowball Edge storage-optimized devices support high-speed, offline bulk data migration of large datasets. Once the data is delivered to S3, FSx for Lustre can be linked to the S3 bucket to populate the Lustre filesystem.

DataSync cannot meet the time constraint over 10 Mbps. Storage Gateway is not designed for large-scale migrations.

Question: 608

A company stores a file in an S3 bucket containing IP allow/deny lists. The file must be accessible via an HTTP endpoint. Firewalls outside AWS must read the file. The company wants to restrict access to only the firewall IP addresses.

The S3 Block Public Access feature is enabled on the account.

Which solution meets these requirements?

- A. Host the bucket as a static website and restrict access by IP.
- B. Create a bucket policy that explicitly allows access only from the firewall IP addresses.
- C. Create a CloudFront distribution with the S3 bucket as the origin. Use an origin access control (OAC) that allows access only from the firewall IP addresses.
- D. Create a Lambda function to validate IP addresses and return the lists.

Answer: B

Explanation:

S3 Block Public Access only blocks public access—not explicitly allowed access.

An S3 bucket policy can explicitly permit access only from specific source IP addresses using the `aws:SourceIp` condition. This allows secure, direct HTTP access to the S3 object from known firewall IP addresses.

Static website hosting (Option A) requires the bucket to be public and is blocked by the enabled S3 Block Public Access setting.

CloudFront with OAC (Option C) is unnecessary and adds cost and complexity.

Lambda (Option D) introduces operational overhead and is not needed since S3 policies can enforce IP restrictions directly.

Question: 609

An insurance company is creating an application to record personal user data. The data includes users' names, ages, and health data. The company wants to run the application in a private subnet on AWS.

Because of data security requirements, the company must have access to the operating system of the compute resources that run the application tier. The company must use a low-latency NoSQL database to store the data.

Which solution will meet these requirements?

- A. Use Amazon EC2 instances for the application tier. Use an Amazon DynamoDB table for the database tier. Create a VPC endpoint for DynamoDB. Assign the instances an instance profile that has permission to access DynamoDB.
- B. Use AWS Lambda functions for the application tier. Use an Amazon DynamoDB table for the database tier. Assign a Lambda function an appropriate IAM role to access the table.
- C. Use AWS Fargate for the application tier. Create an Amazon Aurora PostgreSQL instance inside a private subnet for the database tier.
- D. Use Amazon EC2 instances for the application tier. Use an Amazon S3 bucket to store the data in JSON format. Configure the application to use Amazon Athena to read and write the data to and from the S3 bucket.

Answer: A

Explanation:

The requirement to “have access to the operating system” means the compute layer must be Amazon EC2 (or containers on EC2). Managed runtimes such as Lambda and Fargate do not provide OS-level access.

The requirement for a “low-latency NoSQL database” maps directly to Amazon DynamoDB, which is a fully managed NoSQL key-value and document database that provides single-digit millisecond latency at any scale.

Because the application runs in a private subnet, AWS best practice is to access DynamoDB privately via a VPC endpoint (gateway endpoint for DynamoDB). This avoids traversing the public internet and simplifies security.

An instance profile (EC2 role) is the recommended method to grant EC2 instances permission to access DynamoDB without hardcoding credentials.

Why the other options are not correct:

B: Lambda does not provide OS access, which violates the security requirement.

C: Fargate does not provide OS access, and Aurora PostgreSQL is a relational database, not NoSQL.

D: S3 + Athena is an analytics pattern, not a low-latency NoSQL database solution; query latency is much higher and not suitable for OLTP-style app storage.

Question: 610

A company is building a new web application that serves static and dynamic content from an API. Users will access the application from around the world. The company wants to minimize latency in the most cost-effective way.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy the static content to an Amazon S3 bucket. Use an Amazon API Gateway HTTP API to serve the dynamic content. Create an Amazon CloudFront distribution that uses the S3 bucket and the HTTP API as origins. Enable caching for static content.
- B. Deploy the static content to an Amazon S3 bucket. Provide the bucket website endpoint to users. Use an Amazon API Gateway HTTP API with caching enabled to serve the dynamic content.
- C. Deploy the static content to an Amazon S3 bucket. Use two Amazon EC2 instances as web servers. Deploy an Application Load Balancer to distribute traffic. Create an Amazon CloudFront distribution in front of the S3 bucket to cache static content.
- D. Deploy the static content to an Amazon S3 bucket. Provide the bucket website endpoint to users. Create an Amazon CloudFront distribution in front of the S3 bucket to cache static content.

Answer: A

Explanation:

For global low latency, AWS recommends using Amazon CloudFront as a content delivery network in front of both static and dynamic origins when possible.

Static content on Amazon S3 + CloudFront is the standard cost-effective pattern: S3 for low-cost durable storage, CloudFront for global edge caching.

Dynamic content exposed via Amazon API Gateway HTTP API can also be used as an origin for CloudFront, giving global users reduced latency and the ability to cache appropriate responses (when safe).

This pattern minimizes cost by:

Using S3 instead of EC2 for static hosting.

Using HTTP APIs (cheaper than REST APIs) for dynamic content.

Offloading a significant portion of requests to CloudFront cache.

Why the other options are not correct:

B: Static content is not fronted by CloudFront, so global latency is higher.

C: EC2-based web servers for the core web tier add more operational and cost overhead than needed for a mostly static + API pattern.

D: Only the static content uses CloudFront; the dynamic API still lacks a latency-optimized global entry point.

Question: 611

A retail company is building an order fulfillment system using a microservices architecture on AWS. The system must store incoming orders durably until processing completes successfully. Multiple teams' services process orders according to a defined workflow. Services must be scalable, loosely coupled, and able to handle sudden surges in order volume. The processing steps of each order must be centrally tracked.

Which solution will meet these requirements?

- A. Send incoming orders to an Amazon Simple Notification Service (Amazon SNS) topic. Start an AWS Step Functions workflow for each order that orchestrates the microservices. Use AWS Lambda functions for each microservice.
- B. Send incoming orders to an Amazon Simple Queue Service (Amazon SQS) queue. Start an AWS Step Functions workflow for each order that orchestrates the microservices. Use AWS Lambda functions for each microservice.
- C. Send incoming orders to an Amazon Simple Queue Service (Amazon SQS) queue. Use Amazon EventBridge to distribute events among the microservices. Use AWS Lambda functions for each microservice.
- D. Send incoming orders to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe Amazon EventBridge to the topic to distribute events among the microservices. Use AWS Lambda functions for each microservice.

Answer: B

Explanation:

Durable storage of incoming orders with buffering and ability to handle surges is exactly what Amazon SQS is designed for. SQS provides highly durable, scalable queues that decouple producers from consumers.

Centrally tracking workflow steps is a core use case of AWS Step Functions, which gives a visual workflow and state machine, tracks the state of each order, and can orchestrate calls to multiple microservices (in this case, Lambda functions).

Combining SQS + Step Functions + Lambda gives:

Durable queueing for orders (SQS).

Loose coupling and surge handling (SQS decoupling + auto-scaling Lambda).

Central orchestration and tracking of order-processing steps (Step Functions).

Why the other options are not correct:

A: SNS is a pub/sub service, not a durable work queue, and is not designed for “store-and-retry until processed” workloads in the same way SQS is.

C: SQS + EventBridge provides decoupling but no central, stateful workflow tracking; EventBridge is event routing, not workflow orchestration.

D: SNS + EventBridge still lacks durable order storage and explicit centralized workflow/state tracking.

Question: 612

An application is experiencing performance issues based on increased demand. This increased demand is on read-only historical records that are pulled from an Amazon RDS-hosted database with custom views and queries. A solutions architect must improve performance without changing the database structure.

Which approach will improve performance and MINIMIZE management overhead?

- A. Deploy Amazon DynamoDB, move all the data, and point to DynamoDB.
- B. Deploy Amazon ElastiCache (Redis OSS) and cache the data for the application.
- C. Deploy Memcached on Amazon EC2 and cache the data for the application.
- D. Deploy Amazon DynamoDB Accelerator (DAX) on Amazon RDS to improve cache performance.

Answer: B

Explanation:

AWS recommends using Amazon ElastiCache as an in-memory caching layer in front of relational databases such as Amazon RDS to offload read traffic and significantly improve performance for readheavy workloads. ElastiCache (Redis OSS) provides microsecond latency and can cache the results of frequent or expensive queries without requiring any change to the underlying database schema or engine. This directly addresses the requirement to improve performance without changing the database structure and with minimal operational overhead, because ElastiCache is a fully managed service (patching, failure detection, replacement, etc., are handled by AWS).

Option A (DynamoDB) would require a full data migration and application changes, including schema and query rewrites.

Option C (Memcached on EC2) introduces additional management overhead for EC2 instances (scaling, patching, HA).

Option D (DAX) is a caching layer only for DynamoDB and cannot be used directly with Amazon RDS.

Question: 613

A company is designing a secure solution to grant access to its Amazon RDS for PostgreSQL database.

Applications that run on Amazon EC2 instances must be able to securely authenticate to the database without storing long-term

credentials.

Which solution will meet these requirements?

- A. Enable RDS IAM authentication and configure AWS Secrets Manager to store database credentials. Configure applications to retrieve credentials at runtime.
- B. Configure a custom IAM policy for the database that allows access from the EC2 instances' IP addresses. Configure applications to use a static password to authenticate to the database.
- C. Set up an IAM user for each application. Store the access key ID and secret access key in the EC2 instances' environment variables. Grant the IAM users permission to the database.
- D. Use IAM roles to assign permissions to the EC2 instances. Configure the applications to obtain a token from the RDS database to authenticate by using IAM authentication.

Answer: D

Explanation:

For Amazon RDS for PostgreSQL, AWS provides IAM database authentication. With this feature, applications do not use stored long-term usernames and passwords. Instead, they use temporary authentication tokens that are generated by AWS and validated by the RDS database.

The AWS best practice pattern is:

Attach an IAM role to the EC2 instances (instance profile).

Grant that role the necessary permissions (for example, `rds-db:connect`) to the specific RDS database user.

The application running on the EC2 instance uses the role's temporary credentials to call the RDS token-generation API and obtain a short-lived authentication token.

The application then uses this token as the password when connecting to RDS for PostgreSQL.

This removes the need to store long-term credentials in the application or on the instance and uses IAM roles with temporary credentials, aligning with the security requirement.

Option A still relies on stored credentials (even if in Secrets Manager), which are long-lived and rotated but not token-based per-connection IAM authentication.

Option B uses static passwords and IP-based access, which does not meet the “no long-term credentials” requirement.

Option C stores long-term IAM user keys on the instances, which is explicitly against best practices and does not directly integrate with RDS authentication.

Question: 614

A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities.

The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions.

The company needs to improve application performance and decrease latency for the online game in preparation for user growth.

Which solution will meet these requirements?

- A. Add an Amazon CloudFront distribution in front of the NLBs. Increase the Cache-Control: max-age parameter.
- B. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.
- C. Add AWS Global Accelerator in front of the NLBs. Configure a Global Accelerator endpoint to use the correct listener ports.
- D. Add an Amazon API Gateway endpoint behind the NLBs. Enable API caching. Override method caching for the different stages.

Answer: C

Explanation:

AWS Global Accelerator is designed to improve the availability and performance of applications with global users by using the AWS global network. It provides static anycast IP addresses and routes user traffic over the AWS edge network to the optimal AWS Region and endpoint based on health, geography, and routing policies. Global Accelerator supports both TCP and UDP traffic and can have Network Load Balancers as endpoints.

For latency-sensitive workloads such as multiplayer gaming, Global Accelerator reduces latency and jitter compared to internet-based routing and handles Regional failover quickly.

CloudFront (Option A) is optimized for HTTP/HTTPS content caching and is not appropriate for arbitrary TCP/UDP gaming traffic.

Application Load Balancers (Option B) do not support UDP traffic. API Gateway (Option D) is for HTTP APIs and is not suitable for raw TCP/UDP game traffic.

Question: 615

A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers that a previous employee did not provide details about the resources inventory. The solutions architect needs to build and map the relationship details of the various workloads across all accounts.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report.
- B. Use AWS Step Functions to collect workload details. Build architecture diagrams of the workloads manually.
- C. Use Workload Discovery on AWS to generate architecture diagrams of the workloads.

D. Use AWS X-Ray to view the workload details. Build architecture diagrams with relationships.

Answer: C

Explanation:

Workload Discovery on AWS is a tool that automatically discovers AWS resources across multiple accounts and Regions and builds visual architecture diagrams that show resource relationships. It is specifically designed to help teams understand and document existing workloads with minimal manual effort.

This service provides multi-account visibility, relationship mapping, and exportable diagrams, which directly satisfies the need to “build and map the relationship details of the various workloads” with the least operational overhead.

Systems Manager Inventory (Option A) focuses on collecting configuration and inventory data (primarily for instances and some resources) and does not generate architecture diagrams or full relationship views.

Step Functions (Option B) and X-Ray (Option D) would require custom collection and manual diagramming and are not purpose-built for environment discovery and mapping.

Question: 616

A global ecommerce company is designing a three-tier application on AWS. The application includes a web tier that serves static content, an application tier that handles business logic, and a database tier that stores product information and user data. The application interacts with a relational database.

The company needs a highly available application architecture to serve global users with low latency, with the least operational overhead.

Which solution will meet these requirements?

A. Deploy Amazon EC2 instances in an Auto Scaling group for the application tier and web tier in a single AWS Region. Use an Application Load Balancer to distribute web traffic. Use an Amazon RDS database and Multi-AZ deployments for the database tier.

B. Set up an Amazon CloudFront distribution that uses an Amazon S3 bucket as the origin. Use Amazon Elastic Container Service (Amazon ECS) containers on AWS Fargate to deploy the application tier to each AWS Region where the company operates. Use an Amazon Aurora global database for the database tier.

C. Use an Amazon S3 bucket to store the static web content. Use Amazon EC2 Auto Scaling and EC2 Spot Instances for the application tier. Use Amazon RDS for MySQL with read replicas for the database tier. Use AWS Database Migration Service (AWS DMS) to replicate data to secondary AWS Regions.

D. Use an Amazon S3 bucket to store static web content. Use AWS Lambda functions to handle serverless backend logic in the application tier. Use Amazon API Gateway to invoke the Lambda functions for web requests. Use an Amazon DynamoDB database for the database tier. Deploy the DynamoDB database across multiple AWS Regions.

Answer: B

Explanation:

AWS guidance for global, highly available, low-latency applications using a relational database recommends:

Amazon CloudFront with Amazon S3 for static content to cache data at edge locations globally and minimize latency for static assets.

A regional application tier deployed close to users using managed container services such as Amazon ECS on AWS Fargate, which removes the need to manage servers and scales automatically, reducing operational overhead.

A relational database tier using Amazon Aurora global database, which is purpose-built for globally distributed applications. Aurora global database provides a primary cluster in one Region with low-latency read replicas in secondary Regions and fast cross-Region replication, enabling low-latency reads and high availability for global users.

Option A uses only a single Region, which does not meet the “global users with low latency” requirement.

Option C uses RDS and DMS-based replication, which requires more management and does not provide Aurora's integrated global database features.

Option D replaces the relational database with DynamoDB, which violates the requirement that the application interacts with a relational database.

Question: 617

A company is storing data that will not be frequently accessed in the AWS Cloud. If the company needs to access the data, the data must be retrieved within 12 hours. The company wants a solution that is cost-effective for storage costs per gigabyte.

Which Amazon S3 storage class will meet these requirements?

- A. S3 Standard
- B. S3 Glacier Flexible Retrieval
- C. S3 One Zone-Infrequent Access (S3 One Zone-IA)
- D. S3 Standard-Infrequent Access (S3 Standard-IA)

Answer: B

Explanation:

Amazon S3 Glacier Flexible Retrieval is designed for long-lived, infrequently accessed data where retrieval times of minutes to hours are acceptable. It offers very low storage cost per GB while still supporting faster retrieval than archival tiers that can take many hours or days.

The requirement states that data is rarely accessed and must be retrievable within 12 hours. Glacier Flexible Retrieval meets this retrieval-time requirement and is more cost-effective per GB than Standard and Infrequent Access storage classes for long-term, rarely accessed

data.

S3 Standard, S3 Standard-IA, and S3 One Zone-IA are optimized for more frequent or lower-latency access and have higher storage cost per GB compared to Glacier Flexible Retrieval.

Question: 618

A home security company is expanding its business globally. The company needs to encrypt customer data. The company does not want to manage its own keys. The company needs the keys to be usable in multiple AWS Regions and needs to control access to the keys.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Key Management Service (AWS KMS) to create multi-Region keys. Apply tags to identify each key. Use attribute-based access control (ABAC) condition keys to control access to the keys.
- B. Use AWS Key Management Service (AWS KMS) to create multiple keys by importing key material. Apply tags to identify each key. Use attribute-based access control (ABAC) condition keys to control access to the keys.
- C. Use AWS CloudHSM to create a CloudHSM cluster in the company's primary Region. Synchronize the CloudHSM cluster to additional Regions by using the CloudHSM Management Utility (CMU).
- D. Use AWS CloudHSM to create users. Use the CloudHSM Management Utility (CMU) to share keys with the users. Use the shareKey command to share or unshare the key with additional users in each Region.

Answer: A

Explanation:

AWS Key Management Service (AWS KMS) supports multi-Region keys, which are managed customer master keys that can be replicated to multiple Regions and treated as a single logical key. This allows applications in different Regions to encrypt and decrypt data using equivalent keys while AWS handles key replication, durability, and availability.

KMS is a fully managed key service, so the company does not need to operate hardware or manage low-level key storage. Tags combined with attribute-based access control (ABAC) let you enforce finegrained access to keys by using IAM policies and condition keys, giving centralized and flexible access control with low operational overhead.

CloudHSM (Options C and D) requires managing HSM clusters, scaling, backups, and manual key operations, which significantly increases operational burden. Importing key material in KMS (Option B) adds lifecycle complexity and is unnecessary when native multi-Region keys exist.

Question: 619

A company deployed a three-tier web application in a single Availability Zone in the us-east-1 Region on a single Amazon EC2

instance. Usage of the application is growing.

A solutions architect needs to ensure that the application can handle the growing amount of traffic and that the application is resilient. The solution must be cost-effective.

Which solution will meet these requirements MOST cost-effectively?

- A. Create two additional EC2 instances spread across two separate Availability Zones. Create an Application Load Balancer (ALB). Configure the ALB to route traffic to a target group that contains all three instances. Create an Amazon CloudWatch alarm to scale the EC2 instances vertically to handle the application traffic.
- B. Create eight additional EC2 instances spread across three separate Availability Zones. Create an Application Load Balancer (ALB). Configure the ALB to route traffic to a target group that contains all nine instances. Create an Amazon CloudWatch alarm to scale the EC2 instances horizontally to handle the application traffic.
- C. Create an EC2 Auto Scaling group that contains a minimum of three EC2 instances in the same Availability Zone. Create an Application Load Balancer (ALB). Configure the ALB to route traffic to a target group that contains all the instances. Configure scheduled scaling for the Auto Scaling group.
- D. Create an EC2 Auto Scaling group that contains a minimum of three EC2 instances spread across Availability Zones. Create an Application Load Balancer (ALB). Configure the ALB to route traffic to a target group that contains all the instances. Create an Amazon CloudWatch alarm to scale the EC2 instances horizontally to handle the application traffic.

Answer: D

Explanation:

AWS best practice for a highly available, resilient web application is: Run EC2 instances in an Auto Scaling group across multiple Availability Zones.

Put an Application Load Balancer (ALB) in front of the Auto Scaling group.

Use CloudWatch alarms and scaling policies for horizontal scaling based on demand.

Option D implements exactly this: a minimum of three instances spread across AZs (resilience to AZ failure), behind an ALB, with automatic horizontal scaling. This provides both elasticity and high availability with good cost efficiency (instances scale out/in based on traffic).

Option A uses only three fixed instances and relies on vertical scaling, which does not scale as well and is less resilient.

Option B overprovisions to nine instances from the start, which is unnecessarily expensive.

Option C keeps all instances in a single AZ, which does not meet resilience requirements.

Question: 620

A data science team requires storage for nightly log processing. The size and number of logs is unknown and the logs will persist for

24 hours only.

What is the MOST cost-effective solution?

- A. Amazon S3 Glacier Deep Archive
- B. Amazon S3 Standard
- C. Amazon S3 Intelligent-Tiering
- D. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B

Explanation:

For logs that are:

Written and processed within a short period (24 hours)

Accessed quickly for compute/analytics

With unknown object count and size

Amazon S3 Standard is the most appropriate and cost-effective. Intelligent-Tiering is designed for data stored for longer periods (typically 30+ days) with changing access patterns and charges a per-object monitoring and automation fee that becomes inefficient for very short-lived objects.

S3 Glacier Deep Archive and S3 One Zone-IA are optimized for long-term archival or infrequently accessed data with retrieval time or availability constraints that are not suitable for nightly active processing.

Question: 621

A company has an Amazon S3 data lake that is governed by AWS Lake Formation. The company wants to create a visualization in Amazon QuickSight by joining the data in the data lake with operational data that is stored in an Amazon Aurora MySQL database. The company wants to enforce column-level authorization so that the company's marketing team can access only a subset of columns in the database.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon EMR to ingest the data directly from the database to the QuickSight SPICE engine. Include only the required columns.
- B. Use AWS Glue Studio to ingest the data from the database to the S3 data lake. Attach an IAM policy to the QuickSight users to enforce column-level access control. Use Amazon S3 as the data source in QuickSight.
- C. Use AWS Glue Elastic Views to create a materialized view for the database in Amazon S3. Create an S3 bucket policy to enforce column-

level access control for the QuickSight users. Use Amazon S3 as the data source in QuickSight.

D. Use a Lake Formation blueprint to ingest the data from the database to the S3 data lake. Use Lake Formation to enforce column-level access control for the QuickSight users. Use Amazon Athena as the data source in QuickSight.

Answer: D

Explanation:

AWS Lake Formation provides centralized data access control, including fine-grained (column-level) permissions for data stored in S3 and accessed through services like Amazon Athena.

Using a Lake Formation blueprint to ingest data from Aurora MySQL into the data lake keeps ingestion and governance integrated. When QuickSight uses Athena as the data source, Athena enforces Lake Formation's column-level permissions automatically. This allows the marketing team to see only the authorized subset of columns without custom access-control logic.

Options A, B, and C rely on manually limiting columns at ingestion time or using IAM or S3 bucket policies, which do not provide true column-level authorization for SQL queries and require significantly more manual work and maintenance.

Question: 622

A transaction processing company has weekly scripted batch jobs that run on Amazon EC2 instances. The EC2 instances are in an Auto Scaling group. The number of transactions can vary, but the baseline CPU utilization that is noted on each run is at least 60%. The company needs to provision the capacity 30 minutes before the jobs run.

Currently, engineers complete this task by manually modifying the Auto Scaling group parameters. The company does not have the resources to analyze the required capacity trends for the Auto Scaling group counts. The company needs an automated way to modify the Auto Scaling group's desired capacity.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create a dynamic scaling policy for the Auto Scaling group. Configure the policy to scale based on the CPU utilization metric. Set the target value for the metric to 60%.

B. Create a scheduled scaling policy for the Auto Scaling group. Set the appropriate desired capacity, minimum capacity, and maximum capacity. Set the recurrence to weekly. Set the start time to 30 minutes before the batch jobs run.

C. Create a predictive scaling policy for the Auto Scaling group. Configure the policy to scale based on forecast. Set the scaling metric to CPU utilization. Set the target value for the metric to 60%. In the policy, set the instances to pre-launch 30 minutes before the jobs run.

D. Create an Amazon EventBridge event to invoke an AWS Lambda function when the CPU utilization metric value for the Auto Scaling group reaches 60%. Configure the Lambda function to increase the Auto Scaling group's desired capacity and maximum capacity by 20%.

Explanation:

Predictive scaling in EC2 Auto Scaling uses machine learning to analyze historical load and forecast future capacity needs, then pre-scales capacity before the demand occurs. It is specifically designed for recurring, cyclical workloads such as weekly batch jobs.

In this scenario:

The workload pattern is known and recurring (weekly scripted jobs).

The company explicitly does not want to manually analyze trends.

They need capacity to be ready 30 minutes before jobs start.

Predictive scaling lets you:

Set the metric (CPU utilization) and target (60%).

Configure the policy to pre-launch instances 30 minutes in advance, satisfying the requirement automatically with minimal operational overhead.

Dynamic scaling (Option A) reacts after CPU increases, not before.

Scheduled scaling (Option B) requires manual analysis and ongoing adjustment of capacity values.

EventBridge + Lambda (Option D) introduces unnecessary custom logic and still reacts after CPU reaches 60%, not 30 minutes before.

Question: 623

A company's solutions architect is building a static website to be deployed in Amazon S3 for a production environment. The website integrates with an Amazon Aurora PostgreSQL database by using an AWS Lambda function. The website that is deployed to production will use a Lambda alias that points to a specific version of the Lambda function.

The company must rotate the database credentials every 2 weeks. Lambda functions that the company deployed previously must be able to use the most recent credentials.

Which solution will meet these requirements?

- A. Store the database credentials in AWS Secrets Manager. Turn on rotation. Write code in the Lambda function to retrieve the credentials from Secrets Manager.
- B. Include the database credentials as part of the Lambda function code. Update the credentials periodically and deploy the new Lambda function.
- C. Use Lambda environment variables. Update the environment variables when new credentials are available.

D. Store the database credentials in AWS Systems Manager Parameter Store. Turn on rotation. Write code in the Lambda function to retrieve the credentials from Systems Manager Parameter Store.

Answer: A

Explanation:

AWS Secrets Manager is the managed service for securely storing, rotating, and retrieving database credentials. When you store Aurora credentials in Secrets Manager and enable automatic rotation, Secrets Manager updates the credentials in both the database and the stored secret.

Each Lambda function version or alias can call Secrets Manager at runtime to retrieve the current secret value, so even older deployed Lambda versions that use an alias will always obtain the most recent credentials without redeployment.

Why others are not suitable:

B: Embeds credentials in code, requiring redeployment on every rotation and violating security best practices.

C: Environment variables are version-specific; old aliases would continue using outdated values unless you redeploy or change them.

D: Parameter Store can store and rotate secrets but is less integrated for database credential rotation than Secrets Manager; Secrets Manager is the purpose-built minimal-overhead choice here.

Question: 624

A company has an on-premises volume backup solution that has reached its end of life. The company wants to use AWS as part of a new backup solution and wants to maintain local access to all the data while it is backed up on AWS. The company wants to ensure that the data backed up on AWS is automatically and securely transferred.

Which solution meets these requirements?

A. Use AWS Snowball to migrate data out of the on-premises solution to Amazon S3. Configure on-premises systems to mount the Snowball S3 endpoint to provide local access to the data.

B. Use AWS Snowball Edge to migrate data out of the on-premises solution to Amazon S3. Use the Snowball Edge file interface to provide on-premises systems with local access to the data.

C. Use AWS Storage Gateway and configure a cached volume gateway. Run the Storage Gateway software appliance on premises and configure a percentage of data to cache locally. Mount the gateway storage volumes to provide local access to the data.

D. Use AWS Storage Gateway and configure a stored volume gateway. Run the Storage Gateway software appliance on premises and map the gateway storage volumes to on-premises storage. Mount the gateway storage volumes to provide local access to the data.

Answer: D

Explanation:

AWS Storage Gateway stored volume gateway is designed for use cases where:

The primary data set remains on premises.

You want to maintain full local access to data with low latency.

You need asynchronous, secure replication and backup to AWS.

With stored volumes, you:

Map gateway storage volumes to local on-premises disks.

Present these volumes as iSCSI block devices to on-prem systems.

Storage Gateway then asynchronously backs up snapshots to Amazon S3, providing secure, durable offsite backup automatically.

Cached volume gateway (Option C) keeps the primary copy in AWS and only a subset cached locally, which does not meet the requirement to maintain full local access to all data.

Snowball and Snowball Edge (Options A and B) are primarily for bulk, one-time or periodic offline data transfer and do not provide an ongoing backup solution with continuous local access.

Question: 625

A company is using a loosely coupled serverless architecture on AWS. The architecture consists of multiple web applications and APIs distributed across multiple teams. The company uses AWS Control Tower to provision AWS accounts. The company's development teams use AWS CloudFormation.

The company wants to improve trace monitoring and gain insight into how individual services in application stacks are performing.

Which solution will meet these requirements?

- A. Enable AWS CloudTrail across all accounts by using AWS Control Tower.
- B. Enable AWS X-Ray across all accounts by using AWS Control Tower.
- C. Enable Amazon CloudWatch in the CloudFormation templates.
- D. Enable AWS X-Ray in the CloudFormation templates.

Answer: D

Explanation:

AWS X-Ray is the AWS service designed for distributed tracing, giving end-to-end visibility into how requests flow through microservices, APIs, and serverless components. It provides service maps, latency breakdowns, and performance insight per service, which is exactly what is required.

To use X-Ray effectively, each application stack must be instrumented. For infrastructure deployed through AWS CloudFormation, best practice is to enable and configure X-Ray in the CloudFormation templates (for Lambda, API Gateway, ECS, etc.). That way, all teams and stacks have consistent tracing enabled as part of their IaC.

Control Tower (Option B) cannot “turn on” X-Ray tracing inside applications; it only helps with account setup and guardrails.

CloudTrail (Option A) is for audit logging of API calls, not performance tracing.

CloudWatch (Option C) provides metrics and logs, but not request-level distributed tracing across multiple services.

Question: 626

A financial company is migrating its banking applications to a set of AWS accounts managed by AWS Organizations. The applications will store sensitive customer data on Amazon Elastic Block Store (Amazon EBS) volumes. The company will take regular snapshots for backup purposes.

The company wants to implement controls across all AWS accounts to prevent sharing EBS snapshots publicly.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Enable AWS Config rules for each organizational unit (OU) in Organizations to monitor EBS snapshot permissions.
- B. Enable block public access for EBS snapshots at the organization level.
- C. Create an IAM policy in the root account of the organization that prevents users from modifying snapshot permissions.
- D. Use AWS CloudTrail to track snapshot permission changes.

Answer: B

Explanation:

AWS provides EBS Block Public Access at the AWS Organizations level, which, when enabled, prevents EBS snapshots from being shared publicly across all member accounts. This is an organization-wide control that requires minimal configuration and no ongoing monitoring to prevent public sharing.

This directly satisfies the requirement:

Covers all accounts managed by Organizations.

Explicitly prevents public snapshot sharing.

Has the least operational overhead because it is a single centralized control.

AWS Config (Option A) only monitors and alerts, but does not prevent public sharing.

An IAM policy in the root account (Option C) is harder to enforce consistently across all accounts and may not cover all permission paths.

CloudTrail (Option D) only logs events and does not prevent public access.

Question: 627

A company is building a solution to provide customers with an API that accesses financial data. The API backend needs to compute tax data for each request. The company anticipates greater demand to access the data during the last 3 months of each year.

A solutions architect needs to design a scalable solution that can meet the regular demand and the peak demand at the end of each year.

Which solution will meet these requirements?

- A. Host the API on an Amazon EC2 instance that runs third-party software. Configure the EC2 instance to perform tax computations.
- B. Deploy an Amazon API Gateway REST API. Create an AWS Lambda function to perform tax computations. Integrate the Lambda function with the REST API.
- C. Create an Application Load Balancer (ALB) in front of two Amazon EC2 instances. Configure the EC2 instances to perform tax computations.
- D. Deploy an Amazon API Gateway REST API. Configure an Amazon EC2 instance to perform tax computations. Integrate the EC2 instance with the REST API.

Answer: B

Explanation:

For bursty and seasonal workloads, AWS recommends serverless architectures using Amazon API Gateway + AWS Lambda:

API Gateway exposes a fully managed, scalable API endpoint.

Lambda provides automatic scaling based on request volume, with no need to provision or manage servers.

This is ideal for workloads with variable or seasonal demand, such as end-of-year tax computation bursts.

This combination minimizes operational overhead (no OS patching, no capacity planning) while scaling seamlessly to meet peak demand.

Options A, C, and D rely on EC2 instances. They require capacity planning, scaling configuration, and ongoing management. They do not

scale as effortlessly as Lambda for unpredictable or highly seasonal workloads.

Question: 628

A company hosts dozens of multi-tier applications on AWS. The presentation layer and logic layer are Amazon EC2 Linux instances that use Amazon EBS volumes.

The company needs a solution to ensure that operating system vulnerabilities are not introduced to the EC2 instances when the company deploys new features. The company uses custom AMIs to deploy EC2 instances in an Auto Scaling group. The solution must scale to handle all applications that the company hosts.

Which solution will meet these requirements?

- A. Use Amazon Inspector to patch operating system vulnerabilities. Invoke Amazon Inspector when a new AMI is deployed.
- B. Use AWS Backup to back up the EBS volume of each updated instance. Use the EBS backup volumes to create new AMIs. Use the existing Auto Scaling group to deploy the new AMIs.
- C. Use AWS Systems Manager Patch Manager to patch operating system vulnerabilities in the custom

AMIs.

- D. Use EC2 Image Builder to create new AMIs when the company deploys new features. Include the update-linux component in the build components of the new AMIs. Use the existing Auto Scaling group to deploy the new AMIs.

Answer: D

Explanation:

EC2 Image Builder is the AWS-managed service specifically designed to automate the creation, patching, hardening, and testing of AMIs.

For this scenario, best practice is to:

Use EC2 Image Builder pipelines to generate new golden AMIs whenever features are deployed or on a schedule.

Include build components such as update-linux to automatically apply OS security patches and updates during image creation.

Deploy the new AMIs through existing Auto Scaling groups, ensuring that every newly launched instance is already patched and compliant.

This approach:

Prevents OS vulnerabilities from being introduced at deployment time.

Scales across dozens of applications because AMI pipelines are reusable and automated.

Minimizes manual effort and reduces drift between instances.

Amazon Inspector (Option A) identifies vulnerabilities but does not itself patch AMIs.

AWS Backup (Option B) is for data protection, not vulnerability management.

Patch Manager (Option C) patches running instances, but the requirement is to ensure vulnerabilities are not introduced with new AMIs; golden image pipelines with EC2 Image Builder are the recommended solution.

Question: 629

A company has an organization in AWS Organizations. The company runs Amazon EC2 instances across four AWS accounts in the root organizational unit (OU). There are three nonproduction accounts and one production account. The company wants to prohibit users from launching EC2 instances of a certain size in the nonproduction accounts. The company has created a service control policy (SCP) to deny access to launch instances that use the prohibited types.

Which solutions to deploy the SCP will meet these requirements? (Select TWO.)

- A. Attach the SCP to the root OU for the organization.
- B. Attach the SCP to the three nonproduction Organizations member accounts.
- C. Attach the SCP to the Organizations management account.
- D. Create an OU for the production account. Attach the SCP to the OU. Move the production member account into the new OU.
- E. Create an OU for the required accounts. Attach the SCP to the OU. Move the nonproduction member accounts into the new OU.

Answer: B,E

Explanation:

Service control policies (SCPs) in AWS Organizations apply to the accounts or OUs to which they are attached:

To restrict only nonproduction accounts, you can:

Attach the SCP directly to those specific member accounts (Option B), or

Create a dedicated OU for the nonproduction accounts and attach the SCP to that OU (Option E).

Both methods ensure that the deny policy affects only the nonproduction accounts and does not impact the production account.

Why the others are incorrect:

A: Attaching the SCP to the root OU would apply the deny to all accounts, including production, which does not meet the requirement.

C: SCPs attached to the management account only affect that account, not all member accounts.

D: Creating an OU for the production account and attaching the deny SCP there would block the prohibited instance types in production, which is the opposite of what is required.

Question: 630

A company is building a data analysis platform on AWS by using AWS Lake Formation. The platform will ingest data from different sources such as Amazon S3 and Amazon RDS. The company needs a secure solution to prevent access to portions of the data that contain sensitive information.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM role that includes permissions to access Lake Formation tables.
- B. Create data filters to implement row-level security and cell-level security.
- C. Create an AWS Lambda function that removes sensitive information before Lake Formation ingests the data.
- D. Create an AWS Lambda function that periodically queries and removes sensitive information from Lake Formation tables.

Answer: B

Explanation:

AWS Lake Formation natively supports fine-grained access control, including:

Row-level security

Cell/column-level security

These are implemented through data filters and Lake Formation permissions on governed tables and views. This allows you to centrally define which users or groups can access which rows and which columns, including sensitive data fields, with minimal custom code or maintenance.

Why others are incorrect:

A: IAM alone does not provide row-level or cell-level data security inside Lake Formation tables.

C and D: Using Lambda to scrub or periodically remove sensitive data is complex, error-prone, and high overhead compared to built-in Lake Formation data filters.

Question: 631

A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access

pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours.

Which solution meets these requirements?

- A. Use Amazon DynamoDB with auto scaling. Use on-demand backups and Amazon DynamoDB Streams.
- B. Use Amazon Redshift. Configure concurrency scaling. Activate audit logging. Perform database snapshots every 4 hours.
- C. Use Amazon RDS with Provisioned IOPS. Activate the database auditing parameter. Perform database snapshots every 5 hours.
- D. Use Amazon Aurora MySQL with auto scaling. Activate the database auditing parameter.

Answer: D

Explanation:

Amazon Aurora MySQL provides:

Continuous, incremental backups to Amazon S3 with point-in-time recovery (PITR), allowing recovery to any point within the retention window (typically up to 35 days). This easily achieves an RPO of less than 5 hours, often down to seconds.

Support for database auditing parameters so audit logs can be retained and managed (for example, in database logs or external log services) to meet the 7-day audit requirement.

Auto scaling (via read replicas, Aurora Serverless v2, or capacity management) to handle variable read/write demand throughout the year with minimal operational overhead.

Why others are less suitable:

A: DynamoDB is NoSQL and does not directly satisfy typical relational database + SQL audit requirements; Streams also only retain 24 hours, not 7 days, and on-demand backups do not inherently give an RPO < 5 hours without frequent scheduling.

B: Amazon Redshift is a data warehouse, not a primary transactional application database.

C: Snapshots every 5 hours give an RPO of up to 5 hours, not less than 5 hours, and rely on discrete snapshot points rather than continuous PITR.

Question: 632

A company is developing an application that uses an Amazon Aurora MySQL database. The company plans to regularly make changes to the MySQL database schema to test new features. The tests must not affect the existing production database.

When the company finishes testing, a developer needs to replicate the changes to the production database. The solution must cause minimal downtime.

Which solution will meet these requirements?

- A. Create a new staging Aurora MySQL database cluster based on the existing database. Make the schema changes to the new staging database cluster to test the new features.
- B. Create a read replica based on the existing Aurora MySQL database. Make the schema changes to the read replica. Promote the read replica to primary after successful testing.
- C. Create a blue/green deployment of the Aurora MySQL database. Make schema changes in the staging environment to test new features. Direct traffic from the green environment to the blue environment when testing is complete.
- D. Replicate the Aurora MySQL database to an Amazon DynamoDB table. Make the schema changes to the DynamoDB table to test the new features. Configure the application to use the DynamoDB table when testing is complete.

Answer: C

Explanation:

Aurora blue/green deployments are specifically designed for:

- Safely testing schema changes and database upgrades in an isolated environment (green) without impacting production (blue).
- Performing a fast, low-downtime switchover once testing is complete and validated.

In this pattern:

The blue environment is the current production database.

The green environment is a synchronized copy used for testing changes.

You apply schema changes to the green environment, run tests, and when ready, perform a managed switchover that minimizes downtime and risk.

Why others are not ideal:

A: A separate staging cluster allows testing but does not provide automated, low-downtime synchronization and switchover.

B: Aurora read replicas are for read scaling; schema changes on replicas are not supported in the way described, and promotion alone doesn't solve controlled testing and replication of changes.

D: Moving to DynamoDB changes database engines and data models entirely, and does not match the requirement to keep using Aurora MySQL.

Question: 633

A company deployed an application in two AWS Regions. If the application becomes unavailable in one Region, the application must fail

over to the second Region. The failover process must avoid stale DNS client caches. The company wants to use one endpoint to access both copies of the application.

Which solution will meet these requirements?

- A. Use an Amazon CloudFront distribution that has multiple origins. Correlate each origin with the application in each Region.
- B. Use an Amazon Route 53 weighted routing policy that uses equal weights to route client requests to the second Region if the application becomes unavailable in the original Region.
- C. Use AWS Global Accelerator, and assign a static anycast IP address to the application endpoint.
- D. Use an Amazon Route 53 IP-based routing policy to route requests to the second Region if the application becomes unavailable in the original Region.

Answer: C

Explanation:

AWS Global Accelerator provides static anycast IP addresses that are announced from AWS edge locations. Traffic sent to these IPs is routed over the AWS global network to the nearest healthy endpoint (for example, an ALB or NLB) in any Region. Because clients connect to static IPs instead of DNS names that change, failover does not depend on DNS TTLs or client cache expiry, which avoids stale DNS entries.

You can register endpoints in multiple Regions and let Global Accelerator automatically fail over when health checks fail, all behind a single global endpoint.

CloudFront (Option A) is primarily for HTTP/HTTPS content caching and still uses DNS; it is not intended as a generic multi-Region failover solution for arbitrary applications.

Route 53 (Options B and D) relies on DNS; clients can cache DNS responses and may keep sending traffic to an unhealthy Region until TTLs expire, which does not meet the “avoid stale DNS client

caches” requirement.

Question: 634

A company manages an application that stores data on an Amazon RDS for PostgreSQL Multi-AZ DB instance. High traffic on the application is causing increased latency for many read queries.

A solutions architect must improve the performance of the application.

Which solution will meet this requirement?

- A. Enable Amazon RDS Performance Insights. Configure storage capacity to scale automatically.

- B. Configure the DB instance to use DynamoDB Accelerator (DAX).
- C. Create a read replica of the DB instance. Serve read traffic from the read replica.
- D. Use Amazon Data Firehose between the application and Amazon RDS to increase the concurrency of database requests.

Answer: C

Explanation:

For Amazon RDS relational databases experiencing read-heavy workloads, AWS best practice is to create read replicas and offload read traffic to those replicas. A read replica:

Uses asynchronous replication from the primary.

Can serve read-only queries, thereby reducing load and query latency on the primary.

Can be used behind an application-level read/write splitting mechanism or via endpoints that direct read traffic to replicas.

Performance Insights (Option A) helps diagnose performance problems but does not itself offload traffic.

DAX (Option B) is a cache for DynamoDB, not RDS.

Kinesis Data Firehose (Option D) is a streaming ingestion service and cannot increase RDS query concurrency.

Question: 635 A company is deploying an application in three AWS Regions using an Application Load Balancer. Amazon Route 53 will be used to distribute traffic between these Regions.

Which Route 53 configuration should a solutions architect use to provide the MOST high-performing experience?

- A. Create an A record with a latency policy.
- B. Create an A record with a geolocation policy.
- C. Create a CNAME record with a failover policy.
- D. Create a CNAME record with a geoproximity policy.

Answer: A

Explanation:

Latency-based routing in Amazon Route 53 is designed to route users to the Region that provides the lowest network latency, based on

Amazon's measurements of latency between AWS Regions and users' networks. For applications deployed in multiple Regions, this provides the highest performance experience for global users.

Therefore, creating an A record with a latency routing policy is the correct choice.

Geolocation (Option B) routes based on user location, which may not always correspond to the lowest latency.

Failover (Option C) is for active-passive architectures, not performance optimization.

Geoproximity (Option D) is more complex and focused on directing traffic based on geographic bias rather than measured latency.

Question: 636

A company is creating a mobile financial app that gives users the ability to sign up and store personal information. The app uses an Amazon DynamoDB table to store user details and preferences.

The app generates a credit score report by using the data that is stored in DynamoDB. The app sends credit score reports to users once every month.

The company needs to provide users with an option to remove their data and preferences. The app must delete customer data within one month of receiving a request to delete the data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Lambda function to delete user information. Create an Amazon EventBridge rule that runs when a specified TTL expires. Configure the EventBridge rule to invoke the Lambda function.
- B. Create a DynamoDB stream. Create an AWS Lambda function to delete user information. When a specified TTL expires, write user information to the DynamoDB stream from the DynamoDB table. Configure the DynamoDB stream to invoke the Lambda function to delete user information.
- C. Enable TTL in DynamoDB. Set the expiration date as an attribute. Create an AWS Lambda function to set the TTL based on the expiration date value. Invoke the Lambda function when a user requests to delete personal data.
- D. Enable TTL in DynamoDB. Create an AWS Lambda function to delete user information. Configure AWS Config to detect the DynamoDB state change when TTL expires and to invoke the Lambda function.

Answer: C

Explanation:

Amazon DynamoDB supports Time to Live (TTL), which automatically and asynchronously deletes expired items based on a timestamp attribute. TTL is ideal for automatic data expiration with very low operational overhead.

In this scenario:

When a user requests deletion, the system can calculate an expiration timestamp one month in the future.

A Lambda function is used once per request to write or update the item's TTL attribute to that timestamp (Option C).

DynamoDB TTL then automatically removes the item after the expiration time. Deletion typically occurs within 48 hours of the TTL timestamp, which satisfies the “within one month” requirement.

This approach offloads the actual deletion to DynamoDB and avoids building complex orchestration or periodic cleanup jobs.

Options A, B, and D add unnecessary components (EventBridge, streams, AWS Config, and Lambdabased deletion workflows) on top of TTL or custom logic, increasing operational overhead without providing additional value for the given requirement.

Question: 637

A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS) and the Kubernetes Horizontal Pod Autoscaler. The workload is not consistent throughout the day. A solutions architect notices that the number of nodes does not automatically scale out when the existing nodes have reached maximum capacity in the cluster, which causes performance issues.

Which solution will resolve this issue with the LEAST administrative overhead?

- A. Scale out the nodes by tracking the memory usage.
- B. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
- C. Use an AWS Lambda function to resize the EKS cluster automatically.
- D. Use an Amazon EC2 Auto Scaling group to distribute the workload.

Answer: B

Explanation:

The Kubernetes Horizontal Pod Autoscaler (HPA) scales pods, not nodes. When the cluster runs out of node capacity, the HPA cannot schedule more pods.

On Amazon EKS, AWS recommends using the Kubernetes Cluster Autoscaler to automatically adjust the number of nodes in the cluster's underlying Auto Scaling group based on pending pods.

Cluster Autoscaler watches for pods that cannot be scheduled because of insufficient resources and then scales out nodes automatically with minimal administrative overhead. It also scales in nodes when they are underutilized.

Why others are not correct:

A: Just tracking memory usage does not automatically scale nodes or integrate with Kubernetes scheduling.

C: A custom Lambda-based resizer is unnecessary undifferentiated heavy lifting compared to the native Cluster Autoscaler.

D: An Auto Scaling group is already typically used under EKS, but by itself it does not respond to pod scheduling pressure without Cluster Autoscaler.

Question: 638

A media company needs to migrate its Windows-based video editing environment to AWS. The company's current environment processes 4K video files that require sustained throughput of 2 GB per second across multiple concurrent users.

The company's storage needs increase by 1 TB each week. The company needs a shared file system that supports SMB protocol and can scale automatically based on storage demands.

Which solution will meet these requirements?

- A. Deploy an Amazon FSx for Windows File Server Multi-AZ file system with SSD storage.
- B. Deploy an Amazon Elastic File System (Amazon EFS) file system in Max I/O mode. Provision mount targets in multiple Availability Zones.
- C. Deploy an Amazon FSx for Lustre file system with a Persistent 2 deployment type. Provision the file system with 2 TB of storage.
- D. Deploy Amazon S3 File Gateway by using multiple cached gateway instances. Configure S3 Transfer Acceleration.

Answer: A

Explanation:

The workload is Windows-based and requires a shared file system with SMB support and very high throughput for 4K video editing. Amazon FSx for Windows File Server is a fully managed, highly available native Windows file system that supports the SMB protocol, Active Directory integration, and can deliver high throughput and low latency suitable for media workloads.

FSx for Windows File Server supports automatic storage scaling to grow file system capacity as data increases, matching the requirement of +1 TB per week with minimal admin effort.

A Multi-AZ SSD deployment provides high availability and performance.

Why others are not correct:

B: Amazon EFS is NFS, not SMB, and is optimized for Linux clients, not Windows-based environments.

C: FSx for Lustre is a high-performance POSIX file system typically used with Linux-based HPC workloads, not Windows + SMB.

D: S3 File Gateway is not designed for sustained multi-GB/s shared-edit performance and adds additional complexity; it is more suited for backup/archival and limited shared file access.

Question: 639

A company has built an application that uses an Amazon Simple Queue Service (Amazon SQS) standard queue and an AWS Lambda function. The Lambda function writes messages to the SQS queue. The company needs a solution to ensure that the consumer of the SQS queue never receives duplicate messages.

Which solution will meet this requirement with the FEWEST changes to the current architecture?

- A. Modify the SQS queue to enable long polling for the queue.
- B. Delete the existing SQS queue. Recreate the queue as a FIFO queue. Enable content-based deduplication for the queue.
- C. Modify the SQS queue to enable content-based deduplication for the queue.
- D. Delete the SQS queue. Create an Amazon MQ message broker. Configure the broker to deduplicate messages.

Answer: B

Explanation:

SQS Standard queues provide at-least-once delivery; they can, by design, deliver duplicate messages, and you cannot turn that off.

To guarantee no duplicates to consumers, AWS provides SQS FIFO queues with exactly-once processing semantics in conjunction with deduplication.

The minimal change to achieve this behavior is to:

Recreate the queue as a FIFO queue and

Enable content-based deduplication, so SQS automatically deduplicates messages with identical bodies within the deduplication window.

Why others are not correct:

A: Long polling reduces empty responses and costs but does not eliminate duplicates.

C: Content-based deduplication is available only for FIFO queues, not for standard queues—so you cannot “modify” a standard queue this way.

D: Moving to Amazon MQ is a much bigger architectural change and adds more operational overhead; it's not the “fewest changes” approach.

Question: 640

A company runs an ecommerce platform with a monolithic architecture on Amazon EC2 instances. The platform runs web and API services. The company wants to decouple the architecture and enhance scalability. The company also wants the ability to track orders and reprocess

any failed orders.

Which solution will meet these requirements?

- A. Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure AWS Lambda functions to consume the queue and process orders. Implement an SQS dead-letter queue.
- B. Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure Amazon Elastic Container Service (Amazon ECS) tasks to consume the queue. Implement SQS visibility timeout.
- C. Use Amazon Kinesis Data Streams to queue orders. Use AWS Lambda functions to consume the data stream. Configure Amazon S3 to track and reprocess failed orders.
- D. Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure AWS Lambda functions to consume the queue and process orders. Configure the Lambda functions to use SQS long polling.

Answer: A

Explanation:

To decouple the monolith and enhance scalability, AWS best practice is to introduce an asynchronous message queue, such as Amazon SQS, between the web/API tier and the order-processing logic.

AWS Lambda functions consuming from the SQS queue provide serverless, auto-scaling processing without managing servers.

To track and reprocess failed orders, SQS supports dead-letter queues (DLQs). Messages that cannot be processed successfully after a configurable number of attempts are automatically moved to the DLQ, where operations teams or automated processes can inspect and reprocess them.

Why others are not correct:

B: ECS tasks can consume an SQS queue, but this requires managing container infrastructure and does not inherently provide as simple reprocessing/visibility as combining Lambda with a DLQ.

Visibility timeout is not a tracking or archival mechanism.

C: Kinesis is a streaming service designed for ordered event streams, not primarily for order-queue semantics and DLQs; SQS is simpler and purpose-built for this pattern.

D: Long polling reduces empty responses and API calls but does nothing for tracking or reprocessing failed messages; without a DLQ, failed orders are harder to manage

Question: 641

A company is building a web application. The company needs a load balancing solution that supports HTTPS header-based routing. The

company's security team also requires a rules-based method of blocking specific incoming requests to decrease the effects of malicious activity.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB). Configure an HTTPS listener with mutual TLS enabled.
- B. Create an Application Load Balancer (ALB). Integrate the ALB with AWS WAF. Configure the security team's required rules.
- C. Create an Application Load Balancer (ALB). Integrate the ALB with AWS Config. Apply custom rules to all ALB resources.
- D. Create a Network Load Balancer (NLB). Configure AWS Network Firewall with the security team's required rules.

Answer: B

Explanation:

Application Load Balancer (ALB) supports HTTP/HTTPS layer 7 routing, including header-based routing, path-based routing, and host-based routing.

AWS WAF is designed to provide rules-based filtering (block, allow, count) of HTTP(S) requests to protect applications from common exploits and malicious traffic.

ALB integrates directly with AWS WAF, so you can attach a web ACL with custom rules defined by the security team to block specific patterns while still using header-based routing.

Why the others are not correct:

A: Mutual TLS adds client certificate authentication but does not provide rules-based blocking or WAF-style inspection.

C: AWS Config is for configuration compliance and auditing, not request filtering.

D: NLB operates at layer 4 (TCP/UDP); it does not support HTTP header-based routing.

Question: 642

A media publishing company is building an application on AWS to give users the ability to print their own books. The application frontend runs on a Docker container.

The amount of incoming orders varies significantly. The incoming orders can temporarily exceed the throughput of the company's book printing machines. Order-processing payloads are up to 4 MB in size.

The company needs to develop a solution that can scale to handle incoming orders.

Which solution will meet this requirement?

- A. Use Amazon Simple Queue Service (Amazon SQS) to queue incoming orders. Create an AWS Lambda@Edge function to process orders. Deploy the frontend application on Amazon Elastic Kubernetes Service (Amazon EKS).
- B. Use Amazon Simple Notification Service (Amazon SNS) to queue incoming orders. Create an AWS Lambda function to process orders. Deploy the frontend application on AWS Fargate.
- C. Use Amazon Simple Queue Service (Amazon SQS) to queue incoming orders. Create an AWS Lambda function to process orders. Deploy the frontend application on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type.
- D. Use Amazon Simple Notification Service (Amazon SNS) to queue incoming orders. Create an AWS Lambda@Edge function to process orders. Deploy the frontend application on Amazon EC2 instances.

Answer: C

Explanation:

The workload is bursting and must be able to buffer more orders than the printers can process. The best practice for this is a durable, scalable queue: Amazon SQS.

SQS decouples the frontend from the backend and can handle large, spiky volumes of messages.

AWS Lambda is a good fit for processing messages from SQS because it scales automatically with queue depth and requires no server management.

Deploying the frontend in Docker on Amazon ECS with Fargate provides a fully managed, autoscaling container environment without managing EC2 instances.

Why others are not correct:

A and D: SNS is a pub/sub service, not a queue; it does not provide durable back-pressure buffering when the printers are slower than incoming requests. Lambda@Edge is for CloudFront edge processing, not back-end order processing.

B: Same SNS issue—no durable queue semantics and not ideal for workload that must be buffered and retried.

Question: 643

A company is designing a microservice-based architecture for a new application on AWS. Each microservice will run on its own set of Amazon EC2 instances. Each microservice will need to interact with multiple AWS services.

The company wants to manage permissions for each EC2 instance according to the principle of least privilege.

Which solution will meet this requirement with the LEAST administrative overhead?

- A. Assign an IAM user to each microservice. Use access keys that are stored within the application code to authenticate AWS

service requests.

- B. Create a single IAM role that has permission to access all AWS services. Add the IAM role to an instance profile that is associated with the EC2 instances.
- C. Use AWS Organizations to create a separate account for each microservice. Manage permissions at the account level.
- D. Create individual IAM roles based on the specific needs of each microservice. Add each IAM role to an instance profile that is associated with the appropriate EC2 instance.

Answer: D

Explanation:

AWS best practice is to use IAM roles with instance profiles for EC2 instances so that applications obtain temporary credentials automatically and do not need to store access keys.

To honor the principle of least privilege, each microservice should have an IAM role that grants only the specific permissions it needs.

Therefore, creating individual IAM roles per microservice and attaching them via instance profiles (Option D) both minimizes long-term credential management and applies least privilege cleanly.

Why others are not correct:

- A: Using IAM users with access keys in code is insecure and high-overhead (key rotation, secret management).
- B: A single broad role violates least privilege because every microservice gets more permissions than it needs.
- C: Separate accounts per microservice is extreme over-segmentation and significantly increases operational complexity.

Question: 644

A company uses an Amazon Aurora PostgreSQL DB cluster to store its critical data in the us-east-1 Region. The company wants to develop a disaster recovery plan to recover the database in the us-west-1 Region. The company has a recovery time objective (RTO) of 5 minutes and has a recovery point objective (RPO) of 1 minute.

What should a solutions architect do to meet these requirements?

- A. Create a read replica in us-west-1. Set the DB cluster to automatically fail over to the read replica if the primary instance is not responding.
- B. Create an Aurora global database. Set us-west-1 as the secondary Region. Update connections to use the writer and reader endpoints as appropriate.

- C. Set up a second Aurora DB cluster in us-west-1. Use logical replication to keep the databases synchronized. Create an Amazon EventBridge rule to change the database endpoint if the primary DB cluster does not respond.
- D. Use Aurora automated snapshots to store data in an Amazon S3 bucket. Enable S3 Versioning. Configure S3 Cross-Region Replication to us-west-1. Create a second Aurora DB cluster in us-west-1. Create an Amazon EventBridge rule to restore the snapshot if the primary DB cluster does not respond.

Answer: B

Explanation:

Aurora global database is specifically designed for cross-Region disaster recovery and low-latency global reads.

It provides:

Physical storage-level replication between Regions, typically with lag of under 1 second, easily satisfying RPO 1 minute.

Fast cross-Region failover capabilities, often within minutes or less, which meets the RTO 5 minutes requirement.

You configure the primary Region as the writer cluster and additional Regions (such as us-west-1) as secondary clusters with reader endpoints. In a disaster, you promote the secondary to be the new writer.

Why others are not correct:

A: Standard cross-Region read replicas for Aurora PostgreSQL do not provide automatic multi-Region failover with the same guarantees as Aurora global database and may not meet strict RTO/RPO requirements.

C: Logical replication and custom EventBridge failover is complex, error-prone, and higher overhead compared to the managed global database feature.

D: Restoring from snapshots plus cross-Region replication is too slow to meet RTO 5 minutes and certainly cannot maintain RPO 1 minute.

Question: 645

A company wants to re-architect a large-scale web application to a serverless microservices architecture. The application uses Amazon EC2 instances and is written in Python.

The company selected one component of the web application to test as a microservice. The component supports hundreds of requests per second. The company wants to create and test the microservice on an AWS solution that supports Python. The solution must also scale automatically and require minimal infrastructure and minimal operational support.

Which solution will meet these requirements?

- A. Use a Spot Fleet with Auto Scaling of EC2 instances that run the most recent Amazon Linux operating system.

- B. Use an AWS Elastic Beanstalk web server environment that has high availability configured.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS). Launch Auto Scaling groups of self-managed EC2 instances.
- D. Use an AWS Lambda function that runs custom-developed code.

Answer: D

Explanation:

D is the best match because AWS Lambda is a serverless compute service that runs code without provisioning or managing servers, which directly aligns with the requirement for minimal infrastructure and minimal operational support. Lambda supports Python runtimes, so the team can deploy the selected component as a microservice using the same language. For workloads with hundreds of requests per second, Lambda is designed to handle request-driven execution and automatically scales by running more concurrent instances of the function in response to incoming traffic, without the customer needing to pre-size or manage capacity.

The other options introduce significantly more operational responsibility. Option A still requires managing EC2 instances (AMI lifecycle, patching, capacity planning, scaling policies) and Spot interruptions can add complexity for a web microservice test. Option B (Elastic Beanstalk) reduces some management overhead but still relies on underlying instances and environment management (platform updates, scaling configuration, health monitoring), and it is not a serverless model. Option C (EKS with self-managed nodes) has the highest operational burden here: cluster and node management, scaling node groups, patching, and Kubernetes operations—this conflicts with “minimal operational support.”

Lambda fits the serverless microservices goal: deploy code as small independent services, integrate with other managed services (commonly API-driven invocation patterns), and rely on AWS-managed scaling and availability mechanisms rather than administering fleets or clusters.

Question: 646

A solutions architect needs to save a particular automated database snapshot from an Amazon RDS for Microsoft SQL Server DB instance for longer than the maximum number of days. Which solution will meet these requirements in the MOST operationally efficient way?

- A. Create a manual copy of the snapshot.
- B. Export the contents of the snapshot to an Amazon S3 bucket.
- C. Change the retention period of the snapshot to 45 days.
- D. Create a native SQL Server backup. Save the backup to an Amazon S3 bucket.

Answer: A

Explanation:

Creating a manual copy of the automated snapshot is the most operationally efficient option because it directly meets the requirement—retain a specific snapshot beyond the automated retention window—with the least added process and tooling. In Amazon RDS, automated backups and their snapshots are retained only for the configured backup retention period (up to the service maximum). When that retention period is exceeded, older automated snapshots are removed automatically. However, manual snapshots are retained until you explicitly delete them, so converting (copying) an automated snapshot to a manual snapshot is the standard operational approach to keep a point-in-time backup for long-term retention.

Option C is incorrect because you cannot extend automated snapshot retention beyond the maximum supported retention; also, setting “45 days” may still be within or beyond the service limits depending on the engine and configuration, and it doesn't guarantee indefinite retention. Option B (export to S3) is not the most operationally efficient for the stated goal: exporting is a different workflow (often used for analytics, archiving in open formats, or cross-tool usage) and introduces extra steps, format considerations, and ongoing management in S3. Option D is also heavier operationally: native SQL Server backups require managing backup jobs, storage layout, restores, and permissions, and it shifts responsibility to the customer for operational correctness.

Therefore, A is the simplest and most AWS-native way to preserve an RDS snapshot long-term with minimal operational overhead.

Question: 647

A company runs business applications on AWS. The company uses 50 AWS accounts, thousands of VPCs, and three AWS Regions across the United States and Europe. The company has an existing AWS Direct Connect connection that connects an on-premises data center to a single Region.

A solutions architect needs to establish network connectivity between the on-premises data center and the remaining two Regions. The solutions architect must also establish connectivity between the VPCs. On-premises users and applications must be able to connect to applications that run in the VPCs. The solutions architect creates a transit gateway in each Region and configures the transit gateways as inter-Region peers.

What should the solutions architect do next to meet these requirements?

- A. Create a private virtual interface (VIF) with a gateway type of virtual private gateway. Configure the private VIF to use a virtual private gateway that is associated with one of the VPCs.
- B. Create a private virtual interface (VIF) to a new Direct Connect gateway. Associate the new Direct Connect gateway with a virtual private gateway in each VPC.
- C. Create a transit virtual interface (VIF) with a gateway association to a new Direct Connect gateway. Associate each transit gateway with the new Direct Connect gateway.
- D. Create an AWS Site-to-Site VPN connection that uses a public virtual interface (VIF) for the Direct Connect connection. Attach the Site-to-Site VPN connection to the transit gateways.

Answer: C

Explanation:

The design already uses one transit gateway (TGW) per Region and inter-Region TGW peering, which addresses the multi-Region AWS-side routing. The remaining requirement is to extend on-premises connectivity over Direct Connect so that on-premises networks can reach VPCs attached to the TGWs across all three Regions. The most operationally efficient and scalable approach is to use AWS Direct Connect Gateway (DXGW) with a transit virtual interface (transit VIF) and then associate the Regional transit gateways to that DXGW.

Option C is purpose-built for this: a transit VIF is specifically used to connect a Direct Connect connection to a Direct Connect gateway, and a DXGW can then be associated to multiple transit gateways (and can be used across Regions), enabling centralized connectivity from on-premises to TGW-connected VPCs. With TGW attachments and TGW route tables, you can propagate and control routes across many VPCs and accounts, which fits the “50 accounts, thousands of VPCs” scale. InterRegion TGW peering then allows on-premises routes learned via the DXGW/TGW in one Region to reach workloads in the other Regions through the TGW peering relationships, subject to routing configuration.

Option A is too limited and not scalable because it ties Direct Connect to a virtual private gateway (VGW) associated with a single VPC, which does not meet the multi-VPC, multi-account hub-and-spoke requirement. Option B incorrectly suggests associating a DXGW with a VGW “in each VPC” (VGW is per VPC and would not scale well here, and it doesn't integrate with the TGW hub design you've already built).

Option D is not the intended pattern: Site-to-Site VPN and public VIF do not replace the DXGW + transit VIF architecture for large-scale TGW-based private routing.

Question: 648

A company hosts a web application on Amazon EC2 instances that are part of an Auto Scaling group behind an Application Load Balancer (ALB). The application experiences spikes in requests that come through the ALB throughout each day. The traffic spikes last between 15 and 20 minutes.

The company needs a solution that uses a standard or custom metric to scale the EC2 instances based on the number of requests that come from the ALB.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Amazon CloudWatch alarm to monitor the ALB RequestCount metric. Configure a simple scaling policy to scale the EC2 instances in response to the metric.
- B. Configure a predictive scaling policy based on the ALB RequestCount metric to scale the EC2 instances.
- C. Configure an Amazon CloudWatch alarm to monitor the ALB UnhealthyHostCount metric. Configure a target tracking policy to scale the EC2 instances in response to the metric.
- D. Create an Amazon CloudWatch alarm to monitor a user-defined metric for GET requests. Configure a target tracking policy threshold to scale the EC2 instances.

Answer: D

Explanation:

For short, frequent daily spikes (15-20 minutes), the most cost-effective scaling behavior typically comes from target tracking scaling driven by a request-based metric, because it continuously adjusts capacity to maintain a chosen target value rather than adding/removing capacity in coarse steps. Option D uses a custom CloudWatch metric representing the number of GET requests (sourced from ALB access logs, application instrumentation, or request counting logic) and applies target tracking so the Auto Scaling group scales out and in proportionally to demand. This helps avoid both underprovisioning (bad performance during spikes) and over-provisioning (wasted cost between spikes), which is exactly what cost-optimized elasticity aims for.

Option A (simple scaling) is usually less cost-efficient for bursty traffic because it relies on fixed adjustments and alarm thresholds, often with cooldowns. With 15-20 minute spikes, simple scaling can react too slowly, overshoot, or oscillate, leaving extra instances running after the spike or failing to ramp fast enough at the start. Option B (predictive scaling) is intended for regular, forecastable demand patterns, but it's not always the most cost-effective choice for repeated short spikes because it can pre-scale conservatively and may require more historical data and tuning; it is often paired with dynamic scaling anyway. Option C uses UnhealthyHostCount, which is a health signal—not a demand signal—and scaling on unhealthy hosts can increase cost without addressing request volume; it may also mask underlying application issues rather than scaling to meet load.

Because the requirement explicitly allows standard or custom metrics and focuses on scaling based on request volume, D best matches a cost-efficient, responsive approach that directly tracks incoming workload.

Question: 649

A company is developing a new application that uses Amazon EC2, Amazon S3, and AWS Lambda resources. The company wants to allow employees to access the AWS Management Console by using existing credentials that the company stores and manages in an on-premises Microsoft Active Directory. Each employee must have a specific level of access to the AWS resources that is based on the employee's role.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure AWS Directory Service to create an Active Directory in AWS Managed Microsoft AD. Establish a trust relationship with the on-premises Active Directory. Configure IAM roles and trust policies to give the employees access to the AWS resources.
- B. Use LDAP to directly integrate the on-premises Active Directory with IAM. Map Active Directory groups to IAM roles to control access to AWS resources.
- C. Implement a custom identity broker to authenticate users into the on-premises Active Directory. Configure the identity broker to use AWS STS to grant authorized users IAM role-based access to the AWS resources.
- D. Configure Amazon Cognito to federate users into the on-premises Active Directory. Use Cognito user pools to manage user identities and to manage user access to the AWS resources.

Answer: A

Explanation:

The least operational overhead approach is to use AWS Directory Service with AWS Managed Microsoft AD and establish a trust relationship with the existing on-premises Microsoft Active Directory. This pattern lets employees continue using their existing corporate

credentials while AWS operates the directory infrastructure (patching, monitoring, availability, and domain controller management), reducing the burden compared to building and maintaining custom federation components. With a trust in place, identities and group memberships from the on-premises directory can be used for authentication and authorization workflows in AWS.

To provide role-based access to AWS resources and the AWS Management Console, users should assume IAM roles that grant permissions based on job function. IAM roles with appropriate trust policies enable temporary credentials and least-privilege access without creating long-lived IAM users for every employee. This aligns with AWS best practices of centralized identity management and using roles for access control.

Option B is not a standard direct integration pattern: IAM does not “LDAP bind” directly to onpremises AD for console access. Option C can work (custom identity broker + STS) but increases operational overhead because the company must build, secure, operate, and maintain the broker, including availability, updates, and incident response. Option D is not the typical solution for employee console federation from AD; Amazon Cognito is primarily aimed at customer identity and application sign-in, not as the primary mechanism for workforce console access to AWS accounts.

Therefore, A provides the required federation with the smallest operational footprint while supporting role-based access control.

Question: 650

As part of budget planning, management wants a report of AWS billed items listed by user. The data will be used to create department budgets. A solutions architect needs to determine the most efficient way to obtain this report information.

Which solution meets these requirements?

- A. Run a query with Amazon Athena to generate the report.
- B. Create a report in Cost Explorer and download the report.
- C. Access the bill details from the billing dashboard and download the bill.
- D. Modify a cost budget in AWS Budgets to alert with Amazon Simple Email Service (Amazon SES).

Answer: B

Explanation:

The most efficient way for management-style reporting is AWS Cost Explorer, because it is designed for interactive cost analysis, filtering, grouping, and exporting reports without building a data pipeline. If the organization has activated cost allocation capabilities (for example, by tagging resources or using other allocation methods), Cost Explorer can present costs in a way that supports chargeback/showback and budgeting. From an operational standpoint, generating and downloading a report from Cost Explorer is quick, requires minimal setup, and is repeatable for periodic budget planning.

Option A (Athena) is powerful but typically requires setting up the Cost and Usage Report (CUR) delivery to Amazon S3, defining schemas, and writing/maintaining SQL queries. That is more operational overhead than needed when the ask is “most efficient way to obtain this

report information.” Option C (billing dashboard bill download) provides invoice-style line items, but it is not optimized for slicing/grouping “by user” and may not align with departmental budgeting workflows. Option D is unrelated: AWS Budgets alerts on thresholds and forecasts; it does not generate a billed-items-by-user report.

Important nuance: “by user” reporting in AWS is usually achieved through cost allocation tags, account structure, or other allocation dimensions rather than literal per-person IAM identity billing. In practice, departments/teams/users are mapped via tagging and chargeback structures, which Cost Explorer is intended to analyze and export. As a result, Cost Explorer is the most operationally efficient tool to produce a downloadable report for budget planning.

Therefore, B best meets the requirement with the least effort and fastest path to a usable report.

Question: 651

A company is moving its data management application to AWS. The company wants to transition to an event-driven architecture. The architecture needs to be more distributed and to use serverless concepts while performing the different aspects of the workflow. The company also wants to minimize operational overhead.

Which solution will meet these requirements?

- A. Build out the workflow in AWS Glue. Use AWS Glue to invoke AWS Lambda functions to process the workflow steps.
- B. Build out the workflow in AWS Step Functions. Deploy the application on Amazon EC2 instances. Use Step Functions to invoke the workflow steps on the EC2 instances.
- C. Build out the workflow in Amazon EventBridge. Use EventBridge to invoke AWS Lambda functions on a schedule to process the workflow steps.
- D. Build out the workflow in AWS Step Functions. Use Step Functions to create a state machine. Use the state machine to invoke AWS Lambda functions to process the workflow steps.

Answer: D

Explanation:

To build a distributed, serverless, event-driven workflow with minimal operational overhead, AWS Step Functions orchestrating AWS Lambda is the best fit. Step Functions provides a managed state machine service that coordinates multiple steps, handles retries and error handling patterns, and maintains execution state without requiring the company to run orchestration servers. Lambda provides serverless compute for each workflow task, scaling automatically with demand and eliminating server management.

Option D directly matches “performing different aspects of the workflow” because Step Functions is specifically designed to model multi-step business processes and coordinate activities across services. It also aligns with “minimize operational overhead” because both Step Functions and

Lambda are managed services: there are no instances to patch, no cluster management, and scaling is handled by the platform.

Option B contradicts the serverless goal by deploying the application on EC2 instances, increasing operational overhead (capacity management, patching, scaling, and availability concerns). Option C mixes eventing with scheduling: EventBridge can route events, but invoking Lambda “on a schedule” is not the same as orchestrating a multi-step workflow with branching, waiting, retries, compensation, and state tracking. EventBridge is excellent as an event bus, but Step Functions is the purpose-built workflow orchestrator. Option A is a mismatch because AWS Glue is primarily an ETL/data integration service; while it can run jobs and triggers, it is not the general workflow orchestrator for arbitrary application steps, and using it to invoke Lambda for a broad workflow is less direct and typically less operationally clean than a Step Functions state machine.

Therefore, D is the most aligned solution for serverless, distributed workflow orchestration with low operational burden.

Question: 652

A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.

Which migration solution will meet these requirements?

- A. Use native MySQL tools to migrate the database to Amazon RDS for MySQL. Configure elastic storage scaling.
- B. Migrate the database to Amazon Redshift by using the mysqldump utility. Turn on Auto Scaling for the Amazon Redshift cluster.
- C. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Aurora. Turn on Aurora Auto Scaling.
- D. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoDB. Configure an Auto Scaling policy.

Answer: C

Explanation:

The key requirements are (1) MySQL compatibility for existing applications and (2) the ability to scale automatically during demand spikes for a transactional workload. Amazon Aurora (MySQL-compatible) satisfies compatibility while providing managed database capabilities designed for high performance and scalability. Using AWS Database Migration Service (DMS) is operationally efficient for migration because it supports moving data from on-premises databases to AWS with options for continuous replication to minimize downtime, which is commonly required for transactional systems.

Turning on Aurora Auto Scaling (for Aurora Replicas) allows the database tier to automatically add or remove read replicas based on demand, which helps absorb increased read traffic without manual intervention. This meets the “scale automatically” requirement more directly than simply scaling storage. (Write scaling is addressed through Aurora's architecture and capacity planning; for many workloads, scaling reads is the dominant elasticity need.)

Option A only mentions configuring storage scaling, which does not address compute scaling during load increases; RDS storage autoscaling

helps prevent running out of space but does not automatically add compute capacity in response to demand. Option B is incorrect because Amazon Redshift is a data warehouse designed for analytics/OLAP, not a transactional MySQL OLTP replacement, and mysqldump to Redshift does not preserve application compatibility. Option D changes the data model by migrating to DynamoDB, which is not MySQL-compatible and would require application redesign—violating the compatibility requirement.

Therefore, C best satisfies compatibility and automatic scaling needs while using managed services that reduce operational overhead during and after migration.

Question: 653

A company runs an ecommerce application on premises on Microsoft SQL Server. The company is planning to migrate the application to the AWS Cloud. The application code contains complex T-SQL queries and stored procedures. The company wants to minimize database server maintenance and operating costs after the migration is completed. The company also wants to minimize the need to rewrite code as part of the migration effort.

Which solution will meet these requirements?

- A. Migrate the database to Amazon Aurora PostgreSQL. Turn on Babelfish.
- B. Migrate the database to Amazon S3. Use Amazon Redshift Spectrum for query processing.
- C. Migrate the database to Amazon RDS for SQL Server. Turn on Kerberos authentication.
- D. Migrate the database to an Amazon EMR cluster that includes multiple primary nodes.

Answer: A

Explanation:

The requirement combines two key goals: reduce post-migration database administration/maintenance cost and minimize application rewrites for an existing Microsoft SQL Server application that relies heavily on T-SQL and stored procedures. Amazon Aurora PostgreSQL-Compatible with Babelfish is designed for this exact migration pattern: it helps applications written for SQL Server to run with reduced code changes by providing compatibility for commonly used SQL Server semantics, including T-SQL constructs, procedural logic, and SQL Server-style connectivity patterns (depending on feature usage). Aurora itself is a managed database service that reduces operational overhead compared to self-managed databases by offloading tasks such as provisioning, patching (within service capabilities), backups, and high availability patterns.

Option C (RDS for SQL Server) would indeed minimize rewrites because it keeps the engine the same, but it typically does not optimize operating costs as effectively as migrating off commercial SQL Server licensing/edition costs; it also keeps you on the same engine family, which often carries higher license implications and may not meet the “minimize operating costs” intent as strongly as moving to Aurora PostgreSQL with Babelfish. Option B is not suitable because Redshift Spectrum is for analytics over data in S3, not for running an OLTP ecommerce database with stored procedures. Option D is also a mismatch: EMR is for big data processing frameworks, not a relational OLTP database replacement for SQL Server.

Therefore, A best balances lower ongoing operational cost with reduced refactoring effort by using Aurora's managed capabilities while leveraging Babelfish to ease SQL Server-to-PostgreSQL migration friction.

Question: 654

A company hosts multiple applications on AWS for different product lines. The applications use different compute resources, including Amazon EC2 instances and Application Load Balancers. The applications run in different AWS accounts under the same organization in AWS Organizations across multiple AWS Regions. Teams for each product line have tagged each compute resource in the individual accounts.

The company wants more details about the cost for each product line from the consolidated billing feature in Organizations.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Select a specific AWS generated tag in the AWS Billing console.
- B. Select a specific user-defined tag in the AWS Billing console.
- C. Select a specific user-defined tag in the AWS Resource Groups console.
- D. Activate the selected tag from each AWS account.
- E. Activate the selected tag from the Organizations management account.

Answer: B,E

Explanation:

To break down consolidated billing costs by “product line” using tags, the company must use cost allocation tags and ensure those tags are activated for billing at the payer/management level. Because the teams have already applied user-defined tags to the compute resources across accounts, the correct approach is to (1) use those user-defined tags in the billing tools and (2) activate them so they appear in cost and usage reporting for consolidated billing.

B is required because cost reporting by product line depends on selecting the user-defined cost allocation tag key (for example, ProductLine=...) in AWS billing and cost management tooling (such as Cost Explorer reports). AWS-generated tags are not what teams typically use for chargeback by business dimension, and the scenario explicitly says teams tagged resources (i.e., user-defined).

E is required because in an AWS Organizations consolidated billing setup, cost allocation tag activation is handled centrally from the Organizations management account so that the tags can be used consistently in consolidated cost views and reporting. Activating tags centrally ensures the billing system recognizes and includes those tag dimensions in cost allocation data.

Option C is not relevant to billing breakdown; Resource Groups is for organizing and operating on resources, not for consolidated billing allocation. Option D is unnecessary in this scenario's intended best practice because activation for consolidated billing reporting is done from the management account to standardize reporting. In short, you tag resources (already done), then activate the tag for billing and use it in billing reports to see cost by product line across accounts and Regions.

Therefore, B and E together meet the requirement with the correct billing-focused steps.

Question: 655

A company hosts an end-user application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company needs to configure end-to-end encryption between the ALB and the EC2 instances.

Which solution will meet this requirement with the LEAST operational effort?

A. Deploy AWS CloudHSM. Import a third-party certificate into CloudHSM. Configure the EC2 instances and the ALB to use the CloudHSM imported certificate.

B. Import a third-party certificate bundle into AWS Certificate Manager (ACM). Generate a selfsigned certificate on the EC2 instances. Associate the ACM imported third-party certificate with the ALB.

C. Import a third-party SSL certificate into AWS Certificate Manager (ACM). Install the third-party certificate on the EC2 instances. Associate the ACM imported third-party certificate with the ALB.

D. Use Amazon-issued AWS Certificate Manager (ACM) certificates on the EC2 instances and the ALB.

Answer: D

Explanation:

End-to-end encryption means TLS is used not only from the client to the ALB, but also from the ALB to the EC2 targets. The least operational effort is achieved by using AWS Certificate Manager (ACM) Amazon-issued certificates where possible, because ACM automates key management, certificate provisioning, and renewal for supported use cases. Associating an ACM certificate with the ALB is a standard managed approach for TLS termination at the load balancer, and using managed certificates reduces the overhead of tracking expiration, renewing, and distributing certificates.

For encryption on the backend connection (ALB to instances), the instances must present a certificate and complete TLS. Using Amazon-issued ACM certificates (via supported mechanisms) reduces manual certificate lifecycle work compared with importing and managing third-party certificates. By avoiding third-party cert procurement and manual renewals, the solution minimizes ongoing operations and reduces the risk of outages due to expired certificates. This aligns with the requirement for the least operational effort while meeting the security requirement.

Option A is far more operationally heavy: CloudHSM is intended for scenarios requiring customermanaged HSMs and adds complexity in deployment, scaling, integration, and operations. Option B introduces self-signed certificates on instances, which increases operational friction (distribution, trust, rotation) and is not as clean for managed operations. Option C works but requires managing third-party certificate lifecycle (renewals, re-import, redeploy to instances), which is more overhead than Amazon-issued managed certificates.

Therefore, D best meets end-to-end encryption needs with the lowest ongoing operational burden by using AWS-managed certificate issuance and lifecycle management capabilities.

Question: 656

A company is building new learning management applications on AWS. The company is using Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 to host the applications. The company must ensure that container images are secure. Company administrators must receive notifications of any security vulnerabilities in the images.

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Modify the ECS cluster properties to use privileged mode. Enable host-based logging.
- B. Use the AWS Config conformance pack for Amazon ECS. Use AWS Config to notify administrators if any security vulnerabilities are detected.
- C. Configure AWS WAF to invoke an Amazon CloudWatch alarm when a new security vulnerability is detected.
- D. Use Amazon Inspector to scan container images in Amazon Elastic Container Registry (Amazon ECR).
- E. Use AWS Systems Manager Parameter Store to encrypt container images.

Answer: B,D

Explanation:

To ensure container images are secure and to notify administrators about vulnerabilities, the solution needs (1) a vulnerability scanning capability for container images and (2) a notification mechanism that alerts when findings occur. Amazon Inspector provides automated security assessments and can scan container images stored in Amazon ECR to identify software vulnerabilities and unintended network exposure patterns, producing findings that can be acted upon. Therefore, D addresses the core requirement of detecting vulnerabilities in container images.

To notify administrators with minimal custom work, AWS Config can help by evaluating resources against desired configurations and integrating with notifications through AWS services (for example, via Amazon SNS using Config rules/conformance packs). Using the AWS Config conformance pack for Amazon ECS helps establish a managed set of compliance checks aligned to ECS-related best practices. While Inspector is the system that detects vulnerabilities, Config can be used to enforce and monitor governance controls around the container environment and can trigger notifications when noncompliance is detected. In exam patterns, pairing an Inspector detection capability with a managed governance/notification framework like Config is a common “two-part” answer.

The other options do not meet the requirement: A (privileged mode) can increase risk rather than improve image security; logging does not equal vulnerability detection. C is unrelated because AWS WAF protects web applications at the edge and does not scan container images for CVEs. E is incorrect because Parameter Store stores configuration data and secrets; it does not encrypt container images (ECR encryption at rest is handled by AWS-managed mechanisms and KMS integration, not Parameter Store).

So D provides scanning and B supports managed compliance/notification controls with low operational overhead.

Question: 657

A company is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store.

Which solution will meet these requirements?

- A. Create a new AWS Key Management Service (AWS KMS) key. Use AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS.
- B. Create a new AWS Key Management Service (AWS KMS) key. Enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster.
- C. Create the Amazon EKS cluster with default options. Use the Amazon Elastic Block Store (Amazon EBS) Container Storage Interface (CSI) driver as an add-on.
- D. Create a new AWS Key Management Service (AWS KMS) key with the alias/aws/ebs alias. Enable default Amazon Elastic Block Store (Amazon EBS) volume encryption for the account.

Answer: B

Explanation:

In Amazon EKS, Kubernetes stores objects such as Secrets in the cluster's etcd key-value store. By default, Kubernetes Secrets are base64-encoded and are not automatically encrypted at the application object level unless encryption is configured. Amazon EKS provides a managed capability to encrypt Kubernetes Secrets at rest in etcd using AWS Key Management Service (KMS). The requirement explicitly states that “all secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store,” which maps directly to enabling EKS secrets encryption with a customer-managed KMS key.

Option B is exactly this: create a KMS key and enable EKS KMS secrets encryption on the cluster. With this enabled, EKS uses envelope encryption so that Secrets are encrypted when stored in etcd, and decrypt operations are controlled through KMS permissions. This is the standard, AWS-native method that fulfills the requirement without requiring application changes or external secret stores for the encryption-at-rest requirement in etcd.

Option A (Secrets Manager) is a strong service for secret lifecycle management and rotation, but it does not by itself guarantee that Kubernetes Secrets stored in etcd are encrypted unless EKS secrets encryption is also enabled or the cluster avoids storing secrets in etcd entirely. The question specifically targets etcd encryption. Option C is unrelated: the EBS CSI driver concerns persistent volumes, not etcd secret encryption. Option D concerns EBS volume encryption and a specific AWS- managed key alias used for EBS; it also does not address etcd encryption for Kubernetes Secrets.

Therefore, B is the correct solution because it directly enables encryption of Kubernetes Secrets within etcd using KMS.

Question: 658

A company is developing an ecommerce application that will consist of a load-balanced front end, a container-based application, and a relational database. A solutions architect needs to create a highly available solution that operates with as little manual intervention as possible.

Which solutions meet these requirements? (Select TWO.)

- A. Create an Amazon RDS DB instance in Multi-AZ mode.
- B. Create an Amazon RDS DB instance and one or more replicas in another Availability Zone.
- C. Create an Amazon EC2 instance-based Docker cluster to handle the dynamic application load.
- D. Create an Amazon Elastic Container Service (Amazon ECS) cluster with a Fargate launch type to handle the dynamic application load.
- E. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an Amazon EC2 launch type to handle the dynamic application load.

Answer: A,D

Explanation:

For the relational database tier, Amazon RDS Multi-AZ provides high availability with minimal manual intervention. In Multi-AZ mode, RDS maintains a synchronous standby in a different Availability Zone and provides automatic failover during certain failure scenarios. This reduces operational burden because the service handles replication, health monitoring, and failover orchestration.

For the container-based application tier, Amazon ECS with the Fargate launch type minimizes operational overhead because it removes the need to provision, patch, and scale the underlying EC2 instances that host containers. With Fargate, AWS manages the compute infrastructure, and the team focuses on task definitions, scaling policies, and application configuration. This supports a highly available, load-balanced architecture because ECS services can run tasks across multiple Availability Zones behind an Application Load Balancer, and scaling is managed via ECS Service Auto Scaling.

Option B describes read replicas, which are primarily used for scaling reads and, depending on configuration, may not provide the same automated failover characteristics as Multi-AZ for high availability. Read replicas are not a direct substitute for Multi-AZ HA. Option C (self-managed Docker cluster on EC2) is operationally heavy: you must manage node capacity, patching, cluster lifecycle, and scheduling. Option E (ECS on EC2) is a valid container platform but still requires managing EC2 instances, AMIs, and scaling/patching, which increases manual intervention compared to Fargate.

Therefore, combining RDS Multi-AZ (A) for database resilience and ECS Fargate (D) for serverless container operations meets the HA requirement with the least manual effort.

Question: 659

A company's ecommerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections.

What should a solutions architect do to meet these requirements?

- A. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions inside a VPC.

- B. Point the client driver at an RDS Proxy endpoint. Deploy the Lambda functions inside a VPC.
- C. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions outside a VPC.
- D. Point the client driver at an RDS Proxy endpoint. Deploy the Lambda functions outside a VPC.

Answer: B

Explanation:

Unpredictable Lambda traffic can create a “connection storm” against a relational database because each concurrent Lambda invocation may open its own database connection. Relational engines like PostgreSQL have finite connection capacity, and excessive connections can degrade performance, increase latency, and exhaust database resources. Amazon RDS Proxy is designed to address this by pooling and reusing database connections, multiplexing many application requests over a smaller number of established database connections. This helps keep database performance more predictable under spiky, highly concurrent workloads such as Lambda.

Because the database is a private RDS for PostgreSQL instance, the Lambda functions must run with network access to the VPC subnets and security groups that can reach the database and proxy. Therefore, deploying the Lambda functions inside a VPC and pointing the database client to the RDS Proxy endpoint is the correct operational pattern. RDS Proxy also provides benefits like improved failover behavior handling and centralizing credential management integration patterns (commonly with IAM authentication and/or Secrets Manager), further reducing operational risk.

Option A uses a custom endpoint, which does not solve the fundamental connection scaling problem; it's not a connection pooler. Option C and D place Lambda outside the VPC, which will not work for a private database that has no public connectivity path. Even if connectivity were possible, D would still be incomplete because you'd need private networking to reach the proxy.

Therefore, B best meets the requirement by introducing managed connection pooling via RDS Proxy and ensuring private network reachability by running Lambda in the VPC.

Question: 660

A company runs an on-premises managed file transfer solution to collect images from its clients. The company uses an open source transfer tool to transfer and integrate the images into the company's workflow. The company then runs a custom application to add watermarks to the images.

The company needs to migrate this workload to AWS and wants to use AWS managed services where possible. Uploaded images must be stored as objects. The company wants to automate the watermark addition.

Which solution will meet these requirements?

- A. Use AWS DataSync to automate file transfers. Store the images in an Amazon S3 bucket. Use an application that runs on Amazon EC2 instances to add watermarks.
- B. Use REST APIs to transfer files. Store the images in an Amazon S3 bucket. Use AWS Batch jobs to add watermarks.

- C. Use SFTP with AWS Transfer Family to automate file transfers into Amazon S3 buckets. Configure the Transfer Family workflow to invoke an AWS Lambda function to add watermarks.
- D. Use AWS Transfer Family to transfer images. Store the images in Amazon S3 Glacier Deep Archive. Run an AWS Step Functions state machine to add watermarks.

Answer: C

Explanation:

This workload has two distinct needs: (1) a managed way for external clients to upload files using a familiar managed file transfer protocol, and (2) an automated, low-ops method to run a watermarking step and store the resulting images as objects. AWS Transfer Family provides a fully managed capability to receive files over protocols such as SFTP, while landing those files directly into Amazon S3 as objects. That directly satisfies the requirement that “uploaded images must be stored as objects” with minimal infrastructure management.

To automate watermarking with minimal operational overhead, invoking AWS Lambda is a strong fit. Lambda is serverless and scales automatically with incoming uploads, and it can run the watermarking logic as code without managing servers. Transfer Family supports managed workflows that can orchestrate post-upload actions; using a Transfer Family workflow to invoke a Lambda function provides a clean, managed, event-driven pipeline: clients upload via SFTP objects land in S3 workflow triggers watermark processing output is written back to S3.

Option A adds operational burden by requiring an EC2-based application tier for watermarking, which increases patching and scaling responsibilities. Option B can work (S3 + Batch), but it requires building and operating the upload mechanism (REST APIs) and typically more orchestration than necessary for simple “upload then process” flows; Batch is better when you need large-scale, long-running batch compute rather than lightweight per-object processing. Option D is unsuitable because S3 Glacier Deep Archive is intended for long-term archival with slow retrieval, which conflicts with active workflow processing and watermark automation.

Therefore, C best meets the requirements using managed services end-to-end: Transfer Family for ingestion, S3 for object storage, and Lambda for automated watermarking.

Question: 661

A company uses AWS WAF to protect its web applications. A solutions architect configures a web ACL that uses several rules, including a rule that inspects the HTTP request body for malicious content.

The solutions architect notices that the web ACL is not inspecting large HTTP POST requests properly. As a result, suspicious activities are not being detected. Some large HTTP POST requests are more than 8 MB in size.

The solutions architect must ensure that the web ACL inspects the large HTTP POST requests properly.

Which solution will meet this requirement?

- A. Create two custom AWS WAF rules. Configure one rule to block all oversized requests. Configure the second rule with a higher

priority to allow large requests from legitimate hosts.

B. Enable AWS Shield Advanced. Reconfigure the web ACL to block oversized requests by using Shield Advanced.

C. Verify that the Content-Type header is correctly set in the HTTP requests that AWS WAF rules inspect.

D. Create an AWS Lambda function to preprocess the large requests before AWS rules inspect the requests.

Answer: A

Explanation:

AWS WAF has limits on how much of an HTTP request body it can inspect. When requests exceed the inspectable size, AWS WAF treats the body as oversized relative to the configured inspection limits, which can lead to rules not evaluating the entire body content. If suspicious payloads are embedded beyond the inspected portion of a large POST request (for example, > 8 MB), WAF cannot reliably detect them purely through body inspection rules.

Given this constraint, the most effective way to “ensure” proper protection is to implement an oversize handling strategy using AWS WAF rule logic: block oversized requests by default and then explicitly allow oversized requests only from known legitimate sources. Option A accomplishes this by adding a rule that blocks oversize requests (so attackers cannot bypass inspection by sending very large bodies) and a higher-priority allow rule to permit large requests from trusted hosts (for example, specific known partners, internal CIDRs, or authenticated upstream systems). This design reduces the attack surface and provides deterministic behavior for requests that cannot be fully inspected.

Option B is incorrect because Shield Advanced is for DDoS protection and does not extend WAF's request-body inspection size. Option C is unrelated: Content-Type can influence application parsing, but it will not overcome WAF body-size inspection limitations. Option D is not a practical fit for AWS WAF inspection because WAF evaluates requests at the edge/service layer; it does not natively “call Lambda to rewrite the request body” before WAF evaluates it. Any preprocessing would require a different architectural pattern (such as handling uploads out-of-band), which is beyond the scope and would add operational complexity.

Therefore, A is the correct approach: implement explicit oversized request handling by blocking by default and allowing only vetted large requests.

Question: 662

A company decides to use AWS Key Management Service (AWS KMS) for data encryption operations. The company must create a KMS key and automate the rotation of the key. The company also needs the ability to deactivate the key and schedule the key for deletion.

Which solution will meet these requirements?

A. Create an asymmetric customer managed KMS key. Enable automatic key rotation.

B. Create a symmetric customer managed KMS key. Disable the envelope encryption option.

C. Create a symmetric customer managed KMS key. Enable automatic key rotation.

D. Create an asymmetric customer managed KMS key. Disable the envelope encryption option.

Answer: C

Explanation:

To meet the requirements, the key must support automatic rotation, and the company must be able to disable (deactivate) the key and schedule deletion. In AWS KMS, these lifecycle controls are available for customer managed keys. Automatic key rotation is supported for symmetric customer managed KMS keys used for encryption and decryption operations. When automatic rotation is enabled, AWS KMS generates new cryptographic material for the key on a regular schedule while the key ID remains the same, helping organizations meet compliance and security best practices without manual operational work.

Option C is the only choice that explicitly combines a symmetric customer managed key with automatic rotation enabled, which directly satisfies the rotation requirement. As a customer managed key, it can also be disabled (to prevent use) and scheduled for deletion (with a waiting period), meeting the rest of the lifecycle needs.

Option A is incorrect because automatic rotation is not generally available for asymmetric KMS keys in the same way; asymmetric keys are used for specialized use cases such as signing or asymmetric encryption, and rotation behavior differs. Options B and D mention “disabling envelope encryption,” which is not a configurable “option” you turn off in KMS; envelope encryption is a recommended pattern where KMS protects data keys, and services commonly use it under the hood. Those options therefore do not describe a valid configuration that meets the requirement.

Therefore, the correct solution is to create a symmetric customer managed KMS key and enable automatic rotation, while using standard KMS controls to disable and schedule deletion when required.

Question: 663

A company's HTTP application is behind a Network Load Balancer (NLB). The NLB's target group is configured to use an Amazon EC2 Auto Scaling group with multiple EC2 instances that run the web service.

The company notices that the NLB is not detecting HTTP errors for the application. These errors require a manual restart of the EC2 instances that run the web service. The company needs to improve the application's availability without writing custom scripts or code.

What should a solutions architect do to meet these requirements?

A. Enable HTTP health checks on the NLB, supplying the URL of the company's application.

B. Add a cron job to the EC2 instances to check the local application's logs once each minute. If HTTP errors are detected, the application will restart.

C. Replace the NLB with an Application Load Balancer. Enable HTTP health checks by supplying the URL of the company's application. Configure an Auto Scaling action to replace unhealthy instances.

D. Create an Amazon CloudWatch alarm that monitors the UnhealthyHostCount metric for the NLB. Configure an Auto Scaling action to

replace unhealthy instances when the alarm is in the ALARM state.

Answer: C

Explanation:

A Network Load Balancer operates at Layer 4 (TCP/UDP/TLS) and is optimized for high performance and static IP use cases. While NLB target groups can perform health checks, they are typically oriented around basic reachability and do not provide the same application-layer (Layer 7) visibility as an Application Load Balancer (ALB). The problem statement says the NLB is “not detecting HTTP errors,” which indicates the health signal needs to be based on an HTTP endpoint that can reflect application correctness (for example, returning specific HTTP status codes).

Replacing the NLB with an ALB enables true HTTP/HTTPS health checks against a URL path, including interpretation of HTTP response codes. This is the cleanest managed approach to detect application-layer failure modes that still allow TCP connections but produce bad HTTP responses. Once the ALB detects targets as unhealthy, the target group health status can be used by an Auto Scaling group to take action. With appropriate health check configuration (and, commonly, using ELB health checks as a signal), Auto Scaling can replace unhealthy instances automatically, improving availability without custom scripts.

Option A is misleading: NLB does not provide the same HTTP-aware request routing and rich L7 features; even if an NLB health check is configured, it does not address the broader need for application-layer detection and remediation as directly as ALB. Option B violates the “no custom scripts” requirement. Option D reacts to UnhealthyHostCount, but if the NLB isn't marking hosts unhealthy for HTTP error cases, the metric won't reliably trigger replacement; it also still depends on the NLB's limited visibility into HTTP failures.

Therefore, C best meets the requirement by shifting to ALB for application-layer health checks and using Auto Scaling to replace unhealthy instances automatically.

Question: 664

A security audit reveals that Amazon EC2 instances are not being patched regularly. A solutions architect needs to provide a solution that will run regular security scans across a large fleet of EC2 instances. The solution should also patch the EC2 instances on a regular schedule and provide a report of each instance's patch status.

Which solution will meet these requirements?

- A. Set up Amazon Macie to scan the EC2 instances for software vulnerabilities. Set up a cron job on each EC2 instance to patch the instance on a regular schedule.
- B. Turn on Amazon GuardDuty in the account. Configure GuardDuty to scan the EC2 instances for software vulnerabilities. Set up AWS Systems Manager Session Manager to patch the EC2 instances on a regular schedule.
- C. Set up Amazon Detective to scan the EC2 instances for software vulnerabilities. Set up an Amazon EventBridge scheduled rule to patch the EC2 instances on a regular schedule.
- D. Turn on Amazon Inspector in the account. Configure Amazon Inspector to scan the EC2 instances for software vulnerabilities. Set up AWS Systems Manager Patch Manager to patch the EC2 instances on a regular schedule.

Answer: D

Explanation:

The requirement calls for three capabilities at fleet scale: (1) regular security/vulnerability scanning of EC2 instances, (2) scheduled patching, and (3) reporting on patch compliance/status. The AWS- managed services designed for this are Amazon Inspector for vulnerability scanning and AWS Systems Manager Patch Manager for patch orchestration and compliance reporting.

Amazon Inspector provides automated vulnerability management that can assess EC2 instances for software vulnerabilities and exposure, producing findings that identify missing patches and vulnerable packages. This addresses the security scanning portion without requiring custom tooling on each instance beyond standard agent/configuration requirements.

Systems Manager Patch Manager allows you to define patch baselines, schedule patching operations via maintenance windows, and apply patches across large fleets in a controlled manner. Patch Manager also provides compliance views and reporting so you can see which instances are compliant, which are missing patches, and the results of patch operations. This directly meets the need to patch “on a regular schedule” and to “provide a report of each instance’s patch status.”

The other options are mismatched services: Macie focuses on discovering and protecting sensitive data (especially in S3), not scanning EC2 for software vulnerabilities. GuardDuty is for threat detection based on logs and events; it is not an EC2 vulnerability scanner and does not function as a patching orchestrator. Detective helps investigate security events and relationships; it is not a vulnerability scanning or patch deployment tool. Additionally, cron jobs on each instance create high operational overhead and inconsistent reporting—exactly what the company wants to avoid.

Therefore, D is the correct, operationally excellent approach: use Inspector for vulnerability scanning and Systems Manager Patch Manager for automated, scheduled patching with centralized compliance reporting.

Question: 665

A company runs a container application by using Amazon Elastic Kubernetes Service (Amazon EKS). The application includes microservices that manage customers and place orders. The company needs to route incoming requests to the appropriate microservices.

Which solution will meet this requirement MOST cost-effectively?

- A. Use the AWS Load Balancer Controller to provision a Network Load Balancer.
- B. Use the AWS Load Balancer Controller to provision an Application Load Balancer.
- C. Use an AWS Lambda function to connect the requests to Amazon EKS.
- D. Use Amazon API Gateway to connect the requests to Amazon EKS.

Answer: B

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

Routing requests to the “appropriate microservices” implies Layer 7 (HTTP/HTTPS) routing such as host-based or path-based rules (for example, /customers to one service and /orders to another). In EKS, the most direct and cost-efficient way to provide this style of routing is to use Kubernetes Ingress with the AWS Load Balancer Controller to provision an Application Load Balancer (ALB). An ALB natively supports HTTP/HTTPS request routing features (paths, hosts, headers), which map cleanly to microservices patterns and reduce the need for additional infrastructure components.

A Network Load Balancer (NLB) (Option A) is primarily Layer 4 and is best for TCP/UDP/TLS passthrough and ultra-high performance. It does not provide the same native HTTP path-based routing behavior required to steer requests between microservices based on URL structure, so it often forces you to add another routing layer (like an in-cluster ingress gateway), which can increase cost and operational complexity.

Option C is not an appropriate request routing mechanism; Lambda is compute, not an ingress/routing layer for EKS. Option D (API Gateway) can front services, but for simple microservice routing into EKS it typically adds per-request costs and additional configuration compared with using an ALB that is purpose-built for HTTP load balancing and routing. API Gateway is a great fit for API management features (throttling, usage plans, keys, transformations), but those are not requirements here.

Therefore, B is the most cost-effective and architecturally aligned solution for HTTP routing across EKS-hosted microservices using managed AWS-native ingress integration.

Question: 666

A solutions architect has created an AWS Lambda function that is written in Java. A company will use the Lambda function as a new microservice for its application. The company's customers must be able to call an HTTPS endpoint to reach the microservice. The microservice must use AWS Identity and Access Management (IAM) to authenticate calls.

Which solution will meet these requirements?

- A. Create an Amazon API Gateway REST API. Configure an API method to use the Lambda function. Create a second Lambda function that is configured as an authorizer.
- B. Create an AWS Lambda function URL for the Lambda function. Specify AWS_IAM as the authentication type.
- C. Create an Amazon CloudFront distribution. Deploy the Lambda function to Lambda@Edge. Integrate IAM authentication logic into the Lambda@Edge function.
- D. Create an Amazon CloudFront distribution. Deploy the Lambda function to CloudFront Functions. Specify AWS_IAM as the authentication type.

Answer: B

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

The requirements are straightforward: expose the Lambda microservice through an HTTPS endpoint and authenticate calls using IAM. Lambda function URLs are a built-in feature that provides a dedicated HTTPS endpoint for a Lambda function without requiring API Gateway, ALB, or CloudFront.

When configured with the authentication type `AWS_IAM`, the endpoint requires requests to be signed with AWS Signature Version 4 and authorized by IAM policies. This directly satisfies the “must use IAM to authenticate calls” requirement with the least architectural complexity.

Option A can also secure an endpoint with IAM, but it proposes using a Lambda authorizer, which is typically used for custom authorizers (JWT/OAuth/Cognito/external identity). For IAM authentication in API Gateway, you generally use IAM authorization on the method, not an authorizer function. Also, API Gateway REST APIs introduce additional service configuration and per-request costs when a simpler managed option exists that meets the requirements.

Options C and D are not appropriate. Lambda@Edge and CloudFront Functions run at CloudFront edge locations with different programming and deployment models; they are designed for CDN request/response manipulation, not as the primary mechanism to expose a regional Lambda microservice endpoint with IAM authentication. CloudFront Functions in particular is for lightweight JavaScript at the edge and does not provide a native “`AWS_IAM` authentication type” for invoking an origin Lambda as a microservice endpoint.

Therefore, B is the cleanest and most secure fit: a native HTTPS endpoint backed by Lambda, protected with IAM-based SigV4 authentication.

Question: 667

A company generates SSL certificates from a third-party provider. The company imports the certificates into AWS Certificate Manager (ACM) to use with public web applications.

A solutions architect must implement a solution to notify the company's security team 30 days before an imported certificate expires. The company already has an Amazon Simple Queue Service (Amazon SQS) queue. The company also has an Amazon Simple Notification Service (Amazon SNS) topic that has the security team's email address as a subscriber.

Which solution will provide the security team with the required notification about certificates?

A. Create an AWS Lambda function to scan for expiring certificates. Program the Lambda function to list the certificates in a JSON message and to deliver the message to the SQS queue.

B. Create an AWS Lambda function to scan for expiring certificates. Program the Lambda function to list the certificates in a JSON message and to deliver the message to the SNS topic.

C. Create an Amazon EventBridge rule that specifies the ACM Certificate Approaching Expiration event type. Set the SQS queue as the rule's target.

D. Create an Amazon EventBridge rule that specifies the ACM Certificate Approaching Expiration event type. Set the SNS topic as the rule's target.

Answer: D

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

The requirement is an automated notification 30 days before an imported ACM certificate expires, delivered to the security team via an existing SNS topic with email subscription. The most operationally efficient approach is to use Amazon EventBridge with the managed event type for ACM certificate expiration. ACM publishes an event when a certificate is approaching expiration, and EventBridge can match that event and route it directly to a target service without custom polling logic.

Option D uses an EventBridge rule for the ACM Certificate Approaching Expiration event and sets the SNS topic as the target. This directly delivers an alert to the existing notification channel (email via SNS) and requires minimal code and minimal ongoing maintenance. It also avoids building and scheduling a scanner that must enumerate certificates, calculate dates, handle pagination, and manage failures.

Option C sends the event to SQS. While SQS is useful for decoupling and buffering, the requirement is to notify the security team, and SNS is already configured for email delivery. Using SQS would add an extra consumer component to read from the queue and publish notifications, which is additional operational overhead.

Options A and B require a custom Lambda-based scanning solution. That introduces scheduling (for example, EventBridge schedule), logic to detect "30 days remaining," error handling, and ongoing maintenance. Since ACM already emits a purpose-built event for expiring certificates, polling is unnecessary and less efficient.

Therefore, D is the best solution: it uses a native event from ACM, routes it through EventBridge, and notifies the security team through the existing SNS topic with the least operational effort.

Question: 668

A company hosts a two-tier website that runs on Amazon EC2 instances. The website has a database that runs on Amazon RDS for MySQL. All users are required to log in to the website to see their own customized pages.

The website typically experiences low traffic. Occasionally, the website experiences sudden increases in traffic and becomes unresponsive. During these increases in traffic, the database experiences a heavy write load. A solutions architect must improve the website's availability without changing the application code.

What should the solutions architect do to meet these requirements?

- A. Create an Amazon ElastiCache (Redis OSS) cluster. Configure the application to cache common database queries in the ElastiCache cluster.
- B. Create an Auto Scaling group. Configure Amazon CloudWatch alarms to scale the number of EC2 instances based on the percentage of CPU in use during the traffic increases.
- C. Create an Amazon CloudFront distribution that points to the EC2 instances as the origin. Enable caching of dynamic content, and configure a write throttle from the EC2 instances to the RDS database.

D. Migrate the database to an Amazon Aurora Serverless cluster. Set the maximum Aurora capacity units (ACUs) to a value high enough to respond to the traffic increases. Configure the EC2 instances to connect to the Aurora database.

Answer: D

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

The key constraint is no application code changes, while the observed failure mode during spikes is heavy write load on the database that makes the site unresponsive. Scaling the web tier alone (Option B) may not help if the database is the bottleneck; in fact, adding more EC2 instances can increase concurrent write pressure and worsen the database overload. Caching (Option A) would require modifying the application to use Redis for query caching, which violates the no-code-change requirement. Option C is not a fit: CloudFront caching is effective for cacheable content, but authenticated, personalized pages are typically not cacheable without application changes, and “write throttle” is not a standard CloudFront feature to protect the database.

Migrating from RDS for MySQL to Amazon Aurora Serverless is a managed approach designed to automatically adjust database capacity to match workload demand, which is well suited to “low baseline traffic with occasional sudden spikes.” By setting an appropriate maximum capacity, Aurora Serverless can scale to handle bursts more gracefully than a fixed-size RDS instance, improving availability during unpredictable demand. This reduces operational overhead because capacity management is largely handled by the service rather than by manual instance resizing or complex scaling scripts.

The final step—configuring EC2 instances to connect to the new Aurora endpoint—is an infrastructure configuration change, not an application code rewrite. The application continues to speak the same MySQL-compatible protocol (Aurora MySQL-compatible), preserving compatibility while improving resilience to spikes.

Therefore, D best meets the requirement to improve availability during sudden traffic increases driven by heavy database writes, with minimal operational effort and no application code changes.

Question: 669

A company is migrating its workloads to AWS. The company has sensitive and critical data in onpremises relational databases that run on SQL Server instances. The company wants to use the AWS Cloud to increase security and reduce operational overhead for the databases.

Which solution will meet these requirements?

- A. Migrate the databases to Amazon EC2 instances. Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.
- B. Migrate the databases to a Multi-AZ Amazon RDS for SQL Server DB instance. Use an AWS Key Management Service (AWS KMS) AWS managed key for encryption.
- C. Migrate the data to an Amazon S3 bucket. Use Amazon Macie to ensure data security.

D. Migrate the databases to an Amazon DynamoDB table. Use Amazon CloudWatch Logs to ensure data security.

Answer: B

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

The company wants to increase security and reduce operational overhead for SQL Server relational databases. The most direct AWS managed service for SQL Server is Amazon RDS for SQL Server, which offloads routine database administration tasks such as provisioning, backups, patching (within managed service controls), monitoring integrations, and high availability management compared with running SQL Server on self-managed EC2.

Choosing Multi-AZ enhances availability and resilience by maintaining a synchronous standby in a separate Availability Zone with automated failover. This improves overall security posture and operational reliability because the platform handles common infrastructure failure scenarios without requiring custom clustering or manual intervention.

For encryption of sensitive data, integrating RDS with AWS KMS allows encryption at rest using KMS keys. Using an AWS managed key is operationally simple while still providing strong encryption controls, auditability, and centralized key management.

Option A still requires managing the database on EC2 (OS patching, SQL Server patching, backups, HA design, monitoring, recovery), which increases operational overhead significantly. Option C is not a database migration approach; S3 plus Macie does not provide a relational SQL Server engine for applications. Option D changes the data model by moving to DynamoDB and does not preserve SQL Server relational capabilities; CloudWatch Logs is not a data security control for database content.

Therefore, B best satisfies the goals: a managed relational database service (RDS for SQL Server) with built-in high availability (Multi-AZ) and encryption at rest via KMS, reducing ongoing operational work while improving security controls.

Question: 670

A company's application uses Network Load Balancers, Auto Scaling groups, Amazon EC2 instances, and databases that are deployed in an Amazon VPC. The company wants to capture information about traffic to and from the network interfaces in near real time in its Amazon VPC. The company wants to send the information to Amazon OpenSearch Service for analysis.

Which solution will meet these requirements?

A. Create a log group in Amazon CloudWatch Logs. Configure VPC Flow Logs to send the log data to the log group. Use Amazon Kinesis Data Streams to stream the logs from the log group to OpenSearch Service.

B. Create a log group in Amazon CloudWatch Logs. Configure VPC Flow Logs to send the log data to the log group. Use Amazon Data Firehose to stream the logs from the log group to OpenSearch Service.

C. Create a trail in AWS CloudTrail. Configure VPC Flow Logs to send the log data to the trail. Use Amazon Kinesis Data Streams to stream the logs from the trail to OpenSearch Service.

D. Create a trail in AWS CloudTrail. Configure VPC Flow Logs to send the log data to the trail. Use Amazon Data Firehose to stream the logs from the trail to OpenSearch Service.

Answer: B

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

To capture traffic metadata to and from network interfaces in a VPC, the correct telemetry source is VPC Flow Logs. Flow Logs can publish records to destinations such as CloudWatch Logs, which is a natural first landing zone for near-real-time log ingestion and centralized retention/permissions. The next requirement is to stream that data to Amazon OpenSearch Service for analysis with minimal operational overhead.

Amazon Data Firehose is purpose-built for fully managed delivery of streaming data to destinations including OpenSearch. Firehose handles buffering, batching, retry behavior, and optional transformations, and it minimizes the need to build and operate custom consumers. Therefore, sending VPC Flow Logs to CloudWatch Logs and then using Firehose to deliver to OpenSearch is an operationally efficient and scalable pattern.

Option A uses Kinesis Data Streams, which would require building/operating consumers (or additional integration components) to read from the stream and load into OpenSearch, adding operational complexity compared with Firehose's managed delivery. Options C and D are incorrect because CloudTrail records API activity and governance events, not VPC network flow records; VPC Flow Logs do not publish to CloudTrail trails.

Thus, B best meets the requirement: Flow Logs capture the network interface traffic metadata; CloudWatch Logs provides near-real-time log collection; and Firehose provides managed streaming delivery into OpenSearch for indexing and analytics with minimal operational effort.

Question: 671

A company processes streaming data by using Amazon Kinesis Data Streams and an AWS Lambda function. The streaming data comes from devices that are connected to the internet. The company is experiencing scaling problems and needs to implement shard-level control and custom checkpointing.

Which solution will meet these requirements with the LEAST latency?

- A. Connect Kinesis Data Streams to Amazon Data Firehose to ingest incoming data to an Amazon S3 bucket. Configure S3 Event Notifications to invoke the Lambda function.
- B. Increase the provisioned concurrency settings for the Lambda function. Stream the data from Kinesis Data Streams to an Amazon Simple Queue Service (Amazon SQS) standard queue. Invoke the Lambda function to process the messages.
- C. Run the Lambda function code in an Amazon Elastic Container Service (Amazon ECS) container that runs on AWS Fargate.

Change the code to use the Kinesis Client Library (KCL).

D. Increase the memory and provisioned concurrency settings for the Lambda function. Stream the data from Kinesis Data Streams to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Configure the Lambda function to be invoked by the SQS queue.

Answer: C

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

The requirements “shard-level control” and “custom checkpointing” point directly to the Kinesis Client Library (KCL), which is designed for building consumer applications that coordinate shard processing, perform load balancing across workers, and manage checkpoints in a durable store (commonly DynamoDB) with fine-grained control. While Lambda can consume from Kinesis Data Streams, its event source mapping abstracts shard coordination and checkpoint behavior; it is not the best fit when you explicitly need custom checkpointing logic and shard-level control beyond the managed integration.

Running the existing Lambda processing logic as a long-running consumer in Amazon ECS on AWS Fargate allows the company to operate a scalable KCL-based consumer fleet without managing servers. With Fargate, AWS manages the underlying compute while the application maintains direct control over shard assignment and checkpoint timing/frequency through KCL—exactly what the requirement calls for. Latency is minimized because the consumer reads directly from Kinesis Data Streams and processes records continuously, avoiding additional buffering layers or store-and-forward patterns.

Option A adds significant latency by delivering to S3 through Firehose and then triggering processing from S3 object notifications; this is optimized for delivery and batch-style processing, not low-latency streaming with shard-level control. Options B and D introduce SQS between Kinesis and processing, which adds another hop and does not inherently provide shard-level control or KCL-style checkpointing; FIFO also limits throughput and is not intended for high-scale streaming fan-in. Increasing Lambda provisioned concurrency can reduce cold starts, but it does not solve the need for custom checkpointing at the shard level.

Therefore, C best meets shard-level control and custom checkpointing requirements with the least latency by using a direct KCL consumer on a managed compute platform (Fargate).

Question: 672

A company manages millions of documents in hundreds of Amazon S3 buckets in multiple AWS

Regions. The company must determine whether any of the S3 buckets contain personally identifiable information (PII).

Which solution will meet this requirement with the LEAST operational overhead?

- A. Use Amazon Detective to detect PII in the S3 buckets.
- B. Use AWS Trusted Advisor to generate PII notifications.
- C. Use Amazon Macie to detect PII in the S3 buckets.
- D. Use AWS Lambda functions to review each file in the S3 buckets to identify PII.

Answer: C

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

Amazon Macie is the AWS managed service specifically built to discover, classify, and protect sensitive data in Amazon S3, including many common types of personally identifiable information (PII). Macie uses automated analysis and machine learning to identify sensitive data patterns at scale and produces findings that can be reviewed, prioritized, and integrated into security workflows. Because the environment spans hundreds of buckets, millions of objects, and multiple Regions, the key requirement is to minimize operational overhead while achieving broad coverage.

Option C is therefore the best fit: enabling Macie and configuring it to evaluate the targeted buckets provides a centralized, managed approach without building custom scanners, maintaining parsing logic, or operating distributed processing pipelines. Macie is designed for large-scale S3 estates and reduces ongoing maintenance compared with bespoke solutions.

Option A is incorrect because Amazon Detective is used to investigate and analyze security findings and relationships, not to classify S3 object content for PII. Option B is incorrect because Trusted Advisor provides best-practice checks (cost, security posture items, limits), but it does not inspect S3 object contents to detect PII. Option D would require building and operating custom Lambda scanning across millions of objects, handling pagination, retries, file types, performance tuning, and cost controls—high operational overhead and ongoing maintenance, which the company wants to avoid.

Therefore, C meets the requirement most directly and with the least operational burden by using the AWS service purpose-built for PII discovery in S3.

Question: 673

An ecommerce company is redesigning a web application to run on the AWS Cloud. The application needs to store static website content and must use a Microsoft SQL Server database to store customer data. The company needs to deploy the application in a resilient way across multiple Availability Zones.

Which solution will meet these requirements?

- A. Use an Amazon S3 bucket to store static content. Deploy an Amazon RDS Custom for SQL Server DB instance for the database.
- B. Use an Amazon S3 bucket to store static content. Create an Amazon RDS for SQL Server Multi-AZ deployment for the database.
- C. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach volume to store static content. Deploy an Amazon RDS for SQL Server DB instance for the database.
- D. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach volume to store static content. Deploy SQL Server on two Amazon EC2 instances in separate Availability Zones.

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

To achieve resilience across multiple Availability Zones, the architecture should use services that natively provide Multi-AZ durability and failover with minimal manual intervention. For static website content, Amazon S3 is the natural choice because it stores objects durably and is designed for high availability; it also integrates well with common web delivery patterns. For the database, the requirement is explicitly Microsoft SQL Server and resilience across AZs.

An Amazon RDS for SQL Server Multi-AZ deployment (Option B) is designed to provide high availability by maintaining a standby replica in a different Availability Zone and performing automated failover in certain failure scenarios. This meets the Multi-AZ resiliency requirement while reducing operational overhead compared to self-managed database clustering on EC2.

Option A uses RDS Custom for SQL Server, which is intended for scenarios where you need OS-level and database-level access/customization. That typically increases operational responsibility and is not the simplest “resilient Multi-AZ” answer unless the scenario demands deep customization (it does not here). Options C and D propose using EBS Multi-Attach for static content, which is not a good fit for static website assets and adds complexity; Multi-Attach is limited and does not replace object storage semantics. Option D also adds substantial operational overhead by running and managing SQL Server across EC2 instances in multiple AZs, including patching, backups, and HA configuration.

Therefore, B best meets all requirements: S3 for static object content and RDS for SQL Server MultiAZ for resilient, managed database high availability across Availability Zones.

Question: 674

A company uses AWS CloudFormation to deploy IAM resources within accounts that AWS Control Tower governs. The security team wants to prevent the deployment of IAM roles that include inline policies with the following statements:

"Effect": "Allow", "Action": "*", "Resource": "*"

Which solution will meet this requirement?

- A. Use AWS Control Tower proactive controls to block CloudFormation stacks that match these inline policy statements.
- B. Use AWS Control Tower detective controls to detect and delete IAM inline policies that contain these statements upon deployment.
- C. Use AWS Config to create a rule that detects these statements in any inline IAM policies. Configure the rule to automatically remove these statements by using the AWS-DeleteIAMInlinePolicy remediation.
- D. Use AWS Config to create a rule that detects these statements in inline IAM policies and sends a notification to the security team.

Answer: A

Explanation:

Comprehensive and Detailed 250 to 300 words of Explanation (AWS documentation-based, no links):

The security requirement is preventive: “prevent the deployment” of IAM roles that contain an overly permissive inline policy (Allow on Action:* and Resource:*). Preventive enforcement should occur at deployment time so noncompliant IAM resources are blocked before they exist. In a Control Tower-governed environment where resources are provisioned through CloudFormation, the most appropriate mechanism is a proactive control that evaluates templates and denies noncompliant stack operations.

Option A matches this intent. Control Tower proactive controls are designed to provide preventative guardrails by checking resource configurations (including Infrastructure as Code deployments) against defined rules and blocking deployments that violate policy. This stops risky IAM roles from being created in the first place, which is preferable to after-the-fact detection and cleanup.

Options B, C, and D are detective or reactive patterns. Detective controls (and AWS Config rules) can identify noncompliance after resources are created, but that does not “prevent deployment.” Even if remediation is automated (Option C), there is still a window where an overly permissive role exists and could be assumed or exploited. Option D only notifies, which is insufficient. Option B suggests deleting policies on deployment, but Control Tower detective controls are not intended to automatically delete deployed IAM policies; they primarily detect and report drift/noncompliance.

Therefore, A is the best solution because it enforces the requirement proactively by blocking noncompliant CloudFormation deployments, maintaining a stronger security posture and reducing operational burden associated with detection and remediation.

Question: 675

A company hosts an application on AWS. The application gives users the ability to upload photos and store the photos in an Amazon S3 bucket. The company wants to use Amazon CloudFront and a custom domain name to upload the photo files to the S3 bucket in the eu-west-1 Region.

Which solution will meet these requirements? (Select TWO.)

- A. Use AWS Certificate Manager (ACM) to create a public certificate in the us-east-1 Region. Use the certificate in CloudFront.
- B. Use AWS Certificate Manager (ACM) to create a public certificate in eu-west-1. Use the certificate in CloudFront.
- C. Configure Amazon S3 to allow uploads from CloudFront. Configure S3 Transfer Acceleration.
- D. Configure Amazon S3 to allow uploads from CloudFront origin access control (OAC).
- E. Configure Amazon S3 to allow uploads from CloudFront. Configure an Amazon S3 website endpoint.

Answer: A, D

CloudFront requires ACM certificates to be created in us-east-1. Origin Access Control (OAC) securely allows CloudFront to upload objects to S3 without exposing the bucket publicly. Transfer Acceleration and S3 website endpoints are unnecessary or incompatible for this use case.

Question: 676

A company operates a data lake in Amazon S3. The company wants to query and filter data directly in S3 without downloading objects.

Which solution will meet these requirements?

- A. Use Amazon Athena to query and filter the objects in Amazon S3.
- B. Use Amazon EMR to process and filter the objects.
- C. Use Amazon API Gateway to retrieve filtered results.
- D. Use Amazon ElastiCache to cache the objects.

Answer: A

Explanation:

Amazon Athena enables serverless SQL queries directly on S3 objects, eliminating the need to download or preprocess data. EMR adds operational overhead, and API Gateway/ElastiCache are not data lake query engines.

Question: 677

A company wants DevOps teams to create IAM roles, but no role may have administrative permissions.

Which solution will meet these requirements?

- A. Use SCPs to deny AdministratorAccess policy usage.
- B. Use SCPs to require a permissions boundary when creating IAM roles.
- C. Allow all permissions and auto-delete noncompliant roles.
- D. Attach restrictive permissions boundaries directly to IAM users.

Answer: B

Explanation:

Using Service Control Policies (SCPs) to enforce permissions boundaries ensures that all created roles are constrained, regardless of attached policies. This is the most secure, scalable, and AWS- recommended preventive control.

Question: 678

AWS Lambda functions need shared access to internal libraries and reference data that are updated independently by different teams.

Which solution will meet these requirements?

- A. Use Amazon EBS Multi-Attach.

- B. Store data in the Lambda /tmp directory.
- C. Use Amazon EFS mounted to Lambda.
- D. Use Amazon FSx for Windows File Server.

Answer: C

Explanation:

Amazon EFS integrates natively with Lambda, supports shared, persistent storage, and allows independent updates without redeploying functions. EBS and /tmp are not supported for this use case; FSx is unnecessary.

Question: 679

A company requires centralized auditing for all AWS accounts and compliance monitoring against AWS Foundational Security Best Practices (FSBP) with minimal operational overhead.

Which solution will meet these requirements?

- A. Deploy AWS Control Tower in the management account. Enable AWS Security Hub and Account Factory.
- B. Deploy AWS Control Tower in a member account.
- C. Use AWS Managed Services (AMS) with GuardDuty.
- D. Use AWS Managed Services (AMS) with Security Hub.

Answer: A

Explanation:

AWS Control Tower provides automatic account governance, centralized logging, CloudTrail aggregation, and integrates directly with AWS Security Hub, which evaluates compliance against FSBP standards. This is the lowest operational overhead AWS-native solution.

Question: 680

A company has an application that runs on Amazon EC2 instances and uses an Amazon Aurora database. The EC2 instances connect to the Aurora database by using user names and passwords that the company stores locally in a file.

The company changes the user names and passwords every month. The company wants to minimize the operational overhead of

credential management.

Which solution will meet these requirements?

- A. Store the credentials as a secret within AWS Secrets Manager. Assign IAM permissions to the secret. Reconfigure the application to call the secret. Enable rotation on the secret and configure rotation to occur on a monthly schedule.
- B. Use AWS Systems Manager Parameter Store to create a new parameter for the credentials. Use IAM policies to restrict access to the parameter. Reconfigure the application to access the parameter.
- C. Create an Amazon S3 bucket to store objects. Use an AWS Key Management Service (AWS KMS) key to encrypt the objects. Migrate the credentials file to the S3 bucket. Update the application to retrieve the credentials file from the S3 bucket.
- D. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume for each EC2 instance. Attach the encrypted EBS volumes to the EC2 instances. Migrate the credentials file to the new EBS volumes.

Answer: A

Explanation:

AWS Secrets Manager is purpose-built to securely store, manage, and rotate sensitive credentials such as database user names and passwords. Option A is the most operationally efficient solution because it eliminates manual password rotation, reduces human error, and centralizes secret lifecycle management. Secrets Manager integrates natively with Amazon Aurora, enabling automated credential rotation using AWS-managed or custom Lambda rotation logic. Once rotation is enabled, Secrets Manager updates the database credentials and stores the new values securely without requiring administrators to manually update files on EC2 instances.

By assigning IAM permissions to the secret, access can be tightly controlled using least-privilege principles. The application retrieves credentials at runtime, removing the need to store passwords locally on disk, which significantly improves security posture. Secrets Manager also provides auditing capabilities through AWS CloudTrail, allowing visibility into secret access and changes.

Option B (Systems Manager Parameter Store) can securely store secrets, but automated rotation is not natively supported in the same way as Secrets Manager. Implementing rotation with Parameter Store would require additional custom automation, increasing operational complexity. Option C stores credentials in S3, which is not designed for frequent credential rotation or secure secret access patterns, even when encrypted. Option D only encrypts credentials at rest on individual instances and does not address rotation, distribution, or centralized management, resulting in high operational overhead.

Therefore, A best meets the requirements by providing secure storage, automated monthly rotation, fine-grained access control, and minimal operational effort, aligning with AWS security and operational excellence best practices.

Question: 681

A company is deploying an application that processes streaming data in near-real time. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to provide the lowest possible latency between nodes.

Which networking solution meets these requirements?

- A. Place the EC2 instances in multiple VPCs, and configure VPC peering.
- B. Attach an Elastic Fabric Adapter (EFA) to each EC2 instance.
- C. Run the EC2 instances in a spread placement group.
- D. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

Answer: B

Explanation:

For near-real-time streaming workloads that are highly latency-sensitive, the network layer must provide ultra-low latency and high throughput between compute nodes. Elastic Fabric Adapter (EFA) is specifically designed to meet these requirements. EFA enables EC2 instances to communicate using a high-performance network interface that bypasses traditional TCP/IP networking, providing lower and more consistent latency. This is especially valuable for tightly coupled workloads that require fast inter-node communication.

Option B is the correct choice because EFA supports custom networking stacks and is optimized for applications such as real-time data processing, distributed computing, and high-performance workloads. EFA integrates directly with supported EC2 instance types and allows applications to take advantage of enhanced networking capabilities without requiring complex multi-VPC architectures.

Option A introduces additional network hops and latency due to VPC peering and is not suitable for ultra-low-latency workloads. Option C (spread placement groups) is designed to increase fault tolerance by placing instances on separate hardware, but it does not optimize for low-latency communication and may actually increase network distance between instances. Option D improves storage performance, not network performance, and does not affect inter-instance latency.

Therefore, B is the best solution because Elastic Fabric Adapter is the AWS-recommended approach for achieving the lowest possible network latency between EC2 instances in performance-critical streaming architectures.

Question: 682

A company is migrating a Linux-based web server group to AWS. The web servers must access shared files by using the NFS protocol. The company must not make any changes to the web server application.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket to store the shared files in S3 Standard. Grant the S3 bucket access to the web servers.
- B. Configure an Amazon CloudFront distribution. Set an Amazon S3 bucket as the origin. Store the shared files in the S3 bucket.
- C. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system on the web servers.
- D. Create an Amazon FSx for Windows File Server file system. Configure SMB protocol access for the web servers.

Answer: C

Explanation:

The key requirements are shared file access, NFS protocol compatibility, and no application changes. Amazon Elastic File System (EFS) is a fully managed, scalable file system that natively supports the NFS protocol, making it an ideal drop-in replacement for on-premises shared file systems used by Linux applications.

Option C allows the existing web servers to mount the EFS file system using standard NFS mount commands, preserving application behavior and avoiding code changes. EFS is designed to be accessed concurrently by multiple EC2 instances across Availability Zones, providing high availability and elasticity without manual capacity management. This aligns well with typical web server architectures that rely on shared content or assets.

Option A and B use Amazon S3, which is an object storage service and does not support NFS semantics. Migrating to S3 would require application changes to use object-based APIs instead of file system operations. Option D uses Amazon FSx for Windows File Server, which supports SMB, not NFS, and is intended for Windows-based workloads.

Therefore, C is the correct solution because Amazon EFS provides NFS compatibility, shared access, high availability, and minimal operational overhead while requiring no changes to the existing Linux-based web server applications.

Question: 683

A company runs a latency-sensitive gaming service in the AWS Cloud. The gaming service runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). An Amazon DynamoDB table stores the gaming data. All the infrastructure is in a single AWS Region. The main user base is in that same Region.

A solutions architect needs to update the architecture to support a global expansion of the gaming service. The gaming service must operate with the least possible latency.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution in front of the ALB.
- B. Deploy an Amazon API Gateway regional API endpoint. Integrate the API endpoint with the ALB.
- C. Create an accelerator in AWS Global Accelerator. Add a listener. Configure the endpoint to point to the ALB.
- D. Deploy the ALB and the fleet of EC2 instances to another Region. Use Amazon Route 53 with geolocation routing.

Answer: C

Explanation:

For latency-sensitive, globally distributed applications such as online gaming, minimizing network latency between users and application endpoints is critical. AWS Global Accelerator is designed specifically for this purpose. It uses the AWS global network and Anycast IP addresses to route user traffic to the closest healthy endpoint, reducing latency and improving performance for users worldwide.

Option C is the best solution because Global Accelerator directs traffic over the AWS backbone network instead of the public internet, which significantly reduces jitter and latency. It also provides built-in health checks and automatic failover, improving availability as the service expands globally. Importantly, Global Accelerator works seamlessly with existing Application Load Balancers, allowing the company to enhance performance without redesigning the application stack.

Option A (CloudFront) is optimized for caching static and cacheable content and is not ideal for realtime, stateful gaming traffic. Option B adds unnecessary API abstraction and additional latency. Option D introduces regional duplication and DNS-based routing, which is slower to react to network conditions and does not provide the same low-latency routing as Global Accelerator.

Therefore, C meets the requirement for global expansion with the least possible latency while maintaining a simple and highly performant architecture.

Question: 684

A company uses AWS to run its ecommerce platform. The platform is critical to the company's operations and has a high volume of traffic and transactions. The company configures a multi-factor authentication (MFA) device to secure its AWS account root user credentials. The company wants to ensure that it will not lose access to the root user account if the MFA device is lost.

Which solution will meet these requirements?

- A. Set up a backup administrator account that the company can use to log in if the company loses the MFA device.
- B. Add multiple MFA devices for the root user account to handle the disaster scenario.
- C. Create a new administrator account when the company cannot access the root account.
- D. Attach the administrator policy to another IAM user when the company cannot access the root account.

Answer: B

Explanation:

AWS best practices strongly recommend enabling MFA on the root user, but they also emphasize planning for MFA device loss. AWS allows the configuration of multiple MFA devices for the root user, which provides redundancy and ensures continued access in disaster scenarios.

Option B directly addresses the requirement by enabling a backup authentication method without weakening security controls.

Adding multiple MFA devices ensures that if one device is lost, damaged, or unavailable, the organization can still authenticate securely using the secondary device. This approach maintains strong security while eliminating the risk of permanent root account lockout. It also avoids timeconsuming recovery processes that require contacting AWS Support and proving account ownership.

Option A is not sufficient because IAM administrator accounts do not replace root access and cannot perform certain root-only actions.

Options C and D are not feasible because new administrator accounts or policy changes cannot be made if root access is already lost.

Therefore, B is the correct and AWS-recommended solution to ensure uninterrupted, secure access to the root account while maintaining MFA protection.

Question: 685

A company needs to design a resilient web application to process customer orders. The web application must automatically handle increases in web traffic and application usage without affecting the customer experience or losing customer orders.

Which solution will meet these requirements?

- A. Use a NAT gateway to manage web traffic. Use Amazon EC2 Auto Scaling groups to receive, process, and store processed customer orders. Use an AWS Lambda function to capture and store unprocessed orders.
- B. Use a Network Load Balancer (NLB) to manage web traffic. Use an Application Load Balancer to receive customer orders from the NLB. Use Amazon Redshift with a Multi-AZ deployment to store unprocessed and processed customer orders.
- C. Use a Gateway Load Balancer (GWLB) to manage web traffic. Use Amazon Elastic Container Service (Amazon ECS) to receive and process customer orders. Use the GWLB to capture and store unprocessed orders. Use Amazon DynamoDB to store processed customer orders.
- D. Use an Application Load Balancer to manage web traffic. Use Amazon EC2 Auto Scaling groups to receive and process customer orders. Use Amazon Simple Queue Service (Amazon SQS) to store unprocessed orders. Use Amazon RDS with a Multi-AZ deployment to store processed customer orders.

Answer: D

Explanation:

The primary requirements are automatic scaling, resilience to traffic spikes, and guaranteed order durability without losing customer orders. Option D follows a well-established, AWS-recommended decoupled architecture pattern that meets all these requirements.

Using an Application Load Balancer (ALB) provides Layer 7 request routing and distributes incoming traffic across multiple EC2 instances in Auto Scaling groups, allowing the application layer to scale horizontally as traffic increases. This ensures that spikes in web traffic do not overwhelm individual instances.

The critical resilience feature in this design is Amazon SQS, which acts as a durable buffer for unprocessed customer orders. If the application experiences a sudden traffic surge or if downstream processing slows, SQS stores incoming orders reliably until they can be processed. This decoupling prevents request loss and protects the user experience by absorbing load bursts.

Processed orders are stored in Amazon RDS with a Multi-AZ deployment, which provides synchronous replication to a standby instance in

another Availability Zone and automatic failover. This ensures high availability and data durability for transactional order data.

The other options misuse services or introduce unnecessary complexity. NAT gateways (Option A) do not manage web traffic or provide resiliency for application workloads. Redshift (Option B) is an analytics data warehouse and not suitable for transactional order processing. Gateway Load Balancer (Option C) is designed for network appliances, not application traffic management or order capture.

Therefore, D is the correct solution because it combines load balancing, auto scaling, message queuing, and highly available databases to deliver a resilient, scalable, and fault-tolerant order processing architecture.

Question: 686

A company needs a solution to prevent photos with unwanted content from being uploaded to the company's web application. The solution must not involve training a machine learning (ML) model.

Which solution will meet these requirements?

- A. Create and deploy a model by using Amazon SageMaker Autopilot. Create a real-time endpoint that the web application invokes when new photos are uploaded.
- B. Create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content. Create a Lambda function URL that the web application invokes when new photos are uploaded.
- C. Create an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content. Associate the function with the web application.
- D. Create an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content. Create a Lambda function URL that the web application invokes when new photos are uploaded.

Answer: B

Explanation:

The requirement explicitly states that the solution must not involve training a machine learning model, which immediately eliminates options that require custom ML development. Amazon Rekognition is a fully managed computer vision service that can analyze images and detect inappropriate or unwanted content using pretrained models provided by AWS.

Option B correctly uses Amazon Rekognition's image moderation capabilities to analyze uploaded photos and identify unsafe or unwanted content categories. By invoking Rekognition from an AWS Lambda function, the solution remains serverless, automatically scalable, and operationally efficient. The Lambda function URL provides a secure HTTPS endpoint that the web application can call synchronously during uploads.

Option A violates the requirement because SageMaker Autopilot is designed to build and train ML models. Option C is incorrect because Amazon Comprehend is a natural language processing service for text analysis, not image moderation. CloudFront Functions also do not support calling external AWS services like Rekognition. Option D uses Rekognition Video, which is intended for analyzing video streams, not static images.

Using Lambda and Rekognition together ensures minimal operational overhead, no ML training effort, and immediate enforcement of content moderation policies. The architecture also allows easy extension, such as logging moderation results or blocking uploads based on confidence thresholds.

Therefore, B is the correct solution because it meets the content moderation requirement using AWS-managed ML services without requiring model training.

Question: 687

A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website.

The company needs to reduce coupling within the application and improve website performance. A solutions architect must design the most operationally efficient process for image uploads.

Which combination of actions should the solutions architect take to meet these requirements? (Select TWO.)

- A. Configure the application to upload images to S3 Glacier Flexible Retrieval.
- B. Configure the web server to upload the original images to Amazon S3.
- C. Configure the application to upload images directly from each user's browser to Amazon S3 by using a presigned URL.
- D. Configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded.

Use the function to resize the image.

- E. Create an Amazon EventBridge rule that invokes an AWS Lambda function on a schedule to resize uploaded images.

Answer: C,D

Explanation:

The current architecture tightly couples image upload, resizing, and storage into a single synchronous workflow handled by EC2 instances. This increases upload latency and degrades user experience. The optimal solution is to decouple upload from processing and offload work to managed services.

Option C allows users' browsers to upload images directly to Amazon S3 using presigned URLs, bypassing the EC2 web servers entirely. This significantly reduces server load, network hops, and request latency while improving scalability and performance. Presigned URLs also maintain security by granting time-limited, scoped access to S3.

Option D complements this by using S3 Event Notifications to invoke an AWS Lambda function whenever a new image is uploaded. The Lambda function performs the image resizing asynchronously and stores the processed image back in S3. This event-driven, serverless pattern eliminates tight coupling between upload and processing and ensures automatic scaling with minimal operational overhead.

Option A is unsuitable because Glacier is for archival storage and not designed for active uploads or processing. Option B still routes uploads through EC2, maintaining the bottleneck. Option E introduces unnecessary delays and inefficiency by resizing images on a schedule instead of reacting to upload events.

Therefore, C and D together provide the most performant, scalable, and operationally efficient image upload and processing architecture using AWS-native, event-driven services.

Question: 688

A company uses a general-purpose instance class Amazon RDS for MySQL DB instance in a Multi-AZ configuration. The finance team runs SQL queries to generate reports. Customers experience performance issues during report generation.

A solutions architect needs to minimize the effect of the reporting job on the DB instance.

Which solution will meet these requirements?

- A. Create a proxy in Amazon RDS Proxy. Update the reporting job to query the proxy endpoint.
- B. Update the RDS DB instance configuration to use three Availability Zones.
- C. Add an RDS read replica. Update the reporting job to query the replica endpoint.
- D. Change the RDS configuration to a memory-optimized instance class.

Answer: C

Explanation:

The performance issue occurs because reporting queries compete with production traffic on the same primary database instance. The best-practice AWS solution is to offload read-heavy workloads to a separate database endpoint.

Option C adds an Amazon RDS read replica, which asynchronously replicates data from the primary instance. By redirecting reporting queries to the replica endpoint, the primary database can focus on transactional workloads, significantly improving application performance and customer experience.

Read replicas are specifically designed for this use case: scaling read capacity and isolating reporting or analytics queries. This solution requires minimal changes to the reporting job (endpoint update only) and avoids overprovisioning the primary database.

Option A (RDS Proxy) improves connection management but does not reduce query load or isolate reporting traffic. Option B is invalid because Multi-AZ does not scale reads and is not configurable across three AZs for a single instance. Option D increases instance size but does not address the underlying contention between workloads and increases cost unnecessarily.

Therefore, C is the most efficient and scalable solution to minimize reporting impact while maintaining high performance and availability.

Question: 689

A company uses a general-purpose instance class Amazon RDS for MySQL DB instance. The company has configured the DB instance in a Multi-AZ configuration across two Availability Zones as part of the company's production application.

The company's finance team needs to run SQL queries against the DB instance to generate reports. Customers have reported significant performance issues with the application during report generation.

A solutions architect needs to minimize the effect of the reporting job on the DB instance.

Which solution will meet these requirements?

- A. Create a proxy in Amazon RDS Proxy. Update the reporting job to query the proxy endpoint.
- B. Update the RDS DB instance configuration to use three Availability Zones.
- C. Add an RDS read replica. Update the reporting job to query the replica endpoint.
- D. Change the RDS configuration from a general-purpose instance class to a memory-optimized instance class.

Answer: C

Explanation:

The issue described is a classic case of read-heavy reporting workloads interfering with transactional application traffic. Although the RDS instance is configured for Multi-AZ, Multi-AZ deployments are designed for high availability and failover, not for scaling read performance. All reads and writes still occur on the primary instance, which explains why reporting queries negatively affect customer-facing performance.

Option C, adding an Amazon RDS read replica, is the most effective and AWS-recommended solution. Read replicas asynchronously replicate data from the primary instance and provide a separate endpoint for read-only workloads such as reporting, analytics, and dashboards. By redirecting the finance team's reporting queries to the replica endpoint, the primary database is freed to handle transactional workloads, significantly improving overall application performance.

Option A (RDS Proxy) helps manage database connections efficiently, especially for bursty workloads, but it does not offload query execution or isolate heavy read traffic. Option B is invalid because RDS Multi-AZ does not support spreading a single DB instance across three AZs, and even if it did, it would not reduce read contention. Option D increases instance resources but does not solve the root cause; reporting queries would still compete with production traffic and would also increase cost unnecessarily.

Therefore, C is the optimal solution because it isolates reporting workloads, scales read capacity, improves application responsiveness, and aligns with AWS best practices for high-performance database architectures.

Question: 690

A company has deployed a non-production Amazon EC2 instance by using an Amazon Linux AMI in a private subnet. The company wants to

allow a group of developers to connect to the EC2 instance remotely by using SSH without exposing the EC2 instance to the internet. The developers must be able to connect to the EC2 instance through the AWS Management Console.

Which solution will meet these requirements?

- A. Create a VPC endpoint for AWS Systems Manager in the same subnet as the EC2 instance. Allow inbound access from the endpoint security group to the EC2 instance security group on port 22. Create an IAM role for the EC2 instance and attach the AmazonSSMMangedInstanceCore policy.
- B. Create an EC2 Instance Connect Endpoint in the same subnet as the EC2 instance. Attach a security group to the endpoint that allows inbound connections on port 443. Assign the AmazonEC2InstanceConnect IAM managed policy to the group of developers.
- C. Create an EC2 Instance Connect Endpoint in the same subnet as the EC2 instance. Attach a security group to the endpoint that allows inbound connections on port 22. Assign the AmazonEC2InstanceConnect IAM managed policy to the group of developers.
- D. Create a VPC endpoint for AWS Systems Manager in the same subnet as the EC2 instance. Allow inbound access from the endpoint security group to the EC2 instance security group on port 443. Create an IAM role for the EC2 instance and attach the AmazonSSMReadOnlyAccess policy.

Answer: B

Explanation:

The requirements specify SSH access, no internet exposure, and connection through the AWS Management Console. The AWS-native solution designed specifically for this use case is EC2 Instance Connect Endpoint (EICE).

Option B correctly implements this approach. EC2 Instance Connect Endpoint enables secure SSH access to EC2 instances in private subnets without requiring a bastion host, public IP address, or inbound internet access. Developers authenticate through the AWS Management Console, and the connection is established over HTTPS (port 443), which is why the security group must allow inbound traffic on port 443.

The AmazonEC2InstanceConnect IAM managed policy grants developers permission to push temporary SSH keys to the instance, ensuring short-lived, auditable access that aligns with AWS security best practices. This approach significantly reduces attack surface and operational complexity.

Option A and D incorrectly attempt to use AWS Systems Manager for SSH access on port 22. Systems Manager Session Manager does not use SSH and operates over port 443 without opening inbound ports. Option C is incorrect because EC2 Instance Connect Endpoint does not accept inbound connections on port 22; SSH traffic is tunneled through the endpoint using HTTPS.

Therefore, B is the correct solution because it provides secure, console-based SSH access to private EC2 instances with minimal infrastructure and maximum security.

Question: 691

A company has an application that runs on Amazon EC2 instances in an Auto Scaling group. The application uses hardcoded credentials to access an Amazon RDS database.

To comply with new regulations, the company needs to automatically rotate the database password for the application service account every 90 days.

Which solution will meet these requirements?

- A. Create an AWS Lambda function to generate new passwords and upload them to EC2 instances by using SSH.
- B. Create a secret for the database credentials in AWS Secrets Manager. Enable rotation every 90 days. Modify the application to retrieve credentials from Secrets Manager.
- C. Create an Amazon ECS task to rotate passwords and upload them to EC2 instances.
- D. Create a new EC2 instance that runs a cron job to rotate passwords.

Answer: B

Explanation:

Hardcoded database credentials create security risks and operational challenges, especially when compliance requires regular credential rotation. AWS Secrets Manager is the AWS-recommended service for securely storing, managing, and rotating secrets such as database credentials.

Option B is the correct solution because Secrets Manager natively supports automated credential rotation using AWS-managed or custom Lambda functions. By enabling rotation on a 90-day schedule, Secrets Manager automatically updates the credentials in the RDS database and stores the new values securely. The application retrieves credentials dynamically at runtime, eliminating the need for storing passwords on EC2 instances.

Options A, C, and D all rely on custom scripts, SSH access, and manual distribution of secrets, which significantly increases operational overhead, security risk, and failure potential. These approaches also violate AWS best practices by spreading sensitive credentials across multiple hosts.

Secrets Manager integrates with IAM for fine-grained access control, supports auditing through AWS CloudTrail, and improves overall security posture while reducing operational complexity. Therefore, B best meets the requirements in a secure, scalable, and compliant manner.

Question: 692

A company runs an application that stores and shares photos. Users upload photos to an Amazon S3 bucket. Approximately 150 photos are uploaded daily. The company wants to create a thumbnail for each new photo and store it in a second S3 bucket.

Which solution will meet these requirements MOST cost-effectively?

- A. Use an Amazon EMR cluster and scheduled scripts.
- B. Use an always-on EC2 instance with scheduled scripts.
- C. Configure an S3 event notification to invoke an AWS Lambda function on each upload.

D. Use S3 Storage Lens to invoke a Lambda function.

Answer: C

Explanation:

The workload described is event-driven, low volume, and sporadic, making serverless architecture the most cost-effective choice. AWS Lambda, when triggered by S3 event notifications, runs only when new photos are uploaded and incurs cost only for execution time.

Option C uses S3 event notifications to invoke a Lambda function whenever a new object is created. The Lambda function generates a thumbnail and uploads it to a second S3 bucket. This solution requires no servers, no scheduling, and no idle compute costs, making it extremely cost-efficient for 150 uploads per day.

Options A and B involve running long-lived compute resources that would remain idle most of the time, resulting in unnecessary cost.

Option D is invalid because S3 Storage Lens is designed for storage analytics and reporting, not event-driven processing.

Therefore, C is the most cost-effective and operationally efficient solution.

Question: 693

A company uses Amazon EC2 instances to host a website. The website uses an Amazon S3 bucket to store media files. The company wants to automate infrastructure creation across multiple Regions and securely grant EC2 access to S3 using IAM.

Which solution will meet these requirements MOST securely?

- A. Store IAM access keys in UserData.
- B. Store access keys in S3 and reference them in CloudFormation.
- C. Use an IAM role and instance profile in CloudFormation.
- D. Retrieve access keys dynamically and store them on EC2.

Answer: C

Explanation:

The most secure and AWS-recommended way for EC2 instances to access AWS services is by using IAM roles attached through instance profiles. Option C correctly implements this pattern.

IAM roles provide temporary credentials via the EC2 metadata service, eliminating the need for longterm access keys. Using

CloudFormation ensures consistent, repeatable deployments across Regions while maintaining security best practices.

All other options expose long-term credentials, increasing the risk of compromise and violating AWS security guidance. Therefore, C is the correct and most secure solution.

Question: 694

A company needs to allow a vendor to access CloudWatch Logs in the company's AWS account by using IAM roles for cross-account access.

Which solution will meet these requirements?

- A. Create roles in both accounts and trust the company role.
- B. Create a role in the vendor account and trust the company role.
- C. Create a role in the company account and trust the company role.
- D. Create a role in the company account with permissions and trust the vendor role.

Answer: D

Explanation:

For cross-account access, the standard AWS pattern is to create an IAM role in the resource-owning account (the company's account) and configure the trust policy to allow the external vendor account to assume that role.

Option D correctly follows this model. The company creates an IAM role with the necessary CloudWatch Logs permissions and configures the trust policy to trust the vendor's IAM role ARN. The vendor then assumes this role to access logs securely.

The other options misconfigure trust relationships or place permissions in the wrong account, breaking the cross-account access model. Therefore, D is the correct solution.

Question: 695

A company is designing a web application on AWS. The application uses a VPN connection between the company's on-premises data center and the company's VPC.

The company uses Amazon Route 53 as its DNS service. The application must use private DNS records to communicate with the on-premises services from a VPC.

Which solution will meet these requirements MOST securely?

- A. Create a Route 53 Resolver outbound endpoint. Create a Resolver rule. Associate the Resolver rule with the VPC.

- B. Create a Route 53 Resolver inbound endpoint. Create a Resolver rule. Associate the Resolver rule with the VPC.
- C. Create a Route 53 private hosted zone. Associate the private hosted zone with the on-premises network.
- D. Create a Route 53 public hosted zone. Create a record for each on-premises service.

Answer: A

Explanation:

The requirement is for workloads inside a VPC to resolve private DNS names that are hosted on-premises over a VPN connection. The most secure and AWS-recommended method to achieve this is to use Amazon Route 53 Resolver outbound endpoints with forwarding rules.

Option A enables DNS queries that originate in the VPC to be securely forwarded to on-premises DNS servers through the VPN connection. A Route 53 Resolver outbound endpoint provides elastic, managed DNS forwarding capability without exposing DNS services to the public internet. By creating a Resolver rule, the company can define which domain names (for example, internal corporate domains) should be forwarded to specific on-premises DNS IP addresses. Associating the rule with the VPC ensures only authorized VPCs can use this resolution path.

Option B is used for the opposite scenario: when on-premises systems need to resolve private DNS names hosted in AWS, not when AWS needs to resolve on-premises names. Option C is invalid because Route 53 private hosted zones can only be associated with VPCs, not directly with on-premises networks. Option D is insecure because it exposes internal service records publicly and violates the requirement for private DNS.

Therefore, A is the most secure solution because it keeps DNS resolution private, uses managed AWS infrastructure, integrates cleanly with VPN connectivity, and follows AWS hybrid DNS best practices.

Question: 696

A company runs its workloads on Amazon Elastic Container Service (Amazon ECS). The container images that the ECS task definition uses need to be scanned for Common Vulnerabilities and Exposures (CVEs). New container images that are created also need to be scanned. Which solution will meet these requirements with the FEWEST changes to the workloads?

- A. Use Amazon Elastic Container Registry (Amazon ECR) as a private image repository. Enable scan on push for ECR basic scanning.
- B. Store the container images in an Amazon S3 bucket. Use Amazon Macie to scan the images.
- C. Migrate the workloads to Amazon EKS. Use ECR enhanced scanning.
- D. Store the container images in S3 and trigger Amazon Inspector scans with Lambda.

Answer: A

Explanation:

The requirement is to scan container images for CVEs with the fewest changes to existing ECS workloads. Amazon ECS integrates natively with Amazon Elastic Container Registry (ECR), making ECR the natural place to implement vulnerability scanning without modifying the orchestration platform.

Option A meets the requirement most efficiently. ECR provides built-in vulnerability scanning that can be enabled using scan on push, ensuring that every new image uploaded to the repository is automatically scanned for known CVEs. This capability requires no changes to ECS task definitions beyond referencing ECR images, which is typically already the case. Scan results are maintained by AWS and can be reviewed centrally, providing ongoing security visibility.

Option C introduces a major architectural change by migrating to EKS, which violates the “fewest changes” requirement. Options B and D misuse services: Amazon Macie is designed for sensitive data discovery in S3, not container image vulnerability scanning, and storing images in S3 breaks container-native workflows. Lambda-driven Inspector scans for container images are not required when ECR already provides integrated scanning.

Therefore, A is the optimal solution because it leverages native ECS-ECR integration, requires minimal operational changes, and provides automated CVE scanning aligned with AWS container security best practices.

Question: 697

A company hosts an application on AWS and has generated approximately 2.5 TB of data over 12 years. The data is stored on Amazon EBS volumes.

The company wants a cost-effective backup solution for long-term storage and must be able to retrieve the data within minutes for audits.

Which solution will meet these requirements?

- A. Create EBS snapshots.
- B. Use Amazon S3 Glacier Deep Archive.
- C. Use Amazon S3 Glacier Flexible Retrieval.
- D. Use Amazon Elastic File System (Amazon EFS).

Answer: C

Explanation:

The requirements combine long-term, low-cost storage with retrieval within minutes, which narrows the viable storage classes significantly. Amazon S3 Glacier Flexible Retrieval is designed for exactly this use case: infrequently accessed data that must still be retrievable quickly.

Option C provides substantially lower storage cost than EBS snapshots or EFS while supporting expedited retrievals, which can return data within minutes when required for audits. This meets both the cost and retrieval time constraints. Glacier Flexible Retrieval also integrates seamlessly with S3 lifecycle policies, enabling automated transitions from EBS snapshots or S3 Standard to archival storage.

Option A (EBS snapshots) provides durability but is more expensive for long-term storage at this scale. Option B (Glacier Deep Archive) is cheaper but does not meet the retrieval requirement, as restore times are measured in hours, not minutes. Option D (EFS) is a high-performance file system and is the most expensive option for long-term backups.

Therefore, C is the correct choice because it balances cost efficiency, durability, and audit-friendly retrieval times, aligning with AWS storage best practices for compliance-driven archival data.

Question: 698

A company is building a serverless application to process ecommerce orders. The application must handle bursts of traffic and process orders asynchronously in the order received.

Which solution will meet these requirements?

- A. Use Amazon SNS with AWS Lambda.
- B. Use Amazon SQS FIFO with AWS Lambda.
- C. Use Amazon SQS standard with AWS Batch.
- D. Use Amazon SNS with AWS Batch.

Answer: B

Explanation:

The key requirements are asynchronous processing, high availability, burst handling, and strict message ordering. Amazon SQS FIFO queues are specifically designed to guarantee exactly-once processing and ordered message delivery, making them ideal for transactional workflows like ecommerce order processing.

Option B meets all requirements. SQS FIFO preserves the order of messages within a message group and scales automatically to absorb traffic spikes. Integrating SQS FIFO with AWS Lambda enables serverless, event-driven processing with automatic scaling and no infrastructure management. Lambda processes messages as they arrive while maintaining order guarantees.

Option A (SNS) does not guarantee ordering and is designed for fan-out messaging. Option C (SQS standard) scales well but does not preserve order. Option D introduces unnecessary batch infrastructure and increased latency.

Therefore, B is the best solution because it combines ordered delivery, resilience, elasticity, and serverless compute, ensuring reliable and scalable order processing.

Question: 699 A company must protect sensitive documents in Amazon S3 from deletion or modification for a fixed retention period to meet regulatory requirements.

Which solution will meet these requirements?

- A. Enable S3 Object Lock in governance mode.
- B. Enable S3 Object Lock in compliance mode.
- C. Enable S3 versioning with lifecycle deletion rules.
- D. Transition objects to S3 Glacier Flexible Retrieval.

Answer: B

Explanation:

Regulatory compliance often requires immutability, meaning data cannot be altered or deleted by any user, including administrators, during a retention period. Amazon S3 Object Lock in compliance mode is specifically designed for this requirement.

Option B enforces a write-once-read-many (WORM) model. Once compliance mode is enabled and a retention period is set, no one—including root users—can delete or overwrite the objects until the retention expires. This provides the strongest level of protection and satisfies strict regulatory standards.

Option A (governance mode) allows privileged users to override the lock, which may not satisfy regulatory mandates. Option C does not prevent deletion or modification before retention expiration. Option D addresses storage cost, not immutability.

Therefore, B is the correct solution for regulatory-grade data protection.

Question: 700

A company is deploying a business-critical application that requires durable storage with consistent, low-latency performance.

Which storage option should a solutions architect recommend?

- A. Instance store
- B. Amazon ElastiCache (Memcached)
- C. Provisioned IOPS SSD Amazon EBS volume
- D. Throughput Optimized HDD Amazon EBS volume

Answer: C

Explanation:

Business-critical applications often require predictable, low-latency I/O and high durability. Provisioned IOPS SSD (io1 or io2) Amazon EBS volumes are specifically engineered for these workloads.

Option C provides consistent, high-performance storage with guaranteed IOPS and low latency. EBS volumes are network-attached and persist independently of the EC2 instance lifecycle, ensuring durability and data protection. Provisioned IOPS volumes are commonly used for databases, transactional systems, and latency-sensitive applications.

Option A (instance store) offers low latency but is ephemeral and loses data on instance stop or failure. Option B is in-memory caching and not durable storage. Option D is optimized for large, sequential workloads and does not provide consistent low latency.

Therefore, C is the correct choice because it delivers the required performance, durability, and reliability for mission-critical applications.

Question: 701

A company is using an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The company must ensure that Kubernetes service accounts in the EKS cluster have secure and granular access to specific AWS resources by using IAM roles for service accounts (IRSA).

Which combination of solutions will meet these requirements? (Select TWO.)

- A. Create an IAM policy that defines the required permissions. Attach the policy directly to the IAM role of the EKS nodes.
- B. Implement network policies within the EKS cluster to prevent Kubernetes service accounts from accessing specific AWS services.
- C. Modify the EKS cluster's IAM role to include permissions for each Kubernetes service account. Ensure a one-to-one mapping between IAM roles and Kubernetes roles.
- D. Define an IAM role that includes the necessary permissions. Annotate the Kubernetes service accounts with the Amazon Resource Name (ARN) of the IAM role.
- E. Set up a trust relationship between the IAM roles for the service accounts and an OpenID Connect (OIDC) identity provider.

Answer: D,E

Explanation:

IAM Roles for Service Accounts (IRSA) is the AWS-recommended mechanism to grant fine-grained, pod-level AWS permissions in Amazon EKS. IRSA avoids the security risks of granting permissions to the entire node and aligns with the principle of least privilege.

To implement IRSA, two core components are required. First, an IAM role must be created with a trust policy that references an OIDC identity provider associated with the EKS cluster. This is why Option E is required. The OIDC provider allows AWS STS to authenticate Kubernetes service accounts and issue temporary credentials. Without this trust relationship, service accounts cannot assume IAM roles securely.

Second, the Kubernetes service account must be explicitly annotated with the ARN of the IAM role it is allowed to assume. This is provided by Option D. The annotation creates a direct mapping between a specific service account and a specific IAM role, ensuring granular access control at the pod level.

Option A is incorrect because attaching policies to the node IAM role grants permissions to all pods on the node, which violates least-privilege principles. Option B addresses network-level access but does not control AWS API authorization. Option C is incorrect because IAM roles should not be overloaded with permissions for multiple service accounts, and there is no direct one-to-one mapping between IAM roles and Kubernetes roles in this way.

Therefore, D and E together form the correct and secure IRSA configuration, enabling granular, auditable, and secure access to AWS resources from EKS workloads.

Question: 702

An ecommerce company runs a multi-tier application on AWS. The frontend and backend tiers run on Amazon EC2 instances. The database tier runs on an Amazon RDS for MySQL DB instance.

The application makes frequent calls to return identical datasets from the database. These frequent calls cause performance slowdowns. A solutions architect must improve the performance of the application backend.

Which solution will meet this requirement?

- A. Configure an Amazon Simple Notification Service (Amazon SNS) topic between the EC2 instances and the RDS DB instance.
- B. Configure an Amazon ElastiCache (Redis OSS) cache. Configure the backend EC2 instances to read from the cache.
- C. Configure an Amazon DynamoDB Accelerator (DAX) cluster. Configure the backend EC2 instances to read from the cluster.
- D. Configure Amazon Data Firehose to stream the calls to the database.

Answer: B

Explanation:

The key issue is repeated reads of identical data from an RDS MySQL database, which leads to unnecessary database load and degraded performance. The AWS-recommended solution for this pattern is to introduce an in-memory cache between the application and the

database.

Amazon ElastiCache (Redis OSS) is purpose-built for caching frequently accessed data with microsecond-level latency. By caching identical query results, the backend EC2 instances can serve responses directly from memory instead of repeatedly querying the database. This significantly reduces read pressure on the RDS instance and improves overall application performance and scalability.

Option B is correct because Redis integrates cleanly with EC2-based applications and supports advanced caching patterns such as key expiration, eviction policies, and fine-grained control over cached objects. This approach also improves resilience during traffic spikes by offloading work from the database.

Option A is incorrect because SNS is a messaging service and does not cache or accelerate database queries. Option C is incorrect because DynamoDB Accelerator (DAX) works only with DynamoDB tables, not Amazon RDS. Option D is designed for streaming data ingestion and analytics, not for optimizing synchronous database access.

Therefore, B is the correct solution because it addresses the root cause of the performance problem by caching repeated database reads, following AWS best practices for high-performing application architectures.

Question: 703

A company has an application that uses an Amazon DynamoDB table for storage. A solutions architect discovers that many requests to the table are not returning the latest data.

Users have not reported latency or performance issues.

Which design change should the solutions architect recommend?

- A. Add read replicas to the table.
- B. Use a global secondary index (GSI).
- C. Request strongly consistent reads for the table.
- D. Request eventually consistent reads for the table.

Answer: C

Explanation:

Amazon DynamoDB supports two read consistency models: eventually consistent reads and strongly consistent reads. By default, DynamoDB uses eventual consistency, which means reads might not immediately reflect the most recent writes. This behavior explains why the application is not always receiving the latest data.

Because latency and performance are acceptable, the problem is not throughput or response time — it is data consistency. Therefore, the correct design change is to enable strongly consistent reads, which guarantee that reads always return the most up-to-date data after a successful write.

Option C is correct because strongly consistent reads ensure that the application sees the latest committed value for an item. This is critical for use cases such as inventory systems, user profiles, or transactional workflows where stale data can cause incorrect behavior.

Option A is invalid because DynamoDB does not support traditional read replicas. Option B does not affect consistency; GSIs are used to query data by alternative keys. Option D would worsen the problem because eventual consistency already causes stale reads.

Therefore, C is the appropriate recommendation to ensure correctness of returned data without changing the table schema or introducing additional complexity.

Question: 704

An events company runs a web application on Amazon EKS that uses an Amazon DynamoDB table.

The table has 1,000 RCUs and 500 WCUs provisioned. The application uses eventually consistent reads.

Traffic is usually low but occasionally spikes. During spikes, DynamoDB throttles requests, causing user-facing errors.

What should a solutions architect do to reduce these errors?

- A. Change the DynamoDB table to on-demand capacity mode.
- B. Create a DynamoDB read replica.
- C. Purchase DynamoDB reserved capacity.
- D. Use strongly consistent reads.

Answer: A

Explanation:

The application experiences unpredictable traffic spikes, which exceed the provisioned read and write capacity of the DynamoDB table, resulting in throttling errors. This is a classic scenario where on-demand capacity mode is the most appropriate solution.

Option A is correct because DynamoDB on-demand capacity automatically scales to handle sudden increases in traffic without requiring capacity planning. This eliminates throttling during unpredictable spikes and removes the operational burden of managing RCUs and WCUs. On-demand mode is particularly well-suited for workloads with variable or spiky access patterns.

Option B is incorrect because DynamoDB does not support read replicas in the same way as relational databases. Option C reduces cost for predictable workloads but does not prevent throttling if capacity limits are exceeded. Option D increases read cost and does not address capacity constraints.

Therefore, A is the best solution to improve resilience, eliminate throttling errors, and ensure a smooth user experience during traffic surges.

Question: 705

A company runs an internet-facing web application on AWS and uses Amazon Route 53 with a public hosted zone.

The company wants to log DNS response codes to support future root cause analysis.

Which solution will meet these requirements?

- A. Use Route 53 to configure query logging.
- B. Use AWS CloudTrail to record all Route 53 queries.
- C. Use Amazon CloudWatch metrics for Route 53.
- D. Use AWS Trusted Advisor for root cause analysis.

Answer: A

Explanation:

To capture DNS query and response data, including response codes, Amazon Route 53 provides query logging, which is the most precise and AWS-supported solution for this requirement.

Option A enables Route 53 query logging, which records detailed information about DNS queries, such as the queried domain, record type, source IP, and DNS response code. These logs are delivered to Amazon CloudWatch Logs, where administrators can search, analyze, and retain them for forensic investigation and root cause analysis.

Option B is incorrect because AWS CloudTrail records API calls to AWS services, not DNS query traffic. Option C provides aggregated metrics (such as query counts and health checks) but does not include per-query response codes. Option D offers best-practice recommendations but does not collect or analyze DNS query data.

Therefore, A is the correct solution because Route 53 query logging provides the detailed, low-level DNS visibility required for troubleshooting and operational analysis.

Question: 706

A company is implementing a shared storage solution for a media application that the company hosts on AWS. The company needs the ability to use SMB clients to access stored data.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Create an AWS Storage Gateway Volume Gateway. Create a file share that uses the required client protocol. Connect the application server to the file share.
- B. Create an AWS Storage Gateway Tape Gateway. Configure tapes to use Amazon S3. Connect the application server to the Tape Gateway.

- C. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the application server to the file share.
- D. Create an Amazon FSx for Windows File Server file system. Connect the application server to the file system.

Answer: D

Explanation:

The key requirements are shared storage, SMB protocol support, and least administrative overhead. Amazon FSx for Windows File Server is a fully managed AWS service designed specifically to provide native SMB file shares compatible with Windows-based workloads and SMB clients.

Option D is the optimal solution because FSx for Windows File Server eliminates the need to manage operating systems, patching, backups, file server clustering, and availability. AWS handles infrastructure management, scaling, and availability, while providing high-performance file storage that integrates seamlessly with Active Directory and supports standard Windows file permissions and SMB semantics. This makes it ideal for media applications that rely on shared file access.

Option A (Storage Gateway Volume Gateway) is designed primarily for hybrid environments that extend on-premises storage into AWS, not for cloud-native shared file systems. Option B (Tape Gateway) is intended for backup and archival workflows, not active file access. Option C requires significant operational overhead, including OS maintenance, security patching, scaling, and high availability design.

Therefore, D meets all requirements with the least administrative effort while providing a scalable, secure, and highly available SMB-compatible storage solution.

Question: 707

An analytics application runs on multiple Amazon EC2 Linux instances that use Amazon Elastic File System (Amazon EFS) Standard storage. Files are accessed infrequently after 30 days, but some older files are occasionally retrieved for reporting.

The company wants to reduce storage costs and allow throughput to scale based on file system size. The company will use the EFS lifecycle policy to transition files to Infrequent Access (IA) after 30 days.

Which solution will meet these requirements?

- A. Configure files to transition back to Standard storage on access. Specify provisioned throughput mode.
- B. Specify the provisioned throughput mode only.
- C. Configure files to transition back to Standard storage on access. Specify bursting throughput mode.
- D. Specify the bursting throughput mode only.

Answer: C

Explanation:

The requirements focus on cost optimization, automatic lifecycle management, and throughput scaling based on file system size. Amazon EFS provides two throughput modes: bursting and provisioned. Bursting throughput automatically scales with the size of the file system, which directly aligns with the company's needs.

Option C is correct because it combines two essential features. First, configuring lifecycle management to transition files from Standard to Infrequent Access (IA) after 30 days reduces storage costs for infrequently accessed data. Second, enabling automatic transition back to Standard storage upon access ensures that performance-sensitive reporting workloads are not negatively affected when older files are read.

Using bursting throughput mode allows throughput to grow naturally as the file system size increases, eliminating the need to manually manage or pay for fixed throughput levels. This is ideal for workloads with variable access patterns.

Option A incorrectly specifies provisioned throughput, which is designed for predictable, high-throughput workloads and incurs additional cost. Option B ignores lifecycle transitions. Option D does not account for automatic return to Standard storage, which could lead to suboptimal performance for accessed files.

Therefore, C provides the best balance of performance, cost efficiency, and minimal operational overhead.

Question: 708

A company has a web application with sporadic usage patterns. Usage is heavy at the beginning of each month, moderate weekly, and unpredictable during the week. The application uses a MySQL database and must move to AWS without database modifications.

Which solution will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon RDS for MySQL
- C. MySQL-compatible Amazon Aurora Serverless
- D. MySQL on Amazon EC2 in an Auto Scaling group

Answer: C

Explanation:

The database workload has high variability and unpredictability, which makes fixed-capacity database solutions inefficient and costly. The company also requires MySQL compatibility without schema or application changes.

Amazon Aurora Serverless (MySQL-compatible) is specifically designed for this scenario. Option C automatically scales database capacity up or down based on workload demand, making it highly cost-effective for sporadic usage patterns. During peak periods, Aurora Serverless scales seamlessly to handle increased load, and during idle periods, capacity scales down, significantly reducing cost.

Option A (DynamoDB) is not suitable because it is a NoSQL service and would require significant application and schema changes. Option B

(RDS for MySQL) uses fixed instance sizes and would require overprovisioning to handle peak loads, increasing cost. Option D introduces high operational overhead and does not provide automatic database scaling.

Therefore, C best meets the requirements by providing MySQL compatibility, serverless scaling, and cost optimization for variable workloads.

Question: 709

A company is developing a social media application that must scale rapidly and handle long-running, ordered processes that store large amounts of relational data. Components must scale independently and evolve without downtime.

Which combination of AWS services will meet these requirements?

- A. Amazon ECS with Fargate, Amazon RDS, and Amazon SQS
- B. Amazon ECS with Fargate, Amazon RDS, and Amazon SNS
- C. AWS Lambda, Amazon DynamoDB Streams, and AWS Step Functions
- D. AWS Elastic Beanstalk, Amazon RDS, and Amazon SNS

Answer: A

Explanation:

The application requires independent scaling, ordered asynchronous processing, and relational data storage. Option A provides a microservices-oriented, decoupled architecture using managed services.

Amazon ECS with AWS Fargate enables containerized workloads without server management and allows each component to scale independently. Amazon RDS supports relational data storage for user profiles and content. Amazon SQS ensures reliable, ordered message processing and decouples long-running background tasks, allowing services to evolve independently.

Option B uses SNS, which does not guarantee message ordering. Option C uses DynamoDB, which does not meet the relational data requirement. Option D offers less granular scaling and higher coupling.

Therefore, A best meets scalability, resilience, and modularity requirements.

Question: 710

A website runs on Amazon EC2 behind an ALB with Amazon CloudFront in front. The site is receiving a high rate of unwanted requests from specific IP addresses.

How should the solutions architect address this problem?

- A. Use AWS Shield to configure IP deny rules.
- B. Increase Auto Scaling capacity.
- C. Configure VPC network ACL deny rules.
- D. Use AWS WAF with a rate-based rule on the CloudFront distribution.

Answer: D

Explanation:

The most effective and AWS-recommended way to mitigate unwanted traffic at the edge is to use AWS WAF integrated with CloudFront. Option D allows the architect to block or throttle malicious traffic before it reaches the ALB or EC2 instances.

Rate-based rules automatically block IP addresses that exceed a defined request threshold, making them ideal for mitigating abusive traffic patterns. This reduces load on backend resources and improves performance for legitimate users.

Options A, B, and C are less effective or introduce unnecessary cost and complexity. Shield does not provide granular rate control, Auto Scaling increases cost, and NACLs are difficult to manage dynamically.

Therefore, D is the correct and most secure solution.

Question: 711

A company wants to use AWS Direct Connect to connect on-premises networks to AWS. The company runs many VPCs in a single Region and plans to scale to hundreds of VPCs.

Which service will simplify and scale the network architecture?

- A. VPC endpoints
- B. AWS Transit Gateway
- C. Amazon Route 53
- D. AWS Secrets Manager

Answer: B

Explanation:

As the number of VPCs grows, managing individual VPC connections becomes complex and unscalable. AWS Transit Gateway is specifically

designed to act as a central hub that connects multiple VPCs and on-premises networks through Direct Connect.

Option B simplifies network architecture by allowing hundreds or thousands of VPCs to connect through a single gateway, reducing routing complexity and operational overhead. Transit Gateway supports scalable routing, centralized inspection, and simplified expansion.

The other options do not address network connectivity or scaling challenges. Therefore, B is the correct solution.

Question: 712

A company's expense tracking application gives users the ability to upload images of receipts. The application analyzes the receipts to extract information and stores the raw images in Amazon S3. The application is written in Java and runs on Amazon EC2 On-Demand Instances in an Auto Scaling group behind an Application Load Balancer.

The compute costs and storage costs have increased with the popularity of the application.

Which solution will provide the MOST cost savings without affecting application performance?

- A. Purchase a Compute Savings Plan for the maximum number of necessary EC2 instances. Store the uploaded files in Amazon Elastic File System (Amazon EFS).
- B. Decrease the minimum number of EC2 instances in the Auto Scaling group. Use On-Demand Instances for peak scaling. Store the uploaded files in Amazon Elastic File System (Amazon EFS).
- C. Decrease the maximum number of EC2 instances in the Auto Scaling group. Set up S3 Lifecycle policies to archive the raw images to lower-cost storage tiers after 30 days.
- D. Purchase a Compute Savings Plan for the minimum number of necessary EC2 instances. Use On-Demand Instances for peak scaling. Set up S3 Lifecycle policies to archive the raw images to lower-cost storage tiers after 30 days.

Answer: D

Explanation:

The question asks for the most cost savings while not affecting application performance, and it explicitly calls out both compute costs and storage costs increasing. The best answer is the option that optimizes both cost drivers without reducing capacity needed for peak performance. Option D accomplishes this by combining two well-established cost optimization mechanisms: a Compute Savings Plan for baseline compute usage and S3 Lifecycle policies for lower-cost storage of older receipt images.

A Compute Savings Plan is best applied to the steady, always-needed portion of compute capacity. In an Auto Scaling group, there is typically a predictable minimum footprint that runs continuously. Purchasing a Savings Plan for that baseline delivers substantial discounts compared with On-Demand pricing while keeping performance unchanged. For traffic spikes, continuing to use On-Demand Instances is operationally simple and preserves scaling behavior and customer experience. This avoids the risk that reducing minimum capacity could introduce slower scale-out or insufficient resources during sudden load.

On the storage side, raw receipt images are stored in Amazon S3, and many applications access the newest objects most frequently. Using S3 Lifecycle policies to transition objects after 30 days to lower-cost tiers (such as infrequent access or archival tiers as appropriate for retrieval needs) reduces storage spend without changing how new uploads are handled or processed. It also keeps the application fast for recent receipts, where performance typically matters most.

Options A and B move storage to EFS, which is generally more expensive than S3 for object storage patterns and does not align with “stored as objects.” Option C reduces the maximum Auto Scaling capacity, which can directly harm performance during demand surges. Therefore, D provides the largest cost reduction while preserving performance and scaling characteristics.

Question: 713

A company has hired an external vendor to work in the company's AWS account. The vendor uses an automated tool that the vendor hosts in its own AWS account. The vendor does not have IAM access to the company's AWS account. A solutions architect needs to grant access to the vendor.

Which solution will meet these requirements MOST securely?

- A. Create an IAM role in the company's account to delegate access to the vendor's IAM role. Attach the appropriate IAM policies to the new IAM role to grant the permissions that the vendor requires.
- B. Create an IAM user in the company's account with a password. Attach the appropriate IAM policies to the IAM user.
- C. Create an IAM group in the company's account. Add the IAM user for the vendor's automated tool from the vendor account to the IAM group. Attach policies to the group.
- D. Create a new identity provider (IdP) of provider type AWS account. Supply the vendor's AWS account ID and username. Attach policies to the IdP.

Answer: A

Explanation:

The most secure way to grant an external vendor's AWS-hosted automation tool access into a company AWS account is to use cross-account IAM role assumption. Option A follows the standard AWS pattern: the company creates an IAM role in its own account that has only the permissions the vendor needs, and the role's trust policy allows the vendor's IAM role (in the vendor's AWS account) to assume it. This approach avoids creating long-lived credentials in the company account and supports least privilege, auditing, and easy revocation.

With role assumption, the vendor continues to authenticate in its own AWS account and uses AWS STS to obtain temporary credentials when accessing the company account. Temporary credentials reduce exposure compared to access keys because they expire automatically and can be constrained with session duration, external IDs (where appropriate), and conditions. The company retains full control over permissions through the role policy, and CloudTrail can record AssumeRole events and subsequent API activity for auditing.

Option B is less secure because it creates an IAM user in the company account and would typically require managing long-lived credentials for the vendor tool, increasing the risk of leakage and ongoing operational overhead. Option C is not valid because IAM users are not

shared across accounts; you cannot “add” a user from another account into an IAM group in your account. Option D is incorrect because “AWS account” is not a typical IAM identity provider type for this purpose; crossaccount access is done through role trust policies, not by creating an IdP with an AWS account ID and username.

Therefore, A is the most secure solution because it uses short-lived credentials, enforces least privilege, centralizes control in the company account, and provides strong auditability without issuing permanent access keys.

Question: 714

An ecommerce company is migrating its on-premises workload to the AWS Cloud. The workload consists of a web application and a backend Microsoft SQL Server database. The company expects a high volume of customers during a promotional event. The new AWS infrastructure must be highly available and scalable.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Migrate the web application to two EC2 instances across two Availability Zones behind an Application Load Balancer. Migrate the database to Amazon RDS for Microsoft SQL Server with read replicas in both Availability Zones.
- B. Migrate the web application to an EC2 instance in an Auto Scaling group across two Availability Zones behind an Application Load Balancer. Migrate the database to two EC2 instances across separate Regions with database replication.
- C. Migrate the web application to EC2 instances in an Auto Scaling group across two Availability Zones behind an Application Load Balancer. Migrate the database to Amazon RDS with a Multi-AZ deployment.
- D. Migrate the web application to three EC2 instances across three Availability Zones behind an Application Load Balancer. Migrate the database to three EC2 instances across three Availability Zones.

Answer: C

Explanation:

The requirements are high availability, scalability for traffic spikes, and least administrative overhead. Option C matches established AWS reference patterns: use an Application Load Balancer + Auto Scaling group for the web tier and a managed database service with Multi-AZ for the database tier.

For the web application, placing EC2 instances in an Auto Scaling group across multiple Availability Zones allows the fleet to scale out automatically during the promotional event and to scale in afterward, preserving performance while controlling cost. The Application Load Balancer distributes requests across instances and AZs and removes unhealthy instances from rotation, improving availability.

For the database, Amazon RDS Multi-AZ is the simplest managed HA solution for relational engines, including SQL Server. Multi-AZ provides synchronous replication to a standby instance in another Availability Zone and supports automatic failover. This reduces administrative effort because AWS handles replication, failover orchestration, backups, and underlying infrastructure maintenance within service capabilities. It also ensures database availability during instance or AZ-level disruptions.

Option A is problematic because RDS for SQL Server does not use read replicas as the primary HA mechanism the way Aurora does, and focusing on replicas does not directly address HA in the simplest supported form. Option B and D rely on self-managed SQL Server on EC2, which requires significant operational overhead (patching, backups, replication setup, failover testing, monitoring, and recovery). Option B also suggests multi-Region replication, which adds complexity and is not required to meet “highly available and scalable” for a single-event scale scenario.

Therefore, C provides scalable compute with managed load balancing and managed database high availability, meeting the requirements with the least operational burden.

Question: 715

A company creates a VPC that has one public subnet and one private subnet. The company attaches an internet gateway to the VPC. An Application Load Balancer (ALB) in the public subnet communicates with Amazon EC2 instances in the private subnet.

The EC2 instances in the private subnet must be able to download operating system and application updates from the internet. The instances must not be accessible from the internet.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Associate an Elastic IP address with the NAT gateway.
- B. Add a route of 0.0.0.0/0 to the private subnet route table. Set the NAT gateway as a target.
- C. Deploy a NAT gateway in the public subnet.
- D. Deploy a NAT gateway in the private subnet.
- E. Add a route of 0.0.0.0/0 to the public subnet route table. Set the NAT gateway as a target.
- F. Associate an Elastic IP address with the internet gateway.

Answer: A,B,C

Explanation:

Instances in a private subnet cannot directly reach the internet because they do not have public IP addresses and the private subnet route table typically does not send 0.0.0.0/0 traffic to an internet gateway. However, these instances still need outbound-only internet access to download patches and updates while remaining inaccessible from the internet. The standard AWS design to achieve this is a NAT gateway deployed in a public subnet.

Option C is required because a NAT gateway must reside in a public subnet that has a route to the internet gateway. This allows the NAT gateway to forward outbound traffic from private instances to the internet and return responses, without allowing inbound connections initiated from the internet to those instances.

Option A is required because a NAT gateway uses an Elastic IP address to represent its public-facing identity on the internet. Without an Elastic IP, the NAT gateway cannot communicate with internet endpoints. The Elastic IP is associated with the NAT gateway, not with the internet gateway.

Option B is required because the private subnet needs a route for 0.0.0.0/0 (default route) that targets the NAT gateway. This ensures that outbound internet-bound traffic from the private instances is sent to the NAT gateway rather than being dropped.

Option D is incorrect because a NAT gateway in a private subnet would not have a working route to the internet gateway and would not provide internet egress. Option E is incorrect because the public subnet's default route should point to the internet gateway, not to the NAT gateway. Option F is incorrect because you do not associate Elastic IPs with an internet gateway.

Therefore, A, B, and C correctly implement private subnet outbound internet access while keeping the instances unreachable from the internet.

Question: 716

A company runs a web application on Amazon EC2 instances. The application also uses an Amazon DynamoDB table. The application generates sporadic HTTP 500 errors. The DynamoDB table is operating in on-demand mode, and other applications use the table without any issues.

A solutions architect wants to resolve the HTTP 500 errors without disrupting the web application.

Which solution will meet these requirements?

- A. Configure DynamoDB to support larger write requests for increased throughput.
- B. Enable DynamoDB Streams to monitor changes in the table.
- C. Configure the application to use exponential backoff and retries to query the table.
- D. Configure the application to use strongly consistent reads.

Answer: C

Explanation:

Sporadic HTTP 500 errors in an application that depends on DynamoDB often indicate intermittent downstream failures (for example, transient throttling, temporary network issues, request timeouts, or service-side retryable exceptions). The prompt notes that the DynamoDB table is in on-demand mode and other applications use it without problems, which suggests the issue is not a sustained capacity shortfall or a systemic DynamoDB outage. Instead, the web application likely needs to handle occasional retryable failures more gracefully.

Option C is the best practice for improving reliability when calling AWS services: implement exponential backoff with retries. AWS SDKs commonly include retry behavior, but many systems still require tuning (for example, increasing maximum retries, adding jitter, ensuring idempotent operations where required, and retrying only on retryable errors). Exponential backoff reduces the chance of overwhelming a service during brief contention periods and helps the application recover quickly from transient errors without user-visible failures. This approach is also minimally disruptive because it is a configuration and logic improvement rather than an architectural migration.

Option A is not a standard DynamoDB configuration; DynamoDB has item size limits and throughput is managed via on-demand or provisioned capacity, not by “supporting larger write requests.” Option B (Streams) provides a change data capture feed but does not prevent or mitigate request failures; it is for event processing, replication, or auditing. Option D (strongly consistent reads) addresses stale reads, not HTTP 500 errors, and could increase read cost/RCU consumption without solving the root cause.

Therefore, adding exponential backoff and retries is the most appropriate and operationally sound way to reduce sporadic failures without disrupting the application's architecture or performance.

Question: 717

A security team needs to enforce the rotation of all IAM users' access keys every 90 days. If an access key is found to be older, the key must be made inactive and removed. A solutions architect must create a solution that will check for and remediate any keys older than 90 days.

Which solution meets these requirements with the LEAST operational effort?

- A. Create an AWS Config rule to check for the key age. Configure the AWS Config rule to run an AWS Batch job to remove the key.
- B. Create an Amazon EventBridge rule to check for the key age. Configure the rule to run an AWS Batch job to remove the key.
- C. Create an AWS Config rule to check for the key age. Define an Amazon EventBridge rule to schedule an AWS Lambda function to remove the key.
- D. Create an Amazon EventBridge rule to check for the key age. Define an EventBridge rule to run an AWS Batch job to remove the key.

Answer: C

Explanation:

The requirement has two parts: (1) continuously or regularly evaluate IAM access key age against a 90-day policy, and (2) remediate noncompliant keys by disabling and deleting them. The least operational effort generally comes from using managed compliance evaluation and lightweight serverless remediation.

AWS Config is designed to assess resource configuration compliance over time. Using an AWS Config rule to check access key age fits the governance model because Config maintains a compliance history, supports central reporting, and can evaluate keys across an account (and commonly across an organization with aggregation patterns). After detection, remediation should be automated with minimal infrastructure, and AWS Lambda is the lightest operational tool for directly calling IAM APIs (UpdateAccessKey to Inactive, then DeleteAccessKey) on a schedule.

Option C pairs these appropriately: Config performs compliance evaluation, and an EventBridge scheduled rule triggers a Lambda function to remediate keys older than 90 days. This avoids running and maintaining compute fleets or batch infrastructure. It also provides clear

separation of duties: Config for detection and evidence, Lambda for corrective action.

Options A, B, and D rely on AWS Batch, which introduces additional operational overhead (compute environments, job definitions, queues, scaling, and monitoring) that is unnecessary for a simple IAM housekeeping task. Also, EventBridge by itself is not a compliance evaluation service; it can schedule a job, but it does not inherently track “key age” state or compliance posture the way Config does.

Therefore, C best meets the requirement with the least operational effort by using Config for continuous compliance visibility and a scheduled Lambda for automated key deactivation and deletion.

Question: 718

A media streaming company needs to deploy its video processing application across multiple Availability Zones for high availability. The application consists of containerized microservices that process video files. The microservices must automatically recover from failures.

Which solution meets these requirements with the LEAST operational overhead?

- A. Deploy the containers to Amazon ECS with the EC2 launch type.
- B. Deploy the containers to Amazon EKS with self-managed nodes.
- C. Deploy the containers to Amazon ECS with the Fargate launch type.
- D. Deploy the containers directly to Amazon EC2 instances.

Answer: C

Explanation:

The requirements emphasize high availability across multiple Availability Zones, automatic recovery, and least operational overhead. The simplest way to achieve this for containerized microservices is to use a fully managed container orchestration service with serverless compute. Amazon ECS with the Fargate launch type provides exactly that.

With ECS Fargate, AWS manages the underlying compute infrastructure (no EC2 instances to provision, patch, scale, or secure at the OS level). You define task definitions and services, and ECS schedules tasks across Availability Zones based on your networking configuration (multiple subnets in different AZs). If a container fails, ECS service scheduling automatically replaces it, fulfilling the “recover from failures” requirement. Fargate also integrates with load balancing, CloudWatch logging, and scaling policies, reducing the operational burden of keeping the platform healthy.

Option A (ECS on EC2) is still managed orchestration, but it requires managing the EC2 capacity layer: AMI updates, instance patching, cluster capacity planning, and instance scaling. Option B (EKS with self-managed nodes) typically carries even more operational responsibility because you manage node groups, upgrades, and Kubernetes operational complexity in addition to the underlying instances. Option D is the highest operational overhead because it lacks managed scheduling and self-healing; you would need to build and maintain your own process manager and failover mechanisms.

Therefore, C is the best answer because it provides a managed, multi-AZ capable, self-healing container platform with minimal infrastructure management—matching the operational efficiency goal while maintaining high availability.

Question: 719

A company hosts a public web application on AWS with a three-tier architecture: a frontend Auto Scaling group, an application Auto Scaling group, and an Amazon RDS database.

During unexpected traffic spikes, the company notices long delays in startup time when the frontend and application tiers scale out. The company needs to improve scaling performance without negatively affecting user experience.

Which solution will meet these requirements MOST cost-effectively?

- A. Decrease the minimum number of EC2 instances for both Auto Scaling groups. Increase the desired number of instances to meet peak demand.
- B. Configure the maximum number of instances for both Auto Scaling groups to the number required for peak demand. Create a warm pool.
- C. Increase the maximum number of EC2 instances for both Auto Scaling groups to meet normal demand. Create a warm pool.
- D. Use scheduled scaling. Increase EC2 and RDS instance sizes.

Answer: B

Explanation:

The problem is scale-out latency: when traffic spikes, new instances take time to launch, initialize, and pass health checks. This can temporarily degrade user experience even if Auto Scaling eventually adds enough capacity. The goal is to improve responsiveness to sudden spikes cost-effectively.

A key Auto Scaling feature for this situation is an EC2 Auto Scaling warm pool. A warm pool maintains a set of pre-initialized instances (stopped or running, depending on configuration) that Auto Scaling can rapidly transition into service. This reduces the time needed for instance launch and initialization, enabling faster scale-out during unexpected surges.

Option B is best because it supports the application's peak needs by ensuring Auto Scaling can scale up to the required maximum capacity and keeps warm capacity ready. While a warm pool has some cost (especially if instances are kept running), it is typically more cost-effective than permanently running peak capacity all the time. It improves user experience by minimizing cold start delays during spikes.

Option A is counterproductive: lowering minimum capacity increases the likelihood of scale-out delays. Option C is illogical as written ("increase maximum to meet normal demand"); maximum should reflect peak potential, and a warm pool without sufficient max headroom will not solve the spike problem. Option D uses scheduled scaling, but the spike is unexpected, and increasing instance sizes increases cost and does not address launch delay; it may also reduce horizontal scaling flexibility.

Therefore, B is the most cost-effective way to improve scale-out performance and protect user experience during unexpected bursts.

Question: 720

An online video game company must maintain ultra-low latency for its game servers. The game servers run on Amazon EC2 instances. The company needs a solution that can handle millions of UDP internet traffic requests each second.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Application Load Balancer with the required protocol and ports. Specify the EC2 instances as targets.
- B. Configure a Gateway Load Balancer for the internet traffic. Specify the EC2 instances as targets.
- C. Configure a Network Load Balancer with the required protocol and ports. Specify the EC2 instances as targets.
- D. Launch identical game servers in separate Regions. Route traffic to both sets of servers.

Answer: C

Explanation:

The requirements are UDP support, ultra-low latency, and the ability to handle millions of requests per second in a cost-effective manner. Network Load Balancer (NLB) is the AWS service designed for high-performance Layer 4 load balancing, including TCP, TLS, and UDP. NLB can scale to very high throughput and connection rates while maintaining low latency, which is critical for real-time gaming traffic.

Option C is correct because NLB supports UDP listeners and forwards traffic directly to EC2 targets with minimal processing overhead. This yields lower latency than Layer 7 solutions and is well suited for game protocols that are latency-sensitive and often use UDP for fast, lightweight communication. NLB also provides static IP addresses per Availability Zone and integrates cleanly with Auto Scaling groups.

Option A is incorrect because an Application Load Balancer operates at Layer 7 (HTTP/HTTPS) and does not handle raw UDP traffic; it's optimized for web applications, not gaming protocols. Option B is incorrect because Gateway Load Balancer is intended to deploy and scale third-party virtual appliances (such as firewalls) and uses Geneve encapsulation; it is not meant for distributing general internet game traffic to EC2 game servers. Option D adds significant cost and complexity by deploying multi-Region fleets. It may help global latency in some scenarios, but the question asks for the most cost-effective way to handle massive UDP traffic volume; multi-Region duplication is typically not the first choice for that requirement alone.

Therefore, C (NLB with UDP) best meets the scale and latency requirements with the most efficient, purpose-built AWS load balancing service.

Question: 721

A company is deploying a critical application by using Amazon RDS for MySQL. The application must be highly available and must recover automatically. The company needs to support interactive users (transactional queries) and batch reporting (analytical queries) with no more than a 4-hour lag. The analytical queries must not affect the performance of the transactional queries.

Which solution will meet these requirements?

- A. Configure Amazon RDS for MySQL in a Multi-AZ DB instance deployment with one standby instance. Point the transactional queries to the primary DB instance. Point the analytical queries to a secondary DB instance that runs in a different Availability Zone.
- B. Configure Amazon RDS for MySQL in a Multi-AZ DB cluster deployment with two standby instances. Point the transactional queries to the primary DB instance. Point the analytical queries to the reader endpoint.
- C. Configure Amazon RDS for MySQL to use multiple read replicas across multiple Availability Zones. Point the transactional queries to the primary DB instance. Point the analytical queries to one of the replicas in a different Availability Zone.
- D. Configure Amazon RDS for MySQL as the primary database for the transactional queries with automated backups enabled. Each night, create a read-only database from the most recent snapshot to support the analytical queries. Terminate the previously created database.

Answer: C

Explanation:

The requirement has three key elements: high availability with automatic recovery, separation of transactional and analytical workloads, and acceptable reporting lag up to 4 hours. The clean AWS- native pattern for isolating heavy read/reporting traffic from transactional writes in RDS for MySQL is to use read replicas and direct reporting queries to the replica endpoint(s).

Option C meets these requirements well. The primary RDS for MySQL instance handles transactional traffic (reads/writes). One or more RDS read replicas asynchronously replicate data from the primary. Because replication is asynchronous, some lag is expected; the requirement explicitly tolerates up to a 4-hour lag, which fits the read replica model. Directing batch reporting and analytical queries to a replica prevents those expensive queries from consuming CPU, memory, and I/O on the primary, thereby protecting interactive user performance. Deploying replicas across multiple Availability Zones also improves availability of the reporting tier and reduces the risk that an AZ issue prevents reporting access.

Option A is incorrect because a standard Multi-AZ DB instance uses a synchronous standby that is not readable and is intended for failover, not for serving analytical queries. Option B describes a Multi-AZ DB cluster deployment, which does provide readable standbys in some RDS engines; however, for the classic RDS MySQL exam pattern, the most direct and widely used method to isolate analytics is read replicas. Option D creates a nightly snapshot-based read-only copy; this increases operational complexity, can result in stale data for most of the day, and introduces provisioning delays, which is unnecessary when replicas provide continuous near-real-time copies.

Therefore, C is the best fit because it provides HA for the primary, isolates reporting reads to replicas, and meets the allowed replication lag requirement.

Question: 722

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.

Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

- A. Enable Aurora Auto Scaling for Aurora Replicas. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
- B. Enable Aurora Auto Scaling for Aurora writers. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- D. Enable Aurora Auto Scaling for Aurora writers. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:

The application is described as stateful, which means user sessions or state are likely stored in memory on the application server. When migrating to an Auto Scaling group behind a load balancer, a consistent user experience typically requires session stickiness, so a user's subsequent requests continue to be routed to the same backend instance that holds their session state (unless the application is refactored to externalize session state, which is not stated here). For HTTP/HTTPS web applications, an Application Load Balancer (ALB) is the correct Elastic Load Balancing choice because it operates at Layer 7 and provides built-in support for cookie-based sticky sessions.

On the database side, Aurora PostgreSQL supports scaling read capacity by adding Aurora Replicas, and Aurora Auto Scaling can automatically adjust the number of replicas based on demand. This is the intended mechanism to scale the database tier for growing read traffic while preserving write consistency through a single writer. "Aurora writers" do not scale horizontally the same way; Aurora provides one writer endpoint at a time (with fast failover), so auto-scaling "writers" is not the right construct for demand-based scaling.

Option C combines the correct load balancer type (ALB) with stickiness and the correct database scaling mechanism (Aurora Replicas with auto scaling). Option A uses an NLB, which is Layer 4 and does not provide the same Layer 7 sticky session behavior expected for a web application pattern. Options B and D incorrectly focus on scaling Aurora writers, which does not address scaling demand patterns as effectively as scaling read replicas.

Therefore, C provides the most consistent user experience for a stateful web tier and enables the database tier to scale using Aurora Replica auto scaling.

Question: 723

An insurance company wants to migrate an application that calculates insurance premiums to AWS. The company must run calculations immediately when a customer submits information through the application. The application usually takes 10 seconds to process a calculation.

Which solution will meet this requirement?

- A. Set up an Amazon API Gateway HTTP API to receive the data. Use an AWS Lambda function to process the data immediately.
- B. Upload the customer data to an Amazon S3 bucket. Start an Amazon EC2 Spot Instance to process every data upload.
- C. Set up AWS Transfer Family to receive the customer data. Configure an Amazon EKS job to process the customer data on a schedule.
- D. Upload the data to an Amazon S3 bucket. Invoke an AWS Batch job to process every customer data upload.

Answer: A

Explanation:

The critical requirement is immediate processing upon customer submission, with a typical compute duration of about 10 seconds. This is an ideal fit for a synchronous, request-driven, serverless pattern: API Gateway + AWS Lambda. API Gateway provides a managed HTTPS front door for the application, handles request validation/throttling if needed, and integrates directly with Lambda. Lambda is designed for short-lived computations, scales automatically with request volume, and can complete 10-second tasks well within typical

Lambda execution limits.

Option A therefore meets the requirement with low latency and minimal operational overhead. When a customer submits data, API Gateway invokes the Lambda function immediately, and the function returns the calculated premium response. This preserves an interactive user experience and avoids provisioning or managing servers. It also handles bursts gracefully because Lambda concurrency can scale quickly, and API Gateway can absorb high request rates.

Option B is unsuitable because launching a Spot instance per upload introduces significant startup latency (instance launch time) and is not reliable due to Spot interruptions. Option C and D are not “immediate”: Transfer Family is for file transfers, and scheduled EKS jobs or Batch jobs introduce queuing and scheduling delays. While Batch can be event-driven, it is typically used for longer- running or batch-oriented workloads and still has job-start overhead compared to Lambda.

Therefore, A best matches the requirements for immediate execution, scalability, and low operational burden while maintaining responsive performance for customers.

Question: 724

A global company operates in multiple AWS Regions to meet data residency requirements. The company uses AWS Organizations to manage its accounts. The company wants to restrict IAM roles and access to specific Regions to prevent accidental data operations across geographic boundaries.

Which solution will meet these requirements?

- A. Configure a service control policy (SCP) to deny the `ec2:RunInstances` action in non-compliant Regions.
- B. Configure IAM policies by using the `aws:RequestedRegion` condition.

- C. Configure IAM role trust policies that use the aws:SourceIp condition.
- D. Configure AWS Config to detect unwanted access across Regions.

Answer: B

Explanation:

The requirement is to restrict actions by Region to prevent accidental cross-Region operations that could violate data residency rules. The most direct and flexible way to enforce Region-based restrictions in AWS is to use IAM condition keys, specifically aws:RequestedRegion, in identity-based policies (and commonly in permission boundaries or SCPs when managing an organization).

Option B is correct because it applies Region constraints at authorization time. By adding conditions that allow actions only when aws:RequestedRegion matches an approved list (or denies when it matches disallowed Regions), the company can prevent users and roles from creating, modifying, or accessing resources in the wrong Region. This approach is broad and can be applied across many services, not just EC2, making it suitable for enforcing data residency boundaries across a multiRegion footprint.

Option A is too narrow because it restricts only ec2:RunInstances. Data residency concerns typically apply to many services (S3, RDS, DynamoDB, KMS, etc.), and limiting only EC2 does not prevent accidental data operations elsewhere. Option C controls where requests come from (source IP), not which Region is targeted. Option D (AWS Config) is detective, not preventative; it can alert after noncompliant actions occur, which does not meet the requirement to prevent accidental operations.

Therefore, B best meets the requirement by enforcing Region-level guardrails through IAM authorization conditions, helping ensure workloads remain within approved geographic boundaries.