



**Amazon**

## **Exam Questions AWS-Solution-Architect-Associate**

Amazon AWS Certified Solutions Architect - Associate

### NEW QUESTION 1

- (Topic 4)

A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image. The container needs 50 GB of storage available for temporary files. The infrastructure must be serverless.

Which solution meets these requirements with the LEAST operational overhead?

- A. Create an AWS Lambda function that uses the Docker container image with an Amazon S3 mounted volume that has more than 50 GB of space.
- B. Create an AWS Lambda function that uses the Docker container image with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the AWS Fargate launch type. Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volume.
- D. Create a service with that task definition.
- E. Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the Amazon EC2 launch type with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space. Create a task definition for the container image.
- F. Create a service with that task definition.

**Answer: C**

#### Explanation:

The AWS Fargate launch type is a serverless way to run containers on Amazon ECS, without having to manage any underlying infrastructure. You only pay for the resources required to run your containers, and AWS handles the provisioning, scaling, and security of the cluster. Amazon EFS is a fully managed, elastic, and scalable file system that can be mounted to multiple containers, and provides high availability and durability. By using AWS Fargate and Amazon EFS, you can run your Docker container image with 50 GB of storage available for temporary files, with the least operational overhead. This solution meets the requirements of the question.

References:

- ? AWS Fargate
- ? Amazon Elastic File System
- ? Using Amazon EFS file systems with Amazon ECS

### NEW QUESTION 2

- (Topic 4)

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.
- B. Configure elastic storage scaling.
- C. Migrate the database to Amazon Redshift by using the mysqldump utility.
- D. Turn on Auto Scaling for the Amazon Redshift cluster.
- E. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Aurora.
- F. Turn on Aurora Auto Scaling.
- G. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoDB.
- H. Configure an Auto Scaling policy.

**Answer: C**

#### Explanation:

To migrate a MySQL database to AWS with compatibility and scalability, Amazon Aurora is a suitable option. Aurora is compatible with MySQL and can scale automatically with Aurora Auto Scaling. AWS Database Migration Service (AWS DMS) can be used to migrate the database from on-premises to Aurora with minimal downtime. References:

- ? What Is Amazon Aurora?
- ? Using Amazon Aurora Auto Scaling with Aurora Replicas
- ? What Is AWS Database Migration Service?

### NEW QUESTION 3

- (Topic 4)

A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours.

Which solution meets these requirements with the LEAST operational overhead?

- A. Configure point-in-time recovery for the table.
- B. Use AWS Backup for the table.
- C. Use an AWS Lambda function to make an on-demand backup of the table every hour.
- D. Turn on streams on the table to capture a log of all changes to the table in the last 24 hours. Store a copy of the stream in an Amazon S3 bucket.

**Answer: A**

#### Explanation:

Point-in-time recovery (PITR) for DynamoDB is a feature that enables you to restore your table data to any point in time during the last 35 days. PITR helps protect your table from accidental write or delete operations, such as a test script writing to a production table or a user issuing a wrong command. PITR is easy to use, fully managed, fast, and scalable. You can enable PITR with a single click in the DynamoDB console or with a simple API call. You can restore a table to a new table using the console, the AWS CLI, or the DynamoDB API. PITR does not consume any provisioned table capacity and has no impact on the performance or availability of your production applications. PITR meets the requirements of the company with the least operational overhead, as it does not require any manual backup creation, scheduling, or maintenance. It also provides per-second granularity for restoring the table to any point within the last 24 hours.

References:

- ? Point-in-time recovery for DynamoDB - Amazon DynamoDB
- ? Amazon DynamoDB point-in-time recovery (PITR)
- ? Enable Point-in-Time Recovery (PITR) for Dynamodb global tables

- ? Restoring a DynamoDB table to a point in time - Amazon DynamoDB
- ? Point-in-time recovery: How it works - Amazon DynamoDB

#### NEW QUESTION 4

- (Topic 4)

A company migrated a MySQL database from the company's on-premises data center to an Amazon RDS for MySQL DB instance. The company sized the RDS DB instance to meet the company's average daily workload. Once a month, the database performs slowly when the company runs queries for a report. The company wants to have the ability to run reports and maintain the performance of the daily workloads. Which solution will meet these requirements?

- A. Create a read replica of the databas
- B. Direct the queries to the read replica.
- C. Create a backup of the databas
- D. Restore the backup to another DB instanc
- E. Direct the queries to the new database.
- F. Export the data to Amazon S3. Use Amazon Athena to query the S3 bucket.
- G. Resize the DB instance to accommodate the additional workload.

**Answer: C**

#### Explanation:

Amazon Athena is a service that allows you to run SQL queries on data stored in Amazon S3. It is serverless, meaning you do not need to provision or manage any infrastructure. You only pay for the queries you run and the amount of data scanned<sup>1</sup>.

By using Amazon Athena to query your data in Amazon S3, you can achieve the following benefits:

? You can run queries for your report without affecting the performance of your

Amazon RDS for MySQL DB instance. You can export your data from your DB instance to an S3 bucket and use Athena to query the data in the bucket. This way, you can avoid the overhead and contention of running queries on your DB instance.

? You can reduce the cost and complexity of running queries for your report. You do

not need to create a read replica or a backup of your DB instance, which would incur additional charges and require maintenance. You also do not need to resize your DB instance to accommodate the additional workload, which would increase your operational overhead.

? You can leverage the scalability and flexibility of Amazon S3 and Athena. You can

store large amounts of data in S3 and query them with Athena without worrying about capacity or performance limitations. You can also use different formats, compression methods, and partitioning schemes to optimize your data storage and query performance<sup>1</sup>.

#### NEW QUESTION 5

- (Topic 4)

A company has two VPCs named Management and Production. The Management VPC uses VPNs through a customer gateway to connect to a single device in the data center. The Production VPC uses a virtual private gateway AWS Direct Connect connections. The Management and Production VPCs both use a single VPC peering connection to allow communication between the

What should a solutions architect do to mitigate any single point of failure in this architecture?

- A. Add a set of VPNs between the Management and Production VPCs.
- B. Add a second virtual private gateway and attach it to the Management VPC.
- C. Add a second set of VPNs to the Management VPC from a second customer gateway device.
- D. Add a second VPC peering connection between the Management VPC and the Production VPC.

**Answer: C**

#### Explanation:

This answer is correct because it provides redundancy for the VPN connection between the Management VPC and the data center. If one customer gateway device or one VPN tunnel becomes unavailable, the traffic can still flow over the second customer gateway device and the second VPN tunnel. This way, the single point of failure in the VPN connection is mitigated.

References:

? <https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html>

? <https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/VPC/vpn-tunnel-redundancy.html>

#### NEW QUESTION 6

- (Topic 4)

A company wants to use an AWS CloudFormation stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment. The solution must follow security best practices.

Which solution will meet these requirements?

- A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL
- B. Create an Amazon API Gateway REST API that has the S3 bucket as the target
- C. Configure the CloudFormation stack to use the API Gateway URL
- D. Create a presigned URL for the template object. Configure the CloudFormation stack to use the presigned URL.
- E. Allow public access to the template object in the S3 bucket
- F. Block the public access after the test environment is created

**Answer: C**

#### Explanation:

it allows CloudFormation to access the template in the S3 bucket without granting public access or creating additional resources. A presigned URL is a URL that is signed with the access key of an IAM user or role that has permission to access the object. The presigned URL can be used by anyone who receives it, but it expires after a specified time. By creating a presigned URL for the template object and configuring the CloudFormation stack to use it, the company can grant CloudFormation access to the template based on specific user requests and follow security best practices. References:

? Using Amazon S3 Presigned URLs

? Using Amazon S3 Buckets

### NEW QUESTION 7

- (Topic 4)

An ecommerce company is running a seasonal online sale. The company hosts its website on Amazon EC2 instances spanning multiple Availability Zones. The company wants its website to manage sudden traffic increases during the sale.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Auto Scaling group that is large enough to handle peak traffic load
- B. Stop half of the Amazon EC2 instance
- C. Configure the Auto Scaling group to use the stopped instances to scale out when traffic increases.
- D. Create an Auto Scaling group for the website
- E. Set the minimum size of the Auto Scaling group so that it can handle high traffic volumes without the need to scale out.
- F. Use Amazon CloudFront and Amazon ElastiCache to cache dynamic content with an Auto Scaling group set as the origin
- G. Configure the Auto Scaling group with the instances necessary to populate CloudFront and ElastiCache
- H. Scale in after the cache is fully populated.
- I. Configure an Auto Scaling group to scale out as traffic increases
- J. Create a launch template to start new instances from a preconfigured Amazon Machine Image (AMI).

**Answer:** D

#### Explanation:

The solution that meets the requirements of high availability, resiliency, and minimal operational effort is to use AWS Transfer for SFTP and an Amazon S3 bucket for storage. This solution allows the company to securely transfer files over SFTP to Amazon S3, which is a durable and scalable object storage service. The company can then modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processing. The EC2 instance can be part of an Auto Scaling group with a scheduled scaling policy to run the batch operation only at night. This way, the company can save costs by scaling down the EC2 instances when they are not needed. The other solutions do not meet all the requirements because they either use Amazon EFS or Amazon EBS for storage, which are more expensive and less scalable than Amazon S3, or they do not use a scheduled scaling policy to optimize the EC2 instances usage. References :=

? AWS Transfer for SFTP

? Amazon S3

? Amazon EC2 Auto Scaling

### NEW QUESTION 8

- (Topic 4)

A company has an application that uses Docker containers in its local data center. The application runs on a container host that stores persistent data in a volume on the host. The container instances use the stored persistent data.

The company wants to move the application to a fully managed service because the company does not want to manage any servers or storage infrastructure.

Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS) with self-managed nodes
- B. Create an Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instance
- C. Use the EBS volume as a persistent volume mounted in the containers.
- D. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch type
- E. Create an Amazon Elastic File System (Amazon EFS) volume
- F. Add the EFS volume as a persistent storage volume mounted in the containers.
- G. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch type
- H. Create an Amazon S3 bucket
- I. Map the S3 bucket as a persistent storage volume mounted in the containers.
- J. Use Amazon Elastic Container Service (Amazon ECS) with an Amazon EC2 launch type
- K. Create an Amazon Elastic File System (Amazon EFS) volume
- L. Add the EFS volume as a persistent storage volume mounted in the containers.

**Answer:** B

#### Explanation:

This solution meets the requirements because it allows the company to move the application to a fully managed service without managing any servers or storage infrastructure. AWS Fargate is a serverless compute engine for containers that runs the Amazon ECS tasks. With Fargate, the company does not need to provision, configure, or scale clusters of virtual machines to run containers. Amazon EFS is a fully managed file system that can be accessed by multiple containers concurrently. With EFS, the company does not need to provision and manage storage capacity. EFS provides a simple interface to create and configure file systems quickly and easily. The company can use the EFS volume as a persistent storage volume mounted in the containers to store the persistent data. The company can also use the EFS mount helper to simplify the mounting process. References: Amazon ECS on AWS Fargate, Using Amazon EFS file systems with Amazon ECS, Amazon EFS mount helper.

### NEW QUESTION 9

- (Topic 4)

A company needs to configure a real-time data ingestion architecture for its application. The company needs an API, a process that transforms data as the data is streamed, and a storage solution for the data.

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

- A. Deploy an Amazon EC2 instance to host an API that sends data to an Amazon Kinesis data stream
- B. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data source
- C. Use AWS Lambda functions to transform the data
- D. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
- E. Deploy an Amazon EC2 instance to host an API that sends data to AWS Glue
- F. Stop source/destination checking on the EC2 instance
- G. Use AWS Glue to transform the data and to send the data to Amazon S3.
- H. Configure an Amazon API Gateway API to send data to an Amazon Kinesis data stream
- I. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data source
- J. Use AWS Lambda functions to transform the data
- K. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
- L. Configure an Amazon API Gateway API to send data to AWS Glue
- M. Use AWS Lambda functions to transform the data
- N. Use AWS Glue to send the data to Amazon S3.

**Answer:** C

**Explanation:**

It uses Amazon Kinesis Data Firehose which is a fully managed service for delivering real-time streaming data to destinations such as Amazon S3. This service requires less operational overhead as compared to option A, B, and D. Additionally, it also uses Amazon API Gateway which is a fully managed service for creating, deploying, and managing APIs. These services help in reducing the operational overhead and automating the data ingestion process.

**NEW QUESTION 10**

- (Topic 4)

A company needs to migrate a MySQL database from its on-premises data center to AWS within 2 weeks. The database is 20 TB in size. The company wants to complete the migration with minimal downtime.

Which solution will migrate the database MOST cost-effectively?

- A. Order an AWS Snowball Edge Storage Optimized device
- B. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing change
- C. Send the Snowball Edge device to AWS to finish the migration and continue the ongoing replication.
- D. Order an AWS Snowmobile vehicle
- E. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
- F. Send the Snowmobile vehicle back to AWS to finish the migration and continue the ongoing replication.
- G. Order an AWS Snowball Edge Compute Optimized with GPU device
- H. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
- I. Send the Snowball device to AWS to finish the migration and continue the ongoing replication.
- J. Order a 1 GB dedicated AWS Direct Connect connection to establish a connection with the data center
- K. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing changes.

**Answer:** A

**Explanation:**

This answer is correct because it meets the requirements of migrating a 20 TB MySQL database within 2 weeks with minimal downtime and cost-effectively. The AWS Snowball Edge Storage Optimized device has up to 80 TB of usable storage space, which is enough to fit the database. The AWS Database Migration Service (AWS DMS) can migrate data from MySQL to Amazon Aurora, Amazon RDS for MySQL, or MySQL on Amazon EC2 with minimal downtime by continuously replicating changes from the source to the target. The AWS Schema Conversion Tool (AWS SCT) can convert the source schema and code to a format compatible with the target database. By using these services together, the company can migrate the database to AWS with minimal downtime and cost. The Snowball Edge device can be shipped back to AWS to finish the migration and continue the ongoing replication until the database is fully migrated.

References:

? <https://docs.aws.amazon.com/snowball/latest/developer-guide/device-differences.html>

? [https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Source.MySQL.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.MySQL.html)

? [https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/CHAP\\_Source.MySQL.htm](https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/CHAP_Source.MySQL.htm)

**NEW QUESTION 10**

- (Topic 4)

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
- B. Deploy the Lambda functions inside a VPC.
- C. Point the client driver at an RDS proxy endpoint
- D. Deploy the Lambda functions inside a VPC.
- E. Point the client driver at an RDS custom endpoint
- F. Deploy the Lambda functions outside a VPC.
- G. Point the client driver at an RDS proxy endpoint
- H. Deploy the Lambda functions outside a VPC.

**Answer:** B

**Explanation:**

To maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections, a solutions architect should point the client driver at an RDS proxy endpoint and deploy the Lambda functions inside a VPC. An RDS proxy is a fully managed database proxy that allows applications to share connections to a database, improving database availability and scalability. By using an RDS proxy, the Lambda functions can reuse existing connections, rather than creating new ones for every invocation, reducing the connection overhead and latency. Deploying the Lambda functions inside a VPC allows them to access the private RDS DB instance securely and efficiently, without exposing it to the public internet. References:

? [Using Amazon RDS Proxy with AWS Lambda](#)

? [Configuring a Lambda function to access resources in a VPC](#)

**NEW QUESTION 13**

- (Topic 4)

A company is building a shopping application on AWS. The application offers a catalog that changes once each month and needs to scale with traffic volume. The company wants the lowest possible latency from the application. Data from each user's shopping cart needs to be highly available. User session data must be available even if the user is disconnected and reconnects.

What should a solutions architect do to ensure that the shopping cart data is preserved at all times?

- A. Configure an Application Load Balancer to enable the sticky sessions feature (session affinity) for access to the catalog in Amazon Aurora.
- B. Configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- C. Configure Amazon OpenSearch Service to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- D. Configure an Amazon EC2 instance with Amazon Elastic Block Store (Amazon EBS) storage for the catalog and shopping cart
- E. Configure automated snapshots.

**Answer: B**

**Explanation:**

To ensure that the shopping cart data is preserved at all times, a solutions architect should configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session. This solution has the following benefits:

- ? It offers the lowest possible latency from the application, as ElastiCache for Redis is a blazing fast in-memory data store that provides sub-millisecond latency to power internet-scale real-time applications<sup>1</sup>.
- ? It scales with traffic volume, as ElastiCache for Redis supports horizontal scaling by adding more nodes or shards to the cluster, and vertical scaling by changing the node type<sup>2</sup>.
- ? It is highly available, as ElastiCache for Redis supports replication across multiple Availability Zones and automatic failover in case of a primary node failure<sup>3</sup>.
- ? It preserves user session data even if the user is disconnected and reconnects, as ElastiCache for Redis can store session data, such as user login information and shopping cart contents, in a persistent and durable manner using snapshots or AOF (append-only file) persistence<sup>4</sup>.

References:

- ? 1: <https://aws.amazon.com/elasticache/redis/>
- ? 2: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Scaling.html>
- ? 3: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.html>
- ? 4: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/backups.html>

**NEW QUESTION 15**

- (Topic 4)

A financial company needs to handle highly sensitive data The company will store the data in an Amazon S3 bucket The company needs to ensure that the data is encrypted in transit and at rest The company must manage the encryption keys outside the AWS Cloud Which solution will meet these requirements?

- A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key
- B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key
- C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE)
- D. Encrypt the data at the company's data center before storing the data in the S3 bucket

**Answer: D**

**Explanation:**

This option is the only solution that meets the requirements because it allows the company to encrypt the data with its own encryption keys and tools outside the AWS Cloud. By encrypting the data at the company's data center before storing the data in the S3 bucket, the company can ensure that the data is encrypted in transit and at rest, and that the company has full control over the encryption keys and processes. This option also avoids the need to use any AWS encryption services or features, which may not be compatible with the company's security policies or compliance standards.

\* A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

\* B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

\* C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards. References:

- ? 1 Protecting data with encryption - Amazon Simple Storage Service
- ? 2 Protecting data with server-side encryption - Amazon Simple Storage Service
- ? 3 Protecting data by using client-side encryption - Amazon Simple Storage Service
- ? 4 AWS Key Management Service Concepts - AWS Key Management Service

**NEW QUESTION 17**

- (Topic 4)

An image hosting company uploads its large assets to Amazon S3 Standard buckets The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.

Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

- A. Move assets to S3 Intelligent-Tiering after 30 days.
- B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
- C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
- D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days
- E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

**Answer: AB**

**Explanation:**

S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead<sup>1</sup>. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.

S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle<sup>2</sup>. You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage costs and avoid paying for parts that are not used.

Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs<sup>3</sup>. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.

Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed<sup>1</sup>. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.

Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed<sup>1</sup>. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL: 1: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html> 2:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html> 3: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty-bucket.html#delete-bucket-considerations> : <https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html> : <https://aws.amazon.com/certification/certified-solutions-architect-associate/>

### NEW QUESTION 18

- (Topic 4)

A media company stores movies in Amazon S3. Each movie is stored in a single video file that ranges from 1 GB to 10 GB in size.

The company must be able to provide the streaming content of a movie within 5 minutes of a user purchase. There is higher demand for movies that are less than 20 years old than for movies that are more than 20 years old. The company wants to minimize hosting service costs based on demand.

Which solution will meet these requirements?

- A. Store all media content in Amazon S3. Use S3 Lifecycle policies to move media data into the Infrequent Access tier when the demand for a movie decreases.
- B. Store newer movie video files in S3 Standard Store older movie video files in S3 Standard-Infrequent Access (S3 Standard-IA). When a user orders an older movie, retrieve the video file by using standard retrieval.
- C. Store newer movie video files in S3 Intelligent-Tierin
- D. Store older movie video files inS3 Glacier Flexible Retrieva
- E. When a user orders an older movie, retrieve the video file by using expedited retrieval.
- F. Store newer movie video files in S3 Standar
- G. Store older movie video files in S3 Glacier Flexible Retrieva
- H. When a user orders an older movie, retrieve the video file by using bulk retrieval.

**Answer: C**

#### Explanation:

This solution will meet the requirements of minimizing hosting service costs based on demand and providing the streaming content of a movie within 5 minutes of a user purchase. S3 Intelligent-Tiering is a storage class that automatically optimizes storage costs by moving data to the most cost-effective access tier when access patterns change. It is suitable for data with unknown, changing, or unpredictable access patterns, such as newer movies that may have higher demand<sup>1</sup>. S3 Glacier Flexible Retrieval is a storage class that provides low-cost storage for archive data that is retrieved asynchronously. It offers flexible data retrieval options from minutes to hours, and free bulk retrievals in 5-12 hours. It is ideal for backup, disaster recovery, and offsite data storage needs<sup>2</sup>. By using expedited retrieval, the user can access the older movie video file in 1-5 minutes, which meets the requirement of 5 minutes<sup>3</sup>.

References: 1: Amazon S3 Intelligent-Tiering Storage Class | AWS<sup>4</sup>, Overview section2: Amazon S3 Glacier Flexible Retrieval and Glacier Deep Archive Retrieval ...1, Amazon S3 Glacier Flexible Retrieval section3: Amazon S3 Glacier Flexible Retrieval and Glacier Deep Archive Retrieval ...1, Retrieval Rates section.

### NEW QUESTION 19

- (Topic 4)

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 VPN attachments for the on-premises connections.
- D. Create an AWS Direct Connect connection between the on-premises locations and acentral VP
- E. Connect the central VPC to other VPCs by using peering connections.

**Answer: C**

#### Explanation:

A transit gateway is a network transit hub that enables you to connect your VPCs and on-premises networks in a centralized and scalable way. You can create VPC attachments to connect your VPCs to the transit gateway, and VPN attachments to connect your on-premises networks to the transit gateway over the internet. The transit gateway acts as a router between the attached networks, and simplifies the administration of new connections by reducing the number of peering or VPN connections required. You can also use transit gateway route tables to control the routing of traffic between the attached networks. By creating a transit gateway and using VPC and VPN attachments, you can meet the requirements of the company with the least administrative overhead.

References:

? AWS Transit Gateway

? Transit gateway attachments

? Transit gateway route tables

### NEW QUESTION 21

- (Topic 4)

A company hosts a multi-tier web application on Amazon Linux Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company observes that the Auto Scaling group launches more On-Demand Instances when the application's end users access high volumes of static web content. The company wants to optimize cost.

What should a solutions architect do to redesign the application MOST cost-effectively?

- A. Update the Auto Scaling group to use Reserved Instances instead of On-Demand Instances.
- B. Update the Auto Scaling group to scale by launching Spot Instances instead of On- Demand Instances.
- C. Create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket.
- D. Create an AWS Lambda function behind an Amazon API Gateway API to host the static website contents.

**Answer:** C

**Explanation:**

This answer is correct because it meets the requirements of optimizing cost and reducing the workload on the database. Amazon CloudFront is a content delivery network (CDN) service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the request is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance. You can create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket, which is an origin that you define for CloudFront. This way, you can offload the requests for static web content from your EC2 instances to CloudFront, which can improve the performance and availability of your website, and reduce the cost of running your EC2 instances.

References:

? <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>

**NEW QUESTION 23**

- (Topic 4)

A company wants to securely exchange data between its software as a service (SaaS) application Salesforce account and Amazon S3. The company must encrypt the data at rest by using AWS Key Management Service (AWS KMS) customer managed keys (CMKs). The company must also encrypt the data in transit. The company has enabled API access for the Salesforce account.

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

- A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3.
- B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3.
- C. Create Amazon AppFlow flows to transfer the data securely from Salesforce to Amazon S3.
- D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3.

**Answer:** C

**Explanation:**

Amazon AppFlow is a fully managed integration service that enables users to transfer data securely between SaaS applications and AWS services. It supports Salesforce as a source and Amazon S3 as a destination. It also supports encryption of data at rest using AWS KMS CMKs and encryption of data in transit using SSL/TLS1. By using Amazon AppFlow, the solution can meet the requirements with the least development effort.

\* A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves writing custom code to interact with Salesforce and Amazon S3 APIs, handle authentication, encryption, error handling, and monitoring2.

\* B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves creating a state machine definition to orchestrate the data transfer task, and invoking Lambda functions or other services to perform the actual data transfer3.

\* D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves using the Amazon AppFlow Custom Connector SDK to build and deploy a custom connector for Salesforce, which requires additional configuration and management. Reference URL: <https://aws.amazon.com/appflow/>

**NEW QUESTION 24**

- (Topic 4)

A company has an organization in AWS Organizations that has all features enabled The company requires that all API calls and logins in any existing or new AWS account must be audited The company needs a managed solution to prevent additional work and to minimize costs The company also needs to know when any AWS account is not compliant with the AWS Foundational Security Best Practices (FSBP) standard.

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

- A. Deploy an AWS Control Tower environment in the Organizations management account Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.
- B. Deploy an AWS Control Tower environment in a dedicated Organizations member account Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.
- C. Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ) Submit an RFC to self-service provision Amazon GuardDuty in the MALZ.
- D. Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ) Submit an RFC to self-service provision AWS Security Hub in the MALZ.

**Answer:** A

**Explanation:**

AWS Control Tower is a fully managed service that simplifies the setup and governance of a secure, compliant, multi-account AWS environment. It establishes a landing zone that is based on best-practices blueprints, and it enables governance using controls you can choose from a pre-packaged list. The landing zone is a well-architected, multi-account baseline that follows AWS best practices. Controls implement governance rules for security, compliance, and operations. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts. It aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer, as well as from AWS Partner solutions. AWS Security Hub continuously monitors your environment using automated compliance checks based on the AWS best practices and industry standards, such as the AWS Foundational Security Best Practices (FSBP) standard. AWS Control Tower Account Factory is a feature that automates the provisioning of new AWS accounts that are preconfigured to meet your business, security, and compliance requirements. By deploying an AWS Control Tower environment in the Organizations management account, you can leverage the existing organization structure and policies, and enable AWS Security Hub and AWS Control Tower Account Factory in the environment. This way, you can audit all API calls and logins in any existing or new AWS account, monitor the compliance status of each account with the FSBP standard, and provision new accounts with ease and consistency. This solution meets the requirements with the least operational overhead, as you do not need to manage any infrastructure, perform any data migration, or submit any requests for changes. References:

? AWS Control Tower

? [AWS Security Hub]

? [AWS Control Tower Account Factory]

**NEW QUESTION 29**

- (Topic 4)

A company is deploying a new application to Amazon Elastic Kubernetes Service (Amazon EKS) with an AWS Fargate cluster. The application needs a storage solution for data persistence. The solution must be highly available and fault tolerant. The solution also must be shared between multiple application containers. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create Amazon Elastic Block Store (Amazon EBS) volumes in the same Availability Zones where EKS worker nodes are placed
- B. Register the volumes in a StorageClass object on an EKS cluster
- C. Use EBS Multi-Attach to share the data between containers.
- D. Create an Amazon Elastic File System (Amazon EFS) file system
- E. Register the file system in a StorageClass object on an EKS cluster
- F. Use the same file system for all containers.
- G. Create an Amazon Elastic Block Store (Amazon EBS) volume
- H. Register the volume in a StorageClass object on an EKS cluster
- I. Use the same volume for all containers.
- J. Create Amazon Elastic File System (Amazon EFS) file systems in the same Availability Zones where EKS worker nodes are placed
- K. Register the file systems in a StorageClass object on an EKS cluster
- L. Create an AWS Lambda function to synchronize the data between file systems.

**Answer: B**

**Explanation:**

Amazon EFS is a fully managed, elastic, and scalable file system that can be shared between multiple containers. It provides high availability and fault tolerance by replicating data across multiple Availability Zones. Amazon EFS is compatible with Amazon EKS and AWS Fargate, and can be registered in a StorageClass object on an EKS cluster. Amazon EBS volumes are not supported by AWS Fargate, and cannot be shared between multiple containers without using EBS Multi-Attach, which has limitations and performance implications. EBS Multi-Attach also requires the volumes to be in the same Availability Zone as the worker nodes, which reduces availability and fault tolerance. Synchronizing data between multiple EFS file systems using AWS Lambda is unnecessary, complex, and prone to errors. References:

- ? Amazon EFS Storage Classes
- ? Amazon EKS Storage Classes
- ? Amazon EBS Multi-Attach

**NEW QUESTION 30**

- (Topic 4)

A company is designing a new web service that will run on Amazon EC2 instances behind an Elastic Load Balancing (ELB) load balancer. However, many of the web service clients can only reach IP addresses authorized on their firewalls. What should a solutions architect recommend to meet the clients' needs?

- A. A Network Load Balancer with an associated Elastic IP address.
- B. An Application Load Balancer with an associated Elastic IP address.
- C. An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address.
- D. An EC2 instance with a public IP address running as a proxy in front of the load balancer.

**Answer: A**

**Explanation:**

A Network Load Balancer can be assigned one Elastic IP address for each Availability Zone it uses<sup>1</sup>. This allows the clients to reach the load balancer using a static IP address that can be authorized on their firewalls. An Application Load Balancer cannot be assigned an Elastic IP address<sup>2</sup>. An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address would not work because the load balancer would still use its own IP address as the source of the forwarded requests to the web service. An EC2 instance with a public IP address running as a proxy in front of the load balancer would add unnecessary complexity and cost, and would not provide the same scalability and availability as a Network Load Balancer. References: 1: Network Load Balancers - Elastic Load Balancing<sup>3</sup>, IP address type section<sup>2</sup>: How to assign Elastic IP to Application Load Balancer in AWS<sup>4</sup>, answer section.

**NEW QUESTION 34**

- (Topic 4)

A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture. The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service. Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

- A. Create a new organization in AWS Organizations with all features turned on
- B. Create the new AWS accounts in the organization.
- C. Set up an Amazon Cognito identity pool
- D. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication.
- E. Configure a service control policy (SCP) to manage the AWS account
- F. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service.
- G. Create a new organization in AWS Organization
- H. Configure the organization's authentication mechanism to use AWS Directory Service directly.
- I. Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization
- J. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

**Answer: AE**

**Explanation:**

AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts<sup>1</sup>. By creating a new organization in AWS Organizations with all features turned on, the solution can consolidate and manage the new AWS accounts for different business units.

AWS IAM Identity Center (formerly known as AWS Single Sign-On) is a service that provides single sign-on access for all of your AWS accounts and cloud applications. It connects with Microsoft Active Directory through AWS Directory Service to allow users in that directory to sign in to a personalized AWS access portal using their existing Active Directory user names and passwords. From the AWS access portal, users have access to all the AWS accounts and cloud applications that they have permissions for<sup>2</sup>. By setting up IAM Identity Center in the organization and integrating it with the company's corporate directory

service, the solution can authenticate access to these AWS accounts using a centralized corporate directory service.

\* B. Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication. This solution will not meet the requirement of authenticating access to these AWS accounts by using a centralized corporate directory service, as Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications, not for corporate directory services.

\* C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service. This solution will not work, as SCPs are used to restrict the actions that can be performed by the accounts in an organization, not to manage the accounts themselves. Also, IAM Identity Center cannot be added to AWS Directory Service, as it is a separate service that connects with Microsoft Active Directory through AWS Directory Service.

\* D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly. This solution will not work, as AWS Organizations does not have an authentication mechanism that can use AWS Directory Service directly. AWS Organizations relies on IAM Identity Center to provide single sign-on access for the accounts in an organization.

Reference URL: [https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_integrate\\_services.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_integrate_services.html)

### NEW QUESTION 37

- (Topic 4)

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 for the internet traffic
- B. Specify the EC2 instances as the targets.
- C. Configure a Gateway Load Balancer for the internet traffic
- D. Specify the EC2 instances as the targets.
- E. Configure a Network Load Balancer with the required protocol and ports for the internet traffic
- F. Specify the EC2 instances as the targets.
- G. Launch an identical set of game servers on EC2 instances in separate AWS Region
- H. Route internet traffic to both sets of EC2 instances.

**Answer: C**

#### Explanation:

The most cost-effective solution for the online video game company is to configure a Network Load Balancer with the required protocol and ports for the internet traffic and specify the EC2 instances as the targets. This solution will enable the company to handle millions of UDP requests per second with ultra-low latency and high performance. A Network Load Balancer is a type of Elastic Load Balancing that operates at the connection level (Layer 4) and routes traffic to targets (EC2 instances, microservices, or containers) within Amazon VPC based on IP protocol data. A Network Load Balancer is ideal for load balancing of both TCP and UDP traffic, as it is capable of handling millions of requests per second while maintaining high throughput at ultra-low latency. A Network Load Balancer also preserves the source IP address of the clients to the back-end applications, which can be useful for logging or security purposes.

### NEW QUESTION 38

- (Topic 4)

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:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Condition": {
        "NotIpAddress": {
          "aws:SourceIp": [
            "192.0.2.0/24",
            "203.0.113.0/24"
          ]
        }
      }
    }
  ],
  "Resource": ["*"]
}
```

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

#### NEW QUESTION 40

- (Topic 4)

A recent analysis of a company's IT expenses highlights the need to reduce backup costs. The company's chief information officer wants to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. The company must preserve the existing investment in the on-premises backup applications and workflows.

What should a solutions architect recommend?

- A. Set up AWS Storage Gateway to connect with the backup applications using the NFS interface.
- B. Set up an Amazon EFS file system that connects with the backup applications using the NFS interface.
- C. Set up an Amazon EFS file system that connects with the backup applications using the iSCSI interface.
- D. Set up AWS Storage Gateway to connect with the backup applications using the iSCSI- virtual tape library (VTL) interface.

**Answer:** D

#### Explanation:

it allows the company to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. By setting up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface, the company can store backup data on virtual tapes in S3 or Glacier. This preserves the existing investment in the on-premises backup applications and workflows while leveraging AWS storage services.

References:

- ? AWS Storage Gateway
- ? Tape Gateway

#### NEW QUESTION 43

- (Topic 4)

A company runs an SMB file server in its data center. The file server stores large files that the company frequently accesses for up to 7 days after the file creation date. After 7 days, the company needs to be able to access the files with a maximum retrieval time of 24 hours.

Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to increase the company's storage space.
- C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
- D. Create an Amazon FSx File Gateway to increase the company's storage space.
- E. Create an Amazon S3 Lifecycle policy to transition the data after 7 days.
- F. Configure access to Amazon S3 for each use.
- G. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

**Answer:** B

#### Explanation:

Amazon S3 File Gateway is a service that provides a file-based interface to Amazon S3, which appears as a network file share. It enables you to store and retrieve Amazon S3 objects through standard file storage protocols such as SMB. S3 File Gateway can also cache frequently accessed data locally for low-latency access. S3 Lifecycle policy is a feature that allows you to define rules that automate the management of your objects throughout their lifecycle. You can use S3 Lifecycle policy to transition objects to different storage classes based on their age and access patterns. S3 Glacier Deep Archive is a storage class that offers the lowest cost for long-term data archiving, with a retrieval time of 12 hours or 48 hours. This solution will meet the requirements, as it allows the company to store large files in S3 with SMB file access, and to move the files to S3 Glacier Deep Archive after 7 days for cost savings and compliance.

References:

- ? 1 provides an overview of Amazon S3 File Gateway and its benefits.
- ? 2 explains how to use S3 Lifecycle policy to manage object storage lifecycle.
- ? 3 describes the features and use cases of S3 Glacier Deep Archive storage class.

#### NEW QUESTION 44

- (Topic 4)

A company wants to migrate an on-premises legacy application to AWS. The application ingests customer order files from an on-premises enterprise resource planning (ERP) system. The application then uploads the files to an SFTP server. The application uses a scheduled job that checks for order files every hour. The company already has an AWS account that has connectivity to the on-premises network. The new application on AWS must support integration with the existing ERP system. The new application must be secure and resilient and must use the SFTP protocol to process orders from the ERP system immediately. Which solution will meet these requirements?

- A. Create an AWS Transfer Family SFTP internet-facing server in two Availability Zones.
- B. Use Amazon S3 storage.
- C. Create an AWS Lambda function to process order files.
- D. Use S3 Event Notifications to send s3:ObjectCreated:\* events to the Lambda function.
- E. Create an AWS Transfer Family SFTP internet-facing server in one Availability Zone.
- F. Use Amazon Elastic File System (Amazon EFS) storage.
- G. Create an AWS Lambda function to process order files.
- H. Use a Transfer Family managed workflow to invoke the Lambda function.
- I. Create an AWS Transfer Family SFTP internal server in two Availability Zones.
- J. Use Amazon Elastic File System (Amazon EFS) storage.
- K. Create an AWS Step Functions state machine to process order files.
- L. Use Amazon EventBridge Scheduler to invoke the state machine to periodically check Amazon EFS for order files.
- M. Create an AWS Transfer Family SFTP internal server in two Availability Zones.
- N. Use Amazon S3 storage.
- O. Create an AWS Lambda function to process order files.

P. Use a Transfer Family managed workflow to invoke the Lambda function.

**Answer:** D

**Explanation:**

This solution meets the requirements because it uses the following components and features:

? AWS Transfer Family SFTP internal server: This allows the application to securely transfer order files from the on-premises ERP system to AWS using the SFTP protocol over a private connection. The internal server is deployed in two Availability Zones for high availability and fault tolerance.

? Amazon S3 storage: This provides scalable, durable, and cost-effective object storage for the order files. Amazon S3 also supports encryption at rest and in transit, as well as lifecycle policies and versioning for data protection and compliance.

? AWS Lambda function: This enables the application to process the order files in a serverless manner, without provisioning or managing servers. The Lambda function can perform any custom logic or transformation on the order files, such as validating, parsing, or enriching the data.

? Transfer Family managed workflow: This simplifies the orchestration of the file processing tasks by triggering the Lambda function as soon as a file is uploaded to the SFTP server. The managed workflow also provides error handling, retry policies, and logging capabilities.

**NEW QUESTION 49**

- (Topic 4)

A company is deploying an application that processes large quantities of data in parallel. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to prevent groups of nodes from sharing the same underlying hardware.

Which networking solution meets these requirements?

- A. Run the EC2 instances in a spread placement group.
- B. Group the EC2 instances in separate accounts.
- C. Configure the EC2 instances with dedicated tenancy.
- D. Configure the EC2 instances with shared tenancy.

**Answer:** A

**Explanation:**

it allows the company to deploy an application that processes large quantities of data in parallel and prevent groups of nodes from sharing the same underlying hardware. By running the EC2 instances in a spread placement group, the company can launch a small number of instances across distinct underlying hardware to reduce correlated failures. A spread placement group ensures that each instance is isolated from each other at the rack level. References:

? Placement Groups

? Spread Placement Groups

**NEW QUESTION 50**

- (Topic 4)

To meet security requirements, a company needs to encrypt all of its application data in transit while communicating with an Amazon RDS MySQL DB instance. A recent security audit revealed that encryption at rest is enabled using AWS Key Management Service (AWS KMS), but data in transit is not enabled.

What should a solutions architect do to satisfy the security requirements?

- A. Enable IAM database authentication on the database.
- B. Provide self-signed certificate
- C. Use the certificates in all connections to the RDS instance.
- D. Take a snapshot of the RDS instance
- E. Restore the snapshot to a new instance with encryption enabled.
- F. Download AWS-provided root certificate
- G. Provide the certificates in all connections to the RDS instance.

**Answer:** D

**Explanation:**

To satisfy the security requirements, the solutions architect should download AWS-provided root certificates and provide the certificates in all connections to the RDS instance. This will enable SSL/TLS encryption for data in transit between the application and the RDS instance. SSL/TLS encryption provides a layer of security by encrypting data that moves between the client and the server. Amazon RDS creates an SSL certificate and installs the certificate on the DB instance when the instance is provisioned. The application can use the AWS-provided root certificates to verify the identity of the DB instance and establish a secure connection<sup>1</sup>.

The other options are not correct because they do not enable encryption for data in transit or are not relevant for the use case. Enabling IAM database authentication on the database is not correct because this option only provides a method of authentication, not encryption. IAM database authentication allows users to use AWS Identity and Access Management (IAM) users and roles to access a database, instead of using a database user name and password<sup>2</sup>.

Providing self-signed certificates is not correct because this option is not secure or reliable. Self-signed certificates are certificates that are signed by the same entity that issued them, instead of by a trusted certificate authority (CA). Self-signed certificates can be easily forged or compromised, and are not recognized by most browsers and applications<sup>3</sup>. Taking a snapshot of the RDS instance and restoring it to a new instance with encryption enabled is not correct because this option only enables encryption at rest, not encryption in transit. Encryption at rest protects data that is stored on disk, but does not protect data that is moving between the client and the server<sup>4</sup>.

References:

? Using SSL/TLS to encrypt a connection to a DB instance - Amazon Relational Database Service

? IAM database authentication for MySQL and PostgreSQL - Amazon Relational Database Service

? What are self-signed certificates?

? Encrypting Amazon RDS resources - Amazon Relational Database Service

**NEW QUESTION 55**

- (Topic 4)

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 and enable default Amazon Elastic Block Store (Amazon EBS) volume encryption for the account.

**Answer: B**

**Explanation:**

This option is the most secure and simple way to encrypt the secrets that are stored in Amazon EKS. AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys that can be used to encrypt your data. Amazon EKS KMS secrets encryption is a feature that enables you to use a KMS key to encrypt the secrets that are stored in the Kubernetes etcd key-value store. This provides an additional layer of protection for your sensitive data, such as passwords, tokens, and keys. You can create a new KMS key or use an existing one, and then enable the Amazon EKS KMS secrets encryption on the Amazon EKS cluster. You can also use IAM policies to control who can access or use the KMS key.

Option A is not correct because using AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS is not necessary or efficient. AWS Secrets Manager is a service that helps you securely store, retrieve, and rotate your secrets, such as database credentials, API keys, and passwords. You can use it to manage secrets that are used by your applications or services outside of Amazon EKS, but it is not designed to encrypt the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using AWS Secrets Manager would incur additional costs and complexity, and it would not leverage the native Kubernetes secrets management capabilities.

Option C is not correct because using the Amazon EBS Container Storage Interface (CSI) driver as an add-on does not encrypt the secrets that are stored in Amazon EKS. The Amazon EBS CSI driver is a plugin that allows you to use Amazon EBS volumes as persistent storage for your Kubernetes pods. It is useful for providing durable and scalable storage for your applications, but it does not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the Amazon EBS CSI driver would require additional configuration and resources, and it would not provide the same level of security as using a KMS key.

Option D is not correct because creating a new AWS KMS key with the alias aws/ebs and enabling default Amazon EBS volume encryption for the account does not encrypt the secrets that are stored in Amazon EKS. The alias aws/ebs is a reserved alias that is used by AWS to create a default KMS key for your account. This key is used to encrypt the Amazon EBS volumes that are created in your account, unless you specify a different KMS key. Enabling default Amazon EBS volume encryption for the account is a setting that ensures that all new Amazon EBS volumes are encrypted by default. However, these features do not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the default KMS key or the default encryption setting would not provide the same level of control and security as using a custom KMS key and enabling the Amazon EKS KMS secrets encryption feature. References:

- ? Encrypting secrets used in Amazon EKS
- ? What Is AWS Key Management Service?
- ? What Is AWS Secrets Manager?
- ? Amazon EBS CSI driver
- ? Encryption at rest

**NEW QUESTION 59**

- (Topic 4)

A company stores multiple Amazon Machine Images (AMIs) in an AWS account to launch its Amazon EC2 instances. The AMIs contain critical data and configurations that are necessary for the company's operations. The company wants to implement a solution that will recover accidentally deleted AMIs quickly and efficiently.

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

- A. Create Amazon Elastic Block Store (Amazon EBS) snapshots of the AMI
- B. Store the snapshots in a separate AWS account.
- C. Copy all AMIs to another AWS account periodically.
- D. Create a retention rule in Recycle Bin.
- E. Upload the AMIs to an Amazon S3 bucket that has Cross-Region Replication.

**Answer: C**

**Explanation:**

Recycle Bin is a data recovery feature that enables you to restore accidentally deleted Amazon EBS snapshots and EBS-backed AMIs. When using Recycle Bin, if your resources are deleted, they are retained in the Recycle Bin for a time period that you specify before being permanently deleted. You can restore a resource from the Recycle Bin at any time before its retention period expires. This solution has the least operational overhead, as you do not need to create, copy, or upload any additional resources. You can also manage tags and permissions for AMIs in the Recycle Bin. AMIs in the Recycle Bin do not incur any additional charges.

References:

- ? Recover AMIs from the Recycle Bin
- ? Recover an accidentally deleted Linux AMI

**NEW QUESTION 64**

- (Topic 4)

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 EPS file system in the primary account Use AWS DataSync to copy the contents of the original EPS file system to the new EPS 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 syste
- D. Use the primary account's Lambda function to invoke the secondary account's Lambda function
- E. Move the contents of the file system to a Lambda Layer's Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

**Answer: B**

**Explanation:**

This option is the most cost-effective and scalable way to allow the Lambda function in the primary account to access the EFS file system in the secondary account. VPC peering enables private connectivity between two VPCs without requiring gateways, VPN connections, or dedicated network connections. The Lambda function can use the VPC peering connection to mount the EFS file system as a local file system and access the files as needed. The solution does not

incur additional data transfer or storage costs, and it leverages the existing EFS file system without duplicating or moving the data.

Option A is not cost-effective because it requires creating a new EFS file system and using AWS DataSync to copy the data from the original EFS file system. This would incur additional storage and data transfer costs, and it would not provide real-time access to the files.

Option C is not scalable because it requires creating a second Lambda function in the secondary account and configuring cross-account permissions to invoke it from the primary account. This would add complexity and latency to the solution, and it would increase the Lambda invocation costs.

Option D is not feasible because Lambda layers are not designed to store large amounts of data or provide file system access. Lambda layers are used to share common code or libraries across multiple Lambda functions. Moving the contents of the EFS file system to a Lambda layer would exceed the size limit of 250 MB for a layer, and it would not allow the Lambda function to read or write files to the layer. References:

? What Is VPC Peering?

? Using Amazon EFS file systems with AWS Lambda

? What Are Lambda Layers?

#### NEW QUESTION 66

- (Topic 4)

A company is developing a new machine learning (ML) model solution on AWS. The models are developed as independent microservices that fetch approximately 1 GB of model data from Amazon S3 at startup and load the data into memory. Users access the models through an asynchronous API. Users can send a request or a batch of requests and specify where the results should be sent.

The company provides models to hundreds of users. The usage patterns for the models are irregular. Some models could be unused for days or weeks. Other models could receive batches of thousands of requests at a time.

Which design should a solutions architect recommend to meet these requirements?

- A. Direct the requests from the API to a Network Load Balancer (NLB). Deploy the models as AWS Lambda functions that are invoked by the NLB.
- B. Direct the requests from the API to an Application Load Balancer (ALB). Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from an Amazon Simple Queue Service (Amazon SQS) queue.
- C. Use AWS App Mesh to scale the instances of the ECS cluster based on the SQS queue size.
- D. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue.
- E. Deploy the models as AWS Lambda functions that are invoked by SQS event.
- F. Use AWS Auto Scaling to increase the number of vCPUs for the Lambda functions based on the SQS queue size.
- G. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue.
- H. Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from the queue.
- I. Enable AWS Auto Scaling on Amazon ECS for both the cluster and copies of the service based on the queue size.

**Answer: D**

#### Explanation:

This answer is correct because it meets the requirements of running the ML models as independent microservices that can handle irregular and unpredictable usage patterns. By directing the requests from the API into an Amazon SQS queue, the company can decouple the request processing from the model execution, and ensure that no requests are lost due to spikes in demand. By deploying the models as Amazon ECS services that read from the queue, the company can leverage containers to isolate and package each model as a microservice, and fetch the model data from S3 at startup. By enabling AWS Auto Scaling on Amazon ECS for both the cluster and copies of the service based on the queue size, the company can automatically scale up or down the number of EC2 instances in the cluster and the number of tasks in each service to match the demand and optimize performance.

References:

? <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html>

? <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>

? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/autoscaling-ecs.html>

#### NEW QUESTION 67

- (Topic 4)

A company runs a web application that is deployed on Amazon EC2 instances in the private subnet of a VPC. An Application Load Balancer (ALB) that extends across the public subnets directs web traffic to the EC2 instances. The company wants to implement new security measures to restrict inbound traffic from the ALB to the EC2 instances while preventing access from any other source inside or outside the private subnet of the EC2 instances.

Which solution will meet these requirements?

- A. Configure a route in a route table to direct traffic from the internet to the private IP addresses of the EC2 instances.
- B. Configure the security group for the EC2 instances to only allow traffic that comes from the security group for the ALB.
- C. Move the EC2 instances into the public subnet.
- D. Give the EC2 instances a set of Elastic IP addresses.
- E. Configure the security group for the ALB to allow any TCP traffic on any port.

**Answer: B**

#### Explanation:

To restrict inbound traffic from the ALB to the EC2 instances, the security group for the EC2 instances should only allow traffic that comes from the security group for the ALB. This way, the EC2 instances can only receive requests from the ALB and not from any other source inside or outside the private subnet.

References:

? Security Groups for Your Application Load Balancers

? Security Groups for Your VPC

#### NEW QUESTION 70

- (Topic 4)

A company has deployed a multiplayer game for mobile devices. The game requires live

location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.

The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.

What should a solutions architect do to improve the performance of the data tier?

- A. Take a snapshot of the existing DB instance.
- B. Restore the snapshot with Multi-AZ enabled.
- C. Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
- D. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance.

- E. Modify the game to use DAX.
- F. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance.
- G. Modify the game to use Redis.

**Answer: D**

**Explanation:**

The solution that will improve the performance of the data tier is to deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance and modify the game to use Redis. This solution will enable the game to store and retrieve the location data of the players in a fast and scalable way, as Redis is an in-memory data store that supports geospatial data types and commands. By using ElastiCache for Redis, the game can reduce the load on the RDS for PostgreSQL DB instance, which is not optimized for high-frequency updates and queries of location data. ElastiCache for Redis also supports replication, sharding, and auto scaling to handle the increasing user base of the game. The other solutions are not as effective as the first one because they either do not improve the performance, do not support geospatial data, or do not leverage caching. Taking a snapshot of the existing DB instance and restoring it with Multi-AZ enabled will not improve the performance of the data tier, as it only provides high availability and durability, but not scalability or low latency. Migrating from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards will not improve the performance of the data tier, as OpenSearch Service is mainly designed for full-text search and analytics, not for real-time location tracking. OpenSearch Service also does not support geospatial data types and commands natively, unlike Redis. Deploying Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance and modifying the game to use DAX will not improve the performance of the data tier, as DAX is only compatible with DynamoDB, not with RDS for PostgreSQL. DAX also does not support geospatial data types and commands.

References:

- ? Amazon ElastiCache for Redis
- ? Geospatial Data Support - Amazon ElastiCache for Redis
- ? Amazon RDS for PostgreSQL
- ? Amazon OpenSearch Service
- ? Amazon DynamoDB Accelerator (DAX)

**NEW QUESTION 73**

- (Topic 4)

A company runs demonstration environments for its customers on Amazon EC2 instances. Each environment is isolated in its own VPC. The company's operations team needs to be notified when RDP or SSH access to an environment has been established.

- A. Configure Amazon CloudWatch Application Insights to create AWS Systems Manager OpsItems when RDP or SSH access is detected.
- B. Configure the EC2 instances with an IAM instance profile that has an IAM role with the AmazonSSMManagedInstanceCore policy attached.
- C. Publish VPC flow logs to Amazon CloudWatch Log
- D. Create required metric filter
- E. Create an Amazon CloudWatch metric alarm with a notification action for when the alarm is in the ALARM state.
- F. Configure an Amazon EventBridge rule to listen for events of type EC2 Instance State- change Notification
- G. Configure an Amazon Simple Notification Service (Amazon SNS) topic as a target
- H. Subscribe the operations team to the topic.

**Answer: C**

**Explanation:**

<https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/>

**NEW QUESTION 77**

- (Topic 4)

A company is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company's application. A solutions architect wants to implement a solution that is highly available, fault tolerant, and automatically scalable. What should the solutions architect recommend?

- A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
- B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in different Availability Zones.
- C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
- D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

**Answer: C**

**Explanation:**

If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway's Availability Zone is down, resources in the other Availability Zones lose internet access. To create an Availability Zone-independent architecture, create a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone. <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html#nat-gateway-basics>

**NEW QUESTION 82**

- (Topic 4)

A company has a web application hosted over 10 Amazon EC2 instances with traffic directed by Amazon Route 53. The company occasionally experiences a timeout error when attempting to browse the application. The networking team finds that some DNS queries return IP addresses of unhealthy instances, resulting in the timeout error.

What should a solutions architect implement to overcome these timeout errors?

- A. Create a Route 53 simple routing policy record for each EC2 instance
- B. Associate a health check with each record.
- C. Create a Route 53 failover routing policy record for each EC2 instance
- D. Associate a health check with each record.
- E. Create an Amazon CloudFront distribution with EC2 instances as its origin
- F. Associate a health check with the EC2 instances.
- G. Create an Application Load Balancer (ALB) with a health check in front of the EC2 instance
- H. Route to the ALB from Route 53.

**Answer: D**

**Explanation:**

An Application Load Balancer (ALB) allows you to distribute incoming traffic across multiple backend instances, and can automatically route traffic to healthy instances while removing traffic from unhealthy instances. By using an ALB in front of the EC2 instances and routing traffic to it from Route 53, the load balancer can perform health checks on the instances and only route traffic to healthy instances, which should help to reduce or eliminate timeout errors caused by unhealthy instances.

**NEW QUESTION 86**

- (Topic 4)

A company runs a three-tier application in two AWS Regions. The web tier, the application tier, and the database tier run on Amazon EC2 instances. The company uses Amazon RDS for Microsoft SQL Server Enterprise for the database tier. The database tier is experiencing high load when weekly and monthly reports are run. The company wants to reduce the load on the database tier.

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

- A. Create read replica
- B. Configure the reports to use the new read replicas.
- C. Convert the RDS database to Amazon DynamoDB\_ Configure the reports to use DynamoDB
- D. Modify the existing RDS DB instances by selecting a larger instance size.
- E. Modify the existing ROS DB instances and put the instances into an Auto Scaling group.

**Answer:** A

**Explanation:**

it allows the company to create read replicas of its RDS database and reduce the load on the database tier. By creating read replicas, the company can offload read traffic from the primary database instance to one or more replicas. By configuring the reports to use the new read replicas, the company can improve performance and availability of its database tier. References:

? Working with Read Replicas

? Read Replicas for Amazon RDS for SQL Server

**NEW QUESTION 88**

- (Topic 4)

A company hosts a data lake on Amazon S3. The data lake ingests data in Apache Parquet format from various data sources. The company uses multiple transformation steps to prepare the ingested data. The steps include filtering of anomalies, normalizing of data to standard date and time values, and generation of aggregates for analyses.

The company must store the transformed data in S3 buckets that data analysts access. The company needs a prebuilt solution for data transformation that does not require code. The solution must provide data lineage and data profiling. The company needs to share the data transformation steps with employees throughout the company.

Which solution will meet these requirements?

- A. Configure an AWS Glue Studio visual canvas to transform the dat
- B. Share the transformation steps with employees by using AWS Glue jobs.
- C. Configure Amazon EMR Serverless to transform the dat
- D. Share the transformation steps with employees by using EMR Serverless jobs.
- E. Configure AWS Glue DataBrew to transform the dat
- F. Share the transformation steps with employees by using DataBrew recipes.
- G. Create Amazon Athena tables for the dat
- H. Write Athena SQL queries to transform the dat
- I. Share the Athena SQL queries with employees.

**Answer:** C

**Explanation:**

The most suitable solution for the company's requirements is to configure AWS Glue DataBrew to transform the data and share the transformation steps with employees by using DataBrew recipes. This solution will provide a prebuilt solution for data transformation that does not require code, and will also provide data lineage and data profiling. The company can easily share the data transformation steps with employees throughout the company by using DataBrew recipes.

AWS Glue DataBrew is a visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data for analytics or machine learning by up to 80% faster. Users can upload their data from various sources, such as Amazon S3, Amazon RDS, Amazon Redshift, Amazon Aurora, or Glue Data Catalog, and use a point- and-click interface to apply over 250 built-in transformations. Users can also preview the results of each transformation step and see how it affects the quality and distribution of the data<sup>1</sup>.

A DataBrew recipe is a reusable set of transformation steps that can be applied to one or more datasets. Users can create recipes from scratch or use existing ones from the DataBrew recipe library. Users can also export, import, or share recipes with other users or groups within their AWS account or organization<sup>2</sup>.

DataBrew also provides data lineage and data profiling features that help users understand and improve their data quality. Data lineage shows the source and destination of each dataset and how it is transformed by each recipe step. Data profiling shows various statistics and metrics about each dataset, such as column

**NEW QUESTION 89**

- (Topic 4)

A company runs its applications on Amazon EC2 instances. The company performs periodic financial assessments of its AWS costs. The company recently identified unusual spending.

The company needs a solution to prevent unusual spending. The solution must monitor costs and notify responsible stakeholders in the event of unusual spending. Which solution will meet these requirements?

- A. Use an AWS Budgets template to create a zero spend budget
- B. Create an AWS Cost Anomaly Detection monitor in the AWS Billing and Cost Management console.
- C. Create AWS Pricing Calculator estimates for the current running workload pricing details\_
- D. Use Amazon CloudWatch to monitor costs and to identify unusual spending

**Answer:** B

**Explanation:**

it allows the company to monitor costs and notify responsible stakeholders in the event of unusual spending. By creating an AWS Cost Anomaly Detection monitor in the AWS Billing and Cost Management console, the company can use a machine learning service that automatically detects and alerts on anomalous spend. By

configuring alert thresholds, notification preferences, and root cause analysis, the company can prevent unusual spending and identify its source. References:  
? AWS Cost Anomaly Detection  
? Creating a Cost Anomaly Monitor

#### NEW QUESTION 91

- (Topic 4)

A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance. New company management wants to ensure the application is highly available. What should a solutions architect do to meet this requirement?

- A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer
- B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region.
- C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application.
- D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer

**Answer:** A

#### Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability-zone.html>

#### NEW QUESTION 94

- (Topic 4)

A company has a three-tier application for image sharing. The application uses an Amazon EC2 instance for the front-end layer, another EC2 instance for the application layer, and a third EC2 instance for a MySQL database. A solutions architect must design a scalable and highly available solution that requires the least amount of change to the application.

Which solution meets these requirements?

- A. Use Amazon S3 to host the front-end layer
- B. Use AWS Lambda functions for the application layer
- C. Move the database to an Amazon DynamoDB table
- D. Use Amazon S3 to store and serve users' images.
- E. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application layer
- F. Move the database to an Amazon RDS DB instance with multiple read replicas to serve users' images.
- G. Use Amazon S3 to host the front-end layer
- H. Use a fleet of EC2 instances in an Auto Scaling group for the application layer
- I. Move the database to a memory optimized instance type to store and serve users' images.
- J. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application layer
- K. Move the database to an Amazon RDS Multi-AZ DB instance
- L. Use Amazon S3 to store and serve users' images.

**Answer:** D

#### Explanation:

for "Highly available": Multi-AZ & for "least amount of changes to the application": Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring

#### NEW QUESTION 98

- (Topic 4)

A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will continue to grow and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL.

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

- A. Send activity data to an Amazon Kinesis data stream
- B. Configure the stream to deliver the data to an Amazon S3 bucket.
- C. Send activity data to an Amazon Kinesis Data Firehose delivery stream
- D. Configure the stream to deliver the data to an Amazon Redshift cluster.
- E. Place activity data in an Amazon S3 bucket
- F. Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
- G. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zones
- H. Configure the service to forward data to an Amazon RDS Multi-AZ database.

**Answer:** B

#### Explanation:

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This allows you to use your data to gain new insights for your business and customers. The first step to create a data warehouse is to launch a set of nodes, called an Amazon Redshift cluster. After you provision your cluster, you can upload your data set and then perform data analysis queries. Regardless of the size of the data set, Amazon Redshift offers fast query performance using the same SQL-based tools and business intelligence applications that you use today.

#### NEW QUESTION 102

- (Topic 4)

A company uses Amazon S3 as its data lake. The company has a new partner that must use SFTP to upload data files. A solutions architect needs to implement a highly available SFTP solution that minimizes operational overhead.

Which solution will meet these requirements?

- A. Use AWS Transfer Family to configure an SFTP-enabled server with a publicly accessible endpoint. Choose the S3 data lake as the destination.

- B. Use Amazon S3 File Gateway as an SFTP server Expose the S3 File Gateway endpoint URL to the new partner Share the S3 File Gateway endpoint with the new partner
- C. Launch an Amazon EC2 instance in a private subnet in a VP
- D. Instruct the new partner to upload files to the EC2 instance by using a VP
- E. Run a cron job script on the EC2 instance to upload files to the S3 data lake
- F. Launch Amazon EC2 instances in a private subnet in a VP
- G. Place a Network Load Balancer (NLB) in front of the EC2 instance
- H. Create an SFTP listener port for the NLB Share the NLB hostname with the new partner Run a cron job script on the EC2 instances to upload files to the S3 data lake.

**Answer: A**

**Explanation:**

This option is the most cost-effective and simple way to enable SFTP access to the S3 data lake. AWS Transfer Family is a fully managed service that supports secure file transfers over SFTP, FTPS, and FTP protocols. You can create an SFTP-enabled server with a public endpoint and associate it with your S3 bucket. You can also use AWS Identity and Access Management (IAM) roles and policies to control access to your S3 data lake. The service scales automatically to handle any volume of file transfers and provides high availability and durability. You do not need to provision, manage, or patch any servers or load balancers. Option B is not correct because Amazon S3 File Gateway is not an SFTP server. It is a hybrid cloud storage service that provides a local file system interface to S3. You can use it to store and retrieve files as objects in S3 using standard file protocols such as NFS and SMB. However, it does not support SFTP protocol, and it requires deploying a file gateway appliance on-premises or on EC2.

Option C is not cost-effective or scalable because it requires launching and managing an EC2 instance in a private subnet and setting up a VPN connection for the new partner. This would incur additional costs for the EC2 instance, the VPN connection, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instance to upload files to the S3 data lake, which is not efficient or reliable.

Option D is not cost-effective or scalable because it requires launching and managing multiple EC2 instances in a private subnet and placing a NLB in front of them. This would incur additional costs for the EC2 instances, the NLB, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instances to upload files to the S3 data lake, which is not efficient or reliable. References:

- ? What Is AWS Transfer Family?
- ? What Is Amazon S3 File Gateway?
- ? What Is Amazon EC2?
- ? [What Is Amazon Virtual Private Cloud?]
- ? [What Is a Network Load Balancer?]

**NEW QUESTION 104**

- (Topic 4)

A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers a previous employee did not provide details about the resources invent<sup>^</sup>. 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 (formerly called AWS Perspective) is a tool that visualizes AWS Cloud workloads. It maintains an inventory of the AWS resources across your accounts and Regions, mapping relationships between them, and displaying them in a web UI. It also allows you to query AWS Cost and Usage Reports, search for resources, save and export architecture diagrams, and more<sup>1</sup>. By using Workload Discovery on AWS, the solution can build and map the relationship details of the various workloads across all accounts with the least operational effort.

\* A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report. This solution will not meet the requirement of building and mapping the relationship details of the various workloads across all accounts, as AWS Systems Manager Inventory is a feature that collects metadata from your managed instances and stores it in a central Amazon S3 bucket. It does not provide a map view or architecture diagrams of the workloads<sup>2</sup>.

\* B. Use AWS Step Functions to collect workload details Build architecture diagrams of the workloads manually. This solution will not meet the requirement of the least operational effort, as it involves creating and managing state machines to orchestrate the workload details collection, and building architecture diagrams manually.

\* D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships. This solution will not meet the requirement of the least operational effort, as it involves instrumenting your applications with X-Ray SDKs to collect workload details, and building architecture diagrams manually.

Reference URL: <https://aws.amazon.com/solutions/implementations/workload-discovery-on-aws/>

**NEW QUESTION 108**

- (Topic 4)

A company sends AWS CloudTrail logs from multiple AWS accounts to an Amazon S3 bucket in a centralized account. The company must keep the CloudTrail logs. The company must also be able to query the CloudTrail logs at any time Which solution will meet these requirements?

- A. Use the CloudTrail event history in the centralized account to create an Amazon Athena tabl
- B. Query the CloudTrail logs from Athena.
- C. Configure an Amazon Neptune instance to manage the CloudTrail log
- D. Query the CloudTrail logs from Neptune.
- E. Configure CloudTrail to send the logs to an Amazon DynamoDB tabl
- F. Create a dashboard in Amazon QuIckSight to query the logs in the table.
- G. use Amazon Athena to create an Athena noteboo
- H. Configure CloudTrail to send the logs to the noteboo
- I. Run queries from Athena.

**Answer: A**

**Explanation:**

it allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company

can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. References:

- ? Viewing Events with CloudTrail Event History
- ? Querying AWS CloudTrail Logs
- ? Amazon Athena

#### NEW QUESTION 113

- (Topic 4)

A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a third-party provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead. What should the solutions architect do to meet these requirements?

- A. Deploy a NAT instance in the VP
- B. Route all the internet-based traffic through the NAT instance.
- C. Deploy a NAT gateway in the public subnet
- D. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
- E. Configure an internet gateway and attach it to the VP
- F. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
- G. Configure a virtual private gateway and attach it to the VP
- H. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

**Answer:** B

#### Explanation:

To allow the MySQL database in the private subnets to access the internet without exposing it to the public, a NAT gateway is a suitable solution. A NAT gateway enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. A NAT gateway resides in the public subnets and can handle high throughput of traffic with low latency. A NAT gateway is also a managed service that does not require any operational overhead. References:

- ? NAT Gateways
- ? NAT Gateway Pricing

#### NEW QUESTION 115

- (Topic 4)

A company wants to create an application to store employee data in a hierarchical structured relationship. The company needs a minimum-latency response to high-traffic queries for the employee data and must protect any sensitive data. The company also needs to receive monthly email messages if any financial information is present in the employee data.

Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Use Amazon Redshift to store the employee data in hierarchie
- B. Unload the data to Amazon S3 every month.
- C. Use Amazon DynamoDB to store the employee data in hierarchie
- D. Export the data to Amazon S3 every month.
- E. Configure Amazon fvlacie for the AWS accoun
- F. Integrate Macie with Amazon EventBridge to send monthly events to AWS Lambda.
- G. Use Amazon Athena to analyze the employee data in Amazon S3. Integrate Athena with Amazon QuickSight to publish analysis dashboards and share the dashboards with users.
- H. Configure Amazon Macie for the AWS account Integrate Macie with Amazon EventBridge to send monthly notifications through an Amazon Simple Notification Service (Amazon SNS) subscription.

**Answer:** BE

#### Explanation:

Generally, for building a hierarchical relationship model, a graph database such as Amazon Neptune is a better choice. In some cases, however, DynamoDB is a better choice for hierarchical data modeling because of its flexibility, security, performance, and scale. <https://docs.aws.amazon.com/prescriptive-guidance/latest/dynamodb-hierarchical-data-model/introduction.html>

#### NEW QUESTION 120

- (Topic 4)

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 tabl
- B. Create a proxy application layer to intercept and process the data that each application requests.
- C. Store the data in an Amazon S3 bucke
- D. Process and transform the data by using S3 Object Lambda before returning the data to the requesting application.
- E. Process the data and store the transformed data in three separate Amazon S3 buckets so that each application has its own custom datase
- F. Point each application to its respectiveS3 bucket.
- G. Process the data and store the transformed data in three separate Amazon DynamoDB tables so that each application has its own custom datase
- H. Point each application to its respective DynamoDB table.

**Answer:** B

#### Explanation:

<https://aws.amazon.com/blogs/aws/introducing-amazon-s3-object-lambda-use-your-code-to-process-data-as-it-is-being-retrieved-from-s3/>

S3 Object Lambda is a new feature of Amazon S3 that enables customers to add their own code to process data retrieved from S3 before returning it to the application. By using S3 Object Lambda, the data can be processed and transformed in real-time, without the need to store multiple copies of the data in separate S3 buckets or DynamoDB tables.

In this case, the PII can be removed from the data by the code added to S3 Object Lambda before returning the data to the two applications that do not need to process PII. The one application that requires PII can be pointed to the original S3 bucket where the PII is still stored. Using S3 Object Lambda is the simplest and most cost-effective solution, as it eliminates the need to maintain multiple copies of the same data in different buckets or tables, which can result in additional storage costs and operational overhead.

### NEW QUESTION 123

- (Topic 4)

A company needs to store data from its healthcare application. The application's data frequently changes. A new regulation requires audit z access at all levels of the stored data.

The company hosts the application on an on-premises infrastructure that is running out of storage capacity. A solutions architect must securely migrate the existing data to AWS while satisfying the new regulation.

Which solution will meet these requirements?

- A. Use AWS DataSync to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- B. Use AWS Snowcone to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.
- C. Use Amazon S3 Transfer Acceleration to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- D. Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.

**Answer: A**

#### Explanation:

This answer is correct because it meets the requirements of securely migrating the existing data to AWS and satisfying the new regulation. AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and Amazon S3. DataSync automatically encrypts data in transit and verifies data integrity during transfer. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to Amazon S3. CloudTrail can log data events, which show the resource operations performed on or within a resource in your AWS account, such as S3 object-level API activity. By using CloudTrail to log data events, you can audit access at all levels of the stored data.

References:

? <https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>

? <https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/logging-data-events-with-cloudtrail.html>

### NEW QUESTION 124

- (Topic 4)

A company stores critical data in Amazon DynamoDB tables in the company's AWS account. An IT administrator accidentally deleted a DynamoDB table. The deletion caused a significant loss of data and disrupted the company's operations. The company wants to prevent this type of disruption in the future.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure a trail in AWS CloudTrail
- B. Create an Amazon EventBridge rule for delete action
- C. Create an AWS Lambda function to automatically restore deleted DynamoDB tables.
- D. Create a backup and restore plan for the DynamoDB table
- E. Recover the DynamoDB tables manually.
- F. Configure deletion protection on the DynamoDB tables.
- G. Enable point-in-time recovery on the DynamoDB tables.

**Answer: C**

#### Explanation:

Deletion protection is a feature of DynamoDB that prevents accidental deletion of tables. When deletion protection is enabled, you cannot delete a table unless you explicitly disable it first. This adds an extra layer of security and reduces the risk of data loss and operational disruption. Deletion protection is easy to enable and disable using the AWS Management Console, the AWS CLI, or the DynamoDB API. This solution has the least operational overhead, as you do not need to create, manage, or invoke any additional resources or services. References:

? [Using deletion protection to protect your table](#)

? [Preventing Accidental Table Deletion in DynamoDB](#)

? [Amazon DynamoDB now supports table deletion protection](#)

### NEW QUESTION 126

- (Topic 4)

A social media company wants to allow its users to upload images in an application that is hosted in the AWS Cloud. The company needs a solution that automatically resizes the images so that the images can be displayed on multiple device types. The application experiences unpredictable traffic patterns throughout the day. The company is seeking a highly available solution that maximizes scalability.

What should a solutions architect do to meet these requirements?

- A. Create a static website hosted in Amazon S3 that invokes AWS Lambda functions to resize the images and store the images in an Amazon S3 bucket.
- B. Create a static website hosted in Amazon CloudFront that invokes AWS Step Functions to resize the images and store the images in an Amazon RDS database.
- C. Create a dynamic website hosted on a web server that runs on an Amazon EC2 instance. Configure a process that runs on the EC2 instance to resize the images and store the images in an Amazon S3 bucket.
- D. Create a dynamic website hosted on an automatically scaling Amazon Elastic Container Service (Amazon ECS) cluster that creates a resize job in Amazon Simple Queue Service (Amazon SQS). Set up an image-resizing program that runs on an Amazon EC2 instance to process the resize jobs.

**Answer: A**

#### Explanation:

By using Amazon S3 and AWS Lambda together, you can create a serverless architecture that provides highly scalable and available image resizing capabilities. Here's how the solution would work: Set up an Amazon S3 bucket to store the original images uploaded by users. Configure an event trigger on the S3 bucket to invoke an AWS Lambda function whenever a new image is uploaded. The Lambda function can be designed to retrieve the uploaded image, perform the necessary resizing operations based on device requirements, and store the resized images back in the S3 bucket or a different bucket designated for resized images. Configure the Amazon S3 bucket to make the resized images publicly accessible for serving to users.

### NEW QUESTION 131

- (Topic 4)

A company is designing a new web application that will run on Amazon EC2 Instances. The application will use Amazon DynamoDB for backend data storage. The application traffic will be unpredictable. The company expects that the application read and write throughput to the database will be moderate to high. The company needs to scale in response to application traffic.

Which DynamoDB table configuration will meet these requirements MOST cost-effectively?

- A. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard table class
- B. Set DynamoDB auto scaling to a maximum defined capacity.
- C. Configure DynamoDB in on-demand mode by using the DynamoDB Standard table class.
- D. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class
- E. Set DynamoDB auto scaling to a maximum defined capacity.
- F. Configure DynamoDB in on-demand mode by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class.

**Answer: B**

#### Explanation:

The most cost-effective DynamoDB table configuration for the web application is to configure DynamoDB in on-demand mode by using the DynamoDB Standard table class. This configuration will allow the company to scale in response to application traffic and pay only for the read and write requests that the application performs on the table.

On-demand mode is a flexible billing option that can handle thousands of requests per second without capacity planning. On-demand mode automatically adjusts the table's capacity based on the incoming traffic, and charges only for the read and write requests that are actually performed. On-demand mode is suitable for applications with unpredictable or variable workloads, or applications that prefer the ease of paying for only what they use<sup>1</sup>.

The DynamoDB Standard table class is the default and recommended table class for most workloads. The DynamoDB Standard table class offers lower throughput costs than the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and is more cost-effective for tables where throughput is the dominant cost. The DynamoDB Standard table class also offers the same performance, durability, and availability as the DynamoDB Standard-IA table class<sup>2</sup>. The other options are not correct because they are either not cost-effective or not suitable for the use case. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration requires manual estimation and management of the table's capacity, which adds complexity and cost to the solution. Provisioned mode is a billing option that requires users to specify the amount of read and write capacity units for their tables, and charges for the reserved capacity regardless of usage. Provisioned mode is suitable for applications with predictable or stable workloads, or applications that require finer-grained control over their capacity settings<sup>1</sup>. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration is not cost-effective for tables with moderate to high throughput. The DynamoDB Standard-IA table class offers lower storage costs than the DynamoDB Standard table class, but higher throughput costs. The DynamoDB Standard-IA table class is optimized for tables where storage is the dominant cost, such as tables that store infrequently accessed data<sup>2</sup>. Configuring DynamoDB in on-demand mode by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class is not correct because this configuration is not cost-effective for tables with moderate to high throughput. As mentioned above, the DynamoDB Standard-IA table class has higher throughput costs than the DynamoDB Standard table class, which can offset the savings from lower storage costs.

References:

? Table classes - Amazon DynamoDB

? Read/write capacity mode - Amazon DynamoDB

### NEW QUESTION 132

- (Topic 4)

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
- D. Set the origin as the web application that accepts the appointment requests.
- E. Add an Auto Scaling group for the application that sends meeting invitation
- F. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

**Answer: D**

#### Explanation:

To resolve the issue of longer delivery times for meeting invitations, the solutions architect can recommend adding an Auto Scaling group for the application that sends meeting invitations and configuring the Auto Scaling group to scale based on the depth of the SQS queue. This will allow the application to scale up as the number of appointment requests increases, improving the performance and delivery times of the meeting invitations.

### NEW QUESTION 135

- (Topic 4)

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 NLB
- B. Increase the Cache-Control: max-age parameter.
- C. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.
- D. Add AWS Global Accelerator in front of the NLB
- E. Configure a Global Accelerator endpoint to use the correct listener ports.
- F. Add an Amazon API Gateway endpoint behind the NLB
- G. Enable API caching
- H. Override method caching for the different stages.

**Answer:** C

**Explanation:**

This answer is correct because it improves the application performance and decreases latency for the online game by using AWS Global Accelerator. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your public applications. Global Accelerator provides two global static public IPs that act as a fixed entry point to your application endpoints, such as NLBs, in different AWS Regions. Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. Global Accelerator also terminates TCP and UDP traffic at the edge locations, which reduces the number of hops and improves the network performance. By adding AWS Global Accelerator in front of the NLBs, you can achieve up to 60% improvement in latency for your online game.

References:

? <https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html>

? <https://aws.amazon.com/global-accelerator/>

**NEW QUESTION 140**

- (Topic 4)

A company has NFS servers in an on-premises data center that need to periodically back up small amounts of data to Amazon S3. Which solution meets these requirements and is MOST cost-effective?

- A. Set up AWS Glue to copy the data from the on-premises servers to Amazon S3.
- B. Set up an AWS DataSync agent on the on-premises servers, and sync the data to Amazon S3.
- C. Set up an SFTP sync using AWS Transfer for SFTP to sync data from on premises to Amazon S3.
- D. Set up an AWS Direct Connect connection between the on-premises data center and a VPC, and copy the data to Amazon S3.

**Answer:** B

**Explanation:**

AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and AWS storage services. AWS DataSync can transfer data at speeds up to 10 times faster than open-source tools by using a purpose-built network protocol and parallelizing data transfers. AWS DataSync also handles encryption, data integrity verification, and bandwidth optimization. To use AWS DataSync, users need to deploy a DataSync agent on their on-premises servers, which connects to the NFS servers and syncs the data to Amazon S3. Users can schedule periodic or one-time sync tasks and monitor the progress and status of the transfers.

The other options are not correct because they are either not cost-effective or not suitable for the use case. Setting up AWS Glue to copy the data from the on-premises servers to Amazon S3 is not cost-effective because AWS Glue is a serverless data integration service that is mainly used for extract, transform, and load (ETL) operations, not for simple data backup. Setting up an SFTP sync using AWS Transfer for SFTP to sync data from on premises to Amazon S3 is not cost-effective because AWS Transfer for SFTP is a fully managed service that provides secure file transfer using the SFTP protocol, which is more suitable for exchanging data with third parties than for backing up data. Setting up an AWS Direct Connect connection between the on-premises data center and a VPC, and copying the data to Amazon S3 is not cost-effective because AWS Direct Connect is a dedicated network connection between AWS and the on-premises location, which has high upfront costs and requires additional configuration.

References:

? [AWS DataSync](#)

? [How AWS DataSync works](#)

? [AWS DataSync FAQs](#)

**NEW QUESTION 143**

- (Topic 4)

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing and the company is concerned about a potential increase in cost.

What should a solutions architect do to reduce the cost of the website?

- A. Create an Amazon CloudFront distribution to cache static files at edge locations.
- B. Create an Amazon ElastiCache cluster. Connect the ALB to the ElastiCache cluster to serve cached files.
- C. Create an AWS WAF web ACL and associate it with the ALB.
- D. Add a rule to the web ACL to cache static files.
- E. Create a second ALB in an alternative AWS Region. Route user traffic to the closest Region to minimize data transfer costs.

**Answer:** A

**Explanation:**

Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website.

CloudFront

can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not.

References:

? <https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/>

? <https://nodeployfriday.com/posts/static-website-hosting/>

? <https://aws.amazon.com/cloudfront/>

**NEW QUESTION 144**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### AWS-Solution-Architect-Associate Practice Exam Features:

- \* AWS-Solution-Architect-Associate Questions and Answers Updated Frequently
- \* AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff
- \* AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Solution-Architect-Associate Practice Test Here](#)