



## Amazon

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

Amazon AWS Certified Solutions Architect - Associate

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

- (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 2

- (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 3

- (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 database
- B. Direct the queries to the read replica.
- C. Create a backup of the database
- D. Restore the backup to another DB instance
- 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.

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 4

- (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 5

- (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 6

- (Topic 4)

A company is moving its data and applications to AWS during a multiyear migration project. The company wants to securely access data on Amazon S3 from the company's AWS Region and from the company's on-premises location. The data must not traverse the internet. The company has established an AWS Direct Connect connection between its Region and its on-premises location.

Which solution will meet these requirements?

- A. Create gateway endpoints for Amazon S3. Use the gateway endpoints to securely access the data from the Region and the on-premises location.
- B. Create a gateway in AWS Transit Gateway to access Amazon S3 securely from the Region and the on-premises location.
- C. Create interface endpoints for Amazon S3. Use the interface endpoints to securely access the data from the Region and the on-premises location.
- D. Use an AWS Key Management Service (AWS KMS) key to access the data securely from the Region and the on-premises location.

**Answer: B**

#### Explanation:

A gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service<sup>1</sup>. Amazon S3 does not support gateway endpoints, only interface endpoints<sup>2</sup>. Therefore, option A is incorrect.

An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported service<sup>1</sup>. An interface endpoint can provide secure access to Amazon S3 from within the Region, but not from the on-premises location. Therefore, option C is incorrect.

AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data<sup>3</sup>. AWS KMS does not provide a way to access data on Amazon S3 without traversing the internet. Therefore, option D is incorrect. AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. You can create a gateway in AWS Transit Gateway to access Amazon S3 securely from both the Region and the on-premises location using AWS Direct Connect. Therefore, option B is correct.

#### NEW QUESTION 7

- (Topic 4)

A company wants to analyze and generate reports to track the usage of its mobile app. The app is popular and has a global user base. The company uses a custom report building program to analyze application usage.

The program generates multiple reports during the last week of each month. The program takes less than 10 minutes to produce each report. The company rarely uses the program to generate reports outside of the last week of each month. The company wants to generate reports in the least amount of time when the reports are requested.

Which solution will meet these requirements MOST cost-effectively?

- A. Run the program by using Amazon EC2 On-Demand Instance
- B. Create an Amazon EventBridge rule to start the EC2 instances when reports are requested
- C. Run the EC2 instances continuously during the last week of each month.
- D. Run the program in AWS Lambda
- E. Create an Amazon EventBridge rule to run a Lambda function when reports are requested.
- F. Run the program in Amazon Elastic Container Service (Amazon ECS). Schedule Amazon ECS to run the program when reports are requested.
- G. Run the program by using Amazon EC2 Spot Instance
- H. Create an Amazon EventBridge rule to start the EC2 instances when reports are requested
- I. Run the EC2 instances continuously during the last week of each month.

**Answer: B**

**Explanation:**

This solution meets the requirements most cost-effectively because it leverages the serverless and event-driven capabilities of AWS Lambda and Amazon EventBridge. AWS Lambda allows you to run code without provisioning or managing servers, and you pay only for the compute time you consume. Amazon EventBridge is a serverless event bus service that lets you connect your applications with data from various sources and routes that data to targets such as AWS Lambda. By using Amazon EventBridge, you can create a rule that triggers a Lambda function to run the program when reports are requested, and you can also schedule the rule to run during the last week of each month. This way, you can generate reports in the least amount of time and pay only for the resources you use.

References:

? AWS Lambda

? Amazon EventBridge

**NEW QUESTION 8**

- (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 9**

- (Topic 4)

A company has a multi-tier payment processing application that is based on virtual machines (VMs). The communication between the tiers occurs asynchronously through a third-party middleware solution that guarantees exactly-once delivery.

The company needs a solution that requires the least amount of infrastructure management. The solution must guarantee exactly-once delivery for application messaging

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

- A. Use AWS Lambda for the compute layers in the architecture.
- B. Use Amazon EC2 instances for the compute layers in the architecture.
- C. Use Amazon Simple Notification Service (Amazon SNS) as the messaging component between the compute layers.
- D. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues as the messaging component between the compute layers.
- E. Use containers that are based on Amazon Elastic Kubernetes Service (Amazon EKS) for the compute layers in the architecture.

**Answer: AD**

**Explanation:**

This solution meets the requirements because it requires the least amount of infrastructure management and guarantees exactly-once delivery for application messaging. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You only pay for the compute time you consume. Lambda scales automatically with the size of your workload. Amazon SQS FIFO queues are designed to ensure that messages are processed exactly once, in the exact order that they are sent. FIFO queues have high availability and deliver messages in a strict first-in, first-out order. You can use Amazon SQS to decouple and scale microservices, distributed systems, and serverless applications. References: AWS Lambda, Amazon SQS FIFO queues

**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 10**

- (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 11**

- (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 12

- (Topic 4)

A company needs to use its on-premises LDAP directory service to authenticate its users to the AWS Management Console. The directory service is not compatible with Security Assertion Markup Language (SAML).

Which solution meets these requirements?

- A. Enable AWS IAM Identity Center (AWS Single Sign-On) between AWS and the on-premises LDAP.
- B. Create an IAM policy that uses AWS credentials, and integrate the policy into LDAP.
- C. Set up a process that rotates the IAM credentials whenever LDAP credentials are updated.
- D. Develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials.

**Answer: D**

### Explanation:

The solution that meets the requirements is to develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials. This solution allows the company to use its existing LDAP directory service to authenticate its users to the AWS Management Console, without requiring SAML compatibility. The custom identity broker application or process can act as a proxy between the LDAP directory service and AWS STS, and can request temporary security credentials for the users based on their LDAP attributes and roles. The users can then use these credentials to access the AWS Management Console via a sign-in URL generated by the identity broker. This solution also enhances security by using short-lived credentials that expire after a specified duration.

The other solutions do not meet the requirements because they either require SAML compatibility or do not provide access to the AWS Management Console. Enabling AWS IAM Identity Center (AWS Single Sign-On) between AWS and the on-premises LDAP would require the LDAP directory service to support SAML 2.0, which is not the case for this scenario. Creating an IAM policy that uses AWS credentials and integrating the policy into LDAP would not provide access to the AWS Management Console, but only to the AWS APIs. Setting up a process that rotates the IAM credentials whenever LDAP credentials are updated would also not provide access to the AWS Management Console, but only to the AWS CLI. Therefore, these solutions are not suitable for the given requirements.

## NEW QUESTION 16

- (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 20

- (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 24

- (Topic 4)

A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:

- A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
- Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders

The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.

A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize utilization of the company's AWS resources. Which solution meets these requirements?

- A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
- B. Configure each Auto Scaling group's minimum capacity according to peak workload values.
- C. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
- D. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
- E. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment
- F. Configure the EC2 instances to poll their respective queue
- G. Scale the Auto Scaling groups based on notifications that the queues send.
- H. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment
- I. Configure the EC2 instances to poll their respective queue
- J. Create a metric based on a backlog per instance calculation
- K. Scale the Auto Scaling groups based on this metric.

**Answer:** D

#### Explanation:

The number of instances in your Auto Scaling group can be driven by how long it takes to process a message and the acceptable amount of latency (queue delay). The solution is to use a backlog per instance metric with the target value being the acceptable backlog per instance to maintain.

#### NEW QUESTION 29

- (Topic 4)

A company has an on-premises server that uses an Oracle database to process and store customer information. The company wants to use an AWS database service to achieve higher availability and to improve application performance. The company also wants to offload reporting from its primary database system. Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions. Point the reporting functions toward a separate DB instance from the primary DB instance.
- B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database. Create a read replica in the same zone as the primary DB instance.
- C. Direct the reporting functions to the read replica.
- D. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database. Direct the reporting functions to use the reader instance in the cluster deployment.
- E. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora database.
- F. Direct the reporting functions to the reader instances.

**Answer:** D

#### Explanation:

Amazon Aurora is a fully managed relational database that is compatible with MySQL and PostgreSQL. It provides up to five times better performance than MySQL and

up to three times better performance than PostgreSQL. It also provides high availability and durability by replicating data across multiple Availability Zones and continuously backing up data to Amazon S3<sup>1</sup>. By using Amazon RDS deployed in a Multi-AZ instance deployment

to create an Amazon Aurora database, the solution can achieve higher availability and improve application performance.

Amazon Aurora supports read replicas, which are separate instances that share the same underlying storage as the primary instance. Read replicas can be used to offload read-only queries from the primary instance and improve performance. Read replicas can also be used for reporting functions<sup>2</sup>. By directing the reporting functions to the reader instances, the solution can offload reporting from its primary database system.

\* A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions. Point the reporting functions toward a separate DB instance from the primary DB instance. This solution will not meet the requirement of using an AWS database service, as AWS DMS is a service that helps users migrate databases to AWS, not a database service itself. It also involves creating multiple DB instances in different Regions, which may increase complexity and cost.



- \* B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database. Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica. This solution will not meet the requirement of achieving higher availability, as a Single-AZ deployment does not provide failover protection in case of an Availability Zone outage. It also involves using Oracle as the database engine, which may not provide better performance than Aurora.
- \* C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database. Direct the reporting functions to use the reader instance in the cluster deployment. This solution will not meet the requirement of improving application performance, as Oracle may not provide better performance than Aurora. It also involves using a cluster deployment, which is only supported for Aurora, not for Oracle. Reference URL: <https://aws.amazon.com/rds/aurora/>

### NEW QUESTION 31

- (Topic 4)

A manufacturing company runs its report generation application on AWS. The application generates each report in about 20 minutes. The application is built as a monolith that runs on a single Amazon EC2 instance. The application requires frequent updates to its tightly coupled modules. The application becomes complex to maintain as the company adds new features.

Each time the company patches a software module, the application experiences downtime. Report generation must restart from the beginning after any interruptions. The company wants to redesign the application so that the application can be flexible, scalable, and gradually improved. The company wants to minimize application downtime.

Which solution will meet these requirements?

- A. Run the application on AWS Lambda as a single function with maximum provisioned concurrency.
- B. Run the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy.
- C. Run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling.
- D. Run the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy.

**Answer: C**

#### Explanation:

The solution that will meet the requirements is to run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling. This solution will allow the application to be flexible, scalable, and gradually improved, as well as minimize application downtime. By breaking down the monolithic application into microservices, the company can decouple the modules and update them independently, without affecting the whole application. By running the microservices on Amazon ECS, the company can leverage the benefits of containerization, such as portability, efficiency, and isolation. By enabling service auto scaling, the company can adjust the number of containers running for each microservice based on demand, ensuring optimal performance and cost. Amazon ECS also supports various deployment strategies, such as rolling update or blue/green deployment, that can reduce or eliminate downtime during updates.

The other solutions are not as effective as the first one because they either do not meet the requirements or introduce new challenges. Running the application on AWS Lambda as a single function with maximum provisioned concurrency will not meet the requirements, as it will not break down the monolith into microservices, nor will it reduce the complexity of maintenance. Lambda functions are also limited by execution time (15 minutes), memory size (10 GB), and concurrency quotas, which may not be sufficient for the report generation application. Running the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy will not meet the requirements, as it will introduce the risk of interruptions due to spot price fluctuations. Spot Instances are not guaranteed to be available or stable, and may be reclaimed by AWS at any time with a two-minute warning. This may cause report generation to fail or restart from scratch. Running the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy will not meet the requirements, as it will not break down the monolith into microservices, nor will it minimize application downtime. The all-at-once deployment strategy will deploy updates to all instances simultaneously, causing a brief outage for the application.

References:

- ? Amazon Elastic Container Service
- ? Microservices on AWS
- ? Service Auto Scaling - Amazon Elastic Container Service
- ? AWS Lambda
- ? Amazon EC2 Spot Instances
- ? [AWS Elastic Beanstalk]

### NEW QUESTION 34

- (Topic 4)

A company is using an Application Load Balancer (ALB) to present its application to the internet. The company finds abnormal traffic access patterns across the application. A solutions architect needs to improve visibility into the infrastructure to help the company understand these abnormalities better.

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

- A. Create a table in Amazon Athena for AWS CloudTrail log
- B. Create a query for the relevant information.
- C. Enable ALB access logging to Amazon S3. Create a table in Amazon Athena, and query the logs.
- D. Enable ALB access logging to Amazon S3. Open each file in a text editor, and search each line for the relevant information.
- E. Use Amazon EMR on a dedicated Amazon EC2 instance to directly query the ALB to acquire traffic access log information.

**Answer: B**

#### Explanation:

This solution meets the requirements because it allows the company to improve visibility into the infrastructure by using ALB access logging and Amazon Athena. ALB access logging is a feature that captures detailed information about requests sent to the load balancer, such as the client's IP address, request path, response code, and latency. By enabling ALB access logging to Amazon S3, the company can store the access logs in an S3 bucket as compressed files. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. By creating a table in Amazon Athena for the access logs, the company can query the logs and get results in seconds. This way, the company can better understand the abnormal traffic access patterns across the application.

References:

- ? Access logs for your Application Load Balancer
- ? Querying Application Load Balancer Logs

### NEW QUESTION 38

- (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 40**

- (Topic 4)

A company is reviewing a recent migration of a three-tier application to a VPC. The security team discovers that the principle of least privilege is not being applied to Amazon EC2 security group ingress and egress rules between the application tiers.

What should a solutions architect do to correct this issue?

- A. Create security group rules using the instance ID as the source or destination.
- B. Create security group rules using the security group ID as the source or destination.
- C. Create security group rules using the VPC CIDR blocks as the source or destination.
- D. Create security group rules using the subnet CIDR blocks as the source or destination.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group-rules.html>

**NEW QUESTION 43**

- (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 Zone
- B. Use Amazon S3 storag
- C. Create an AWS Lambda function to process order file
- 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 Zon
- F. Use Amazon Elastic File System (Amazon EFS) storag
- G. Create an AWS Lambda function to process order file
- H. Use a Transfer Family managed workflow to invoke the Lambda function.
- I. Create an AWS Transfer Family SFTP internal server in two Availability Zone
- J. Use Amazon Elastic File System (Amazon EFS) storag
- K. Create an AWS Step Functions state machine to process order file
- 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 Zone
- N. Use Amazon S3 storag
- O. Create an AWS Lambda function to process order file
- 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 45**

- (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 46**

- (Topic 4)

A company has a mobile chat application with a data store based in Amazon DynamoDB. Users would like new messages to be read with as little latency as possible. A solutions architect needs to design an optimal solution that requires minimal application changes.

Which method should the solutions architect select?

- A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages table.
- B. Update the code to use the DAX endpoint.
- C. Add DynamoDB read replicas to handle the increased read load.
- D. Update the application to point to the read endpoint for the read replicas.
- E. Double the number of read capacity units for the new messages table in DynamoDB.
- F. Continue to use the existing DynamoDB endpoint.
- G. Add an Amazon ElastiCache for Redis cache to the application stack.
- H. Update the application to point to the Redis cache endpoint instead of DynamoDB.

**Answer:** A

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-high-latency/>

Amazon DynamoDB Accelerator (DAX) is a fully managed in-memory cache for DynamoDB that improves the performance of DynamoDB tables by up to 10 times and

provides a microsecond level of response time at any scale. It is compatible with DynamoDB API operations and requires minimal code changes to use<sup>1</sup>. By configuring DAX for the

new messages table, the solution can reduce the latency for reading new messages with minimal application changes.

\* B. Add DynamoDB read replicas to handle the increased read load. Update the application to point to the read endpoint for the read replicas. This solution will not work, as DynamoDB does not support read replicas as a feature. Read replicas are available for Amazon RDS, not for DynamoDB<sup>2</sup>.

\* C. Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint. This solution will not meet the requirement of reading new messages with as little latency as possible, as increasing the read capacity units will only increase the throughput of DynamoDB, not the performance or latency<sup>3</sup>.

\* D. Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB. This solution will not meet the requirement of minimal application changes, as adding ElastiCache for Redis will require significant code changes to implement caching logic, such as querying cache first, updating cache after writing to DynamoDB, and invalidating cache when needed. Reference URL:

<https://aws.amazon.com/dynamodb/dax/>

**NEW QUESTION 49**

- (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 51**

- (Topic 4)

A solutions architect is designing a new API using Amazon API Gateway that will receive

requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made.

Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

- A. An AWS Glue job



- B. An AWS Lambda function
- C. A containerized service hosted in Amazon Elastic Kubernetes Service (Amazon EKS)
- D. A containerized service hosted in Amazon ECS with Amazon EC2

**Answer:** B

**Explanation:**

API Gateway + Lambda is the perfect solution for modern applications with serverless architecture.

**NEW QUESTION 55**

- (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 56**

- (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 57

- (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 Notificatio
- G. Configure an Amazon Simple Notification Service (Amazon SNS) topic as a targe
- 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 60

- (Topic 4)

A retail company uses a regional Amazon API Gateway API for its public REST APIs. The API Gateway endpoint is a custom domain name that points to an Amazon Route 53 alias record. A solutions architect needs to create a solution that has minimal effects on customers and minimal data loss to release the new version of APIs.

Which solution will meet these requirements?

- A. Create a canary release deployment stage for API Gatewa
- B. Deploy the latest API versio
- C. Point an appropriate percentage of traffic to the canary stag
- D. After API verification, promote the canary stage to the production stage.
- E. Create a new API Gateway endpoint with a new version of the API in OpenAPI YAMLfile forma
- F. Use the import-to-update operation in merge mode into the API in API Gatewa
- G. Deploy the new version of the API to the production stage.
- H. Create a new API Gateway endpoint with a new version of the API in OpenAPI JSON file forma
- I. Use the import-to-update operation in overwrite mode into the API in API Gatewa
- J. Deploy the new version of the API to the production stage.
- K. Create a new API Gateway endpoint with new versions of the API definition
- L. Create a custom domain name for the new API Gateway AP
- M. Point the Route 53 alias record to the new API Gateway API custom domain name.

**Answer:** A

#### Explanation:

This answer is correct because it meets the requirements of releasing the new version of APIs with minimal effects on customers and minimal data loss. A canary release deployment is a software development strategy in which a new version of an API is deployed for testing purposes, and the base version remains deployed as a production release for normal operations on the same stage. In a canary release deployment, total API traffic is separated at random into a production release and a canary release with a pre- configured ratio. Typically, the canary release receives a small percentage of API traffic and the production release takes up the rest. The updated API features are only visible to API traffic through the canary. You can adjust the canary traffic percentage to optimize test coverage or performance. By keeping canary traffic small and the selection random, most users are not adversely affected at any time by potential bugs in the new version, and no single user is adversely affected all the time. After the test metrics pass your requirements, you can promote the canary release to the production release and disable the canary from the deployment. This makes the new features available in the production stage. References:

? <https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>

#### NEW QUESTION 63

- (Topic 4)

A solutions architect is designing a REST API in Amazon API Gateway for a cash payback service The application requires 1 GB of memory and 2 GB of storage for its computation resources. The application will require that the data is in a relational format.

Which additional combination of AWS services will meet these requirements with the LEAST administrative effort? {Select TWO.)

- A. Amazon EC2
- B. AWS Lambda
- C. Amazon RDS
- D. Amazon DynamoDB
- E. Amazon Elastic Kubernetes Services (Amazon EKS)

**Answer:** BC

#### Explanation:

AWS Lambda is a service that lets users run code without provisioning or managing servers. It automatically scales and manages the underlying compute resources for the code. It supports multiple languages, such as Java, Python, Node.js, and G1o. By using AWS Lambda for the REST API, the solution can meet the requirements of 1 GB of memory and minimal administrative effort.

Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups. It supports multiple database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server2. By using Amazon RDS for the data store, the solution can meet the requirements of 2 GB of storage and a relational format.

\* A. Amazon EC2. This solution will not meet the requirement of minimal administrative effort, as Amazon EC2 is a service that provides virtual servers in the cloud that users have to configure and manage themselves. It requires users to choose an instance type, an operating system, a security group, and other options3.

\* D. Amazon DynamoDB. This solution will not meet the requirement of a relational format, as Amazon DynamoDB is a service that provides a key-value and document database that delivers single-digit millisecond performance at any scale. It is a non-relational or NoSQL database that does not support joins or transactions.

\* E. Amazon Elastic Kubernetes Services (Amazon EKS). This solution will not meet the requirement of minimal administrative effort, as Amazon EKS is a service

that provides a fully managed Kubernetes service that users have to configure and manage themselves. It requires users to create clusters, nodes groups, pods, services, and other Kubernetes resources.

Reference URL: <https://aws.amazon.com/lambda/>

#### NEW QUESTION 64

- (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 66

- (Topic 4)

A company wants to host a scalable web application on AWS. The application will be accessed by users from different geographic regions of the world. Application users will be able to download and upload unique data up to gigabytes in size. The development team wants a cost-effective solution to minimize upload and download latency and maximize performance.

What should a solutions architect do to accomplish this?

- A. Use Amazon S3 with Transfer Acceleration to host the application.
- B. Use Amazon S3 with CacheControl headers to host the application.
- C. Use Amazon EC2 with Auto Scaling and Amazon CloudFront to host the application.
- D. Use Amazon EC2 with Auto Scaling and Amazon ElastiCache to host the application.

**Answer:** C

#### Explanation:

This answer is correct because it meets the requirements of hosting a scalable web application that can handle large data transfers from different geographic regions. Amazon EC2 provides scalable compute capacity for hosting web applications. Auto Scaling can automatically adjust the number of EC2 instances based on the demand and traffic patterns. Amazon CloudFront is a content delivery network (CDN) that can cache static and dynamic content at edge locations closer to the users, reducing latency and improving performance. CloudFront can also use S3 Transfer Acceleration to speed up the transfers between S3 buckets and CloudFront edge locations.

References:

? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html>

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

? <https://aws.amazon.com/s3/transfer-acceleration/>

#### NEW QUESTION 69

- (Topic 4)

A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account.

The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.

Which authentication option will meet these requirements MOST securely?

- A. Integrate DynamoDB with AWS Secrets Manager in the inventory application account
- B. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB table
- C. Schedule secret rotation for every 30 days.
- D. In every business account, create an IAM user that has programmatic access
- E. Configure the application to use the correct IAM user access key ID and secret access key to authenticate and read the DynamoDB table
- F. Manually rotate IAM access keys every 30 days.
- G. In every business account, create an IAM role named BU\_ROLE with a policy that gives the role access to the DynamoDB table and a trust policy to trust a specific role in the inventory application account
- H. In the inventory account, create a role named APP\_ROLE that allows access to the STS AssumeRole API operation
- I. Configure the application to use APP\_ROLE and assume the cross-account role BU\_ROLE to read the DynamoDB table.
- J. Integrate DynamoDB with AWS Certificate Manager (ACM). Generate identity certificates to authenticate DynamoDB
- K. Configure the application to use the correct certificate to authenticate and read the DynamoDB table.

**Answer:** C

#### Explanation:

This solution meets the requirements most securely because it uses IAM roles and the STS AssumeRole API operation to authenticate and authorize the inventory application to access the DynamoDB tables in different accounts. IAM roles are more secure than IAM users or certificates because they do not require long-term credentials or passwords. Instead, IAM roles provide temporary security credentials that are automatically rotated and can be configured with a limited

duration. The STS AssumeRole API operation enables you to request temporary credentials for a role that you are allowed to assume. By using this operation, you can delegate access to resources that are in different AWS accounts that you own or that are owned by third parties. The trust policy of the role defines which entities can assume the role, and the permissions policy of the role defines which actions can be performed on the resources. By using this solution, you can avoid hard-coding credentials or certificates in the inventory application, and you can also avoid storing them in Secrets Manager or ACM. You can also leverage the built-in security features of IAM and STS, such as MFA, access logging, and policy conditions.

References:

? IAM Roles

? STS AssumeRole

? Tutorial: Delegate Access Across AWS Accounts Using IAM Roles

### NEW QUESTION 73

- (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 RDS 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 74

- (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 data
- B. Share the transformation steps with employees by using AWS Glue jobs.
- C. Configure Amazon EMR Serverless to transform the data
- D. Share the transformation steps with employees by using EMR Serverless jobs.
- E. Configure AWS Glue DataBrew to transform the data
- F. Share the transformation steps with employees by using DataBrew recipes.
- G. Create Amazon Athena tables for the data
- H. Write Athena SQL queries to transform the data
- 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.

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.

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 75

- (Topic 4)

A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users. What should a solutions architect recommend?

- A. Deploy Amazon Inspector and associate it with the ALB.
- B. Deploy AWS WAF, associate it with the ALB, and configure a rate-limiting rule.
- C. Deploy rules to the network ACLs associated with the ALB to block the incoming traffic.
- D. Deploy Amazon GuardDuty and enable rate-limiting protection when configuring GuardDuty.

**Answer:** B



**Explanation:**

This answer is correct because it meets the requirements of blocking the illegitimate incoming requests in a way that has a minimal impact on legitimate users. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define. You can associate AWS WAF with an ALB to protect the web application from malicious requests. You can configure a rate-limiting rule in AWS WAF to track the rate of requests for each originating IP address and block requests from an IP address that exceeds a certain limit within a five-minute period. This way, you can mitigate potential DDoS attacks and improve the performance of your website.

References:

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

? <https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type- rate-based.html>

**NEW QUESTION 77**

- (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 82**

- (Topic 4)

A gaming company wants to launch a new internet-facing application in multiple AWS Regions. The application will use the TCP and UDP protocols for communication. The company needs to provide high availability and minimum latency for global users.

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

- A. Create internal Network Load Balancers in front of the application in each Region.
- B. Create external Application Load Balancers in front of the application in each Region.
- C. Create an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region.
- D. Configure Amazon Route 53 to use a geolocation routing policy to distribute the traffic.
- E. Configure Amazon CloudFront to handle the traffic and route requests to the application in each Region.

**Answer:** BC

**Explanation:**

This combination of actions will provide high availability and minimum latency for global users by using AWS Global Accelerator and Application Load Balancers. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your internet-facing applications by using the AWS global network. It provides two global static public IPs that act as a fixed entry point to your application endpoints, such as Application Load Balancers, in multiple Regions<sup>1</sup>. Global Accelerator uses the AWS backbone network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. It also offers TCP and UDP support, traffic encryption, and DDoS protection<sup>2</sup>. Application Load Balancers are external load balancers that distribute incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. They support both HTTP and HTTPS (SSL/TLS) protocols, and offer advanced features such as content-based routing, health checks, and integration with other AWS services<sup>3</sup>. By creating external Application Load Balancers in front of the application in each Region, you can ensure that the application can handle varying load patterns and scale on demand. By creating an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region, you can leverage the performance, security, and availability of the AWS global network to deliver the best possible user experience.

References: 1: What is AWS Global Accelerator? - AWS Global Accelerator<sup>4</sup>, Overview section<sup>2</sup>: Network Acceleration Service - AWS Global Accelerator - AWS<sup>5</sup>, Why AWS Global Accelerator? section. 3: What is an Application Load Balancer? - Elastic Load Balancing<sup>6</sup>, Overview section.

**NEW QUESTION 85**

- (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 88

- (Topic 4)

A company runs multiple workloads in its on-premises data center. The company's data center cannot scale fast enough to meet the company's expanding business needs. The company wants to collect usage and configuration data about the on-premises servers and workloads to plan a migration to AWS. Which solution will meet these requirements?

- A. Set the home AWS Region in AWS Migration Hu
- B. Use AWS Systems Manager to collect data about the on-premises servers.
- C. Set the home AWS Region in AWS Migration Hu
- D. Use AWS Application Discovery Service to collect data about the on-premises servers.
- E. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant template
- F. Use AWS Trusted Advisor to collect data about the on-premises servers.
- G. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates. Use AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers.

**Answer: B**

#### Explanation:

The most suitable solution for the company's requirements is to set the home AWS Region in AWS Migration Hub and use AWS Application Discovery Service to collect data about the on-premises servers. This solution will enable the company to gather usage and configuration data of its on-premises servers and workloads, and plan a migration to AWS.

AWS Migration Hub is a service that simplifies and accelerates migration tracking by aggregating migration status information into a single console. Users can view the discovered servers, group them into applications, and track the migration status of each application from the Migration Hub console in their home Region. The home Region is the AWS Region where users store their migration data, regardless of which Regions they migrate into<sup>1</sup>.

AWS Application Discovery Service is a service that helps users plan their migration to AWS by collecting usage and configuration data about their on-premises servers and databases. Application Discovery Service is integrated with AWS Migration Hub and supports two methods of performing discovery: agentless discovery and agent-based discovery. Agentless discovery can be performed by deploying the Application Discovery Service Agentless Collector through VMware vCenter, which collects static configuration data and utilization data for virtual machines (VMs) and databases. Agent-based discovery can be performed by deploying the AWS Application Discovery Agent on each of the VMs and physical servers, which collects static configuration data, detailed time-series system-performance information, inbound and outbound network connections, and processes that are running<sup>2</sup>.

The other options are not correct because they do not meet the requirements or are not relevant for the use case. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Trusted Advisor to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another, such as from Oracle to PostgreSQL<sup>3</sup>. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits<sup>4</sup>. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. As mentioned above, AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another. AWS DMS is a service that helps users migrate relational databases, non-relational databases, and other types of data stores to

AWS with minimal downtime<sup>5</sup>. References:

? Home Region - AWS Migration Hub

? What is AWS Application Discovery Service? - AWS Application Discovery Service

? AWS Schema Conversion Tool - Amazon Web Services

? What Is Trusted Advisor? - Trusted Advisor

? What Is AWS Database Migration Service? - AWS Database Migration Service

### NEW QUESTION 91

- (Topic 4)

A company has a production workload that is spread across different AWS accounts in various AWS Regions. The company uses AWS Cost Explorer to continuously monitor costs and usage. The company wants to receive notifications when the cost and usage spending of the workload is unusual. Which combination of steps will meet these requirements? (Select TWO.)

- A. In the AWS accounts where the production workload is running, create a linked account budget by using Cost Explorer in the AWS Cost Management console
- B. In the AWS accounts where the production workload is running, create a linked account monitor by using AWS Cost Anomaly Detection in the AWS Cost Management console
- C. In the AWS accounts where the production workload is running, create a Cost and Usage Report by using Cost Anomaly Detection in the AWS Cost Management console.
- D. Create a report and send email messages to notify the company on a weekly basis.
- E. Create a subscription with the required threshold and notify the company by using weekly summaries.

**Answer: BE**

#### Explanation:

AWS Cost Anomaly Detection allows you to create monitors that track the cost and usage of your AWS resources and alert you when there is an unusual spending pattern. You can create monitors based on different dimensions, such as AWS services, accounts, tags, or cost categories. You can also create alert subscriptions that notify you by email or Amazon SNS when an anomaly is detected. You can specify the threshold and frequency of the alerts, and choose to receive weekly summaries of your anomalies. Reference URLs:

1 <https://aws.amazon.com/aws-cost-management/aws-cost-anomaly-detection/>

2 <https://docs.aws.amazon.com/cost-management/latest/userguide/getting-started-ad.html>

3 <https://docs.aws.amazon.com/cost-management/latest/userguide/manage-ad.html>

### NEW QUESTION 93

- (Topic 4)

A company wants to manage Amazon Machine Images (AMIs). The company currently copies AMIs to the same AWS Region where the AMIs were created. The company needs to design an application that captures AWS API calls and sends alerts whenever the Amazon EC2 CreateImage API operation is called within the company's account.

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

- A. Create an AWS Lambda function to query AWS CloudTrail logs and to send an alert when a CreateImage API call is detected.
- B. Configure AWS CloudTrail with an Amazon Simple Notification Service (Amazon SNS) notification that occurs when updated logs are sent to Amazon S3. Use Amazon Athena to create a new table and to query on CreateImage when an API call is detected.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API call.
- D. Configure the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected.
- E. Configure an Amazon Simple Queue Service (Amazon SQS) FIFO queue as a target for AWS CloudTrail log.
- F. Create an AWS Lambda function to send an alert to an Amazon Simple Notification Service (Amazon SNS) topic when a CreateImage API call is detected.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/monitor-ami-events.html#:~:text=For%20example%2C%20you%20can%20create%20an%20EventBridge%20rule%20that%20detects%20when%20the%20AMI%20creation%20process%20has%20completed%20and%20then%20invokes%20an%20Amazon%20SNS%20topic%20to%20send%20an%20email%20notification%20to%20you.>

**NEW QUESTION 97**

- (Topic 4)

A company runs a container application on a Kubernetes cluster in the company's data center. The application uses Advanced Message Queuing Protocol (AMQP) to communicate with a message queue. The data center cannot scale fast enough to meet the company's expanding business needs. The company wants to migrate the workloads to AWS.

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

- A. Migrate the container application to Amazon Elastic Container Service (Amazon ECS). Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages.
- B. Migrate the container application to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon MQ to retrieve the messages.
- C. Use highly available Amazon EC2 instances to run the application. Use Amazon MQ to retrieve the messages.
- D. Use AWS Lambda functions to run the application. Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages.

**Answer:** B

**Explanation:**

This option is the best solution because it allows the company to migrate the container application to AWS with minimal changes and leverage a managed service to run the Kubernetes cluster and the message queue. By using Amazon EKS, the company can run the container application on a fully managed Kubernetes control plane that is compatible with the existing Kubernetes tools and plugins. Amazon EKS handles the provisioning, scaling, patching, and security of the Kubernetes cluster, reducing the operational overhead and complexity. By using Amazon MQ, the company can use a fully managed message broker service that supports AMQP and other popular messaging protocols. Amazon MQ handles the administration, maintenance, and scaling of the message broker, ensuring high availability, durability, and security of the messages.

\* A. Migrate the container application to Amazon Elastic Container Service (Amazon ECS). Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages. This option is not optimal because it requires the company to change the container orchestration platform from Kubernetes to ECS, which can introduce additional complexity and risk. Moreover, it requires the company to change the messaging protocol from AMQP to SQS, which can also affect the application logic and performance. Amazon ECS and Amazon SQS are both fully managed services that simplify the deployment and management of containers and messages, but they may not be compatible with the existing application architecture and requirements.

\* C. Use highly available Amazon EC2 instances to run the application. Use Amazon MQ to retrieve the messages. This option is not ideal because it requires the company to manage the EC2 instances that host the container application. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, the company would need to install and maintain the Kubernetes software on the EC2 instances, which can also add complexity and risk. Amazon MQ is a fully managed message broker service that supports AMQP and other popular messaging protocols, but it cannot compensate for the lack of a managed Kubernetes service.

\* D. Use AWS Lambda functions to run the application. Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages. This option is not feasible because AWS Lambda does not support running container applications directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run container applications on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the container API, which can introduce additional complexity and overhead. Moreover, Lambda functions have limitations in terms of available CPU, memory, and runtime, which may not suit the application needs. Amazon SQS is a fully managed message queue service that supports asynchronous communication, but it does not support AMQP or other messaging protocols.

References:

- ? 1 Amazon Elastic Kubernetes Service - Amazon Web Services
- ? 2 Amazon MQ - Amazon Web Services
- ? 3 Amazon Elastic Container Service - Amazon Web Services
- ? 4 AWS Lambda FAQs - Amazon Web Services

**NEW QUESTION 99**

- (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 hierarchy.
- B. Unload the data to Amazon S3 every month.
- C. Use Amazon DynamoDB to store the employee data in hierarchy.
- D. Export the data to Amazon S3 every month.
- E. Configure Amazon IAM for the AWS account.
- F. Integrate Amazon 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 104

- (Topic 4)

A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming data. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting.

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

- A. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.
- B. Use AWS Glue to deliver streaming data to Amazon S3.
- C. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.
- D. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.

**Answer:** A

#### Explanation:

To ingest and process high-volume streaming data with the least operational overhead, Amazon Kinesis Data Firehose is a suitable solution. Amazon Kinesis Data Firehose can capture, transform, and deliver streaming data to Amazon S3 or other destinations. Amazon Kinesis Data Firehose can scale automatically to match the throughput of the data and handle any amount of data. Amazon Kinesis Data Firehose is also a fully managed service that does not require any servers to provision or manage. References:

? [What Is Amazon Kinesis Data Firehose?](#)

? [Amazon Kinesis Data Firehose Pricing](#)

#### NEW QUESTION 105

- (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/awscloudtrail/latest/userguide/logging-data-events-with-cloudtrail.html>

#### NEW QUESTION 109

- (Topic 4)

A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance. The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.

Which solution will allow the node to join the cluster?

- A. Grant the required permission in AWS Identity and Access Management (IAM) to the AmazonEKSNodeRole IAM role.
- B. Create interface VPC endpoints to allow nodes to access the control plane.
- C. Recreate nodes in the public subnet. Restrict security groups for EC2 nodes.
- D. Allow outbound traffic in the security group of the nodes.

**Answer:** B

#### Explanation:

Kubernetes API requests within your cluster's VPC (such as node to control plane communication) use the private VPC endpoint.

<https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html>

#### NEW QUESTION 113

- (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 distributio
- 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 118**

- (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 cachin
- 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 123**

- (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 126**

- (Topic 4)

A company is deploying an application in three AWS Regions using an Application Load Balancer Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST high-performing experience?

- A. Create an A record with a latency policy.



- B. Create an A record with a geolocation policy.
- C. Create a CNAME record with a failover policy.
- D. Create a CNAME record with a geoproximity policy.

**Answer:** A

**Explanation:**

To provide the most high-performing experience for the users of the application, a solutions architect should use a latency routing policy for the Route 53 A record. This policy allows Route 53 to route traffic to the AWS Region that provides the lowest possible latency for the users<sup>1</sup>. A latency routing policy can also improve the availability of the application, as Route 53 can automatically route traffic to another Region if the primary Region becomes unavailable<sup>2</sup>.

References:

? 1: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-latency>

? 2: [https://aws.amazon.com/route53/faqs/#Latency\\_Based\\_Routing](https://aws.amazon.com/route53/faqs/#Latency_Based_Routing)

**NEW QUESTION 130**

- (Topic 4)

A solutions architect must provide an automated solution for a company's compliance policy that states security groups cannot include a rule that allows SSH from 0.0.0.0/0. The company needs to be notified if there is any breach in the policy. A solution is needed as soon as possible.

What should the solutions architect do to meet these requirements with the LEAST operational overhead?

- A. Write an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one.
- B. Enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created.
- C. Create an IAM role with permissions to globally open security groups and network ACL
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic to generate a notification every time the role is assumed by a user.
- E. Configure a service control policy (SCP) that prevents non-administrative users from creating or editing security group
- F. Create a notification in the ticketing system when a user requests a rule that needs administrator permissions.

**Answer:** B

**Explanation:**

The most suitable solution for the company's compliance policy is to enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created. This solution has the least operational overhead because it uses a predefined rule that is already available in AWS Config, which is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. The restricted-ssh rule checks whether security groups that are in use have inbound rules that allow SSH from 0.0.0.0/0 addresses, and reports them as noncompliant<sup>1</sup>. Users can configure the rule to send notifications to an Amazon SNS topic when a noncompliant change occurs, and subscribe to the topic to receive alerts via email, SMS, or other methods<sup>2</sup>.

The other options are not correct because they either have more operational overhead or do not meet the requirements. Writing an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one is not correct because it requires custom code development and maintenance, which adds complexity and cost to the solution. Creating an IAM role with permissions to globally open security groups and network ACLs, and creating an Amazon SNS topic to generate a notification every time the role is assumed by a user is not correct because it does not prevent or detect the creation of noncompliant rules by other users or roles, and it does not address the existing rules that may violate the policy. Configuring a service control policy (SCP) that prevents non-administrative users from creating or editing security groups, and creating a notification in the ticketing system when a user requests a rule that needs administrator permissions is not correct because it does not provide an automated solution for the policy enforcement and notification, and it may limit the flexibility and productivity of the users.

References:

? restricted-ssh - AWS Config

? Getting Notifications When Your Resources Change - AWS Config

**NEW QUESTION 133**

- (Topic 4)

A company runs a real-time data ingestion solution on AWS. The solution consists of the most recent version of Amazon Managed Streaming for Apache Kafka (Amazon MSK). The solution is deployed in a VPC in private subnets across three Availability Zones.

A solutions architect needs to redesign the data ingestion solution to be publicly available over the internet. The data in transit must also be encrypted.

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

- A. Configure public subnets in the existing VPC
- B. Deploy an MSK cluster in the public subnet
- C. Update the MSK cluster security settings to enable mutual TLS authentication.
- D. Create a new VPC that has public subnet
- E. Deploy an MSK cluster in the public subnet
- F. Update the MSK cluster security settings to enable mutual TLS authentication.
- G. Deploy an Application Load Balancer (ALB) that uses private subnet
- H. Configure an ALB security group inbound rule to allow inbound traffic from the VPC CIDR block for HTTPS protocol.
- I. Deploy a Network Load Balancer (NLB) that uses private subnet
- J. Configure an NLB listener for HTTPS communication over the internet.

**Answer:** A

**Explanation:**

The solution that meets the requirements with the most operational efficiency is to configure public subnets in the existing VPC and deploy an MSK cluster in the public subnets. This solution allows the data ingestion solution to be publicly available over the internet without creating a new VPC or deploying a load balancer. The solution also ensures that the data in transit is encrypted by enabling mutual TLS authentication, which requires both the client and the server to present certificates for verification. This solution leverages the public access feature of Amazon MSK, which is available for clusters running Apache Kafka 2.6.0 or later versions<sup>1</sup>.

The other solutions are not as efficient as the first one because they either create unnecessary resources or do not encrypt the data in transit. Creating a new VPC with public subnets would incur additional costs and complexity for managing network resources and routing. Deploying an ALB or an NLB would also add more costs and latency for the data ingestion solution. Moreover, an ALB or an NLB would not encrypt the data in transit by itself, unless they are configured with HTTPS listeners and certificates, which would require additional steps and maintenance. Therefore, these solutions are not optimal for the given requirements.

References:

? Public access - Amazon Managed Streaming for Apache Kafka

### NEW QUESTION 137

- (Topic 4)

A company has hired a solutions architect to design a reliable architecture for its application. The application consists of one Amazon RDS DB instance and two manually provisioned Amazon EC2 instances that run web servers. The EC2 instances are located in a single Availability Zone.

An employee recently deleted the DB instance, and the application was unavailable for 24 hours as a result. The company is concerned with the overall reliability of its environment.

What should the solutions architect do to maximize reliability of the application's infrastructure?

- A. Delete one EC2 instance and enable termination protection on the other EC2 instance
- B. Update the DB instance to be Multi-AZ, and enable deletion protection.
- C. Update the DB instance to be Multi-AZ, and enable deletion protection
- D. Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling group across multiple Availability Zones.
- E. Create an additional DB instance along with an Amazon API Gateway and an AWS Lambda function
- F. Configure the application to invoke the Lambda function through API Gateway
- G. Have the Lambda function write the data to the two DB instances.
- H. Place the EC2 instances in an EC2 Auto Scaling group that has multiple subnets located in multiple Availability Zones
- I. Use Spot Instances instead of On-Demand Instance
- J. Set up Amazon CloudWatch alarms to monitor the health of the instance
- K. Update the DB instance to be Multi-AZ, and enable deletion protection.

**Answer: B**

#### Explanation:

This answer is correct because it meets the requirements of maximizing the reliability of the application's infrastructure. You can update the DB instance to be Multi-AZ, which means that Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance. It can also help protect your databases against DB instance failure and Availability Zone disruption. You can also enable deletion protection on the DB instance, which prevents the DB instance from being deleted by any user. You can place the EC2 instances behind an Application Load Balancer, which distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This increases the availability and fault tolerance of your applications. You can run the EC2 instances in an EC2 Auto Scaling group across multiple Availability Zones, which ensures that you have the correct number of EC2 instances available to handle the load for your application. You can use scaling policies to adjust the number of instances in your Auto Scaling group in response to changing demand.

References:

? <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZSingleStandby.html>

? [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_DeleteInstance.html#USER\\_DeleteInstance.DeletionProtection](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html#USER_DeleteInstance.DeletionProtection)

? <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html>

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

### NEW QUESTION 141

- (Topic 4)

A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.

However, the company's security policy states that any external service cannot initiate a connection to the EC2 instances.

What should a solutions architect recommend to resolve this issue?

- A. Create a NAT gateway and make it the destination of the subnet's route table.
- B. Create an internet gateway and make it the destination of the subnet's route table
- C. Create a virtual private gateway and make it the destination of the subnet's route table.
- D. Create an egress-only internet gateway and make it the destination of the subnet's route table.

**Answer: D**

#### Explanation:

An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in your VPC to the internet, and prevents the internet from initiating an IPv6 connection with your instances. This meets the company's security policy and requirements. To use an egress-only internet gateway, you need to add a route in the subnet's route table that routes IPv6 internet traffic (::/0) to the egress-only internet gateway.

Reference URLs:

1 <https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

2 <https://dev.to/aws-builders/what-is-an-egress-only-internet-gateways-in-aws-7gp>

3 <https://docs.aws.amazon.com/vpc/latest/userguide/route-table-options.html>

### NEW QUESTION 144

- (Topic 4)

A company is storing 700 terabytes of data on a large network-attached storage (NAS) system in its corporate data center. The company has a hybrid environment with a 10 Gbps AWS Direct Connect connection.

After an audit from a regulator, the company has 90 days to move the data to the cloud. The company needs to move the data efficiently and without disruption.

The company still needs to be able to access and update the data during the transfer window.

Which solution will meet these requirements?

- A. Create an AWS DataSync agent in the corporate data center
- B. Create a data transfer task
- C. Start the transfer to an Amazon S3 bucket.
- D. Back up the data to AWS Snowball Edge Storage Optimized device
- E. Ship the devices to an AWS data center
- F. Mount a target Amazon S3 bucket on the on-premises file system.
- G. Use rsync to copy the data directly from local storage to a designated Amazon S3 bucket over the Direct Connect connection.
- H. Back up the data on tape
- I. Ship the tapes to an AWS data center
- J. Mount a target Amazon S3 bucket on the on-premises file system.

**Answer:** A

**Explanation:**

This answer is correct because it meets the requirements of moving the data efficiently and without disruption, and still being able to access and update the data during the transfer window. AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS and helps you move data quickly and securely between on-premises storage, edge locations, other clouds, and AWS Storage. You can create an AWS DataSync agent in the corporate data center to connect your NAS system to AWS over the Direct Connect connection. You can create a data transfer task to specify the source location, destination location, and options for transferring the data. You can start the transfer to an Amazon S3 bucket and monitor the progress of the task. DataSync automatically encrypts data in transit and verifies data integrity during transfer. DataSync also supports incremental transfers, which means that only files that have changed since the last transfer are copied. This way, you can ensure that your data is synchronized between your NAS system and S3 bucket, and you can access and update the data during the transfer window.

References:

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

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

**NEW QUESTION 147**

- (Topic 4)

A company needs a solution to prevent photos with unwanted content from being uploaded to the company's web application. The solution must not involve training a machine learning (ML) model. Which solution will meet these requirements?

- A. Create and deploy a model by using Amazon SageMaker Autopilo
- B. Create a real-time endpoint that the web application invokes when new photos are uploaded.
- C. Create an AWS Lambda function that uses Amazon Rekognition to detect unwanted conten
- D. Create a Lambda function URL that the web application invokes when new photos are uploaded.
- E. Create an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted conten
- F. Associate the function with the web application.
- G. Create an AWS Lambda function that uses Amazon Rekognition Video to detectunwanted conten
- H. Create a Lambda function URL that the web application invokes when new photos are uploaded.

**Answer:** B

**Explanation:**

The solution that will meet the requirements is to create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content, and create a Lambda function URL that the web application invokes when new photos are uploaded. This solution does not involve training a machine learning model, as Amazon Rekognition is a fully managed service that provides pre-trained computer vision models for image and video analysis. Amazon Rekognition can detect unwanted content such as explicit or suggestive adult content, violence, weapons, drugs, and more. By using AWS Lambda, the company can create a serverless function that can be triggered by an HTTP request from the web application. The Lambda function can use the Amazon Rekognition API to analyze the uploaded photos and return a response indicating whether they contain unwanted content or not.

The other solutions are not as effective as the first one because they either involve training a machine learning model, do not support image analysis, or do not work with photos. Creating and deploying a model by using Amazon SageMaker Autopilot involves training a machine learning model, which is not required for the scenario. Amazon SageMaker Autopilot is a service that automatically creates, trains, and tunes the best machine learning models for classification or regression based on the data provided by the user. Creating an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content does not support image analysis, as Amazon Comprehend is a natural language processing service that analyzes text, not images. Amazon Comprehend can extract insights and relationships from text such as language, sentiment, entities, topics, and more. Creating an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content does not work with photos, as Amazon Rekognition Video is designed for analyzing video streams, not static images. Amazon Rekognition Video can detect activities, objects, faces, celebrities, text, and more in video streams.

References:

? Amazon Rekognition

? AWS Lambda

? Detecting unsafe content - Amazon Rekognition

? Amazon SageMaker Autopilot

? Amazon Comprehend

**NEW QUESTION 152**

.....

## Relate Links

**100% Pass Your AWS-Solution-Architect-Associate Exam with ExamBible Prep Materials**

<https://www.exambible.com/AWS-Solution-Architect-Associate-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>