

# Amazon

## Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate



**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 is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket. During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.

Which solution will meet these requirements?

- A. Enable S3 Intelligent-Tiering for the S3 bucket.
- B. Enable S3 Transfer Acceleration for the S3 bucket.
- C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC.
- D. Create an interface endpoint for Amazon S3 in the VPC.
- E. Associate this endpoint with all route tables in the VPC.

**Answer: C**

#### Explanation:

A gateway VPC endpoint for Amazon S3 enables private connections between the VPC and Amazon S3 that do not require an internet gateway or NAT device.

This minimizes costs and prevents traffic from traversing the internet. A gateway VPC endpoint uses a prefix list as the route target in a VPC route table to route traffic privately to Amazon S3<sup>1</sup>. Associating the endpoint with all route tables in the VPC ensures that all subnets can access Amazon S3 through the endpoint.

Option A is incorrect because S3 Intelligent-Tiering is a storage class that optimizes storage costs by automatically moving objects between two access tiers based on changing access patterns. It does not affect the network traffic between the VPC and Amazon S3<sup>2</sup>.

Option B is incorrect because S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between clients and an S3 bucket. It does not prevent traffic from traversing the internet<sup>3</sup>.

Option D is incorrect because an interface VPC endpoint for Amazon S3 is powered by AWS PrivateLink, which requires an elastic network interface (ENI) with a private IP address in each subnet. This adds complexity and cost to the solution. Moreover, an interface VPC endpoint does not support cross-Region access to Amazon S3. Reference URL: 1: <https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html> 2:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-dynamic-data-access> 3:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html> : <https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/>

#### NEW QUESTION 5

- (Topic 4)

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

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

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

**Answer: C**

#### Explanation:

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

References:

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

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

#### NEW QUESTION 6

- (Topic 4)

A company 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 solutions architect wants to use the following JSON text as an identity-based policy to grant specific permissions:

```
{  "Statement": [
    {
        "Action": [
            "ssm:ListDocuments",
            "ssm:GetDocument"
        ],
        "Effect": "Allow",
        "Resource": "*",
        "Sid": ""
    }
  ],
  "Version": "2012-10-17"
}
```

Which IAM principals can the solutions architect attach this policy to? (Select TWO.)

- A. Role
- B. Group
- C. Organization
- D. Amazon Elastic Container Service (Amazon ECS) resource
- E. Amazon EC2 resource

**Answer:** AB

**Explanation:**

This JSON text is an identity-based policy that grants specific permissions. The IAM principals that the solutions architect can attach this policy to are Role and Group. This is because the policy is written in JSON and is an identity-based policy, which can be attached to IAM principals such as users, groups, and roles. Identity-based policies are permissions policies that you attach to IAM identities (users, groups, or roles) and explicitly state what that identity is allowed (or denied) to do<sup>1</sup>. Identity-based policies are different from resource-based policies, which define the permissions around the specific resource<sup>1</sup>. Resource-based policies are attached to a resource, such as an Amazon S3 bucket or an Amazon EC2 instance<sup>1</sup>. Resource-based policies can also specify a principal, which is the entity that is allowed or denied access to the resource<sup>1</sup>. Organization is not an IAM principal, but a feature of AWS Organizations that allows you to manage multiple AWS accounts centrally<sup>2</sup>. Amazon ECS resource and Amazon EC2 resource are not IAM principals, but AWS resources that can have resource-based policies attached to them<sup>3,4</sup>. References:

- ? Identity-based policies and resource-based policies
- ? AWS Organizations
- ? Amazon ECS task role
- ? Amazon EC2 instance profile

**NEW QUESTION 8**

- (Topic 4)

A company containerized a Windows job that runs on .NET 6 Framework under a Windows container. The company wants to run this job in the AWS Cloud. The job runs every 10 minutes. The job's runtime varies between 1 minute and 3 minutes.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Lambda function based on the container image of the job
- B. Configure Amazon EventBridge to invoke the function every 10 minutes.
- C. Use AWS Batch to create a job that uses AWS Fargate resource
- D. Configure the job scheduling to run every 10 minutes.
- E. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job
- F. Create a scheduled task based on the container image of the job to run every 10 minutes.
- G. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job
- H. Create a standalone task based on the container image of the job
- I. Use Windows task scheduler to run the job every 10 minutes.

**Answer:** A

**Explanation:**

AWS Lambda supports container images as a packaging format for functions. You can use existing container development workflows to package and deploy Lambda functions as container images of up to 10 GB in size. You can also use familiar tools such as Docker CLI to build, test, and push your container images to Amazon Elastic Container Registry (Amazon ECR). You can then create an AWS Lambda function based on the container image of your job and configure Amazon EventBridge to invoke the function every 10 minutes using a cron expression. This solution will be cost-effective as you only pay for the compute time you consume when your function runs. References: <https://docs.aws.amazon.com/lambda/latest/dg/images-create.html>  
<https://docs.aws.amazon.com/eventbridge/latest/userguide/run-lambda-schedule.html>



**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 uses on-premises servers to host its applications. The company is running out of storage capacity. The applications use both block storage and NFS storage. The company needs a high-performing solution that supports local caching without re-architecting its existing applications.

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

- A. Mount Amazon S3 as a file system to the on-premises servers.
- B. Deploy an AWS Storage Gateway file gateway to replace NFS storage.
- C. Deploy AWS Snowball Edge to provision NFS mounts to on-premises servers.
- D. Deploy an AWS Storage Gateway volume gateway to replace the block storage.
- E. Deploy Amazon Elastic File System (Amazon EFS) volumes and mount them to on-premises servers.

**Answer:** BD

**Explanation:**

<https://aws.amazon.com/storagegateway/file/>

File Gateway provides a seamless way to connect to the cloud in order to store application data files and backup images as durable objects in Amazon S3 cloud storage. File Gateway offers SMB or NFS-based access to data in Amazon S3 with local caching. It can be used for on-premises applications, and for Amazon EC2-based applications that need file protocol access to S3 object storage.

<https://aws.amazon.com/storagegateway/volume/>

Volume Gateway presents cloud-backed iSCSI block storage volumes to your on-premises applications. Volume Gateway stores and manages on-premises data in Amazon S3 on your behalf and operates in either cache mode or stored mode. In the cached Volume Gateway mode, your primary data is stored in Amazon S3, while retaining your frequently accessed data locally in the cache for low latency access.

**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 14**

- (Topic 4)

A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.

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

- A. Install an external image management library on an EC2 instance
- B. Use the image management library to process the images.
- C. Create a CloudFront origin request policy
- D. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
- E. Use a Lambda@Edge function with an external image management library
- F. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
- G. Create a CloudFront response headers policy
- H. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

**Answer: C**

**Explanation:**

Lambda@Edge is a service that allows you to run Lambda functions at CloudFront edge locations. It can be used to modify requests and responses that flow through CloudFront. CloudFront origin request policy is a policy that controls the values (URL query strings, HTTP headers, and cookies) that are included in requests that CloudFront sends to the origin. It can be used to collect additional information at the origin or to customize the origin response. CloudFront response headers policy is a policy that specifies the HTTP headers that CloudFront removes or adds in responses that it sends to viewers. It can be used to add security or custom headers to responses.

Based on these definitions, the solution that will meet the requirements with the least operational overhead is:

\* C. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.

This solution would allow the application to use a Lambda@Edge function to resize the images dynamically and serve appropriate formats to clients based on the User-Agent HTTP header in the request. The Lambda@Edge function would run at the edge locations, reducing latency and load on the origin. The application code would only need to include an external image management library that can perform image manipulation tasks.

**NEW QUESTION 17**

- (Topic 4)

A company is creating an application. The company stores data from tests of the application in multiple on-premises locations.

The company needs to connect the on-premises locations to VPCs in an AWS Region in the AWS Cloud. The number of accounts and VPCs will increase during the next year. The network architecture must simplify the administration of new connections and must provide the ability to scale.

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

- A. Create a peering connection between the VPCs. Create a VPN connection between the VPCs and the on-premises locations.
- B. Launch an Amazon EC2 instance. On the instance, include VPN software that uses a VPN connection to connect all VPCs and on-premises locations.
- C. Create a transit gateway. Create VPC attachments for the VPC connections. Create VPN attachments for the on-premises connections.
- D. Create an AWS Direct Connect connection between the on-premises locations and a central VPC.
- E. Connect the central VPC to other VPCs by using peering connections.

**Answer: C**

**Explanation:**

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

References:

? [AWS Transit Gateway](#)

? [Transit gateway attachments](#)

? [Transit gateway route tables](#)

**NEW QUESTION 22**

- (Topic 4)

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

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

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

**Answer: C**

**Explanation:**

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

References:

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

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

**NEW QUESTION 27**

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

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



- I. Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization  
J. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

**Answer:** AE

**Explanation:**

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

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

\* B. Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication. This solution will not meet the requirement of authenticating access to these AWS accounts by using a centralized corporate directory service, as Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications, not for corporate directory services<sup>3</sup>.

\* C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service. This solution will not work, as SCPs are used to restrict the actions that can be performed by the accounts in an organization, not to manage the accounts themselves<sup>1</sup>. Also, IAM Identity Center cannot be added to AWS Directory Service, as it is a separate service that connects with Microsoft Active Directory through AWS Directory Service<sup>2</sup>.

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

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

**NEW QUESTION 35**

- (Topic 4)

An online video game company must maintain ultra-low latency for its game servers. The game servers run on Amazon EC2 instances. The company needs a solution that can

handle millions of UDP internet traffic requests each second.

Which solution will meet these requirements MOST cost-effectively?

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

**Answer:** C

**Explanation:**

The most cost-effective solution for the online video game company is to configure a Network Load Balancer with the required protocol and ports for the internet traffic and specify the EC2 instances as the targets. This solution will enable the company to handle millions of UDP requests per second with ultra-low latency and high performance. A Network Load Balancer is a type of Elastic Load Balancing that operates at the connection level (Layer 4) and routes traffic to targets (EC2 instances, microservices, or containers) within Amazon VPC based on IP protocol data. A Network Load Balancer is ideal for load balancing of both TCP and UDP traffic, as it is capable of handling millions of requests per second while maintaining high throughput at ultra-low latency. A Network Load Balancer also preserves the source IP address of the clients to the back-end applications, which can be useful for logging or security purposes<sup>1</sup>.

**NEW QUESTION 40**

- (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 44**

- (Topic 4)



A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL. The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time. Which solution will meet these requirements with the LEAST operational overhead?

- A. Migrate the web tier and the application tier to Amazon EC2 instances in private subnet
- B. Migrate the database tier to Amazon RDS for MySQL in private subnets.
- C. Migrate the web tier to Amazon EC2 instances in public subnet
- D. Migrate the application tier to EC2 instances in private subnet
- E. Migrate the database tier to Amazon Aurora MySQL in private subnets.
- F. Migrate the web tier to Amazon EC2 instances in public subnet
- G. Migrate the application tier to EC2 instances in private subnet
- H. Migrate the database tier to Amazon RDS for MySQL in private subnets.
- I. Migrate the web tier and the application tier to Amazon EC2 instances in public subnet
- J. Migrate the database tier to Amazon Aurora MySQL in public subnets.

**Answer: C**

**Explanation:**

The solution that meets the requirements with the least operational overhead is to migrate the web tier to Amazon EC2 instances in public subnets, migrate the application tier to EC2 instances in private subnets, and migrate the database tier to Amazon RDS for MySQL in private subnets. This solution allows the company to migrate its three-tier application to AWS by making minimal changes to the architecture, as it preserves the same web, application, and database tiers and uses the same MySQL database engine. The solution also provides a database solution that can restore data to a specific point in time, as Amazon RDS for MySQL supports automated backups and point-in-time recovery. This solution also reduces the operational overhead by using managed services such as Amazon EC2 and Amazon RDS, which handle tasks such as provisioning, patching, scaling, and monitoring.

The other solutions do not meet the requirements as well as the first one because they either involve more changes to the architecture, do not provide point-in-time recovery, or do not follow best practices for security and availability. Migrating the database tier to Amazon Aurora MySQL would require changing the database engine and potentially modifying the application code to ensure compatibility. Migrating the web tier and the application tier to public subnets would expose them to more security risks and reduce their availability in case of a subnet failure. Migrating the database tier to public subnets would also compromise its security and performance. References:

- ? Migrate Your Application Database to Amazon RDS
- ? Amazon RDS for MySQL
- ? Amazon Aurora MySQL
- ? Amazon VPC

**NEW QUESTION 49**

- (Topic 4)

A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent accidental deletion of the documents and ensure that all versions of the documents are available. Users must be able to download, modify, and upload documents. Which combination of actions should be taken to meet these requirements? (Choose two.)

- A. Enable a read-only bucket ACL.
- B. Enable versioning on the bucket.
- C. Attach an IAM policy to the bucket.
- D. Enable MFA Delete on the bucket.
- E. Encrypt the bucket using AWS KMS.

**Answer: BD**

**Explanation:**

Versioning is a feature of Amazon S3 that allows users to keep multiple versions of the same object in a bucket. It can help prevent accidental deletion of the documents and ensure that all versions of the documents are available<sup>1</sup>. MFA Delete is a feature of Amazon S3 that adds an extra layer of security by requiring two forms of authentication to delete a version or change the versioning state of a bucket. It can help prevent unauthorized or accidental deletion of the documents<sup>2</sup>. By enabling both versioning and MFA Delete on the bucket, the solution can meet the requirements.

- \* A. Enable a read-only bucket ACL. This solution will not meet the requirement of allowing users to download, modify, and upload documents, as a read-only bucket ACL will prevent write access to the bucket<sup>3</sup>.
  - \* C. Attach an IAM policy to the bucket. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as an IAM policy is used to grant or deny permissions to users or roles, not to enable versioning or MFA Delete<sup>4</sup>.
  - \* E. Encrypt the bucket using AWS KMS. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as encrypting the bucket using AWS KMS is a method of protecting data at rest, not enabling versioning or MFA Delete.
- Reference URL: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html>

**NEW QUESTION 51**

- (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 56

- (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 60

- (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 63

- (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 67

- (Topic 4)

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

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

**Answer:** D

#### Explanation:

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

The other options are not correct because they do not enable encryption for data in transit or are not relevant for the use case. Enabling IAM database authentication on the database is not correct because this option only provides a method of authentication, not encryption. IAM database authentication allows users to use AWS Identity and Access Management (IAM) users and roles to access a database, instead of using a database user name and password<sup>2</sup>. Providing self-signed certificates is not correct because this option is not secure or reliable. Self-signed certificates are certificates that are signed by the same entity that issued them, instead of by a trusted certificate authority (CA). Self-signed certificates can be easily forged or compromised, and are not recognized by most browsers and applications<sup>3</sup>. Taking a snapshot of the RDS instance and restoring it to a new instance with encryption enabled is not correct because this option only enables encryption at rest, not encryption in transit. Encryption at rest protects data that is stored on disk, but does not protect data that is moving between the client and the server<sup>4</sup>.

References:

- ? Using SSL/TLS to encrypt a connection to a DB instance - Amazon Relational Database Service
- ? IAM database authentication for MySQL and PostgreSQL - Amazon Relational Database Service
- ? What are self-signed certificates?
- ? Encrypting Amazon RDS resources - Amazon Relational Database Service

### NEW QUESTION 69

- (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 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 73

- (Topic 4)

A company operates a two-tier application for image processing. The application uses two Availability Zones, each with one public subnet and one private subnet. An Application Load Balancer (ALB) for the web tier uses the public subnets. Amazon EC2 instances for the application tier use the private subnets. Users report that the application is running more slowly than expected. A security audit of the web server log files shows that the application is receiving millions of illegitimate requests from a small number of IP addresses. A solutions architect needs to resolve the immediate performance problem while the company investigates a more permanent solution. What should the solutions architect recommend to meet this requirement?



- A. Modify the inbound security group for the web tier
- B. Add a deny rule for the IP addresses that are consuming resources.
- C. Modify the network ACL for the web tier subnet
- D. Add an inbound deny rule for the IP addresses that are consuming resources
- E. Modify the inbound security group for the application tier
- F. Add a deny rule for the IP addresses that are consuming resources.
- G. Modify the network ACL for the application tier subnet
- H. Add an inbound deny rule for the IP addresses that are consuming resources

**Answer: B**

**Explanation:**

Deny the request from the first entry at the public subnet, don't allow it to cross and get to the private subnet.

In this scenario, the security audit reveals that the application is receiving millions of illegitimate requests from a small number of IP addresses. To address this issue, it is recommended to modify the network ACL (Access Control List) for the web tier subnets. By adding an inbound deny rule specifically targeting the IP addresses that are consuming resources, the network ACL can block the illegitimate traffic at the subnet level before it reaches the web servers. This will help alleviate the excessive load on the web tier and improve the application's performance.

**NEW QUESTION 77**

- (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 78**

- (Topic 4)

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

Which solution will meet these requirements?

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

**Answer: B**

**Explanation:**

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

References:

? [Security Groups for Your Application Load Balancers](#)

? [Security Groups for Your VPC](#)

**NEW QUESTION 79**

- (Topic 4)

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

What should the solutions architect recommend?

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 83**

- (Topic 4)

A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company's AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (PII) or financial information, including passport numbers and credit card numbers.

Which solution will meet these requirements?

- A. Configure AWS Audit Manager on the account.
- B. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
- C. Configure Amazon S3 Inventory on the S3 bucket.
- D. Configure Amazon Athena to query the inventory.
- E. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
- F. Use Amazon S3 Select to run a report across the S3 bucket.

**Answer: C**

**Explanation:**

Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS. Macie can run data discovery jobs that use managed identifiers for various types of PII or financial information, such as passport numbers and credit card numbers. Macie can also generate findings that alert you to potential issues or risks with your data. References: <https://docs.aws.amazon.com/macie/latest/userguide/macie-identifiers.html>

**NEW QUESTION 86**

- (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 87**

- (Topic 4)

A solutions architect has created two IAM policies: Policy1 and Policy2. Both policies are attached to an IAM group.

## Policy 1

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:Get*",
        "iam:List*",
        "kms:List*",
        "ec2:*",
        "ds:*",
        "logs:Get*",
        "logs:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

## Policy 2

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "ds:Delete*",
      "Resource": "*"
    }
  ]
}
```

A cloud engineer is added as an IAM user to the IAM group. Which action will the cloud engineer be able to perform?

- A. Deleting IAM users
- B. Deleting directories
- C. Deleting Amazon EC2 instances
- D. Deleting logs from Amazon CloudWatch Logs

**Answer: C**

### Explanation:

<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/ds/index.html>

## NEW QUESTION 90

- (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 95

- (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 97

- (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<sup>1</sup>.

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

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

### NEW QUESTION 102

- (Topic 4)

A company wants to use high-performance computing and artificial intelligence to improve its fraud prevention and detection technology. The company requires distributed processing to complete a single workload as quickly as possible.

Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS) and multiple containers.
- B. Use AWS ParallelCluster and the Message Passing Interface (MPI) libraries.
- C. Use an Application Load Balancer and Amazon EC2 instances.
- D. Use AWS Lambda functions.

**Answer:** B

#### Explanation:

AWS ParallelCluster is a service that allows you to create and manage high- performance computing (HPC) clusters on AWS. It supports multiple schedulers, including AWS Batch, which can run distributed workloads across multiple EC2 instances<sup>1</sup>. MPI is a standard for message passing between processes in parallel computing. It provides functions for sending and receiving data, synchronizing processes, and managing communication groups<sup>2</sup>. By using AWS ParallelCluster and MPI libraries, you can take advantage of the following benefits:

- ? You can easily create and configure HPC clusters that meet your specific requirements, such as instance type, number of nodes, network configuration, and storage options<sup>1</sup>.
- ? You can leverage the scalability and elasticity of AWS to run large-scale parallel workloads without worrying about provisioning or managing servers<sup>1</sup>.
- ? You can use MPI libraries to optimize the performance and efficiency of your parallel applications by enabling inter-process communication and data exchange<sup>2</sup>.
- ? You can choose from a variety of MPI implementations that are compatible with AWS ParallelCluster, such as Open MPI, Intel MPI, and MPICH3.

#### NEW QUESTION 104

- (Topic 4)

A company runs a container application by using Amazon Elastic Kubernetes Service (Amazon EKS). The application includes microservices that manage customers and place orders. The company needs to route incoming requests to the appropriate microservices. Which solution will meet this requirement MOST cost-effectively?

- A. Use the AWS Load Balancer Controller to provision a Network Load Balancer.
- B. Use the AWS Load Balancer Controller to provision an Application Load Balancer.
- C. Use an AWS Lambda function to connect the requests to Amazon EKS.
- D. Use Amazon API Gateway to connect the requests to Amazon EKS.

**Answer:** B

#### Explanation:

An Application Load Balancer is a type of Elastic Load Balancer that operates at the application layer (layer 7) of the OSI model. It can distribute incoming traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. It can also route requests based on the content of the request, such as the host name, path, or query parameters<sup>1</sup>.

The AWS Load Balancer Controller is a controller that helps you manage Elastic Load Balancers for your Kubernetes cluster. It can provision Application Load Balancers or Network Load Balancers when you create Kubernetes Ingress or Service resources<sup>2</sup>.

By using the AWS Load Balancer Controller to provision an Application Load Balancer for your Amazon EKS cluster, you can achieve the following benefits:

- ? You can route incoming requests to the appropriate microservices based on the rules you define in your Ingress resource. For example, you can route requests with different host names or paths to different microservices that handle customers and orders<sup>2</sup>.
- ? You can improve the performance and availability of your container applications by distributing the load across multiple targets and enabling health checks and automatic scaling<sup>1</sup>.
- ? You can reduce the cost and complexity of managing your load balancers by using a single controller that integrates with Amazon EKS and Kubernetes. You do not need to manually create or configure load balancers or update them when your cluster changes<sup>2</sup>.

#### NEW QUESTION 107

- (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 108

- (Topic 4)

A company wants to use an event-driven programming model with AWS Lambda. The company wants to reduce startup latency for Lambda functions that run on Java 11. The company does not have strict latency requirements for the applications. The company wants to reduce cold starts and outlier latencies when a function scales up.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure Lambda provisioned concurrency.
- B. Increase the timeout of the Lambda functions.
- C. Increase the memory of the Lambda functions.
- D. Configure Lambda SnapStart.

**Answer:** D

#### Explanation:

To reduce startup latency for Lambda functions that run on Java 11, Lambda SnapStart is a suitable solution. Lambda SnapStart is a feature that enables faster cold starts and lower outlier latencies for Java 11 functions. Lambda SnapStart uses a pre- initialized Java Virtual Machine (JVM) to run the functions, which reduces the initialization time and memory footprint. Lambda SnapStart does not incur any additional charges. References:

? Lambda SnapStart for Java 11 Functions

? Lambda SnapStart FAQs

#### NEW QUESTION 113

- (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 117**

- (Topic 4)

A company hosts an application used to upload files to an Amazon S3 bucket. Once uploaded, the files are processed to extract metadata which takes less than 5 seconds. The volume and frequency of the uploads varies from a few files each hour to hundreds of concurrent uploads. The company has asked a solutions architect to design a cost-effective architecture that will meet these requirements.

What should the solutions architect recommend?

- A. Configure AWS CloudTrail trails to log S3 API calls. Use AWS AppSync to process the files.
- B. Configure an object-created event notification within the S3 bucket to invoke an AWS Lambda function to process the files.
- C. Configure Amazon Kinesis Data Streams to process and send data to Amazon S3. Invoke an AWS Lambda function to process the files.
- D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to process the files uploaded to Amazon S3. Invoke an AWS Lambda function to process the files.

**Answer:** B

**Explanation:**

This option is the most cost-effective and scalable way to process the files uploaded to S3. AWS CloudTrail is used to log API calls, not to trigger actions based on them. AWS AppSync is a service for building GraphQL APIs, not for processing files. Amazon Kinesis Data Streams is used to ingest and process streaming data, not to send data to S3. Amazon SNS is a pub/sub service that can be used to notify subscribers of events, not to process files. References:

? Using AWS Lambda with Amazon S3

? AWS CloudTrail FAQs

? What Is AWS AppSync?

? [What Is Amazon Kinesis Data Streams?]

? [What Is Amazon Simple Notification Service?]

**NEW QUESTION 122**

- (Topic 4)

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

Which solution will meet these requirements?

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

**Answer:** A

**Explanation:**

It allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. References:

? Viewing Events with CloudTrail Event History

? Querying AWS CloudTrail Logs

? Amazon Athena

**NEW QUESTION 123**

- (Topic 4)

A company previously migrated its data warehouse solution to AWS. The company also has an AWS Direct Connect connection. Corporate office users query the



data warehouse using a visualization tool. The average size of a query returned by the data warehouse is 50 MB and each webpage sent by the visualization tool is approximately 500 KB. Result sets returned by the data warehouse are not cached. Which solution provides the LOWEST data transfer egress cost for the company?

- A. Host the visualization tool on premises and query the data warehouse directly over the internet.
- B. Host the visualization tool in the same AWS Region as the data warehouse
- C. Access it over the internet.
- D. Host the visualization tool on premises and query the data warehouse directly over a Direct Connect connection at a location in the same AWS Region.
- E. Host the visualization tool in the same AWS Region as the data warehouse and access it over a Direct Connect connection at a location in the same Region.

**Answer: D**

**Explanation:**

<https://aws.amazon.com/directconnect/pricing/> <https://aws.amazon.com/blogs/aws/aws-data-transfer-prices-reduced/>

**NEW QUESTION 124**

- (Topic 4)

A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company designed the application to work with session affinity (sticky sessions) for a better user experience.

The application must be available publicly over the internet as an endpoint\_ A WAF must be applied to the endpoint for additional security. Session affinity (sticky sessions) must be configured on the endpoint

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

- A. Create a public Network Load Balancer Specify the application target group.
- B. Create a Gateway Load Balancer Specify the application target group.
- C. Create a public Application Load Balancer Specify the application target group.
- D. Create a second target group
- E. Add Elastic IP addresses to the EC2 instances
- F. Create a web ACL in AWS WAF Associate the web ACL with the endpoint

**Answer: CE**

**Explanation:**

C and E are the correct answers because they allow the company to create a public endpoint for its web application that supports session affinity (sticky sessions) and has a WAF applied for additional security. By creating a public Application Load Balancer, the company can distribute incoming traffic across multiple EC2 instances in an Auto Scaling group and specify the application target group. By creating a web ACL in AWS WAF and associating it with the Application Load Balancer, the company can protect its web application from common web exploits. By enabling session stickiness on the Application Load Balancer, the company can ensure that subsequent requests from a user during a session are routed to the same target. References:

? Application Load Balancers

? AWS WAF

? Target Groups for Your Application Load Balancers

? How Application Load Balancer Works with Sticky Sessions

**NEW QUESTION 126**

- (Topic 4)

A global marketing company has applications that run in the ap-southeast-2 Region and the eu-west-1 Region. Applications that run in a VPC in eu-west-1 need to communicate securely with databases that run in a VPC in ap-southeast-2.

Which network design will meet these requirements?

- A. Create a VPC peering connection between the eu-west-1 VPC and the ap-southeast-2 VPC
- B. Create an inbound rule in the eu-west-1 application security group that allows traffic from the database server IP addresses in the ap-southeast-2 security group.
- C. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VPC
- D. Update the subnet route table
- E. Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1.
- F. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VPC
- G. Update the subnet route tables Create an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses.
- H. Create a transit gateway with a peering attachment between the eu-west-1 VPC and the ap-southeast-2 VPC
- I. After the transit gateways are properly peered and routing is configured, create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1.

**Answer: C**

**Explanation:**

"You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC."

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>

**NEW QUESTION 128**

- (Topic 4)

A company has multiple Windows file servers on premises. The company wants to migrate and consolidate its files into an Amazon FSx for Windows File Server file system. File permissions must be preserved to ensure that access rights do not change.

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

- A. Deploy AWS DataSync agents on premise
- B. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- C. Copy the shares on each file server into Amazon S3 buckets by using the AWS CLI Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- D. Remove the drives from each file server Ship the drives to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system

- E. Order an AWS Snowcone device
- F. Connect the device to the on-premises network
- G. Launch AWS DataSync agents on the device
- H. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system,
- I. Order an AWS Snowball Edge Storage Optimized device
- J. Connect the device to the on-premises network
- K. Copy data to the device by using the AWS CLI
- L. Ship the device back to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.

**Answer:** AD

**Explanation:**

A This option involves deploying DataSync agents on your on-premises file servers and using DataSync to transfer the data directly to the FSx for Windows File Server. DataSync ensures that file permissions are preserved during the migration process. D This option involves using an AWS Snowcone device, a portable data transfer device. You would connect the Snowcone device to your on-premises network, launch DataSync agents on the device, and schedule DataSync tasks to transfer the data to FSx for Windows File Server. DataSync handles the migration process while preserving file permissions.

**NEW QUESTION 133**

- (Topic 4)

A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities. The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions. The company needs to improve application performance and decrease latency for the online game in preparation for user growth.

Which solution will meet these requirements?

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

**Answer:** C

**Explanation:**

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

References:

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

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

**NEW QUESTION 134**

- (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 135**

- (Topic 4)

A company uses AWS Organizations. The company wants to operate some of its AWS accounts with different budgets. The company wants to receive alerts and

automatically prevent provisioning of additional resources on AWS accounts when the allocated budget threshold is met during a specific period. Which combination of solutions will meet these requirements? (Select THREE.)

- A. Use AWS Budgets to create a budget
- B. Set the budget amount under the Cost and Usage Reports section of the required AWS accounts.
- C. Use AWS Budgets to create a budget
- D. Set the budget amount under the Billing dashboards of the required AWS accounts.
- E. Create an IAM user for AWS Budgets to run budget actions with the required permissions.
- F. Create an IAM role for AWS Budgets to run budget actions with the required permissions.
- G. Add an alert to notify the company when each account meets its budget threshold
- H. Add a budget action that selects the IAM identity created with the appropriate config rule to prevent provisioning of additional resources.
- I. Add an alert to notify the company when each account meets its budget threshold
- J. Add a budget action that selects the IAM identity created with the appropriate service control policy (SCP) to prevent provisioning of additional resources.

**Answer:** BDF

**Explanation:**

To use AWS Budgets to create and manage budgets for different AWS accounts, the company needs to do the following steps:

1. Use AWS Budgets to create a budget for each AWS account that needs a different budget amount. The budget can be based on cost or usage metrics, and can have different time periods, filters, and thresholds. The company can set the budget amount under the Billing dashboards of the required AWS accounts.

2. Create an IAM role for AWS Budgets to run budget actions with the required permissions. A budget action is a response that AWS Budgets initiates when a budget exceeds a specified threshold. The IAM role allows AWS Budgets to perform actions on behalf of the company, such as applying an IAM policy or a service control policy (SCP) to restrict the provisioning of additional resources.

3. Add an alert to notify the company when each account meets its budget threshold.

The alert can be sent via email or Amazon SNS. The company can also add a budget action that selects the IAM role created and the appropriate SCP to prevent provisioning of additional resources. An SCP is a type of policy that can be applied to an AWS account or an organizational unit (OU) within AWS Organizations.

An SCP can limit the actions that users and roles can perform in the account or OU.

References:

1. <https://aws.amazon.com/budgets/>

2. <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html>

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

4.

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html)

**NEW QUESTION 140**

- (Topic 4)

A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 bucket
- B. Allow access from all the EC2 instances in the VPC.
- C. Create an Amazon Elastic File System (Amazon EFS) file system
- D. Mount the EFS file system from each EC2 instance.
- E. Create a file system on a Provisioned IOPS SSD (PIOPS) Amazon Elastic Block Store (Amazon EBS) volume
- F. Attach the EBS volume to all the EC2 instances.
- G. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance
- H. Synchronize the EBS volumes across the different EC2 instances.

**Answer:** B

**Explanation:**

It allows the EC2 instances to read and write rapidly and concurrently to shared storage across two Availability Zones. Amazon EFS provides a scalable, elastic, and highly available file system that can be mounted from multiple EC2 instances. Amazon EFS supports high levels of throughput and IOPS, and consistent low latencies. Amazon EFS also supports NFSv4 lock upgrading and downgrading, which enables high levels of concurrency. References:

1. [Amazon EFS Features](#)

2. [Using Amazon EFS with Amazon EC2](#)

**NEW QUESTION 143**

- (Topic 4)

A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month.

What is the MOST cost-effective solution to connect these VPCs?

- A. Implement AWS Transit Gateway to connect the VPC
- B. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
- C. Implement an AWS Site-to-Site VPN tunnel between the VPC
- D. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
- E. Set up a VPC peering connection between the VPC
- F. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
- G. Set up a 1 GB AWS Direct Connect connection between the VPC
- H. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

**Answer:** C

**Explanation:**

To connect two VPCs in the same Region within the same AWS account, VPC peering is the most cost-effective solution. VPC peering allows direct network traffic between the VPCs without requiring a gateway, VPN connection, or AWS Transit Gateway. VPC peering also does not incur any additional charges for data transfer between the VPCs.



References:

- ? What Is VPC Peering?
- ? VPC Peering Pricing

#### NEW QUESTION 147

- (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 Zone
- 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 151

- (Topic 4)

A company is running its production and nonproduction environment workloads in multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to design a solution that will prevent the modification of cost usage tags. Which solution will meet these requirements?

- A. Create a custom AWS Config rule to prevent tag modification except by authorized principals.
- B. Create a custom trail in AWS CloudTrail to prevent tag modification
- C. Create a service control policy (SCP) to prevent tag modification except by authorized principals.
- D. Create custom Amazon CloudWatch logs to prevent tag modification.

**Answer: C**

#### Explanation:

This solution meets the requirements because it uses SCPs to restrict the actions that can be performed on cost usage tags in the organization. SCPs are a type of organization policy that you can use to manage permissions in your organization. SCPs specify the maximum permissions for an organization, organizational unit (OU), or account. You can use SCPs to enforce consistent tag policies across your organization and prevent unauthorized or accidental changes to your tags. You can also create exceptions for authorized principals, such as administrators or auditors, who need to modify tags for legitimate purposes.

References:

- ? Service control policies (SCPs) - AWS Organizations
- ? Tag policies - AWS Organizations

#### NEW QUESTION 154

- (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 155

- (Topic 4)

An ecommerce company runs applications in AWS accounts that are part of an organization in AWS Organizations. The applications run on Amazon Aurora PostgreSQL databases across all the accounts. The company needs to prevent malicious activity and must identify abnormal failed and incomplete login attempts to the databases.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts.
- B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization.
- C. Publish the Aurora general logs to a log group in Amazon CloudWatch Logs. Export the log data to a central Amazon S3 bucket.
- D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket.

**Answer: C**

#### Explanation:

This option is the most operationally efficient way to meet the requirements because it allows the company to monitor and analyze the database login activity across all the accounts in the organization. By publishing the Aurora general logs to a log group in Amazon CloudWatch Logs, the company can enable the logging of the database connections, disconnections, and failed authentication attempts. By exporting the log data to a central Amazon S3 bucket, the company can store the log data in a durable and cost-effective way and use other AWS services or tools to perform further analysis or alerting on the log data. For example, the company can use Amazon Athena to query the log data in Amazon S3, or use Amazon SNS to send notifications based on the log data.

\* A. Attach service control policies (SCPs) to the root of the organization to identify the failed login attempts. This option is not effective because SCPs are not designed to identify the failed login attempts, but to restrict the actions that the users and roles can perform in the member accounts of the organization. SCPs are applied to the AWS API calls, not to the database login attempts. Moreover, SCPs do not provide any logging or analysis capabilities for the database activity.

\* B. Enable the Amazon RDS Protection feature in Amazon GuardDuty for the member accounts of the organization. This option is not optimal because the Amazon RDS Protection feature in Amazon GuardDuty is not available for Aurora PostgreSQL databases, but only for Amazon RDS for MySQL and Amazon RDS for MariaDB databases. Moreover, the Amazon RDS Protection feature does not monitor the database login attempts, but the network and API activity related to the RDS instances.

\* D. Publish all the Aurora PostgreSQL database events in AWS CloudTrail to a central Amazon S3 bucket. This option is not sufficient because AWS CloudTrail does not capture the database login attempts, but only the AWS API calls made by or on behalf of the Aurora PostgreSQL database. For example, AWS CloudTrail can record the events such as creating, modifying, or deleting the database instances, clusters, or snapshots, but not the events such as connecting, disconnecting, or failing to authenticate to the database. References:

- ? 1 Working with Amazon Aurora PostgreSQL - Amazon Aurora
- ? 2 Working with log groups and log streams - Amazon CloudWatch Logs
- ? 3 Exporting Log Data to Amazon S3 - Amazon CloudWatch Logs
- ? [4] Amazon GuardDuty FAQs
- ? [5] Logging Amazon RDS API Calls with AWS CloudTrail - Amazon Relational Database Service

#### NEW QUESTION 157

- (Topic 4)

A company needs to provide customers with secure access to its data. The company processes customer data and stores the results in an Amazon S3 bucket. All the data is subject to strong regulations and security requirements. The data must be encrypted at rest. Each customer must be able to access only their data from their AWS account. Company employees must not be able to access the data.

Which solution will meet these requirements?

- A. Provision an AWS Certificate Manager (ACM) certificate for each customer.
- B. Encrypt the data client-side.
- C. In the private certificate policy, deny access to the certificate for all principals except an IAM role that the customer provides.
- D. Provision a separate AWS Key Management Service (AWS KMS) key for each customer.
- E. Encrypt the data server-side.
- F. In the S3 bucket policy, deny decryption of data for all principals except an IAM role that the customer provides.
- G. Provision a separate AWS Key Management Service (AWS KMS) key for each customer.
- H. Encrypt the data server-side.
- I. In each KMS key policy, deny decryption of data for all principals except an IAM role that the customer provides.
- J. Provision an AWS Certificate Manager (ACM) certificate for each customer.
- K. Encrypt the data client-side.
- L. In the public certificate policy, deny access to the certificate for all principals except an IAM role that the customer provides.

**Answer: C**

#### Explanation:

The correct solution is to provision a separate AWS KMS key for each customer and encrypt the data server-side. This way, the company can use the S3 encryption feature to protect the data at rest and delegate the control of the encryption keys to the customers. The customers can then use their own IAM roles to access and decrypt their data. The company employees will not be able to access the data because they are not authorized by the KMS key policies. The other options are incorrect because:

? Option A and D are using ACM certificates to encrypt the data client-side. This is

not a recommended practice for S3 encryption because it adds complexity and overhead to the encryption process. Moreover, the company will have to manage the certificates and their policies for each customer, which is not scalable and secure.

? Option B is using a separate KMS key for each customer, but it is using the S3

bucket policy to control the decryption access. This is not a secure solution because the bucket policy applies to the entire bucket, not to individual objects.

Therefore, the customers will be able to access and decrypt each other's data if they have the permission to list the bucket contents. The bucket policy also overrides the KMS key policy, which means the company employees can access the data if they have the permission to use the KMS key.

References:

- ? S3 encryption
- ? KMS key policies

? ACM certificates

#### NEW QUESTION 159

- (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 Autopilot
- 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 content
- 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 content
- F. Associate the function with the web application.
- G. Create an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content
- 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 164

- (Topic 4)

A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS). The company's workload is not consistent throughout the day. The company wants Amazon EKS to scale in and out according to the workload.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. Use an AWS Lambda function to resize the EKS cluster
- B. Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.
- C. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
- D. Use Amazon API Gateway and connect it to Amazon EKS
- E. Use AWS App Mesh to observe network activity.

**Answer: BC**

#### Explanation:

<https://docs.aws.amazon.com/eks/latest/userguide/horizontal-pod-autoscaler.html> <https://docs.aws.amazon.com/eks/latest/userguide/autoscaling.html>

Horizontal pod autoscaling is a feature of Kubernetes that automatically scales the number of pods in a deployment, replication controller, or replica set based on that resource's CPU utilization. It requires a metrics source such as the Kubernetes Metrics Server to provide CPU usage data<sup>1</sup>. Cluster autoscaling is a feature of Kubernetes that automatically adjusts the number of nodes in a cluster when pods fail or are rescheduled onto other nodes. It requires an integration with AWS Auto Scaling groups to manage the EC2 instances that join the cluster<sup>2</sup>. By using both horizontal pod autoscaling and cluster autoscaling, the solution can ensure that Amazon EKS scales in and out according to the workload.

#### NEW QUESTION 166

- (Topic 4)

A company uses high concurrency AWS Lambda functions to process a constantly increasing number of messages in a message queue during marketing events. The Lambda functions use CPU intensive code to process the messages. The company wants to reduce the compute costs and to maintain service latency for its customers.

Which solution will meet these requirements?

- A. Configure reserved concurrency for the Lambda function
- B. Decrease the memory allocated to the Lambda functions.
- C. Configure reserved concurrency for the Lambda function
- D. Increase the memory according to AWS Compute Optimizer recommendations.
- E. Configure provisioned concurrency for the Lambda function
- F. Decrease the memory allocated to the Lambda functions.
- G. Configure provisioned concurrency for the Lambda function
- H. Increase the memory according to AWS Compute Optimizer recommendations.

**Answer: D**



**Explanation:**

The company wants to reduce the compute costs and maintain service latency for its Lambda functions that process a constantly increasing number of messages in a message queue. The Lambda functions use CPU intensive code to process the messages. To meet these requirements, a solutions architect should recommend the following solution:

? Configure provisioned concurrency for the Lambda functions. Provisioned

concurrency is the number of pre-initialized execution environments that are allocated to the Lambda functions. These execution environments are prepared to respond immediately to incoming function requests, reducing the cold start latency. Configuring provisioned concurrency also helps to avoid throttling errors due to reaching the concurrency limit of the Lambda service.

? Increase the memory according to AWS Compute Optimizer recommendations.

AWS Compute Optimizer is a service that provides recommendations for optimal AWS resource configurations based on your utilization data. By increasing the memory allocated to the Lambda functions, you can also increase the CPU power and improve the performance of your CPU intensive code. AWS Compute Optimizer can help you find the optimal memory size for your Lambda functions based on your workload characteristics and performance goals.

This solution will reduce the compute costs by avoiding unnecessary over-provisioning of memory and CPU resources, and maintain service latency by using provisioned concurrency and optimal memory size for the Lambda functions.

References:

? Provisioned Concurrency

? AWS Compute Optimizer

**NEW QUESTION 170**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

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