# AWS-Solution-Architect-Associate Dumps

# Amazon AWS Certified Solutions Architect - Associate

## https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html

**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 MySQ
B. Configureelastic storage scaling.
C. Migrate the database to Amazon Redshift by using the mysqldump utilit
D. Turn on Auto Scaling for the Amazon Redshift cluster.
E. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Auror
F. Turn on Aurora Auto Scaling.
G. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoD
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 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 3**
- (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 service1. Amazon S3 does not support gateway endpoints, only interface endpoints2. 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 service1. 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 data3. 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 4**
- (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 do1. Identity-based policies are different from resource-based policies, which define the permissions around the specific resource1. Resource-based policies are attached to a resource, such as an Amazon S3 bucket or an Amazon EC2 instance1. Resource-based policies can also specify a principal, which is the entity that is allowed or denied access to the resource1. Organization is not an IAM principal, but a feature of AWS Organizations that allows you to manage multiple AWS accounts centrally2. Amazon ECS resource and Amazon EC2 resource are not IAM principals, but AWS resources that can have resource-based policies attached to them34. References:
? Identity-based policies and resource-based policies
? AWS Organizations
? Amazon ECS task role
? Amazon EC2 instance profile

**NEW QUESTION 5**
- (Topic 4)
A company needs to configure a real-time data ingestion architecture for its application. The company needs an API. a process that transforms data as the data is streamed, and a storage solution for the data.
Which solution will meet these requirements with the LEAST operational overhead?

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

**Answer:** C

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

**NEW QUESTION 6**
- (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 jo
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 jo
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 jo
H. Create a standalone task based on the container image of the jo
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 7**
- (Topic 4)
A company needs to migrate a MySQL database from its on-premises data center to AWS within 2 weeks. The database is 20 TB in size. The company wants to complete the migration with minimal downtime.
Which solution will migrate the database MOST cost-effectively?

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

**Answer:** A

**Explanation:**
 This answer is correct because it meets the requirements of migrating a 20 TB MySQL database within 2 weeks with minimal downtime and cost-effectively. The AWS Snowball Edge Storage Optimized device has up to 80 TB of usable storage space, which is enough to fit the database. The AWS Database Migration Service (AWS DMS) can migrate data from MySQL to Amazon Aurora, Amazon RDS for MySQL, or MySQL on Amazon EC2 with minimal downtime by continuously replicating changes from the source to the target. The AWS Schema Conversion Tool (AWS SCT) can convert the source schema and code to a format compatible with the target database. By using these services together, the company can migrate the database to AWS with minimal downtime and cost. The Snowball Edge device can be shipped back to AWS to finish the migration and continue the ongoing replication until the database is fully migrated.
References:
? https://docs.aws.amazon.com/snowball/latest/developer-guide/device- differences.html
? https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.MySQL.html
? https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/CHAP_So urce.MySQL.htm


**NEW QUESTION 8**
- (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 endpoin
B. Deploy the Lambda functions inside a VPC.
C. Point the client driver at an RDS proxy endpoin
D. Deploy the Lambda functions inside a VPC.
E. Point the client driver at an RDS custom endpoin
F. Deploy the Lambda functions outside a VPC.
G. Point the client driver at an RDS proxy endpoin
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 9**
- (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 instanc
B. Use the image management library to process the images.
C. Create a CloudFront origin request polic
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 librar
F. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
G. Create a CloudFront response headers polic
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 tasks1.

**NEW QUESTION 10**
- (Topic 4)
An image hosting company uploads its large assets to Amazon S3 Standard buckets The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.
Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

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

**Answer:** AB

**Explanation:**
 S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead1. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.
S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle2. You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage costs and avoid paying for parts that are not used.
Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs3. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.
Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.
Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed1. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL: 1: https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html 2: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html 3: https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty- bucket.html#delete-bucket-considerations : https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html : https://aws.amazon.com/certification/certified-solutions-architect-associate/

**NEW QUESTION 10**
- (Topic 4)
A company is creating an application The company stores data from tests of the application in multiple on-premises locations
The company needs to connect the on-premises locations to VPCs in an AWS Region in the AWS Cloud The number of accounts and VPCs will increase during the next year The network architecture must simplify the administration of new connections and must provide the ability to scale.
Which solution will meet these requirements with the LEAST administrative overhead'?

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

**Answer:** C

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

**NEW QUESTION 12**
- (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/Introducti on.html
? https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html

**NEW QUESTION 13**
- (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 traffi
B. Specify the EC2 instances as the targets.
C. Configure a Gateway Load Balancer for the internet traffi
D. Specify the EC2 instances as the targets.
E. Configure a Network Load Balancer with the required protocol and ports for the internet traffi
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 purposes1.

**NEW QUESTION 14**
- (Topic 4)
A company is implementing new data retention policies for all databases that run on Amazon RDS DB instances. The company must retain daily backups for a minimum period of 2 years. The backups must be consistent and restorable.
Which solution should a solutions architect recommend to meet these requirements?

A. Create a backup vault in AWS Backup to retain RDS backup
B. Create a new backup plan with a daily schedule and an expiration period of 2 years after creatio
C. Assign the RDS DB instances to the backup plan.

D. Configure a backup window for the RDS DB instances for daily snapshot
E. Assign a snapshot retention policy of 2 years to each RDS DB instanc
F. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule snapshot deletions.
G. Configure database transaction logs to be automatically backed up to Amazon CloudWatch Logs with an expiration period of 2 years.
H. Configure an AWS Database Migration Service (AWS DMS) replication tas
I. Deploy a replication instance, and configure a change data capture (CDC) task to stream database changes to Amazon S3 as the targe
J. Configure S3 Lifecycle policies to delete the snapshots after 2 years.

**Answer:** A

**Explanation:**
AWS Backup is a fully managed service that enables users to centralize and automate the backup of data across AWS services. It can create and manage backup plans that specify the frequency and retention period of backups. It can also assign backup resources to backup vaults, which are containers that store backup data1. By using AWS Backup, the solution can ensure that the RDS backups are consistent, restorable, and retained for a minimum period of 2 years.
* B. Configure a backup window for the RDS DB instances for daily snapshots. Assign a snapshot retention policy of 2 years to each RDS DB instance. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule snapshot deletions. This solution will not meet the requirement of ensuring that the backups are consistent and restorable, as Amazon DLM is not compatible with RDS snapshots and cannot be used to schedule snapshot deletions2.
* C. Configure database transaction logs to be automatically backed up to Amazon CloudWatch Logs with an expiration period of 2 years. This solution will not meet the requirement of ensuring that the backups are consistent and restorable, as database transaction logs are not sufficient to restore a database to a point in time. They only capture the changes made to the database, not the full state of the database3.
* D. Configure an AWS Database Migration Service (AWS DMS) replication task. Deploy a replication instance, and configure a change data capture (CDC) task to stream database changes to Amazon S3 as the target. Configure S3 Lifecycle policies to delete the snapshots after 2 years. This solution will not meet the requirement of ensuring that the backups are consistent and restorable, as AWS DMS is a service that helps users migrate databases to AWS, not back up databases. It also requires additional resources and configuration, such as replication instances and CDC tasks.
Reference URL: https://docs.aws.amazon.com/aws- backup/latest/devguide/whatisbackup.html

**NEW QUESTION 17**
- (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 22**
- (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 1AM 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 instanc
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 connection1.

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 password2. 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 applications3. 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 server4.

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 24**
- (Topic 4)
A company used an Amazon RDS for MySQL DB instance during application testing. Before terminating the DB instance at the end of the test cycle, a solutions architect created two backups. The solutions architect created the first backup by using the mysqldump utility to create a database dump. The solutions architect created the second backup by enabling the final DB snapshot option on RDS termination.

The company is now planning for a new test cycle and wants to create a new DB instance from the most recent backup. The company has chosen a MySQL-compatible edition of Amazon Aurora to host the DB instance.
Which solutions will create the new DB instance? (Select TWO.)

A. Import the RDS snapshot directly into Aurora.
B. Upload the RDS snapshot to Amazon S3. Then import the RDS snapshot into Aurora.
C. Upload the database dump to Amazon S3. Then import the database dump into Aurora.
D. Use AWS Database Migration Service (AWS DMS) to import the RDS snapshot into Aurora.
E. Upload the database dump to Amazon S3. Then use AWS Database Migration Service (AWS DMS) to import the database dump into Aurora.

**Answer:** AC

**Explanation:**
These answers are correct because they meet the requirements of creating a new DB instance from the most recent backup and using a MySQL-compatible edition of Amazon Aurora to host the DB instance. You can import the RDS snapshot directly into Aurora if the MySQL DB instance and the Aurora DB cluster are running the same version of MySQL. For example, you can restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.6, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is simple and requires the fewest number of steps. You can upload the database dump to Amazon S3 and then import the database dump into Aurora if the MySQL DB instance and the Aurora DB cluster are running different versions of MySQL. For example, you can import a MySQL version 5.6 database dump into Aurora MySQL version 5.7, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is more flexible and allows you to migrate across different versions of MySQL.
References:
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Import.html
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Dump.html


**NEW QUESTION 29**
- (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 tie
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 tie
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, dont 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 32**
- (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 subne
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 37**
- (Topic 4)
A company has deployed a multiplayer game for mobile devices. The game requires live
location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.
The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.
What should a solutions architect do to improve the performance of the data tier?

A. Take a snapshot of the existing DB instanc
B. Restore the snapshot with Multi-AZ enabled.
C. Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
D. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instanc
E. Modify the game to use DAX.
F. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instanc
G. Modify the game to use Redis.

**Answer:** D

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

**NEW QUESTION 41**
- (Topic 4)
A company runs demonstration environments for its customers on Amazon EC2 instances. Each environment is isolated in its own VPC. The company's
operations team needs to be notified when RDP or SSH access to an environment has been established.

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

**Answer:** C

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

**NEW QUESTION 44**
- (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 differentAvailability 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 48**
- (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 49**
- (Topic 4)
A company built an application with Docker containers and needs to run the application in the AWS Cloud The company wants to use a managed sen/ice to host the application
The solution must scale in and out appropriately according to demand on the individual container services The solution also must not result in additional operational overhead or infrastructure to manage
Which solutions will meet these requirements? (Select TWO)

A. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
B. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate.

C. Provision an Amazon API Gateway API Connect the API to AWS Lambda to run thecontainers.
D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes.
E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes.

**Answer:** AB

**Explanation:**
 These options are the best solutions because they allow the company to run the application with Docker containers in the AWS Cloud using a managed service that scales automatically and does not require any infrastructure to manage. By using AWS Fargate, the company can launch and run containers without having to provision, configure, or scale clusters of EC2 instances. Fargate allocates the right amount of compute resources for each container and scales them up or down as needed. By using Amazon ECS or Amazon EKS, the company can choose the container orchestration platform that suits its needs. Amazon ECS is a fully managed service that integrates with other AWS services and simplifies the deployment and management of containers. Amazon EKS is a managed service that runs Kubernetes on AWS and provides compatibility with existing Kubernetes tools and plugins.
* C. Provision an Amazon API Gateway API Connect the API to AWS Lambda to run the containers. This option is not feasible because AWS Lambda does not support running Docker containers directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run Docker containers on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the Docker API, which can introduce additional complexity and overhead.
* D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes. This option is not optimal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.
* E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.
References:
? 1 AWS Fargate - Amazon Web Services
? 2 Amazon Elastic Container Service - Amazon Web Services
? 3 Amazon Elastic Kubernetes Service - Amazon Web Services
? 4 AWS Lambda FAQs - Amazon Web Services

**NEW QUESTION 50**
- (Topic 4)
A company runs a three-tier application in two AWS Regions. The web tier, the application tier, and the database tier run on Amazon EC2 instances. The company uses Amazon RDS for Microsoft SQL Server Enterprise for the database tier The database tier is experiencing high load when weekly and monthly reports are run. The company wants to reduce the load on the database tier.
Which solution will meet these requirements with the LEAST administrative effort?

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

**Answer:** A

**Explanation:**
 it allows the company to create read replicas of its RDS database and reduce the load on the database tier. By creating read replicas, the company can offload read traffic from the primary database instance to one or more replicas. By configuring the reports to use the new read replicas, the company can improve performance and availability of its database tier. References:
? Working with Read Replicas
? Read Replicas for Amazon RDS for SQL Server

**NEW QUESTION 52**
- (Topic 4)
A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users.
What should a solutions architect recommend?

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

**Answer:** B

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

**NEW QUESTION 53**
- (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 parameters1.
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 resources2.
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 orders2.
? You can improve the performance and availability of your container applications by
distributing the load across multiple targets and enabling health checks and automatic scaling1.
? 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 changes2.

**NEW QUESTION 57**
- (Topic 4)
A company needs to minimize the cost of its 1 Gbps AWS Direct Connect connection. The company's average connection utilization is less than 10%. A solutions architect must recommend a solution that will reduce the cost without compromising security.
Which solution will meet these requirements?

A. Set up a new 1 Gbps Direct Connect connectio
B. Share the connection with another AWS account.
C. Set up a new 200 Mbps Direct Connect connection in the AWS Management Console.
D. Contact an AWS Direct Connect Partner to order a 1 Gbps connectio
E. Share the connection with another AWS account.
F. Contact an AWS Direct Connect Partner to order a 200 Mbps hosted connection for an existing AWS account.

**Answer:** D

**Explanation:**
 company need to setup a cheaper connection (200 M) but B is incorrect because you can only order port speeds of 1, 10, or 100 Gbps for more flexibility you can go with hosted connection, You can order port speeds between 50 Mbps and 10 Gbps. https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct- connect.html

**NEW QUESTION 60**
- (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 62**
- (Topic 4)
A gaming company wants to launch a new internet-facing application in multiple AWS Regions The application will use the TCP and UDP protocols for communication. The company needs to provide high availability and minimum latency for global users.
Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

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

**Answer:** BC

**Explanation:**
 This combination of actions will provide high availability and minimum latency for global users by using AWS Global Accelerator and Application Load Balancers. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your internet-facing applications by using the AWS global network. It provides two global static public IPs that act as a fixed entry point to your application endpoints, such as Application Load Balancers, in

multiple Regions1. Global Accelerator uses the AWS backbone network to route traffic to the optimal regional
endpoint based on health, client location, and policies that you configure. It also offers TCP and UDP support, traffic encryption, and DDoS protection2. Application Load Balancers are external load balancers that distribute incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. They support both HTTP and HTTPS (SSL/TLS) protocols, and offer advanced features such as content-based routing, health checks, and integration with other AWS services3. By creating external Application Load Balancers in front of the application in each Region, you can ensure that the application can handle varying load patterns and scale on demand. By creating an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region, you can leverage the performance, security, and availability of the AWS global network to deliver the best possible user experience.
References: 1: What is AWS Global Accelerator? - AWS Global Accelerator4, Overview section2: Network Acceleration Service - AWS Global Accelerator - AWS5, Why AWS Global Accelerator? section. 3: What is an Application Load Balancer? - Elastic Load Balancing6, Overview section.

**NEW QUESTION 67**
- (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 laye
B. Use AWS Lambda functions for the application laye
C. Move the database to an Amazon DynamoDB tabl
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 laye
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 laye
H. Use a fleet of EC2 instances in an Auto Scaling group for the application laye
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 laye
K. Move the database to an Amazon RDS Multi-AZ DB instanc
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 69**
- (Topic 4)
A company runs multiple workloads in its on-premises data center. The company's data center cannot scale fast enough to meet the company's expanding business needs. The company wants to collect usage and configuration data about the on-premises servers and workloads to plan a migration to AWS.
Which solution will meet these requirements?

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

**Answer:** B

**Explanation:**
The most suitable solution for the company's requirements is to set the home AWS Region in AWS Migration Hub and use AWS Application Discovery Service to collect data about the on-premises servers. This solution will enable the company to gather usage and configuration data of its on-premises servers and workloads, and plan a migration to AWS.
AWS Migration Hub is a service that simplifies and accelerates migration tracking by aggregating migration status information into a single console. Users can view the discovered servers, group them into applications, and track the migration status of each application from the Migration Hub console in their home Region. The home Region is the AWS Region where users store their migration data, regardless of which Regions they migrate into1.
AWS Application Discovery Service is a service that helps users plan their migration to AWS by collecting usage and configuration data about their on-premises servers and databases. Application Discovery Service is integrated with AWS Migration Hub and supports two methods of performing discovery: agentless discovery and agent-based discovery. Agentless discovery can be performed by deploying the Application Discovery Service Agentless Collector through VMware vCenter, which collects static configuration data and utilization data for virtual machines (VMs) and databases. Agent-based discovery can be performed by deploying the AWS Application Discovery Agent on each of the VMs and physical servers, which collects static configuration data, detailed time-series system-performance information, inbound and outbound network connections, and processes that are running2.
The other options are not correct because they do not meet the requirements or are not relevant for the use case. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Trusted Advisor to collect data about the on- premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another, such as from Oracle to PostgreSQL3. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits4. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. As mentioned above, AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another. AWS DMS is a service that helps users migrate relational databases, non-relational databases, and other types of data stores to
AWS with minimal downtime5. References:
? Home Region - AWS Migration Hub
? What is AWS Application Discovery Service? - AWS Application Discovery Service
? AWS Schema Conversion Tool - Amazon Web Services
? What Is Trusted Advisor? - Trusted Advisor
? What Is AWS Database Migration Service? - AWS Database Migration Service

**NEW QUESTION 70**
- (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 tog 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 75**
- (Topic 4)
A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers a previous employee did not provide details about the resources invent^. The solutions architect needs to build and map the relationship details of the various workloads across all accounts.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report.
B. Use AWS Step Functions to collect workload details Build architecture diagrams of theworkloads manually.
C. Use Workload Discovery on AWS to generate architecture diagrams of the workloads.
D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships

**Answer:** C

**Explanation:**
Workload Discovery on AWS (formerly called AWS Perspective) is a tool that visualizes AWS Cloud workloads. It maintains an inventory of the AWS resources across your accounts and Regions, mapping relationships between them, and displaying them in a web UI. It also allows you to query AWS Cost and Usage Reports, search for resources, save and export architecture diagrams, and more1. By using Workload Discovery on AWS, the solution can build and map the relationship details of the various workloads across all accounts with the least operational effort.
* A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report. This solution will not meet the requirement of building and mapping the relationship details of the various workloads across all accounts, as AWS Systems Manager Inventory is a feature that collects metadata from your managed instances and stores it in a central Amazon S3 bucket. It does not provide a map view or architecture diagrams of the workloads2.
* B. Use AWS Step Functions to collect workload details Build architecture diagrams of the work-loads manually. This solution will not meet the requirement of the least operational effort, as it involves creating and managing state machines to orchestrate the workload details collection, and building architecture diagrams manually.
* D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships. This solution will not meet the requirement of the least operational effort, as it involves instrumenting your applications with X-Ray SDKs to collect workload details, and building architecture diagrams manually.
Reference URL: https://aws.amazon.com/solutions/implementations/workload-discovery- on-aws/

**NEW QUESTION 78**
- (Topic 4)
A company is deploying an application that processes streaming data in near-real time The company plans to use Amazon EC2 instances for the workload The network architecture must be configurable to provide the lowest possible latency between nodes
Which combination of network solutions will meet these requirements? (Select TWO)

A. Enable and configure enhanced networking on each EC2 instance
B. Group the EC2 instances in separate accounts
C. Run the EC2 instances in a cluster placement group
D. Attach multiple elastic network interfaces to each EC2 instance
E. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

**Answer:** AC

**Explanation:**
 These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth, lower latency, and lower jitter. Enhanced networking uses single root I/O virtualization (SR-IOV) or Elastic Fabric Adapter (EFA) to provide direct access to the network hardware. You can enable and configure enhanced networking by choosing a supported instance type and a compatible operating system, and installing the required drivers. Option C runs the EC2 instances in a cluster placement group, which is a logical grouping of instances within a single Availability Zone that are placed close together on the same underlying hardware. Cluster placement groups provide the lowest network latency and the highest network throughput among the placement group options. You can run the EC2 instances in a cluster placement group by creating a placement group and launching the instances into it. Option B is not suitable because grouping the EC2 instances in separate accounts does not provide the lowest possible latency between nodes. Separate accounts are used to isolate and organize resources for different purposes, such as security, billing, or compliance. However, they do not affect the network performance or proximity of the instances. Moreover, grouping the EC2 instances in separate accounts would incur additional costs and complexity, and it would require setting up cross-account networking and permissions.
Option D is not suitable because attaching multiple elastic network interfaces to each EC2 instance does not provide the lowest possible latency between nodes.

Elastic network interfaces are virtual network interfaces that can be attached to EC2 instances to provide additional network capabilities, such as multiple IP addresses, multiple subnets, or enhanced security. However, they do not affect the network performance or proximity of the instances. Moreover, attaching multiple elastic network interfaces to each EC2 instance would consume additional resources and limit the instance type choices.
Option E is not suitable because using Amazon EBS optimized instance types does not provide the lowest possible latency between nodes. Amazon EBS optimized instance types are instances that provide dedicated bandwidth for Amazon EBS volumes, which are block storage volumes that can be attached to EC2 instances. EBS optimized instance types improve the performance and consistency of the EBS volumes, but they do not affect the network performance or proximity of the instances. Moreover, using EBS optimized instance types would incur additional costs and may not be necessary for the streaming data workload.
References:
? Enhanced networking on Linux
? Placement groups
? Elastic network interfaces
? Amazon EBS-optimized instances

## NEW QUESTION 79
- (Topic 4)
A company has a production workload that is spread across different AWS accounts in various AWS Regions. The company uses AWS Cost Explorer to continuously monitor costs and usage. The company wants to receive notifications when the cost and usage spending of the workload is unusual.
Which combination of steps will meet these requirements? (Select TWO.)

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

**Answer:** BE

**Explanation:**
AWS Cost Anomaly Detection allows you to create monitors that track the cost and usage of your AWS resources and alert you when there is an unusual spending pattern. You can create monitors based on different dimensions, such as AWS services, accounts, tags, or cost categories. You can also create alert subscriptions that notify you by email or Amazon SNS when an anomaly is detected. You can specify the threshold and frequency of the alerts, and choose to receive weekly summaries of your anomalies. Reference URLs:
1 https://aws.amazon.com/aws-cost-management/aws-cost-anomaly-detection/
2 https://docs.aws.amazon.com/cost-management/latest/userguide/getting-started-ad.html
3 https://docs.aws.amazon.com/cost-management/latest/userguide/manage-ad.html

## NEW QUESTION 84
- (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 warehous
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 88
- (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 VP
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 VP
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 VP
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 VP
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 92**
- (Topic 4)
A company is running a photo hosting service in the us-east-1 Region. The service enables users across multiple countries to upload and view photos. Some photos are heavily viewed for months, and others are viewed for less than a week. The application allows uploads of up to 20 MB for each photo. The service uses the photo metadata to determine which photos to display to each user.
Which solution provides the appropriate user access MOST cost-effectively?

A. Store the photos in Amazon DynamoD
B. Turn on DynamoDB Accelerator (DAX) to cache frequently viewed items.
C. Store the photos in the Amazon S3 Intelligent-Tiering storage clas
D. Store the photo metadata and its S3 location in DynamoDB.
E. Store the photos in the Amazon S3 Standard storage clas
F. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage clas
G. Use the object tags to keep track of metadata.
H. Store the photos in the Amazon S3 Glacier storage clas
I. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Glacier Deep Archive storage clas
J. Store the photo metadata and its S3 location in Amazon OpenSearch Service.

**Answer:** B

**Explanation:**
This solution provides the appropriate user access most cost-effectively
because it uses the Amazon S3 Intelligent-Tiering storage class, which automatically optimizes storage costs by moving data to the most cost-effective access tier when access patterns change, without performance impact or operational overhead1. This storage class is ideal for data with unknown, changing, or unpredictable access patterns, such as photos that are heavily viewed for months or less than a week. By storing the photo metadata and its S3 location in DynamoDB, the application can quickly query and retrieve the relevant photos for each user. DynamoDB is a fast, scalable, and fully managed NoSQL database service that supports key-value and document data models2.
References: 1: Amazon S3 Intelligent-Tiering Storage Class | AWS3, Overview section2: Amazon DynamoDB - NoSQL Cloud Database Service4, Overview section.

**NEW QUESTION 95**
- (Topic 4)
A company stores critical data in Amazon DynamoDB tables in the company's AWS account. An IT administrator accidentally deleted a DynamoDB table. The deletion caused a significant loss of data and disrupted the company's operations. The company wants to prevent this type of disruption in the future.
Which solution will meet this requirement with the LEAST operational overhead?

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

**Answer:** C

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

**NEW QUESTION 96**
- (Topic 4)
A social media company wants to allow its users to upload images in an application that is hosted in the AWS Cloud. The company needs a solution that automatically resizes the images so that the images can be displayed on multiple device types. The application experiences unpredictable traffic patterns throughout the day. The company is seeking a highly available solution that maximizes scalability.
What should a solutions architect do to meet these requirements?

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

**Answer:** A

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

**NEW QUESTION 101**
- (Topic 4)
A company is designing a new web application that will run on Amazon EC2 Instances. The application will use Amazon DynamoDB for backend data storage. The application traffic will be unpredictable. T company expects that the application read and write throughput to the database will be moderate to high. The company needs to scale in response to application traffic.
Which DynamoDB table configuration will meet these requirements MOST cost-effectively?

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

**Answer:** B

**Explanation:**
The most cost-effective DynamoDB table configuration for the web application is to configure DynamoDB in on-demand mode by using the DynamoDB Standard table class. This configuration will allow the company to scale in response to application traffic and pay only for the read and write requests that the application performs on the table.
On-demand mode is a flexible billing option that can handle thousands of requests per second without capacity planning. On-demand mode automatically adjusts the table's capacity based on the incoming traffic, and charges only for the read and write requests that are actually performed. On-demand mode is suitable for applications with unpredictable or variable workloads, or applications that prefer the ease of paying for only what they use1.
The DynamoDB Standard table class is the default and recommended table class for most workloads. The DynamoDB Standard table class offers lower throughput costs than the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and is more cost-effective for tables where throughput is the dominant cost. The DynamoDB Standard table class also offers the same performance, durability, and availability as the DynamoDB Standard-IA table class2.
The other options are not correct because they are either not cost-effective or not suitable for the use case. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration requires manual estimation and management of the table's capacity, which adds complexity and cost to the solution. Provisioned mode is a billing option that requires users to specify the amount of read and write capacity units for their tables, and charges for the reserved capacity regardless of usage. Provisioned mode is suitable for applications with predictable or stable workloads, or applications that require finer-grained control over their capacity settings1. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration is not cost-effective for tables with moderate to high throughput. The DynamoDB Standard-IA table class offers lower storage costs than the DynamoDB Standard table class, but higher throughput costs. The DynamoDB Standard-IA table class is optimized for tables where storage is the dominant cost, such as tables that store infrequently accessed data2. Configuring DynamoDB in on-demand mode by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class is not correct because this configuration is not cost- effective for tables with moderate to high throughput. As mentioned above, the DynamoDB Standard-IA table class has higher throughput costs than the DynamoDB Standard table class, which can offset the savings from lower storage costs.
References:
? Table classes - Amazon DynamoDB
? Read/write capacity mode - Amazon DynamoDB

**NEW QUESTION 104**
- (Topic 4)
A company is moving its on-premises Oracle database to Amazon Aurora PostgreSQL. The database has several applications that write to the same tables. The applications need to be migrated one by one with a month in between each migration. Management has expressed concerns that the database has a high number of reads and writes. The data must be kept in sync across both databases throughout the migration.
What should a solutions architect recommend?

A. Use AWS DataSync for the initial migratio
B. Use AWS Database Migration Service (AWS DMS) to create a change data capture (CDC) replication task and a table mapping to select all tables.
C. Use AWS DataSync for the initial migratio
D. Use AWS Database Migration Service (AWS DMS) to create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
E. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a memory optimized replication instanc
F. Create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
G. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a compute optimized replication instanc
H. Create a full load plus change data capture (CDC) replication task and a table mapping to select the largest tables.

**Answer:** C

**Explanation:**
https://aws.amazon.com/ko/premiumsupport/knowledge-center/dms- memory-optimization/

**NEW QUESTION 105**
- (Topic 4)
A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB) The website serves static content Website traffic is increasing and the company is concerned about a potential increase in cost.
What should a solutions architect do to reduce the cost of the website?

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

**Answer:** A

**Explanation:**
Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website. CloudFront
can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the

origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not. References:
? https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/
? https://nodeployfriday.com/posts/static-website-hosting/
? https://aws.amazon.com/cloudfront/

## NEW QUESTION 110
- (Topic 4)
A company is deploying an application in three AWS Regions using an Application Load Balancer Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST high- performing experience?

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

**Answer:** A

**Explanation:**
To provide the most high-performing experience for the users of the application, a solutions architect should use a latency routing policy for the Route 53 A record. This policy allows Route 53 to route traffic to the AWS Region that provides the lowest possible latency for the users1. A latency routing policy can also improve the availability of the application, as Route 53 can automatically route traffic to another Region if the primary Region becomes unavailable2.
References:
? 1: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing- policy.html#routing-policy-latency
? 2: https://aws.amazon.com/route53/faqs/#Latency_Based_Routing

## NEW QUESTION 112
- (Topic 4)
A company has a popular gaming platform running on AWS. The application is sensitive to latency because latency can impact the user experience and introduce unfair advantages to some players. The application is deployed in every AWS Region. It runs on Amazon EC2 instances that are part of Auto Scaling groups configured behind Application Load Balancers (ALBs). A solutions architect needs to implement a mechanism to monitor the health of the application and redirect traffic to healthy endpoints.
Which solution meets these requirements?

A. Configure an accelerator in AWS Global Accelerato
B. Add a listener for the port that the application listens on, and attach it to a Regional endpoint in each Regio
C. Add the ALB as the endpoint.
D. Create an Amazon CloudFront distribution and specify the ALB as the origin serve
E. Configure the cache behavior to use origin cache header
F. Use AWS Lambda functions to optimize the traffic.
G. Create an Amazon CloudFront distribution and specify Amazon S3 as the origin serve
H. Configure the cache behavior to use origin cache header
I. Use AWS Lambda functions to optimize the traffic.
J. Configure an Amazon DynamoDB database to serve as the data store for the applicatio
K. Create a DynamoDB Accelerator (DAX) cluster to act as the in-memory cache for DynamoDB hosting the application data.

**Answer:** A

**Explanation:**
AWS Global Accelerator directs traffic to the optimal healthy endpoint based on health checks, it can also route traffic to the closest healthy endpoint based on geographic location of the client. By configuring an accelerator and attaching it to a Regional endpoint
in each Region, and adding the ALB as the endpoint, the solution will redirect traffic to healthy endpoints, improving the user experience by reducing latency and ensuring that the application is running optimally. This solution will ensure that traffic is directed to the closest healthy endpoint and will help to improve the overall user experience.

## NEW QUESTION 113
- (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 115**
- (Topic 4)
A company wants to rearchitect a large-scale web application to a serverless microservices architecture. The application uses Amazon EC2 instances and is written in Python.
The company selected one component of the web application to test as a microservice. The component supports hundreds of requests each second. The company wants to create and test the microservice on an AWS solution that supports Python. The solution must also scale automatically and require minimal infrastructure and minimal operational support.
Which solution will meet these requirements?

A. Use a Spot Fleet with auto scaling of EC2 instances that run the most recent Amazon Linux operating system.
B. Use an AWS Elastic Beanstalk web server environment that has high availability configured.
C. Use Amazon Elastic Kubernetes Service (Amazon EKS). Launch Auto Scaling groups of self-managed EC2 instances.
D. Use an AWS Lambda function that runs custom developed code.

**Answer:** D

**Explanation:**
 AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You can use Lambda to create and test microservices that are written in Python or other supported languages. Lambda scales automatically to handle the number of requests per second. You only pay for the compute time you consume. Lambda also integrates with other AWS services, such as Amazon API Gateway, Amazon S3, Amazon DynamoDB, and Amazon SQS, to enable event-driven architectures. Lambda has minimal infrastructure and operational overhead, as you do not need to manage servers, operating systems, patches, or scaling policies.
The other options are not serverless solutions and require more infrastructure and operational support. They also do not scale automatically to handle the number of requests per second. A Spot Fleet is a collection of EC2 instances that run on spare capacity at low prices. However, Spot Instances can be interrupted by AWS at any time, which can affect the availability and performance of your microservice. AWS Elastic Beanstalk is a service that automates the deployment and management of web applications on EC2 instances. However, you still need to provision, configure, and monitor the underlying EC2 instances and load balancers. Amazon EKS is a service that runs Kubernetes on AWS. However, you still need to create, configure, and manage the EC2 instances that form the Kubernetes cluster and nodes. You also need to install and update the Kubernetes software and tools. References:
? What is AWS Lambda?
? Building Lambda functions with Python
? Create a layer for a Lambda Python function
? AWS Lambda – Function in Python
? How do I call my AWS Lambda function from a local python script?

**NEW QUESTION 118**
- (Topic 4)
A company runs a real-time data ingestion solution on AWS. The solution consists of the most recent version of Amazon Managed Streaming for Apache Kafka (Amazon MSK). The solution is deployed in a VPC in private subnets across three Availability Zones.
A solutions architect needs to redesign the data ingestion solution to be publicly available over the internet. The data in transit must also be encrypted.
Which solution will meet these requirements with the MOST operational efficiency?

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

**Answer:** A

**Explanation:**
 The solution that meets the requirements with the most operational efficiency is to configure public subnets in the existing VPC and deploy an MSK cluster in the public subnets. This solution allows the data ingestion solution to be publicly available over the internet without creating a new VPC or deploying a load balancer. The solution also ensures that the data in transit is encrypted by enabling mutual TLS authentication, which requires both the client and the server to present certificates for verification. This solution leverages the public access feature of Amazon MSK, which is available for clusters running Apache Kafka 2.6.0 or later versions1.
The other solutions are not as efficient as the first one because they either create unnecessary resources or do not encrypt the data in transit. Creating a new VPC with public subnets would incur additional costs and complexity for managing network resources and routing. Deploying an ALB or an NLB would also add more costs and latency for the data ingestion solution. Moreover, an ALB or an NLB would not encrypt the data in transit by itself, unless they are configured with HTTPS listeners and certificates, which would require additional steps and maintenance. Therefore, these solutions are not optimal for the given requirements. References:
? Public access - Amazon Managed Streaming for Apache Kafka

**NEW QUESTION 123**
- (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 authonzed 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 128**
- (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 132**
- (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 137**
- (Topic 4)
A company has multiple AWS accounts with applications deployed in the us-west-2 Region Application logs are stored within Amazon S3 buckets in each account The company wants
to build a centralized log analysis solution that uses a single S3 bucket Logs must not leave us-west-2, and the company wants to incur minimal operational overhead
Which solution meets these requirements and is MOST cost-effective?

A. Create an S3 Lifecycle policy that copies the objects from one of the application S3 buckets to the centralized S3 bucket
B. Use S3 Same-Region Replication to replicate logs from the S3 buckets to another S3 bucket in us-west-2 Use this S3 bucket for log analysis.
C. Write a script that uses the PutObject API operation every day to copy the entire contents of the buckets to another S3 bucket in us-west-2 Use this S3 bucket for log analysis.
D. Write AWS Lambda functions in these accounts that are triggered every time logs are delivered to the S3 buckets (s3 ObjectCreated a event) Copy the logs to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.

**Answer:** B

**Explanation:**
 This solution meets the following requirements:
? It is cost-effective, as it only charges for the storage and data transfer of the replicated objects, and does not require any additional AWS services or custom scripts. S3 Same-Region Replication (SRR) is a feature that automatically replicates objects across S3 buckets within the same AWS Region. SRR can help you aggregate logs from multiple sources to a single destination for analysis and auditing. SRR also preserves the metadata, encryption, and access control of the source objects.
? It is operationally efficient, as it does not require any manual intervention or scheduling. SRR replicates objects as soon as they are uploaded to the source bucket, ensuring that the destination bucket always has the latest log data. SRR also handles any updates or deletions of the source objects, keeping the destination bucket in sync. SRR can be enabled with a few clicks in the S3 console or with a simple API call.
? It is secure, as it does not allow the logs to leave the us-west-2 Region. SRR only replicates objects within the same AWS Region, ensuring that the data sovereignty and compliance requirements are met. SRR also supports encryption of the source and destination objects, using either server-side encryption with AWS KMS or S3- managed keys, or client-side encryption.
References:
? Same-Region Replication - Amazon Simple Storage Service
? How do I replicate objects across S3 buckets in the same AWS Region?
? Centralized Logging on AWS | AWS Solutions | AWS Solutions Library


**NEW QUESTION 140**
- (Topic 4)
A company wants to run its experimental workloads in the AWS Cloud. The company has a budget for cloud spending. The company's CFO is concerned about cloud spending accountability for each department. The CFO wants to receive notification when the spending threshold reaches 60% of the budget.
Which solution will meet these requirements?

A. Use cost allocation tags on AWS resources to label owner
B. Create usage budgets in AWS Budget
C. Add an alert threshold to receive notification when spending exceeds 60% of the budget.
D. Use AWS Cost Explorer forecasts to determine resource owner
E. Use AWS Cost Anomaly Detection to create alert threshold notifications when spending exceeds 60% of the budget.
F. Use cost allocation tags on AWS resources to label owner
G. Use AWS Support API on AWS Trusted Advisor to create alert threshold notifications when spending exceeds 60% of the budget
H. Use AWS Cost Explorer forecasts to determine resource owner
I. Create usage budgets in AWS Budget
J. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

**Answer:** A

**Explanation:**
 This solution meets the requirements because it allows the company to track and manage its cloud spending by using cost allocation tags to assign costs to different departments, creating usage budgets to set spending limits, and adding alert thresholds to receive notifications when the spending reaches a certain percentage of the budget. This way, the company can monitor its experimental workloads and avoid overspending on the cloud.
References:
? Using Cost Allocation Tags
? Creating an AWS Budget
? Creating an Alert for an AWS Budget


**NEW QUESTION 142**
- (Topic 4)
A company is building a solution that will report Amazon EC2 Auto Scaling events across all the applications in an AWS account. The company needs to use a serverless solution to store the EC2 Auto Scaling status data in Amazon S3. The company then will use the data in Amazon S3 to provide near-real-time updates in a dashboard. The solution must not affect the speed of EC2 instance launches.
How should the company move the data to Amazon S3 to meet these requirements?

A. Use an Amazon CloudWatch metric stream to send the EC2 Auto Scaling status data to Amazon Kinesis Data Firehos
B. Store the data in Amazon S3.
C. Launch an Amazon EMR cluster to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehos
D. Store the data in Amazon S3.
E. Create an Amazon EventBridge rule to invoke an AWS Lambda function on a schedul
F. Configure the Lambda function to send the EC2 Auto Scaling status data directly to Amazon S3.
G. Use a bootstrap script during the launch of an EC2 instance to install Amazon Kinesis Agen
H. Configure Kinesis Agent to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehos
I. Store the data in Amazon S3.

**Answer:** A

**Explanation:**
You can use metric streams to continually stream CloudWatch metrics to a destination of your choice, with near-real-time delivery and low latency. One of the use cases is Data Lake: create a metric stream and direct it to an Amazon Kinesis Data Firehose delivery stream that delivers your CloudWatch metrics to a data lake such as Amazon S3. https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Metric- Streams.html

**NEW QUESTION 143**
- (Topic 4)
A company runs a three-tier web application in the AWS Cloud that operates across three Availability Zones. The application architecture has an Application Load Balancer, an Amazon EC2 web server that hosts user session states, and a MySQL database that runs on an EC2 instance. The company expects sudden increases in application traffic. The company wants to be able to scale to meet future application capacity demands and to ensure high availability across all three Availability Zones.
Which solution will meet these requirements?

A. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deploymen
B. Use Amazon ElastiCache for Redis with high availability to store session data and to cache read
C. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
D. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deploymen
E. Use Amazon ElastiCache for Memcached with high availability to store session data and to cache read
F. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
G. Migrate the MySQL database to Amazon DynamoD
H. Use DynamoDB Accelerator (DAX) to cache read
I. Store the session data in DynamoD
J. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
K. Migrate the MySQL database to Amazon RDS for MySQL in a single Availability Zon
L. Use Amazon ElastiCache for Redis with high availability to store session data and to cache read
M. Migrate the web server to an Auto Scaling group that is in three Availability Zones.

**Answer:** A

**Explanation:**
 This answer is correct because it meets the requirements of scaling to meet future application capacity demands and ensuring high availability across all three Availability Zones. By migrating the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment, the company can benefit from automatic failover, backup, and patching of the database across multiple Availability Zones. By using Amazon ElastiCache for Redis with high availability, the company can store session data and cache reads in a fast, in-memory data store that can also fail over across Availability Zones. By migrating the web server to an Auto Scaling group that is in three Availability Zones, the company can automatically scale the web server capacity based on the demand and traffic patterns. References:
? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.ht
ml
? https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html
? https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2- auto-scaling.html

**NEW QUESTION 145**
- (Topic 4)
A company wants to use artificial intelligence (AI) to determine the quality of its customer service calls. The company currently manages calls in four different languages, including English. The company will offer new languages in the future. The company does not have the resources to regularly maintain machine learning (ML) models.
The company needs to create written sentiment analysis reports from the customer service call recordings. The customer service call recording text must be translated into English.
Which combination of steps will meet these requirements? (Select THREE.)

A. Use Amazon Comprehend to translate the audio recordings into English.
B. Use Amazon Lex to create the written sentiment analysis reports.
C. Use Amazon Polly to convert the audio recordings into text.
D. Use Amazon Transcribe to convert the audio recordings in any language into text.
E. Use Amazon Translate to translate text in any language to English.
F. Use Amazon Comprehend to create the sentiment analysis reports.

**Answer:** DEF

**Explanation:**
 These answers are correct because they meet the requirements of creating written sentiment analysis reports from the customer service call recordings in any language and translating them into English. Amazon Transcribe is a service that uses advanced machine learning technologies to recognize speech in audio files and transcribe them into text. You can use Amazon Transcribe to convert the audio recordings in any language into text, and specify the language code of the source audio. Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. You can use Amazon Translate to translate text in any language to English, and specify the source and target language codes. Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. You can use Amazon Comprehend to create the sentiment analysis reports, which determine if the text is positive, negative, neutral, or mixed.
References:
? https://docs.aws.amazon.com/transcribe/latest/dg/what-is-transcribe.html
? https://docs.aws.amazon.com/translate/latest/dg/what-is.html
? https://docs.aws.amazon.com/comprehend/latest/dg/how-sentiment.html

**NEW QUESTION 146**
- (Topic 4)
A company runs a three-tier web application in a VPC across multiple Availability Zones. Amazon EC2 instances run in an Auto Scaling group for the application tier.
The company needs to make an automated scaling plan that will analyze each resource's daily and weekly historical workload trends. The configuration must scale resources appropriately according to both the forecast and live changes in utilization.
Which scaling strategy should a solutions architect recommend to meet these requirements?

A. Implement dynamic scaling with step scaling based on average CPU utilization from the EC2 instances.
B. Enable predictive scaling to forecast and scal
C. Configure dynamic scaling with target tracking.
D. Create an automated scheduled scaling action based on the traffic patterns of the web application.
E. Set up a simple scaling polic
F. Increase the cooldown period based on the EC2 instance startup time

**Answer:** B

**Explanation:**
This solution meets the requirements because it allows the company to use both predictive scaling and dynamic scaling to optimize the capacity of its Auto Scaling group. Predictive scaling uses machine learning to analyze historical data and forecast future traffic patterns. It then adjusts the desired capacity of the group in advance of the predicted changes. Dynamic scaling uses target tracking to maintain a specified metric (such as CPU utilization) at a target value. It scales the group in or out as needed to keep the metric close to the target. By using both scaling methods, the company can benefit from faster, simpler, and more accurate scaling that responds to both forecasted and live changes in utilization. References:
? Predictive scaling for Amazon EC2 Auto Scaling
? [Target tracking scaling policies for Amazon EC2 Auto Scaling]


**NEW QUESTION 149**
- (Topic 4)
A gaming company uses Amazon DynamoDB to store user information such as geographic location, player data, and leaderboards. The company needs to configure continuous backups to an Amazon S3 bucket with a minimal amount of coding. The backups must not affect availability of the application and must not affect the read capacity units (RCUs) that are defined for the table
Which solution meets these requirements?

A. Use an Amazon EMR cluste
B. Create an Apache Hive job to back up the data to Amazon S3.
C. Export the data directly from DynamoDB to Amazon S3 with continuous backup
D. Turn on point-in-time recovery for the table.
E. Configure Amazon DynamoDB Stream
F. Create an AWS Lambda function to consume the stream and export the data to an Amazon S3 bucket.
G. Create an AWS Lambda function to export the data from the database tables to Amazon S3 on a regular basi
H. Turn on point-in-time recovery for the table.

**Answer:** B

**Explanation:**
https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and- design-patterns/
https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/


**NEW QUESTION 150**
......

# Thank You for Trying Our Product

\* 100% Pass or Money Back

     All our products come with a 90-day Money Back Guarantee.

\* One year free update

     You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

     We currently serve more than 30,000,000 customers.

\* Shop Securely

     All transactions are protected by VeriSign!

**100% Pass Your AWS-Solution-Architect-Associate Exam with Our Prep Materials Via below:**

https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html