



# Amazon-Web-Services

## Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

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

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

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

### NEW QUESTION 1

- (Topic 1)

A company has a three-tier web application that is deployed on AWS. The web servers are deployed in a public subnet in a VPC. The application servers and database servers are deployed in private subnets in the same VPC. The company has deployed a third-party virtual firewall appliance from AWS Marketplace in an inspection VPC. The appliance is configured with an IP interface that can accept IP packets. A solutions architect needs to integrate the web application with the appliance to inspect all traffic to the application before the traffic reaches the web server. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a Network Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection
- B. Create an Application Load Balancer in the public subnet of the application's VPC to route the traffic to the appliance for packet inspection
- C. Deploy a transit gateway in the inspection VPC. Configure route tables to route the incoming packets through the transit gateway
- D. Deploy a Gateway Load Balancer in the inspection VPC. Create a Gateway Load Balancer endpoint to receive the incoming packets and forward the packets to the appliance

**Answer:** D

#### Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/scaling-network-traffic-inspection-using-aws-gateway-load-balancer/>

### NEW QUESTION 2

- (Topic 1)

A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning. Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

- A. Configure the application to send the data to Amazon Kinesis Data Firehose.
- B. Use Amazon Simple Email Service (Amazon SES) to format the data and to send the report by email.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Glue job to query the application's API for the data.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data.
- E. Store the application data in Amazon S3. Create an Amazon Simple Notification Service (Amazon SNS) topic as an S3 event destination to send the report by email.

**Answer:** BD

#### Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/send-email-formatted.html>

\* D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data. This step can be done using AWS Lambda to extract the shipping statistics and organize the data into an HTML format.

\* B. Use Amazon Simple Email Service (Amazon SES) to format the data and send the report by email. This step can be done by using Amazon SES to send the report to multiple email addresses at the same time every morning.

Therefore, options D and B are the correct choices for this question. Option A is incorrect because Kinesis Data Firehose is not necessary for this use case. Option C is incorrect because AWS Glue is not required to query the application's API. Option E is incorrect because S3 event notifications cannot be used to send the report by email.

### NEW QUESTION 3

- (Topic 1)

A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted.

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

- A. Create AWS Secrets Manager secrets for encrypted certificates
- B. Manually update the certificates as needed
- C. Control access to the data by using fine-grained IAM access.
- D. Create an AWS Lambda function that uses the Python cryptography library to receive and perform encryption operations
- E. Store the function in an Amazon S3 bucket.
- F. Create an AWS Key Management Service (AWS KMS) customer managed key
- G. Allow the EC2 role to use the KMS key for encryption operations
- H. Store the encrypted data on Amazon S3.
- I. Create an AWS Key Management Service (AWS KMS) customer managed key
- J. Allow the EC2 role to use the KMS key for encryption operations
- K. Store the encrypted data on Amazon Elastic Block Store (Amazon EBS) volumes.

**Answer:** D

### NEW QUESTION 4

- (Topic 1)

A company has created an image analysis application in which users can upload photos and add photo frames to their images. The users upload images and metadata to indicate which photo frames they want to add to their images. The application uses a single Amazon EC2 instance and Amazon DynamoDB to store the metadata.

The application is becoming more popular, and the number of users is increasing. The company expects the number of concurrent users to vary significantly depending on the time of day and day of week. The company must ensure that the application can scale to meet the needs of the growing user base.

Which solution meets these requirements?

- A. Use AWS Lambda to process the photos
- B. Store the photos and metadata in DynamoDB.
- C. Use Amazon Kinesis Data Firehose to process the photos and to store the photos and metadata.
- D. Use AWS Lambda to process the photos

- E. Store the photos in Amazon S3. Retain DynamoDB to store the metadata.
- F. Increase the number of EC2 instances to three
- G. Use Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes to store the photos and metadata.

**Answer: C**

**Explanation:**

<https://www.quora.com/How-can-I-use-DynamoDB-for-storing-metadata-for-Amazon-S3-objects>

This solution meets the requirements of scalability, performance, and availability. AWS Lambda can process the photos in parallel and scale up or down automatically depending on the demand. Amazon S3 can store the photos and metadata reliably and durably, and provide high availability and low latency. DynamoDB can store the metadata efficiently and provide consistent performance. This solution also reduces the cost and complexity of managing EC2 instances and EBS volumes.

Option A is incorrect because storing the photos in DynamoDB is not a good practice, as it can increase the storage cost and limit the throughput. Option B is incorrect because Kinesis Data Firehose is not designed for processing photos, but for streaming data to destinations such as S3 or Redshift. Option D is incorrect because increasing the number of EC2 instances and using Provisioned IOPS SSD volumes does not guarantee scalability, as it depends on the load balancer and the application code. It also increases the cost and complexity of managing the infrastructure.

References:

? <https://aws.amazon.com/certification/certified-solutions-architect-professional/>

? <https://www.examtips.com/discussions/amazon/view/7193-exam-aws-certified-solutions-architect-professional-topic-1/>

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

**NEW QUESTION 5**

- (Topic 1)

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

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

**Answer: B**

**Explanation:**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3>

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>

**NEW QUESTION 6**

- (Topic 1)

A company hosts an application on multiple Amazon EC2 instances. The application processes messages from an Amazon SQS queue, writes to an Amazon RDS table, and deletes the message from the queue. Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.

What should a solutions architect do to ensure messages are being processed once only?

- A. Use the CreateQueue API call to create a new queue
- B. Use the AddPermission API call to add appropriate permissions
- C. Use the ReceiveMessage API call to set an appropriate wait time
- D. Use the ChangeMessageVisibility API call to increase the visibility timeout

**Answer: D**

**Explanation:**

The visibility timeout begins when Amazon SQS returns a message. During this time, the consumer processes and deletes the message. However, if the consumer fails before deleting the message and your system doesn't call the DeleteMessage action for that message before the visibility timeout expires, the message becomes visible to other consumers and the message is received again. If a message must be received only once, your consumer should delete it within the duration of the visibility timeout. <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html>

Keyword: SQS queue writes to an Amazon RDS. From this, Option D is the best choice; other options are ruled out [Option A - You can't introduce one more Queue in the existing one; Option B - only Permission; Option C - Only Retrieves Messages]. FIFO queues are designed to never introduce duplicate messages. However, your message producer might introduce duplicates in certain scenarios: for example, if the producer sends a message, does not receive a response, and then resends the same message. Amazon SQS APIs provide deduplication functionality that prevents your message producer from sending duplicates. Any duplicates introduced by the message producer are removed within a 5-minute deduplication interval. For standard queues, you might occasionally receive a duplicate copy of a message (at-least- once delivery). If you use a standard queue, you must design your applications to be idempotent (that is, they must not be affected adversely when processing the same message more than once).

**NEW QUESTION 7**

- (Topic 1)

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

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

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

- J. Use AWS Lambda functions to transform the data
- K. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
- L. Configure an Amazon API Gateway API to send data to AWS Glue
- M. Use AWS Lambda functions to transform the data
- N. Use AWS Glue to send the data to Amazon S3.

**Answer:** C

#### NEW QUESTION 8

- (Topic 1)

A company has an Amazon S3 bucket that contains critical data. The company must protect the data from accidental deletion.

Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

- A. Enable versioning on the S3 bucket.
- B. Enable MFA Delete on the S3 bucket.
- C. Create a bucket policy on the S3 bucket.
- D. Enable default encryption on the S3 bucket.
- E. Create a lifecycle policy for the objects in the S3 bucket.

**Answer:** AB

#### Explanation:

To protect data in an S3 bucket from accidental deletion, versioning should be enabled, which enables you to preserve, retrieve, and restore every version of every object in an S3 bucket. Additionally, enabling MFA (multi-factor authentication) Delete on the S3 bucket adds an extra layer of protection by requiring an authentication token in addition to the user's access keys to delete objects in the bucket.

Reference:

AWS S3 Versioning documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>

AWS S3 MFA Delete documentation: <https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html>

#### NEW QUESTION 9

- (Topic 1)

A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes. The application data must be stored in a standard file system structure. The company wants a solution that scales automatically, is highly available, and requires minimum operational overhead.

Which solution will meet these requirements?

- A. Migrate the application to run as containers on Amazon Elastic Container Service (Amazon ECS). Use Amazon S3 for storage.
- B. Migrate the application to run as containers on Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon Elastic Block Store (Amazon EBS) for storage.
- C. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling group. Use Amazon Elastic File System (Amazon EFS) for storage.
- D. Use Amazon Elastic File System (Amazon EFS) for storage.
- E. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling group. Use Amazon Elastic Block Store (Amazon EBS) for storage.
- F. Use Amazon Elastic Block Store (Amazon EBS) for storage.

**Answer:** C

#### Explanation:

EFS is a standard file system, it scales automatically and is highly available.

#### NEW QUESTION 10

- (Topic 1)

A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week.

What should the company do to guarantee the EC2 capacity?

- A. Purchase Reserved instances that specify the Region needed.
- B. Create an On-Demand Capacity Reservation that specifies the Region needed.
- C. Purchase Reserved instances that specify the Region and three Availability Zones needed.
- D. Create an On-Demand Capacity Reservation that specifies the Region and three Availability Zones needed.

**Answer:** D

#### Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>

Reserve instances: You will have to pay for the whole term (1 year or 3 years) which is not cost effective.

#### NEW QUESTION 10

- (Topic 1)

A company recently migrated to AWS and wants to implement a solution to protect the traffic that flows in and out of the production VPC. The company had an inspection server in its on-premises data center. The inspection server performed specific operations such as traffic flow inspection and traffic filtering. The company wants to have the same functionalities in the AWS Cloud.

Which solution will meet these requirements?

- A. Use Amazon GuardDuty for traffic inspection and traffic filtering in the production VPC.
- B. Use Traffic Mirroring to mirror traffic from the production VPC for traffic inspection and filtering.
- C. Use AWS Network Firewall to create the required rules for traffic inspection and traffic filtering for the production VPC.
- D. Use AWS Firewall Manager to create the required rules for traffic inspection and traffic filtering for the production VPC.

**Answer:** C

**Explanation:**

AWS Network Firewall supports both inspection and filtering as required

**NEW QUESTION 13**

- (Topic 1)

An application development team is designing a microservice that will convert large images to smaller, compressed images. When a user uploads an image through the web interface, the microservice should store the image in an Amazon S3 bucket, process and compress the image with an AWS Lambda function, and store the image in its compressed form in a different S3 bucket.

A solutions architect needs to design a solution that uses durable, stateless components to process the images automatically.

Which combination of actions will meet these requirements? (Choose two.)

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure the S3 bucket to send a notification to the SQS queue when an image is uploaded to the S3 bucket.
- B. Configure the Lambda function to use the Amazon Simple Queue Service (Amazon SQS) queue as the invocation source. When the SQS message is successfully processed, delete the message in the queue.
- C. Configure the Lambda function to monitor the S3 bucket for new uploads. When an uploaded image is detected, write the file name to a text file in memory and use the text file to keep track of the images that were processed.
- D. Launch an Amazon EC2 instance to monitor an Amazon Simple Queue Service (Amazon SQS) queue. When items are added to the queue, log the file name in a text file on the EC2 instance and invoke the Lambda function.
- E. Configure an Amazon EventBridge (Amazon CloudWatch Events) event to monitor the S3 bucket. When an image is uploaded, send an alert to an Amazon Simple Notification Service (Amazon SNS) topic with the application owner's email address for further processing.
- F. Send an alert to an Amazon Simple Notification Service (Amazon SNS) topic with the application owner's email address for further processing.

**Answer: AB**

**Explanation:**

? Creating an Amazon Simple Queue Service (SQS) queue and configuring the S3 bucket to send a notification to the SQS queue when an image is uploaded to the S3 bucket will ensure that the Lambda function is triggered in a stateless and durable manner.

? Configuring the Lambda function to use the SQS queue as the invocation source, and deleting the message in the queue after it is successfully processed will ensure that the Lambda function processes the image in a stateless and durable manner.

Amazon SQS is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating message-oriented middleware, and empowers developers to focus on differentiating work. When new images are uploaded to the S3 bucket, SQS will trigger the Lambda function to process the image and compress it. Once the image is processed, the SQS message is deleted, ensuring that the Lambda function is stateless and durable.

**NEW QUESTION 17**

- (Topic 1)

A survey company has gathered data for several years from areas in the United States. The company hosts the data in an Amazon S3 bucket that is 3 TB in size and growing. The company has started to share the data with a European marketing firm that has S3 buckets. The company wants to ensure that its data transfer costs remain as low as possible.

Which solution will meet these requirements?

- A. Configure the Requester Pays feature on the company's S3 bucket.
- B. Configure S3 Cross-Region Replication from the company's S3 bucket to one of the marketing firm's S3 buckets.
- C. Configure cross-account access for the marketing firm so that the marketing firm has access to the company's S3 bucket.
- D. Configure the company's S3 bucket to use S3 Intelligent-Tiering. Sync the S3 bucket to one of the marketing firm's S3 buckets.

**Answer: A**

**Explanation:**

"Typically, you configure buckets to be Requester Pays buckets when you want to share data but not incur charges associated with others accessing the data. For example, you might use Requester Pays buckets when making available large datasets, such as zip code directories, reference data, geospatial information, or web crawling data." <https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysBuckets.html>

**NEW QUESTION 19**

- (Topic 1)

A company has an automobile sales website that stores its listings in a database on Amazon RDS. When an automobile is sold, the listing needs to be removed from the website, and the data must be sent to multiple target systems.

Which design should a solutions architect recommend?

- A. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) queue for the targets to consume.
- B. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) FIFO queue for the targets to consume.
- C. Subscribe to an RDS event notification and send an Amazon Simple Queue Service (Amazon SQS) queue fanned out to multiple Amazon Simple Notification Service (Amazon SNS) topics. Use AWS Lambda functions to update the targets.
- D. Subscribe to an RDS event notification and send an Amazon Simple Notification Service (Amazon SNS) topic fanned out to multiple Amazon Simple Queue Service (Amazon SQS) queues. Use AWS Lambda functions to update the targets.

**Answer: D**

**Explanation:**

<https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html> <https://docs.aws.amazon.com/lambda/latest/dg/with-sns.html>

**NEW QUESTION 24**

- (Topic 1)

A company has an application that runs on Amazon EC2 instances and uses an Amazon

Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.

What should a solutions architect do to accomplish this goal?

- A. Use AWS Secrets Manager
- B. Turn on automatic rotation.
- C. Use AWS Systems Manager Parameter Store
- D. Turn on automatic rotation.
- E. Create an Amazon S3 bucket to store objects that are encrypted with an AWS Key Management Service (AWS KMS) encryption key
- G. Migrate the credential file to the S3 bucket
- H. Point the application to the S3 bucket.
- I. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instance)
- J. Attach the new EBS volume to each EC2 instance
- K. Migrate the credential file to the new EBS volume
- L. Point the application to the new EBS volume.

**Answer:** A

**Explanation:**

<https://aws.amazon.com/cn/blogs/security/how-to-connect-to-aws-secrets-manager-service-within-a-virtual-private-cloud/>  
<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>

**NEW QUESTION 25**

- (Topic 1)

A company runs its Infrastructure on AWS and has a registered base of 700,000 users for a document management application. The company intends to create a product that converts large PDF files to JPEG image files. The PDF files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over time.

Which solution meets these requirements MOST cost-effectively?

- A. Save the PDF files to Amazon S3. Configure an S3 PUT event to invoke an AWS Lambda function to convert the files to JPEG format and store them back in Amazon S3.
- B. Save the PDF files to Amazon DynamoDB.
- C. Use the DynamoDB Streams feature to invoke an AWS Lambda function to convert the files to JPEG format and store them back in DynamoDB.
- D. Upload the PDF files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances.
- E. Amazon Elastic Block Store (Amazon EBS) storage and an Auto Scaling group.
- F. Use a program in the EC2 instances to convert the files to JPEG format. Save the PDF files and the JPEG files in the EBS store.
- G. Upload the PDF files to an AWS Elastic Beanstalk application that includes Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) storage, and an Auto Scaling group.
- H. Use a program in the EC2 instances to convert the files to JPEG format. Save the PDF files and the JPEG files in the EBS store.

**Answer:** A

**Explanation:**

Elastic Beanstalk is expensive, and DocumentDB has a 400KB max to upload files. So Lambda and S3 should be the one.

**NEW QUESTION 27**

- (Topic 1)

A hospital recently deployed a RESTful API with Amazon API Gateway and AWS Lambda. The hospital uses API Gateway and Lambda to upload reports that are in PDF format and JPEG format. The hospital needs to modify the Lambda code to identify protected health information (PHI) in the reports. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use existing Python libraries to extract the text from the reports and to identify the PHI from the extracted text.
- B. Use Amazon Textract to extract the text from the reports. Use Amazon SageMaker to identify the PHI from the extracted text.
- C. Use Amazon Textract to extract the text from the reports. Use Amazon Comprehend Medical to identify the PHI from the extracted text.
- D. Use Amazon Rekognition to extract the text from the reports. Use Amazon Comprehend Medical to identify the PHI from the extracted text.

**Answer:** C

**Explanation:**

To meet the requirements of the company to have access to both AWS and on-premises file storage with minimum latency, a hybrid cloud architecture can be used. One solution is to deploy and configure Amazon FSx for Windows File Server on AWS, which provides fully managed Windows file servers. The on-premises file data can be moved to the FSx File Gateway, which can act as a bridge between on-premises and AWS file storage. The cloud workloads can be configured to use FSx for Windows File Server on AWS, while the on-premises workloads can be configured to use the FSx File Gateway. This solution minimizes operational overhead and requires no significant changes to the existing file access patterns. The connectivity between on-premises and AWS can be established using an AWS Site-to-Site VPN connection.

Reference:

AWS FSx for Windows File Server: <https://aws.amazon.com/fsx/windows/> AWS FSx File Gateway: <https://aws.amazon.com/fsx/file-gateway/>  
AWS Site-to-Site VPN: <https://aws.amazon.com/vpn/site-to-site-vpn/>

**NEW QUESTION 28**

- (Topic 2)

A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.

What should a solutions architect do to meet this requirement?

- A. Configure AWS Storage Gateway in volume gateway mode.
- B. Mount the volume to each Windows instance.
- C. Configure Amazon FSx for Windows File Server.
- D. Mount the Amazon FSx file system to each Windows instance.
- E. Configure a file system by using Amazon Elastic File System (Amazon EFS). Mount the EFS file system to each Windows instance.
- F. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required size.
- G. Attach each EC2 instance to the volume.
- H. Mount the file system within the volume to each Windows instance.

**Answer:** B

**Explanation:**

This solution meets the requirement of migrating a Windows-based application that requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. Amazon FSx for Windows File Server provides fully managed shared storage built on Windows Server, and delivers a wide range of data access, data management, and administrative capabilities. It supports the Server Message Block (SMB) protocol and can be mounted to EC2 Windows instances across multiple Availability Zones.

Option A is incorrect because AWS Storage Gateway in volume gateway mode provides cloud-backed storage volumes that can be mounted as iSCSI devices from on-premises application servers, but it does not support SMB protocol or EC2 Windows instances. Option C is incorrect because Amazon Elastic File System (Amazon EFS) provides a scalable and elastic NFS file system for Linux-based workloads, but it does not support SMB protocol or EC2 Windows instances.

Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume.

References:

? <https://aws.amazon.com/fsx/windows/>

? <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/using-file-shares.html>

**NEW QUESTION 32**

- (Topic 2)

A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead.

Which solution will meet these requirements?

- A. Update the Route 53 records to use a latency routing policy
- B. Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
- C. Set up a Route 53 active-passive failover configuration
- D. Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
- E. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoint
- F. Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
- G. Update the Route 53 records to use a multivalue answer routing policy
- H. Create a health check
- I. Direct traffic to the website if the health check passes
- J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

**Answer:** B

**Explanation:**

This solution meets the requirements of directing users to a backup static error page if the primary website is unavailable, minimizing changes and infrastructure overhead. Route 53 active-passive failover configuration can route traffic to a primary resource when it is healthy or to a secondary resource when the primary resource is unhealthy. Route 53 health checks can monitor the health of the ALB endpoint and trigger the failover when needed. The static error page can be hosted in an S3 bucket that is configured as a website, which is a simple and cost-effective way to serve static content.

Option A is incorrect because using a latency routing policy can route traffic based on the lowest network latency for users, but it does not provide failover functionality. Option C is incorrect because using an active-active configuration with the ALB and an EC2 instance can increase the infrastructure overhead and complexity, and it does not guarantee that the EC2 instance will always be healthy. Option D is incorrect because using a multivalue answer routing policy can return multiple values for a query, but it does not provide failover functionality.

References:

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

? <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

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

**NEW QUESTION 35**

- (Topic 2)

A company has a multi-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs to modify the infrastructure to be highly available without modifying the application.

Which architecture should the solutions architect choose that provides high availability?

- A. Create an Auto Scaling group that uses three instances across each of two Regions.
- B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.
- C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.
- D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

**Answer:** B

**Explanation:**

High availability can be enabled for this architecture quite simply by modifying the existing Auto Scaling group to use multiple availability zones. The ASG will automatically balance the load so you don't actually need to specify the instances per AZ.

**NEW QUESTION 37**

- (Topic 2)

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs.

The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).

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

- A. Use Amazon Athena for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- B. Use Amazon Kinesis Data Analytics for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster
- D. Use an AWS Glue extract, transform, and load (ETL) job to convert the data into JSON format Load the data into multiple Amazon OpenSearch Service

(Amazon Elasticsearch Service) clusters

E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake Use AWS Glue to crawl the source extract the data and load the data into Amazon S3 in Apache Parquet format

**Answer:** AE

**Explanation:**

Amazon Athena is the best choice for running one-time queries on streaming data. Although Amazon Kinesis Data Analytics provides an easy and familiar standard SQL language to analyze streaming data in real-time, it is designed for continuous queries rather than one-time queries[1]. On the other hand, Amazon Athena is a serverless interactive query service that allows querying data in Amazon S3 using SQL. It is optimized for ad-hoc querying and is ideal for running one-time queries on streaming data[2]. AWS Lake Formation uses as a central place to have all your data for analytics purposes (E). Athena integrate perfect with S3 and can makes queries (A).

**NEW QUESTION 42**

- (Topic 2)

A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.

The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.

Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

- A. Use Spot Instances for the data ingestion layer
- B. Use On-Demand Instances for the data ingestion layer
- C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
- D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
- E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

**Answer:** AC

**Explanation:**

EC2 instance Savings Plan saves 72% while Compute Savings Plans saves 66%. But according to link, it says "Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy, and also apply to Fargate and Lambda usage." EC2 instance Savings Plans are not applied to Fargate or Lambda

**NEW QUESTION 45**

- (Topic 2)

A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.

How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

- A. Configure the web application to send an order message to Amazon Kinesis Data Firehos
- B. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- C. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request Use Lambda to query the database, call the payment service, and pass in the order information.
- D. Store the order in the databas
- E. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to pollAmazon SN
- F. retrieve the message, and process the order.
- G. Store the order in the databas
- H. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queu
- I. Set the payment service to retrieve the message and process the orde
- J. Delete the message from the queue.

**Answer:** D

**Explanation:**

This approach ensures that the order creation and payment processing steps are separate and atomic. By sending the order information to an SQS FIFO queue, the payment service can process the order one at a time and in the order they were received. If the payment service is unable to process an order, it can be retried later, preventing the creation of multiple orders. The deletion of the message from the queue after it is processed will prevent the same message from being processed multiple times.

**NEW QUESTION 46**

- (Topic 2)

A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after li is deployed. The company ls unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.

Which solution will meet these requirements?

- A. Store container images In an Amazon Elastic Container Registry (Amazon ECR) repositor
- B. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the container
- C. Use target tracking to scale automatically based on demand.
- D. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repositor
- E. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the container
- F. Use target tracking to scale automatically based on demand.
- G. Store container images in a repository that runs on an Amazon EC2 instanc
- H. Run the containers on EC2 instances that are spread across multiple Availability Zone
- I. Monitor the average CPU utilization in Amazon CloudWate
- J. Launch new EC2 instances as needed
- K. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image Launch EC2 Instances in an Auto Scaling group across multiple Availability Zone

L. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

**Answer:** A

**Explanation:**

AWS Fargate is a serverless experience for user applications, allowing the user to concentrate on building applications instead of configuring and managing servers. Fargate also automates resource management, allowing users to easily scale their applications in response to demand.

**NEW QUESTION 47**

- (Topic 2)

A company is running an online transaction processing (OLTP) workload on AWS. This workload uses an unencrypted Amazon RDS DB instance in a Multi-AZ deployment. Daily database snapshots are taken from this instance.

What should a solutions architect do to ensure the database and snapshots are always encrypted moving forward?

- A. Encrypt a copy of the latest DB snapshot
- B. Replace existing DB instance by restoring the encrypted snapshot
- C. Create a new encrypted Amazon Elastic Block Store (Amazon EBS) volume and copy the snapshots to it. Enable encryption on the DB instance
- D. Copy the snapshots and enable encryption using AWS Key Management Service (AWS KMS). Restore encrypted snapshot to an existing DB instance
- E. Copy the snapshots to an Amazon S3 bucket that is encrypted using server-side encryption with AWS Key Management Service (AWS KMS) managed keys (SSE-KMS)

**Answer:** A

**Explanation:**

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_RestoreFromSnapshot.html#USER\\_RestoreFromSnapshot.CON](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#USER_RestoreFromSnapshot.CON)  
Under "Encrypt unencrypted resources" - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

**NEW QUESTION 48**

- (Topic 2)

A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year.

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

- A. Move the data to the S3 bucket
- B. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use the built-in key rotation behavior of SSE-S3 encryption keys.
- C. Create an AWS Key Management Service (AWS KMS) customer managed key
- D. Enable automatic key rotation
- E. Set the S3 bucket's default encryption behavior to use the customer managed KMS key
- F. Move the data to the S3 bucket.
- G. Create an AWS Key Management Service (AWS KMS) customer managed key
- H. Set the S3 bucket's default encryption behavior to use the customer managed KMS key
- I. Move the data to the S3 bucket
- J. Manually rotate the KMS key every year.
- K. Encrypt the data with customer key material before moving the data to the S3 bucket
- L. Create an AWS Key Management Service (AWS KMS) key without key material
- M. Import the customer key material into the KMS key
- N. Enable automatic key rotation.

**Answer:** B

**Explanation:**

SSE-S3 - is free and uses AWS owned CMKs (CMK = Customer Master Key). The encryption key is owned and managed by AWS, and is shared among many accounts. Its rotation is automatic with time that varies as shown in the table here. The time is not explicitly defined.

SSE-KMS - has two flavors:

AWS managed CMK. This is free CMK generated only for your account. You can only view its policies and audit usage, but not manage it. Rotation is automatic - once per 1095 days (3 years),

Customer managed CMK. This uses your own key that you create and can manage. Rotation is not enabled by default. But if you enable it, it will be automatically rotated every 1 year. This variant can also use an imported key material by you. If you create such key with an imported material, there is no automated rotation. Only manual rotation.

SSE-C - customer provided key. The encryption key is fully managed by you outside of AWS. AWS will not rotate it.

This solution meets the requirements of moving data to an Amazon S3 bucket, encrypting the data when it is stored in the S3 bucket, and automatically rotating the encryption key every year with the least operational overhead. AWS Key Management Service (AWS KMS) is a service that enables you to create and manage encryption keys for your data. A customer managed key is a symmetric encryption key that you create and manage in AWS KMS. You can enable automatic key rotation for a customer managed key, which means that AWS KMS generates new cryptographic material for the key every year. You can set the S3 bucket's default encryption behavior to use the customer managed KMS key, which means that any object that is uploaded to the bucket without specifying an encryption method will be encrypted with that key.

Option A is incorrect because using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not allow you to control or manage the encryption keys. SSE-S3 uses a unique key for each object, and encrypts that key with a master key that is regularly rotated by S3. However, you cannot enable or disable key rotation for SSE-S3 keys, or specify the rotation interval. Option C is incorrect because manually rotating the KMS key every year can increase the operational overhead and complexity, and it may not meet the requirement of rotating the key every year if you forget or delay the rotation process. Option D is incorrect because encrypting the data with customer key material before moving the data to the S3 bucket can increase the operational overhead and complexity, and it may not provide consistent encryption for all objects in the bucket. Creating a KMS key without key material and importing the customer key material into the KMS key can enable you to use your own source of random bits to generate your KMS keys, but it does not support automatic key rotation.

References:

? <https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html>

? <https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html>

**NEW QUESTION 53**

- (Topic 2)

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new S3 bucket
- B. Load the data into the new S3 bucket
- C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- D. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon Athena to query the data.
- E. Create a new S3 bucket
- F. Load the data into the new S3 bucket
- G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
- I. Load the data into the existing S3 bucket
- J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- K. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
- L. Load the data into the existing S3 bucket
- M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

**Answer:** A

**Explanation:**

This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.

Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi-Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html>

? <https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/>

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

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

**NEW QUESTION 58**

- (Topic 2)

A company runs its two-tier ecommerce website on AWS. The web tier consists of a load balancer that sends traffic to Amazon EC2 instances. The database tier uses an Amazon RDS DB instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available. Which combination of configuration options will meet these requirements? (Choose two.)

- A. Use an Auto Scaling group to launch the EC2 instances in private subnet
- B. Deploy an RDS Multi-AZ DB instance in private subnets.
- C. Configure a VPC with two private subnets and two NAT gateways across two Availability Zone
- D. Deploy an Application Load Balancer in the private subnets.
- E. Use an Auto Scaling group to launch the EC2 instances in public subnets across two Availability Zone
- F. Deploy an RDS Multi-AZ DB instance in private subnets.
- G. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zone
- H. Deploy an Application Load Balancer in the public subnet.
- I. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zone
- J. Deploy an Application Load Balancer in the public subnets.

**Answer:** AE

**Explanation:**

Before you begin: Decide which two Availability Zones you will use for your EC2 instances. Configure your virtual private cloud (VPC) with at least one public subnet in each of these Availability Zones. These public subnets are used to configure the load balancer. You can launch your EC2 instances in other subnets of these Availability Zones instead.

**NEW QUESTION 62**

- (Topic 2)

An entertainment company is using Amazon DynamoDB to store media metadata. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application. What should a solutions architect recommend to meet this requirement?

- A. Use Amazon ElastiCache for Redis.
- B. Use Amazon DynamoDB Accelerator (DAX).
- C. Replicate data by using DynamoDB global tables.
- D. Use Amazon ElastiCache for Memcached with Auto Discovery enabled.

**Answer:** B

**Explanation:**

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

**NEW QUESTION 67**

- (Topic 2)

A company runs workloads on AWS. The company needs to connect to a service from an external provider. The service is hosted in the provider's VPC. According to the company's security team, the connectivity must be private and must be restricted

to the target service. The connection must be initiated only from the company's VPC.  
Which solution will meet these requirements?

- A. Create a VPC peering connection between the company's VPC and the provider's VPC
- B. Update the route table to connect to the target service.
- C. Ask the provider to create a virtual private gateway in its VPC
- D. Use AWS PrivateLink to connect to the target service.
- E. Create a NAT gateway in a public subnet of the company's VPC
- F. Update the route table to connect to the target service.
- G. Ask the provider to create a VPC endpoint for the target service
- H. Use AWS PrivateLink to connect to the target service.

**Answer: D**

**Explanation:**

**\*\*AWS PrivateLink provides private connectivity between VPCs, AWS services, and your on-premises networks, without exposing your traffic to the public internet\*\*.** AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify your network architecture. Interface **\*\*VPC endpoints\*\***, powered by AWS PrivateLink, connect you to services hosted by AWS Partners and supported solutions available in AWS Marketplace.  
<https://aws.amazon.com/privatelink/>

**NEW QUESTION 71**

- (Topic 2)

A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.

Which solution meets these requirements **MOST** cost-effectively?

- A. Use Spot Instances exclusively to handle the maximum capacity required.
- B. Use Reserved Instances exclusively to handle the maximum capacity required.
- C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
- D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

**Answer: D**

**Explanation:**

We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted.  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand-instances.html>

**NEW QUESTION 76**

- (Topic 2)

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

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

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

**Answer: C**

**Explanation:**

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

Creating an Amazon EventBridge (Amazon CloudWatch Events) rule for the CreateImage API call and configuring the target as an Amazon Simple Notification Service (Amazon SNS) topic to send an alert when a CreateImage API call is detected will meet the requirements with the least operational overhead. Amazon EventBridge is a serverless event bus that makes it easy to connect applications together using data from your own applications, integrated Software as a Service (SaaS) applications, and AWS services. By creating an EventBridge rule for the CreateImage API call, the company can set up alerts whenever this operation is called within their account. The alert can be sent to an SNS topic, which can then be configured to send notifications to the company's email or other desired destination.

**NEW QUESTION 77**

- (Topic 2)

A gaming company hosts a browser-based application on AWS. The users of the application consume a large number of videos and images that are stored in Amazon S3. This content is the same for all users.

The application has increased in popularity, and millions of users worldwide are accessing these media files. The company wants to provide the files to the users while reducing the load on the origin.

Which solution meets these requirements **MOST** cost-effectively?

- A. Deploy an AWS Global Accelerator accelerator in front of the web servers.
- B. Deploy an Amazon CloudFront web distribution in front of the S3 bucket.
- C. Deploy an Amazon ElastiCache for Redis instance in front of the web servers.
- D. Deploy an Amazon ElastiCache for Memcached instance in front of the web servers.

**Answer: B**

**Explanation:**

ElastiCache, enhances the performance of web applications by quickly retrieving information from fully-managed in-memory data stores. It utilizes Memcached and Redis, and manages to considerably reduce the time your applications would, otherwise, take to read data from disk-based databases. Amazon CloudFront supports dynamic content from HTTP and WebSocket protocols, which are based on the Transmission Control Protocol (TCP) protocol. Common use cases include dynamic API calls, web pages and web applications, as well as an application's static files such as audio and images. It also supports on-demand media streaming over HTTP. AWS Global Accelerator supports both User Datagram Protocol (UDP) and TCP-based protocols. It is commonly used for non- HTTP use cases, such as gaming, IoT and voice over IP. It is also good for HTTP use cases that need static IP addresses or fast regional failover

**NEW QUESTION 82**

- (Topic 2)

A company runs an Oracle database on premises. As part of the company's migration to AWS, the company wants to upgrade the database to the most recent available version. The company also wants to set up disaster recovery (DR) for the database. The company needs to minimize the operational overhead for normal operations and DR setup. The company also needs to maintain access to the database's underlying operating system.

Which solution will meet these requirements?

- A. Migrate the Oracle database to an Amazon EC2 instance
- B. Set up database replication to a different AWS Region.
- C. Migrate the Oracle database to Amazon RDS for Oracle
- D. Activate Cross-Region automated backups to replicate the snapshots to another AWS Region.
- E. Migrate the Oracle database to Amazon RDS Custom for Oracle
- F. Create a read replica for the database in another AWS Region.
- G. Migrate the Oracle database to Amazon RDS for Oracle
- H. Create a standby database in another Availability Zone.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-custom.html> and <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/working-with-custom-oracle.html>

**NEW QUESTION 84**

- (Topic 2)

A solutions architect is optimizing a website for an upcoming musical event. Videos of the performances will be streamed in real time and then will be available on demand. The event is expected to attract a global online audience.

Which service will improve the performance of both the real-time and on-demand streaming?

- A. Amazon CloudFront
- B. AWS Global Accelerator
- C. Amazon Route 53
- D. Amazon S3 Transfer Acceleration

**Answer: A**

**Explanation:**

You can use CloudFront to deliver video on demand (VOD) or live streaming video using any HTTP origin. One way you can set up video workflows in the cloud is by using CloudFront together with AWS Media Services. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand-streaming-video.html>

**NEW QUESTION 85**

- (Topic 2)

A global company is using Amazon API Gateway to design REST APIs for its loyalty club

users in the us-east-1 Region and the ap-southeast-2 Region. A solutions architect must design a solution to protect these API Gateway managed REST APIs across multiple accounts from SQL injection and cross-site scripting attacks.

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

- A. Set up AWS WAF in both Region
- B. Associate Regional web ACLs with an API stage.
- C. Set up AWS Firewall Manager in both Region
- D. Centrally configure AWS WAF rules.
- E. Set up AWS Shield in both Region
- F. Associate Regional web ACLs with an API stage.
- G. Set up AWS Shield in one of the Region
- H. Associate Regional web ACLs with an API stage.

**Answer: A**

**Explanation:**

Using AWS WAF has several benefits. Additional protection against web attacks using criteria that you specify. You can define criteria using characteristics of web requests such as the following: Presence of SQL code that is likely to be malicious (known as SQL injection). Presence of a script that is likely to be malicious (known as cross-site scripting). AWS Firewall Manager simplifies your administration and maintenance tasks across multiple accounts and resources for a variety of protections. <https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html>

**NEW QUESTION 86**

- (Topic 2)

A company wants to migrate its existing on-premises monolithic application to AWS.

The company wants to keep as much of the front-end code and the backend code as possible. However, the company wants to break the application into smaller applications. A different team will manage each application. The company needs a highly scalable solution that minimizes operational overhead.

Which solution will meet these requirements?

- A. Host the application on AWS Lambda Integrate the application with Amazon API Gateway.
- B. Host the application with AWS Amplif
- C. Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
- D. Host the application on Amazon EC2 instance
- E. Set up an Application Load Balancer with EC2 instances in an Auto Scaling group as targets.
- F. Host the application on Amazon Elastic Container Service (Amazon ECS) Set up an Application Load Balancer with Amazon ECS as the target.

**Answer:** D

**Explanation:**

<https://aws.amazon.com/blogs/compute/microservice-delivery-with-amazon-ecs-and-application-load-balancers/>

**NEW QUESTION 87**

- (Topic 2)

A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable. Which solution will meet these requirements?

- A. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive immediately.
- B. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 2 years.
- C. Use S3 Intelligent-Tierin
- D. Activate the archiving option to ensure that data is archived in S3 Glacier Deep Archive.
- E. Set up an S3 Lifecycle policy to transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately and to S3 Glacier Deep Archive after 2 years.

**Answer:** B

**Explanation:**

[https://aws.amazon.com/about-aws/whats-new/2018/04/announcing-s3-one-zone-infrequent-access-a-new-amazon-s3-storage-class/?nc1=h\\_ls](https://aws.amazon.com/about-aws/whats-new/2018/04/announcing-s3-one-zone-infrequent-access-a-new-amazon-s3-storage-class/?nc1=h_ls)

**NEW QUESTION 91**

- (Topic 2)

A company wants to run applications in containers in the AWS Cloud. These applications are stateless and can tolerate disruptions within the underlying infrastructure. The company needs a solution that minimizes cost and operational overhead. What should a solutions architect do to meet these requirements?

- A. Use Spot Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- B. Use Spot Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.
- C. Use On-Demand Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- D. Use On-Demand Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.

**Answer:** A

**Explanation:**

<https://aws.amazon.com/cn/blogs/compute/cost-optimization-and-resilience-eks-with-spot-instances/>

**NEW QUESTION 96**

- (Topic 2)

A solutions architect needs to securely store a database user name and password that an application uses to access an Amazon RDS DB instance. The application that accesses the database runs on an Amazon EC2 instance. The solutions architect wants to create a secure parameter in AWS Systems Manager Parameter Store.

What should the solutions architect do to meet this requirement?

- A. Create an IAM role that has read access to the Parameter Store paramete
- B. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
- C. Assign this IAM role to the EC2 instance.
- D. Create an IAM policy that allows read access to the Parameter Store paramete
- E. Allow Decrypt access to an AWS Key Management Service (AWS KMS) key that is used to encrypt the paramete
- F. Assign this IAM policy to the EC2 instance.
- G. Create an IAM trust relationship between the Parameter Store parameter and the EC2 instanc
- H. Specify Amazon RDS as a principal in the trust policy.
- I. Create an IAM trust relationship between the DB instance and the EC2 instanc
- J. Specify Systems Manager as a principal in the trust policy.

**Answer:** B

**Explanation:**

[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_aws-services-that-work-with-iam.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_aws-services-that-work-with-iam.html)

**NEW QUESTION 97**

- (Topic 2)

A company's website provides users with downloadable historical performance reports. The website needs a solution that will scale to meet the company's website demands globally. The solution should be cost-effective, limit the provisioning of infrastructure resources, and provide the fastest possible response time. Which combination should a solutions architect recommend to meet these requirements?

- A. Amazon CloudFront and Amazon S3
- B. AWS Lambda and Amazon DynamoDB
- C. Application Load Balancer with Amazon EC2 Auto Scaling
- D. Amazon Route 53 with internal Application Load Balancers

**Answer:** A

**Explanation:**

Cloudfront for rapid response and s3 to minimize infrastructure.

**NEW QUESTION 101**

- (Topic 2)

A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.

Which solution will meet these requirements?

- A. Use AWS DataSync to connect the S3 buckets to the web application.
- B. Deploy AWS Global Accelerator to connect the S3 buckets to the web application.
- C. Deploy Amazon CloudFront to connect the S3 buckets to CloudFront edge servers.
- D. Use Amazon Simple Queue Service (Amazon SQS) to connect the S3 buckets to the web application.

**Answer:** C

**Explanation:**

CloudFront uses a local cache to provide the response, AWS Global accelerator proxies requests and connects to the application all the time for the response.  
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html#private-content-granting-permissions-to-oai>

**NEW QUESTION 102**

- (Topic 3)

A media company hosts its website on AWS. The website application's architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora. The company's cyber security team reports that the application is vulnerable to SQL injection.

How should the company resolve this issue?

- A. Use AWS WAF in front of the ALB Associate the appropriate web ACLs with AWS WAF.
- B. Create an ALB listener rule to reply to SQL injection with a fixed response
- C. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.
- D. Set up Amazon Inspector to block all SQL injection attempts automatically

**Answer:** A

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/waf-block-common-attacks/#:~:text=To%20protect%20your%20applications%20against,%2C%20query%20string%2C%20or%20URI.> -----

----- Protect against SQL injection and cross-site scripting To protect your applications against SQL injection and cross-site scripting (XSS) attacks, use the built-in SQL injection and cross-site scripting engines. Remember that attacks can be performed on different parts of the HTTP request, such as the HTTP header, query string, or URI. Configure the AWS WAF rules to inspect different parts of the HTTP request against the built-in mitigation engines.

**NEW QUESTION 105**

- (Topic 3)

An ecommerce company is experiencing an increase in user traffic. The company's store is deployed on Amazon EC2 instances as a two-tier web application consisting of a web tier and a separate database tier. As traffic increases, the company notices that the architecture is causing significant delays in sending timely marketing and order confirmation email to users. The company wants to reduce the time it spends resolving complex email delivery issues and minimize operational overhead.

What should a solutions architect do to meet these requirements?

- A. Create a separate application tier using EC2 instances dedicated to email processing.
- B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES).
- C. Configure the web instance to send email through Amazon Simple Notification Service (Amazon SNS)
- D. Create a separate application tier using EC2 instances dedicated to email processing
- E. Place the instances in an Auto Scaling group.

**Answer:** B

**Explanation:**

Amazon SES is a cost-effective and scalable email service that enables businesses to send and receive email using their own email addresses and domains. Configuring the web instance to send email through Amazon SES is a simple and effective solution that can reduce the time spent resolving complex email delivery issues and minimize operational overhead.

**NEW QUESTION 109**

- (Topic 3)

A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.

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

- A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge.
- B. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
- C. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
- D. Store the server-side code on Amazon FSx for Windows File Share
- E. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
- F. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volume
- G. Mount the EBS volume on each EC2 instance to share the files.

**Answer:** AD

**Explanation:**

A because Elasticache, despite being ideal for leaderboards per Amazon, doesn't cache at edge locations. D because FSx has higher performance for low latency needs. <https://www.techtarget.com/searchaws/tip/Amazon-FSx-vs-EFS-Compare-the-AWS-file-services> "FSx is built for high performance and submillisecond latency using solid-state drive storage volumes. This design enables users to select storage capacity and latency independently. Thus, even a subterabyte file system can have 256 Mbps or higher throughput and support volumes up to 64 TB."

Amazon S3 is an object storage service that can store static files such as images, videos, documents, etc. Amazon EFS is a file storage service that can store files in a hierarchical structure and supports NFS protocol. Amazon FSx for Windows File Server is a file storage service that can store files in a hierarchical structure and supports SMB protocol. Amazon EBS is a block storage service that can store data in fixed-size blocks and attach to EC2 instances.

Based on these definitions, the combination of steps that should be taken to meet the requirements are:

\* A. Store the static files on Amazon S3. Use Amazon CloudFront to cache objects at the edge. D. Store the server-side code on Amazon FSx for Windows File Server. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.

**NEW QUESTION 111**

- (Topic 3)

A company wants to use high performance computing (HPC) infrastructure on AWS for financial risk modeling. The company's HPC workloads run on Linux. Each HPC workflow runs on hundreds of Amazon EC2 Spot Instances, is short-lived, and generates thousands of output files that are ultimately stored in persistent storage for analytics and long-term future use.

The company seeks a cloud storage solution that permits the copying of on-premises data to long-term persistent storage to make data available for processing by all EC2 instances. The solution should also be a high performance file system that is integrated with persistent storage to read and write datasets and output files. Which combination of AWS services meets these requirements?

- A. Amazon FSx for Lustre integrated with Amazon S3
- B. Amazon FSx for Windows File Server integrated with Amazon S3
- C. Amazon S3 Glacier integrated with Amazon Elastic Block Store (Amazon EBS)
- D. Amazon S3 bucket with a VPC endpoint integrated with an Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume

**Answer:** A

**Explanation:**

<https://aws.amazon.com/fsx/lustre/>

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Many workloads such as machine learning, high performance computing (HPC), video rendering, and financial simulations depend on compute instances accessing the same set of data through high-performance shared storage.

**NEW QUESTION 113**

- (Topic 3)

A company hosts a marketing website in an on-premises data center. The website consists of static documents and runs on a single server. An administrator updates the website content infrequently and uses an SFTP client to upload new documents.

The company decides to host its website on AWS and to use Amazon CloudFront. The company's solutions architect creates a CloudFront distribution. The solutions architect

must design the most cost-effective and resilient architecture for website hosting to serve as the CloudFront origin.

Which solution will meet these requirements?

- A. Create a virtual server by using Amazon Lightsail
- B. Configure the web server in the Lightsail instance
- C. Upload website content by using an SFTP client.
- D. Create an AWS Auto Scaling group for Amazon EC2 instance
- E. Use an Application Load Balance
- F. Upload website content by using an SFTP client.
- G. Create a private Amazon S3 bucket
- H. Use an S3 bucket policy to allow access from a CloudFront origin access identity (OAI). Upload website content by using the AWS CLI.
- I. Create a public Amazon S3 bucket
- J. Configure AWS Transfer for SFTP
- K. Configure the S3 bucket for website hosting
- L. Upload website content by using the SFTP client.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/cli/latest/reference/transfer/describe-server.html>

**NEW QUESTION 115**

- (Topic 3)

A telemarketing company is designing its customer call center functionality on AWS. The company needs a solution that provides multiple speaker recognition and generates transcript files. The company wants to query the transcript files to analyze the business patterns. The transcript files must be stored for 7 years for auditing purposes.

Which solution will meet these requirements?

- A. Use Amazon Recognition for multiple speaker recognition
- B. Store the transcript files in Amazon S3. Use machine learning models for transcript file analysis
- C. Use Amazon Transcribe for multiple speaker recognition
- D. Use Amazon Athena for transcript file analysis
- E. Use Amazon Translate for multiple speaker recognition
- F. Store the transcript files in Amazon Redshift. Use SQL queries for transcript file analysis
- G. Use Amazon Recognition for multiple speaker recognition
- H. Store the transcript files in Amazon S3. Use Amazon Textract for transcript file analysis

**Answer:** B

**Explanation:**

Amazon Transcribe now supports speaker labeling for streaming transcription. Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for you to convert speech-to-text. In live audio transcription, each stream of audio may contain multiple speakers. Now you can conveniently turn on the ability to label speakers, thus helping to identify who is saying what in the output transcript. <https://aws.amazon.com/about-aws/whats-new/2020/08/amazon-transcribe-supports-speaker-labeling-streaming-transcription/>

**NEW QUESTION 116**

- (Topic 3)

An Amazon EC2 instance is located in a private subnet in a new VPC. This subnet does not have outbound internet access, but the EC2 instance needs the ability to download monthly security updates from an outside vendor.

What should a solutions architect do to meet these requirements?

- A. Create an internet gateway, and attach it to the VP
- B. Configure the private subnet route table to use the internet gateway as the default route.
- C. Create a NAT gateway, and place it in a public subne
- D. Configure the private subnet route table to use the NAT gateway as the default route.
- E. Create a NAT instance, and place it in the same subnet where the EC2 instance is locate
- F. Configure the private subnet route table to use the NAT instance as the default route.
- G. Create an internet gateway, and attach it to the VP
- H. Create a NAT instance, and place it in the same subnet where the EC2 instance is locate
- I. Configure the private subnet route table to use the internet gateway as the default route.

**Answer: B**

**Explanation:**

This approach will allow the EC2 instance to access the internet and download the monthly security updates while still being located in a private subnet. By creating a NAT gateway and placing it in a public subnet, it will allow the instances in the private subnet to access the internet through the NAT gateway. And then, configure the private subnet route table to use the NAT gateway as the default route. This will ensure that all outbound traffic is directed through the NAT gateway, allowing the EC2 instance to access the internet while still maintaining the security of the private subnet.

**NEW QUESTION 117**

- (Topic 3)

A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task. The developer already has an IAM user with valid IAM credentials required for Amazon S3.

What should a solutions architect do to grant the permissions?

- A. Add required IAM permissions in the resource policy of the Lambda function
- B. Create a signed request using the existing IAM credentials in the Lambda function
- C. Create a new IAM user and use the existing IAM credentials in the Lambda function.
- D. Create an IAM execution role with the required permissions and attach the IAM role to the Lambda function

**Answer: D**

**Explanation:**

To grant the necessary permissions to an AWS Lambda function to upload files to Amazon S3, a solutions architect should create an IAM execution role with the required permissions and attach the IAM role to the Lambda function. This approach follows the principle of least privilege and ensures that the Lambda function can only access the resources it needs to perform its specific task.

**NEW QUESTION 120**

- (Topic 3)

A company has a multi-tier application deployed on several Amazon EC2 instances in an Auto Scaling group. An Amazon RDS for Oracle instance is the application's data layer that uses Oracle-specific

PL/SQL functions. Traffic to the application has been steadily increasing. This is causing the EC2 instances to become overloaded and the RDS instance to run out of storage. The Auto Scaling group does not have any scaling metrics and defines the minimum healthy instance count only. The company predicts that traffic will continue to increase at a steady but unpredictable rate before levelling off.

What should a solutions architect do to ensure the system can automatically scale for the increased traffic? (Select TWO.)

- A. Configure storage Auto Scaling on the RDS for Oracle Instance.
- B. Migrate the database to Amazon Aurora to use Auto Scaling storage.
- C. Configure an alarm on the RDS for Oracle Instance for low free storage space
- D. Configure the Auto Scaling group to use the average CPU as the scaling metric
- E. Configure the Auto Scaling group to use the average free memory as the seeing metric

**Answer: AD**

**Explanation:**

Auto scaling storage RDS will ease storage issues and migrating Oracle Pl/Sql to Aurora is cumbersome. Also Aurora has auto storage scaling by default. [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_PIOPS.StorageTypes.html#USER\\_PIOPS.Autoscaling](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.Autoscaling)

**NEW QUESTION 124**

- (Topic 3)

A solution architect needs to assign a new microservice for a company's application. Clients must be able to call an HTTPS endpoint to reach the microservice. The microservice also must use AWS identity and Access Management (IAM) to authentication calls. The solutions architect will write the logic for this microservice by using a single AWS Lambda function that is written in Go 1.x.

Which solution will deploy the function in the in the MOST operationally efficient way?

- A. Create an Amazon API Gateway REST AP
- B. Configure the method to use the Lambda functio
- C. Enable IAM authentication on the API.

- D. Create a Lambda function URL for the function.
- E. Specify AWS\_IAM as the authentication type.
- F. Create an Amazon CloudFront distribution.
- G. Deploy the function to Lambda@Edge.
- H. Integrate IAM authentication logic into the Lambda@Edge function.
- I. Create an Amazon CloudFront distribution.
- J. Deploy the function to CloudFront Function.
- K. Specify AWS\_IAM as the authentication type.

**Answer:** A

**Explanation:**

A. Create an Amazon API Gateway REST API. Configure the method to use the Lambda function. Enable IAM authentication on the API. This option is the most operationally efficient as it allows you to use API Gateway to handle the HTTPS endpoint and also allows you to use IAM to authenticate the calls to the microservice. API Gateway also provides many additional features such as caching, throttling, and monitoring, which can be useful for a microservice.

**NEW QUESTION 126**

- (Topic 3)

An IAM user made several configuration changes to AWS resources in their company's account during a production deployment last week. A solutions architect learned that a couple of security group rules are not configured as desired. The solutions architect wants to confirm which IAM user was responsible for making changes.

Which service should the solutions architect use to find the desired information?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS CloudTrail
- D. AWS Config

**Answer:** C

**Explanation:**

The best option is to use AWS CloudTrail to find the desired information. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of AWS account activities. CloudTrail can be used to log all changes made to resources in an AWS account, including changes made by IAM users, EC2 instances, AWS management console, and other AWS services. By using CloudTrail, the solutions architect can identify the IAM user who made the configuration changes to the security group rules.

**NEW QUESTION 131**

- (Topic 3)

A company is hosting a three-tier ecommerce application in the AWS Cloud. The company hosts the website on Amazon S3 and integrates the website with an API that handles sales requests. The company hosts the API on three Amazon EC2 instances behind an Application Load Balancer (ALB). The API consists of static and dynamic front-end content along with backend workers that process sales requests asynchronously.

The company is expecting a significant and sudden increase in the number of sales requests during events for the launch of new products.

What should a solutions architect recommend to ensure that all the requests are processed successfully?

- A. Add an Amazon CloudFront distribution for the dynamic content.
- B. Increase the number of EC2 instances to handle the increase in traffic.
- C. Add an Amazon CloudFront distribution for the static content.
- D. Place the EC2 instances in an Auto Scaling group to launch new instances based on network traffic.
- E. Add an Amazon CloudFront distribution for the dynamic content.
- F. Add an Amazon ElastiCache instance in front of the ALB to reduce traffic for the API to handle.
- G. Add an Amazon CloudFront distribution for the static content.
- H. Add an Amazon Simple Queue Service (Amazon SQS) queue to receive requests from the website for later processing by the EC2 instances.

**Answer:** B

**Explanation:**

This option is the most efficient because it uses Amazon CloudFront, which is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. It also uses a CloudFront distribution for the static content, which reduces the load on the EC2 instances and improves the performance and availability of the website. It also uses an Auto Scaling group to launch new instances based on network traffic, which automatically adjusts the compute capacity of your EC2 instances based on load or a schedule. This solution meets the requirement of ensuring that all the requests are processed successfully during events for the launch of new products. Option A is less efficient because it uses a CloudFront distribution for the dynamic content, which is not necessary as the dynamic content is already handled by the API on the EC2 instances. It also increases the number of EC2 instances to handle the increase in traffic, which could incur higher costs and complexity than using an Auto Scaling group. Option C is less efficient because it uses an Amazon ElastiCache instance in front of the ALB to reduce traffic for the API to handle, which is a way to provide a fully managed in-memory data store service that provides sub-millisecond latency for caching and data processing. However, this could introduce additional complexity and latency, and does not scale automatically based on network traffic. Option D is less efficient because it uses an Amazon Simple Queue Service (Amazon SQS) queue to receive requests from the website for later processing by the EC2 instances, which is a way to send, store, and receive messages between software components at any volume. However, this does not provide a faster response time to the users as they have to wait for their requests to be processed by the EC2 instances.

**NEW QUESTION 136**

- (Topic 3)

A company has a three-tier application on AWS that ingests sensor data from its users' devices. The traffic flows through a Network Load Balancer (NLB) then to Amazon EC2 instances for the web tier and finally to EC2 instances for the application tier. The application tier makes calls to a database.

What should a solutions architect do to improve the security of the data in transit?

- A. Configure a TLS listener. Deploy the server certificate on the NLB.
- B. Configure AWS Shield Advanced. Enable AWS WAF on the NLB.
- C. Change the load balancer to an Application Load Balancer (ALB). Enable AWS WAF on the ALB.
- D. Encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances by using AWS Key Management Service (AWS KMS).

**Answer:** A

**Explanation:**

The best option to improve the security of the data in transit is to configure a TLS listener and deploy the server certificate on the NLB. This will ensure that the data is encrypted and secure as it travels through the network. Additionally, you could also configure AWS Shield Advanced and enable AWS WAF on the NLB to further protect the network from malicious attacks. Alternatively, you could also change the load balancer to an Application Load Balancer (ALB) and enable AWS WAF on the ALB. Finally, you could also encrypt the Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instances by using AWS Key Management Service (AWS KMS).

You must specify an SSL certificate for a TLS listener. The load balancer uses the certificate to terminate the connection and decrypt requests from clients before routing them to targets. <https://docs.aws.amazon.com/elasticloadbalancing/latest/network/create-listener.html>

**NEW QUESTION 141**

- (Topic 3)

A company must migrate 20 TB of data from a data center to the AWS Cloud within 30 days. The company's network bandwidth is limited to 15 Mbps and cannot exceed 70% utilization. What should a solutions architect do to meet these requirements?

- A. Use AWS Snowball.
- B. Use AWS DataSync.
- C. Use a secure VPN connection.
- D. Use Amazon S3 Transfer Acceleration.

**Answer:** A

**Explanation:**

AWS Snowball is a secure data transport solution that accelerates moving large amounts of data into and out of the AWS cloud. It can move up to 80 TB of data at a time, and provides a network bandwidth of up to 50 Mbps, so it is well-suited for the task. Additionally, it is secure and easy to use, making it the ideal solution for this migration.

**NEW QUESTION 146**

- (Topic 3)

An application runs on Amazon EC2 instances in private subnets. The application needs to access an Amazon DynamoDB table. What is the MOST secure way to access the table while ensuring that the traffic does not leave the AWS network?

- A. Use a VPC endpoint for DynamoDB.
- B. Use a NAT gateway in a public subnet.
- C. Use a NAT instance in a private subnet.
- D. Use the internet gateway attached to the VPC.

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

A VPC endpoint for DynamoDB enables Amazon EC2 instances in your VPC to use their private IP addresses to access DynamoDB with no exposure to the public internet. Your EC2 instances do not require public IP addresses, and you don't need an internet gateway, a NAT device, or a virtual private gateway in your VPC. You use endpoint policies to control access to DynamoDB. Traffic between your VPC and the AWS service does not leave the Amazon network.

**NEW QUESTION 148**

- (Topic 3)

A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time.

The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy. Which solution meets these requirements?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon S3 Glacier Deep Archive
- D. AWS Backup

**Answer:** A

**Explanation:**

Amazon EFS provides a simple, scalable, fully managed file system that can be simultaneously accessed from multiple EC2 instances and provides built-in redundancy. It is optimized for multiple EC2 instances to access the same files, and it is designed to be highly available, durable, and secure. It can scale up to petabytes of data and can handle thousands of concurrent connections, and is a cost-effective solution for storing and accessing large amounts of data.

**NEW QUESTION 150**

- (Topic 3)

A company wants to configure its Amazon CloudFront distribution to use SSL/TLS certificates. The company does not want to use the default domain name for the distribution. Instead, the company wants to use a different domain name for the distribution.

Which solution will deploy the certificate with incurring any additional costs?

- A. Request an Amazon issued private certificate from AWS Certificate Manager (ACM) in the us-east-1 Region
- B. Request an Amazon issued private certificate from AWS Certificate Manager (ACM) in the us-west-1 Region.
- C. Request an Amazon issued public certificate from AWS Certificate Manager (ACU) in the us-east-1 Region
- D. Request an Amazon issued public certificate from AWS Certificate Manager (ACU) in the us-west-1 Region.

**Answer:** C

**Explanation:**

This option is the most efficient because it requests an Amazon issued public certificate from AWS Certificate Manager (ACM), which is a service that lets you easily provision, manage, and deploy public and private SSL/TLS certificates for use with AWS services and your internal connected resources<sup>1</sup>. It also requests the certificate in the us-east-1 Region, which is required for using an ACM certificate with CloudFront<sup>2</sup>. It also meets the requirement of deploying the certificate without incurring any additional costs, as ACM does not charge for certificates that are used with supported AWS services<sup>3</sup>. This solution meets the requirement of configuring its CloudFront distribution to use SSL/TLS certificates and using a different domain name for the distribution. Option A is less efficient because it requests an Amazon issued private certificate from ACM, which is a type of certificate that can be used only within your organization or virtual private cloud (VPC). However, this does not meet the requirement of configuring its CloudFront distribution to use SSL/TLS certificates, as CloudFront requires a public certificate. It also requests the certificate in the us-east-1 Region, which is correct. Option B is less efficient because it requests an Amazon issued private certificate from ACM, which is incorrect for the same reason as option A. It also requests the certificate in the us-west-1 Region, which is incorrect as CloudFront requires a certificate in the us-east-1 Region. Option D is less efficient because it requests an Amazon issued public certificate from ACM, which is correct. However, it requests the certificate in the us-west-1 Region, which is incorrect as CloudFront requires a certificate in the us-east-1 Region.

#### NEW QUESTION 152

- (Topic 3)

A company provides an API to its users that automates inquiries for tax computations based on item prices. The company experiences a larger number of inquiries during the holiday season only that cause slower response times. A solutions architect needs to design a solution that is scalable and elastic. What should the solutions architect do to accomplish this?

- A. Provide an API hosted on an Amazon EC2 instance
- B. The EC2 instance performs the required computations when the API request is made.
- C. Design a REST API using Amazon API Gateway that accepts the item name
- D. API Gateway passes item names to AWS Lambda for tax computations.
- E. Create an Application Load Balancer that has two Amazon EC2 instances behind it
- F. The EC2 instances will compute the tax on the received item names.
- G. Design a REST API using Amazon API Gateway that connects with an API hosted on an Amazon EC2 instance
- H. API Gateway accepts and passes the item names to the EC2 instance for tax computations.

**Answer: B**

#### Explanation:

Lambda server-less is scalable and elastic than EC2 api gateway solution

#### NEW QUESTION 155

- (Topic 3)

A company uses a payment processing system that requires messages for a particular payment ID to be received in the same order that they were sent. Otherwise, the payments might be processed incorrectly.

Which actions should a solutions architect take to meet this requirement? (Select TWO.)

- A. Write the messages to an Amazon DynamoDB table with the payment ID as the partition key
- B. Write the messages to an Amazon Kinesis data stream with the payment ID as the partition key.
- C. Write the messages to an Amazon ElastiCache for Memcached cluster with the payment ID as the key
- D. Write the messages to an Amazon Simple Queue Service (Amazon SQS) queue. Set the message attribute to use the payment ID
- E. Write the messages to an Amazon Simple Queue Service (Amazon SQS) FIFO queue
- F. Set the message group to use the payment ID.

**Answer: BE**

#### Explanation:

1) SQS FIFO queues guarantee that messages are received in the exact order they are sent. Using the payment ID as the message group ensures all messages for a payment ID are received sequentially. 2) Kinesis data streams can also enforce ordering on a per partition key basis. Using the payment ID as the partition key will ensure strict ordering of messages for each payment ID.

#### NEW QUESTION 157

- (Topic 3)

A company is launching a new application deployed on an Amazon Elastic Container Service (Amazon ECS) cluster and is using the Fargate launch type for ECS tasks. The company is monitoring CPU and memory usage because it is expecting high traffic to the application upon its launch. However, the company wants to reduce costs when utilization decreases.

What should a solutions architect recommend?

- A. Use Amazon EC2 Auto Scaling to scale at certain periods based on previous traffic patterns
- B. Use an AWS Lambda function to scale Amazon ECS based on metric breaches that trigger an Amazon CloudWatch alarm
- C. Use Amazon EC2 Auto Scaling with simple scaling policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm
- D. Use AWS Application Auto Scaling with target tracking policies to scale when ECS metric breaches trigger an Amazon CloudWatch alarm

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 162

- (Topic 3)

A company plans to use Amazon ElastiCache for its multi-tier web application. A solutions architect creates a Cache VPC for the ElastiCache cluster and an App VPC for the application's Amazon EC2 instances. Both VPCs are in the us-east-1 Region.

The solutions architect must implement a solution to provide the application's EC2 instances with access to the ElastiCache cluster.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a peering connection between the VPCs. Add a route table entry for the peering connection in both VPCs. Configure an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group.
- B. Create a Transit VPC. Update the VPC route tables in the Cache VPC and the App VPC to route traffic through the Transit VPC. Configure an inbound rule for

the ElastiCache cluster's security group to allow inbound connection from the application's security group  
C. Create a peering connection between the VPCs Add a route table entry for the peering connection in both VPCs Configure an inbound rule for the peering connection's security group to allow inbound connection from the application's security group  
D. Create a Transit VPC Update the VPC route tables in the Cache VPC and the App VPC to route traffic through the Transit VPC Configure an inbound rule for the Transit VPCs security group to allow inbound connection from the application's security group

**Answer:** A

**Explanation:**

Creating a peering connection between the two VPCs and configuring an inbound rule for the ElastiCache cluster's security group to allow inbound connection from the application's security group is the most cost-effective solution. Peering connections are free and you only incur the cost of configuring the security group rules. The Transit VPC solution requires additional VPCs and associated resources, which would incur additional costs.

Before Testing | AWS Certification Information and Policies | AWS <https://aws.amazon.com/certification/policies/before-testing/>

**NEW QUESTION 164**

- (Topic 3)

A development team has launched a new application that is hosted on Amazon EC2 instances inside a development VPC. A solution architect needs to create a new VPC in the same account. The new VPC will be peered with the development VPC. The VPC CIDR block for the development VPC is 192.168.0.0/24. The solutions architect needs to create a CIDR block for the new VPC. The CIDR block must be valid for a VPC peering connection to the development VPC. What is the SMALLEST CIDR block that meets these requirements?

- A. 10.0.1.0/32
- B. 192.168.0.0/24
- C. 192.168.1.0/32
- D. 10.0.1.0/24

**Answer:** D

**Explanation:**

The allowed block size is between a /28 netmask and /16 netmask. The CIDR block must not overlap with any existing CIDR block that's associated with the VPC. <https://docs.aws.amazon.com/vpc/latest/userguide/configure-your-vpc.html>

**NEW QUESTION 169**

- (Topic 3)

A company's application runs on AWS. The application stores large documents in an Amazon S3 bucket that uses the S3 Standard-infrequent Access (S3 Standard-IA) storage class. The company will continue paying to store the data but wants to save on its total S3 costs. The company wants authorized external users to have the ability to access the documents in milliseconds. Which solution will meet these requirements MOST cost-effectively?

- A. Configure the S3 bucket to be a Requester Pays bucket
- B. Change the storage tier to S3 Standard for all existing and future objects.
- C. Turn on S3 Transfer Acceleration for the S3 bucket
- D. Use Amazon CloudFront to handle all the requests to the S3 bucket

**Answer:** D

**Explanation:**

This option is the most efficient because it uses Amazon CloudFront, which is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users<sup>1</sup>. It also uses CloudFront to handle all the requests to the S3 bucket, which reduces the S3 costs by caching the content at the edge locations and serving it from there. It also allows authorized external users to access the documents in milliseconds, as CloudFront delivers the content with low latency and high data transfer rates. This solution meets the requirement of continuing paying to store the data but saving on its total S3 costs. Option A is less efficient because it configures the S3 bucket to be a Requester Pays bucket, which is a way to shift the cost of data transfer and requests from the bucket owner to the requester<sup>2</sup>. However, this does not reduce the total S3 costs, as the company still has to pay for storing the data and for any requests made by its own users. Option B is less efficient because it changes the storage tier to S3 Standard for all existing and future objects, which is a way to store frequently accessed data with high durability and availability<sup>3</sup>. However, this does not reduce the total S3 costs, as S3 Standard has higher storage costs than S3 Standard-IA. Option C is less efficient because it turns on S3 Transfer Acceleration for the S3 bucket, which is a way to speed up transfers into and out of an S3 bucket by routing requests through CloudFront edge locations<sup>4</sup>. However, this does not reduce the total S3 costs, as S3 Transfer Acceleration has additional charges for data transfer and requests.

**NEW QUESTION 171**

- (Topic 3)

A company hosts a three-tier web application that includes a PostgreSQL database The database stores the metadata from documents The company searches the metadata for key terms to retrieve documents that the company reviews in a report each month The documents are stored in Amazon S3 The documents are usually written only once, but they are updated frequently The reporting process takes a few hours with the use of relational queries The reporting process must not affect any document modifications or the addition of new documents.

What are the MOST operationally efficient solutions that meet these requirements? (Select TWO )

- A. Set up a new Amazon DocumentDB (with MongoDB compatibility) cluster that includes a read replica Scale the read replica to generate the reports.
- B. Set up a new Amazon RDS for PostgreSQL Reserved Instance and an On-Demand read replica Scale the read replica to generate the reports
- C. Set up a new Amazon Aurora PostgreSQL DB cluster that includes a Reserved Instance and an Aurora Replica issue queries to the Aurora Replica to generate the reports.
- D. Set up a new Amazon RDS for PostgreSQL Multi-AZ Reserved Instance Configure the reporting module to query the secondary RDS node so that the reporting module does not affect the primary node
- E. Set up a new Amazon DynamoDB table to store the documents Use a fixed write capacity to support new document entries Automatically scale the read capacity to support the reports

**Answer:** BC

**Explanation:**

These options are operationally efficient because they use Amazon RDS read replicas to offload the reporting workload from the primary DB instance and avoid affecting any document modifications or the addition of new documents<sup>1</sup>. They also use Reserved Instances for the primary DB instance to reduce costs and On-Demand or Aurora Replicas for the read replicas to scale as needed. Option A is less efficient because it uses Amazon S3 Glacier Flexible Retrieval, which is a cold storage class that has higher retrieval costs and longer retrieval times than Amazon S3 Standard. It also uses EventBridge rules to invoke the job nightly, which does not meet the requirement of processing incoming data files as soon as possible. Option D is less efficient because it uses AWS Lambda to process the files, which has a maximum execution time of 15 minutes per invocation, which might not be enough for processing each file that needs 3-8 minutes. It also uses S3 event notifications to invoke the Lambda function when the files arrive, which could cause concurrency issues if there are thousands of small data files arriving periodically. Option E is less efficient because it uses Amazon DynamoDB, which is a NoSQL database service that does not support relational queries, which are needed for generating the reports. It also uses fixed write capacity, which could cause throttling or underutilization depending on the incoming data files.

#### NEW QUESTION 175

- (Topic 3)

A company stores confidential data in an Amazon Aurora PostgreSQL database in the ap-southeast-3 Region. The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key. The company was recently acquired and must securely share a backup of the database with the acquiring company's AWS account in ap-southeast-3.

What should a solutions architect do to meet these requirements?

- A. Create a database snapshot. Copy the snapshot to a new unencrypted snapshot. Share the new snapshot with the acquiring company's AWS account.
- B. Create a database snapshot. Add the acquiring company's AWS account to the KMS key policy. Share the snapshot with the acquiring company's AWS account.
- C. Create a database snapshot that uses a different AWS managed KMS key. Add the acquiring company's AWS account to the KMS key policy.
- D. Share the snapshot with the acquiring company's AWS account.
- E. Create a database snapshot. Download the database snapshot. Upload the database snapshot to an Amazon S3 bucket. Update the S3 bucket policy to allow access from the acquiring company's AWS account.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html>

There's no need to create another custom AWS KMS key. <https://aws.amazon.com/premiumsupport/knowledge-center/aurora-share-encrypted-snapshot/> Give target account access to the custom AWS KMS key within the source account. 1. Log in to the source account, and go to the AWS KMS console in the same Region as the DB cluster snapshot. 2. Select Customer-managed keys from the navigation pane. 3. Select your custom AWS KMS key (ALREADY CREATED). 4. From the Other AWS accounts section, select Add another AWS account, and then enter the AWS account number of your target account. Then: Copy and share the DB cluster snapshot.

#### NEW QUESTION 178

- (Topic 3)

A company hosts a multi-tier web application that uses an Amazon Aurora MySQL DB cluster for storage. The application tier is hosted on Amazon EC2 instances. The company's IT security guidelines mandate that the database credentials be encrypted and rotated every 14 days.

What should a solutions architect do to meet this requirement with the LEAST operational effort?

- A. Create a new AWS Key Management Service (AWS KMS) encryption key. Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Associate the secret with the Aurora DB cluster. Configure a custom rotation period of 14 days.
- B. Create two parameters in AWS Systems Manager Parameter Store: one for the username as a string parameter and one that uses the SecureString type for the password. Select AWS Key Management Service (AWS KMS) encryption for the password parameter, and load these parameters in the application tier. Implement an AWS Lambda function that rotates the password every 14 days.
- C. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system in all EC2 instances of the application tier.
- D. Restrict the access to the file on the file system so that the application can read the file and that only super users can modify the file. Implement an AWS Lambda function that rotates the key in Aurora every 14 days and writes new credentials into the file.
- E. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon S3 bucket that the application uses to load the credentials. Download the file to the application regularly to ensure that the correct credentials are used. Implement an AWS Lambda function that rotates the Aurora credentials every 14 days and uploads these credentials to the file in the S3 bucket.

**Answer: A**

#### Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types-oracle/>

#### NEW QUESTION 182

- (Topic 3)

A company has implemented a self-managed DNS service on AWS. The solution consists of the following:

- Amazon EC2 instances in different AWS Regions
- Endpoints of a standard accelerator in AWS Global Accelerator

The company wants to protect the solution against DDoS attacks. What should a solutions architect do to meet this requirement?

- A. Subscribe to AWS Shield Advanced. Add the accelerator as a resource to protect.
- B. Subscribe to AWS Shield Advanced. Add the EC2 instances as resources to protect.
- C. Create an AWS WAF web ACL that includes a rate-based rule. Associate the web ACL with the accelerator.
- D. Create an AWS WAF web ACL that includes a rate-based rule. Associate the web ACL with the EC2 instances.

**Answer: A**

#### Explanation:

AWS Shield is a managed service that provides protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. AWS Shield Standard is automatically enabled to all AWS customers at no additional cost. AWS Shield Advanced is an optional paid service. AWS Shield Advanced provides additional protections against more sophisticated and larger attacks for your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Route 53. <https://docs.aws.amazon.com/waf/latest/developerguide/ddos-event-mitigation-logic-gax.html>

#### NEW QUESTION 187

- (Topic 3)

A company runs an application on Amazon EC2 Linux instances across multiple Availability Zones. The application needs a storage layer that is highly available and Portable Operating System Interface (POSIX) compliant. The storage layer must provide maximum data durability and must be shareable across the EC2 instances. The data in the storage layer will be accessed frequently for the first 30 days and will be accessed infrequently after that time. Which solution will meet these requirements MOST cost-effectively?

- A. Use the Amazon S3 Standard storage class. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Glacier.
- B. Use the Amazon S3 Standard storage class.
- C. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Standard-Infrequent Access (S3 Standard-IA).
- D. Use the Amazon Elastic File System (Amazon EFS) Standard storage class.
- E. Create a Lifecycle management policy to move infrequently accessed data to EFS Standard- Infrequent Access (EFS Standard-IA).
- F. Use the Amazon Elastic File System (Amazon EFS) One Zone storage class.
- G. Create a Lifecycle management policy to move infrequently accessed data to EFS One Zone- Infrequent Access (EFS One Zone-IA).

**Answer: C**

**Explanation:**

<https://aws.amazon.com/efs/features/infrequent-access/>

#### NEW QUESTION 190

- (Topic 3)

A solutions architect is designing a multi-tier application for a company. The application's users upload images from a mobile device. The application generates a thumbnail of each image and returns a message to the user to confirm that the image was uploaded successfully. The thumbnail generation can take up to 60 seconds, but the company wants to provide a faster response time to its users to notify them that the original image was received. The solutions architect must design the application to asynchronously dispatch requests to the different application tiers. What should the solutions architect do to meet these requirements?

- A. Write a custom AWS Lambda function to generate the thumbnail and alert the user.
- B. Use the image upload process as an event source to invoke the Lambda function.
- C. Create an AWS Step Functions workflow. Configure Step Functions to handle the orchestration between the application tiers and alert the user when thumbnail generation is complete.
- D. Create an Amazon Simple Queue Service (Amazon SQS) message queue.
- E. As images are uploaded, place a message on the SQS queue for thumbnail generation.
- F. Alert the user through an application message that the image was received.
- G. Create Amazon Simple Notification Service (Amazon SNS) notification topics and subscriptions. Use one subscription with the application to generate the thumbnail after the image upload is complete.
- H. Use a second subscription to message the user's mobile app by way of a push notification after thumbnail generation is complete.

**Answer: C**

**Explanation:**

This option is the most efficient because it uses Amazon SQS, which is a fully managed message queuing service that lets you send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be available<sup>1</sup>. It also uses an SQS message queue to asynchronously dispatch requests to the different application tiers, which decouples the image upload process from the thumbnail generation process and enables scalability and reliability. It also alerts the user through an application message that the image was received, which provides a faster response time to the user than waiting for the thumbnail generation to complete. Option A is less efficient because it uses a custom AWS Lambda function to generate the thumbnail and alert the user, which is a way to run code without provisioning or managing servers. However, this does not use an asynchronous dispatch mechanism to separate the image upload process from the thumbnail generation process. It also uses the image upload process as an event source to invoke the Lambda function, which could cause concurrency issues if there are many images uploaded at once. Option B is less efficient because it uses AWS Step Functions, which is a fully managed service that provides a graphical console to arrange and visualize the components of your application as a series of steps<sup>2</sup>. However, this does not use an asynchronous dispatch mechanism to separate the image upload process from the thumbnail generation process. It also uses Step Functions to handle the orchestration between the application tiers and alert the user when thumbnail generation is complete, which could introduce additional complexity and latency. Option D is less efficient because it uses Amazon SNS, which is a fully managed messaging service that enables you to send messages or notifications directly to users with SMS text messages or email<sup>3</sup>. However, this does not use an asynchronous dispatch mechanism to separate the image upload process from the thumbnail generation process. It also uses SNS notification topics and subscriptions to generate the thumbnail after the image upload is complete and message the user's mobile app by way of a push notification after thumbnail generation is complete, which could introduce additional complexity and latency.

#### NEW QUESTION 191

- (Topic 3)

A company wants to use Amazon S3 for the secondary copy of its on-premises dataset. The company would rarely need to access this copy. The storage solution's cost should be minimal. Which storage solution meets these requirements?

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

**Answer: D**

**Explanation:**

S3 One Zone-IA is a storage class that is designed for data that is accessed less frequently, but requires rapid access when needed. Unlike other S3 Storage Classes which store data in a minimum of three Availability Zones (AZs), S3 One Zone-IA stores data in a single AZ and costs 20% less than S3 Standard-IA. This storage class meets the requirements of the company because it provides a low-cost solution for the secondary copy of its on-premises dataset that would rarely need to be accessed. The other storage classes are either more expensive or not suitable for infrequently accessed data.  
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html>

#### NEW QUESTION 192

- (Topic 3)

A company recently migrated its entire IT environment to the AWS Cloud. The company discovers that users are provisioning oversized Amazon EC2 instances

and modifying security group rules without using the appropriate change control process A solutions architect must devise a strategy to track and audit these inventory and configuration changes.

Which actions should the solutions architect take to meet these requirements? (Select TWO )

- A. Enable AWS CloudTrail and use it for auditing
- B. Use data lifecycle policies for the Amazon EC2 instances
- C. Enable AWS Trusted Advisor and reference the security dashboard
- D. Enable AWS Config and create rules for auditing and compliance purposes
- E. Restore previous resource configurations with an AWS CloudFormation template

**Answer:** AD

**Explanation:**

A) Enable AWS CloudTrail and use it for auditing. AWS CloudTrail provides a record of API calls and can be used to audit changes made to EC2 instances and security groups. By analyzing CloudTrail logs, the solutions architect can track who provisioned oversized instances or modified security groups without proper approval. D) Enable AWS Config and create rules for auditing and compliance purposes. AWS Config can record the configuration changes made to resources like EC2 instances and security groups. The solutions architect can create AWS Config rules to monitor for non-compliant changes, like launching certain instance types or opening security group ports without permission. AWS Config would alert on any violations of these rules.

**NEW QUESTION 193**

- (Topic 3)

A company wants to implement a disaster recovery plan for its primary on-premises file storage volume. The file storage volume is mounted from an Internet Small Computer Systems Interface (iSCSI) device on a local storage server. The file storage volume holds hundreds of terabytes (TB) of data.

The company wants to ensure that end users retain immediate access to all file types from the on-premises systems without experiencing latency.

Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

- A. Provision an Amazon S3 File Gateway as a virtual machine (VM) that is hosted on premise
- B. Set the local cache to 10 T
- C. Modify existing applications to access the files through the NFS protocol
- D. To recover from a disaster, provision an Amazon EC2 instance and mount the S3 bucket that contains the files.
- E. Provision an AWS Storage Gateway tape gateway
- F. Use a data backup solution to back up all existing data to a virtual tape library
- G. Configure the data backup solution to run nightly after the initial backup is complete
- H. To recover from a disaster, provision an Amazon EC2 instance and restore the data to an Amazon Elastic Block Store (Amazon EBS) volume from the volumes in the virtual tape library.
- I. Provision an AWS Storage Gateway Volume Gateway cached volume
- J. Set the local cache to 10 T
- K. Mount the Volume Gateway cached volume to the existing file server by using iSCSI
- L. and copy all files to the storage volume
- M. Configure scheduled snapshots of the storage volume
- N. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.
- O. Provision an AWS Storage Gateway Volume Gateway stored volume with the same amount of disk space as the existing file storage volume
- P. Mount the Volume Gateway stored volume to the existing file server by using iSCSI, and copy all files to the storage volume
- Q. Configure scheduled snapshots of the storage volume
- R. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.

**Answer:** D

**Explanation:**

"The company wants to ensure that end users retain immediate access to all file types from the on-premises systems" - Cached volumes: low latency access to most recent data - Stored volumes: entire dataset is on premise, scheduled backups to S3 Hence Volume Gateway stored volume is the apt choice.

**NEW QUESTION 196**

- (Topic 3)

A company uses a 100 GB Amazon RDS for Microsoft SQL Server Single-AZ DB instance in the us-east-1 Region to store customer transactions. The company needs high availability and automate recovery for the DB instance.

The company must also run reports on the RDS database several times a year. The report process causes transactions to take longer than usual to post to the customer's accounts.

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

- A. Modify the DB instance from a Single-AZ DB instance to a Multi-AZ deployment.
- B. Take a snapshot of the current DB instance
- C. Restore the snapshot to a new RDS deployment in another Availability Zone.
- D. Create a read replica of the DB instance in a different Availability Zone
- E. Point All requests for reports to the read replica.
- F. Migrate the database to RDS Custom.
- G. Use RDS Proxy to limit reporting requests to the maintenance window.

**Answer:** AC

**Explanation:**

<https://medium.com/awesome-cloud/aws-difference-between-multi-az-and-read-replicas-in-amazon-rds-60fe848ef53a>

**NEW QUESTION 201**

- (Topic 3)

A company hosts a web application on multiple Amazon EC2 instances The EC2 instances are in an Auto Scaling group that scales in response to user demand The company wants to optimize cost savings without making a long-term commitment

Which EC2 instance purchasing option should a solutions architect recommend to meet these requirements'?

- A. Dedicated Instances only
- B. On-Demand Instances only
- C. A mix of On-Demand instances and Spot Instances
- D. A mix of On-Demand instances and Reserved instances

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-mixed-instances-groups.html>

**NEW QUESTION 202**

- (Topic 3)

A company runs a fleet of web servers using an Amazon RDS for PostgreSQL DB instance. After a routine compliance check, the company sets a standard that requires a recovery point objective (RPO) of less than 1 second for all its production databases.

Which solution meets these requirements?

- A. Enable a Multi-AZ deployment for the DB Instance
- B. Enable auto scaling for the DB instance in one Availability Zone.
- C. Configure the DB instance in one Availability Zone and create multiple read replicas in a separate Availability Zone
- D. Configure the DB instance in one Availability Zone, and configure AWS Database Migration Service (AWS DMS) change data capture (CDC) tasks

**Answer:** A

**Explanation:**

This option is the most efficient because it uses a Multi-AZ deployment for the DB instance, which provides enhanced availability and durability for RDS database instances by automatically replicating the data to a standby instance in a different Availability Zone<sup>1</sup>. It also provides a recovery point objective (RPO) of less than 1 second for all its production databases, as the standby instance is kept in sync with the primary instance using synchronous physical replication<sup>2</sup>. This solution meets the requirement of requiring a RPO of less than 1 second for all its production databases. Option B is less efficient because it uses auto scaling for the DB instance in one Availability Zone, which is a way to automatically adjust the compute capacity of your DB instance based on load or a schedule<sup>3</sup>. However, this does not provide a RPO of less than 1 second for all its production databases, as it does not replicate the data to another Availability Zone. Option C is less efficient because it uses read replicas in a separate Availability Zone, which are read-only copies of your primary database that can serve read traffic and support scaling. However, this does not provide a RPO of less than 1 second for all its production databases, as read replicas use asynchronous replication and can lag behind the primary database. Option D is less efficient because it uses AWS Database Migration Service (AWS DMS) change data capture (CDC) tasks, which are tasks that capture changes made to source data and apply them to target data. However, this does not provide a RPO of less than 1 second for all its production databases, as AWS DMS uses asynchronous replication and can lag behind the source database.

**NEW QUESTION 205**

.....

## Relate Links

**100% Pass Your SAA-C03 Exam with Exambible Prep Materials**

<https://www.exambible.com/SAA-C03-exam/>

## Contact us

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

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