

# Amazon-Web-Services

## Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



#### NEW QUESTION 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:

Explanation

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>: "When you create a Capacity Reservation, you specify: The Availability Zone in which to reserve the capacity"

#### NEW QUESTION 2

A solutions architect is designing a new hybrid architecture to extend a company's on-premises infrastructure to AWS. The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.

What should the solutions architect do to meet these requirements?

- A. Provision an AWS Direct Connect connection to a Region. Provision a VPN connection as a backup if the primary Direct Connect connection fails.
- B. Provision a VPN tunnel connection to a Region for private connectivity.
- C. Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.
- D. Provision an AWS Direct Connect connection to a Region. Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
- E. Provision an AWS Direct Connect connection to a Region. Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

**Answer: A**

#### Explanation:

Explanation

"In some cases, this connection alone is not enough. It is always better to guarantee a fallback connection as the backup of DX. There are several options, but implementing it with an AWS Site-To-Site VPN is a real

cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX."

<https://www.proud2becloud.com/hybrid-cloud-networking-backup-aws-direct-connect-network-connection-with>

#### NEW QUESTION 3

A solutions architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack, and access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

**Answer: C**

#### Explanation:

Explanation

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html>

"With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it."

#### NEW QUESTION 4

A solutions architect is designing the architecture of a new application being deployed to the AWS Cloud. The application will run on Amazon EC2 On-Demand Instances and will automatically scale across multiple Availability Zones. The EC2 instances will scale up and down frequently throughout the day. An Application Load Balancer (ALB) will handle the load distribution. The architecture needs to support distributed session data management. The company is willing to make changes to code if needed.

What should the solutions architect do to ensure that the architecture supports distributed session data management?

- A. Use Amazon ElastiCache to manage and store session data.
- B. Use session affinity (sticky sessions) of the ALB to manage session data.
- C. Use Session Manager from AWS Systems Manager to manage the session.
- D. Use the GetSessionToken API operation in AWS Security Token Service (AWS STS) to manage the session.

**Answer: A**

#### Explanation:

Explanation

<https://aws.amazon.com/vi/caching/session-management/>

In order to address scalability and to provide a shared data storage for sessions that can be accessible from any individual web server, you can abstract the HTTP sessions from the web servers themselves. A common solution for this is to leverage an In-Memory Key/Value store such as Redis and Memcached.

ElastiCache offerings for In-Memory key/value stores include ElastiCache for Redis, which can support replication, and ElastiCache for Memcached which does

not support replication.

#### NEW QUESTION 5

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost. Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instance
- B. Stop the RDS database when it is not in use.
- C. Purchase EC2 Instance Savings Plans to cover five EC2 instance
- D. Purchase an RDS Reserved DB Instance
- E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as needed
- F. Stop the RDS database when it is not in use.
- G. Purchase EC2 Instance Savings Plans to cover two EC2 instance
- H. Use up to three additional EC2 On-Demand Instances as needed
- I. Purchase an RDS Reserved DB Instance.

**Answer:** D

#### NEW QUESTION 6

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:

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 suite & other Options 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

A company hosts a containerized web application on a fleet of on-premises servers that process incoming requests. The number of requests is growing quickly. The on-premises servers cannot handle the increased number of requests. The company wants to move the application to AWS with minimum code changes and minimum development effort.

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

- A. Use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling
- B. Use an Application Load Balancer to distribute the incoming requests.
- C. Use two Amazon EC2 instances to host the containerized web application
- D. Use an Application Load Balancer to distribute the incoming requests
- E. Use AWS Lambda with a new code that uses one of the supported languages
- F. Create multiple Lambda functions to support the load
- G. Use Amazon API Gateway as an entry point to the Lambda functions.
- H. Use a high performance computing (HPC) solution such as AWS ParallelCluster to establish an HPC cluster that can process the incoming requests at the appropriate scale.

**Answer:** A

#### NEW QUESTION 8

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.

What is the FASTEST way to aggregate data from all of these global sites?

- A. Enable Amazon S3 Transfer Acceleration on the destination bucket
- B. Use multipart uploads to directly upload site data to the destination bucket.
- C. Upload site data to an Amazon S3 bucket in the closest AWS Region
- D. Use S3 cross-Region replication to copy objects to the destination bucket.
- E. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region
- F. Use S3 cross-Region replication to copy objects to the destination bucket.
- G. Upload the data to an Amazon EC2 instance in the closest Region
- H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume

- I. Once a day take an EBS snapshot and copy it to the centralized Region
- J. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

**Answer:** A

**Explanation:**

Explanation

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:

You have customers that upload to a centralized bucket from all over the world.

You transfer gigabytes to terabytes of data on a regular basis across continents.

You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html>

[https://aws.amazon.com/s3/transferacceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications:](https://aws.amazon.com/s3/transferacceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications:)

"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much

as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile

applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"

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

"Improved throughput - You can upload parts in parallel to improve throughput."

**NEW QUESTION 9**

A company needs the ability to analyze the log files of its proprietary application. The logs are stored

in JSON format in an Amazon S3 bucket. Queries will be simple and will run on-demand. A solutions

architect needs to perform the analysis with minimal changes to the existing architecture.

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

- A. Use Amazon Redshift to load all the content into one place and run the SQL queries as needed
- B. Use Amazon CloudWatch Logs to store the logs. Run SQL queries as needed from the Amazon CloudWatch console
- C. Use Amazon Athena directly with Amazon S3 to run the queries as needed
- D. Use AWS Glue to catalog the logs. Use a transient Apache Spark cluster on Amazon EMR to run the SQL queries as needed

**Answer:** C

**Explanation:**

Explanation

Amazon Athena can be used to query JSON in S3.

**NEW QUESTION 10**

The management account has an Amazon S3 bucket that contains project reports. The company

wants to limit access to this S3 bucket to only users of accounts within the organization in AWS

Organizations.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Add the `aws:PrincipalOrgID` global condition key with a reference to the organization ID to the S3 bucket policy.
- B. Create an organizational unit (OU) for each department.
- C. Add the `aws:PrincipalOrgPaths` global condition key to the S3 bucket policy.
- D. Use AWS CloudTrail to monitor the `CreateAccount`, `InviteAccountToOrganization`, `LeaveOrganization`, and `RemoveAccountFromOrganization` events.
- E. Update the S3 bucket policy accordingly.
- F. Tag each user that needs access to the S3 bucket.
- G. Add the `aws:PrincipalTag` global condition key to the S3 bucket policy.

**Answer:** A

**Explanation:**

Explanation

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

The `aws:PrincipalOrgID` global key provides an alternative to listing all the account IDs for all AWS

accounts in an organization. For example, the following Amazon S3 bucket policy allows members of

any account in the XXX organization to add an object into the examtopics bucket.

```
{"Version": "2020-09-10",
```

```
"Statement": {
```

```
"Sid": "AllowPutObject",
```

```
"Effect": "Allow",
```

```
"Principal": "*",
```

```
"Action": "s3:PutObject",
```

```
"Resource": "arn:aws:s3:::examtopics/*",
```

```
"Condition": {"StringEquals":
```

```
{"aws:PrincipalOrgID":["XXX"]}}}
```

[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_condition-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html)

**NEW QUESTION 10**

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company

duplicate the architecture and created a second EC2 instance and EBS volume in another Availability

Zone, placing both behind an Application Load Balancer. After completing this change, users reported

that, each time they refreshed the website, they could see one subset of their documents or the

other, but never all of the documents at the same time.

What should a solutions architect propose to ensure users see all of their documents at once?

- A. Copy the data so both EBS volumes contain all the documents.
- B. Configure the Application Load Balancer to direct a user to the server with the documents
- C. Copy the data from both EBS volumes to Amazon EFS Modify the application to save newdocuments to Amazon EFS
- D. Configure the Application Load Balancer to send the request to both servers Return eachdocument from the correct server.

**Answer: C**

**Explanation:**

Explanation

Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System versions 4.0 and 4.1 (NFSv4) protocol. We recommend using a current generation Linux NFSv4.1 client, such as those found in the latest Amazon Linux, Redhat, and Ubuntu

AMIs, in conjunction with the Amazon EFS Mount Helper. For instructions, see Using the amazon-efsutils Tools.

For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client.

You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.

**NEW QUESTION 11**

A company observes an increase in Amazon EC2 costs in its most recent bill

The billing team notices unwanted vertical scaling of instance types for a couple of EC2 instances

A solutions architect needs to create a graph comparing the last 2 months of EC2 costs and perform an in-depth analysis to identify the root cause of the vertical scaling

How should the solutions architect generate the information with the LEAST operational overhead?

- A. Use AWS Budgets to create a budget report and compare EC2 costs based on instance types
- B. Use Cost Explorer's granular filtering feature to perform an in-depth analysis of EC2 costs based on instance types
- C. Use graphs from the AWS Billing and Cost Management dashboard to compare EC2 costs based on instance types for the last 2 months
- D. Use AWS Cost and Usage Reports to create a report and send it to an Amazon S3 bucket Use Amazon QuickSight with Amazon S3 as a source to generate an interactive graph based on instance types.

**Answer: B**

**Explanation:**

Explanation

AWS Cost Explorer is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your costs. <https://docs.aws.amazon.com/costmanagement/latest/userguide/ce-what-is.html>

**NEW QUESTION 12**

A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database.

During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort.

Which solution will meet these requirements?

- A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instances. Connect the database by using native Java Database Connectivity (JDBC) drivers.
- B. Change the platform from Aurora to Amazon DynamoD
- C. Provision a DynamoDB Accelerator (DAX) cluste
- D. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
- E. Set up two Lambda function
- F. Configure one function to receive the informatio
- G. Configure the other function to load the information into the databas
- H. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
- I. Set up two Lambda function
- J. Configure one function to receive the informatio
- K. Configure the other function to load the information into the databas
- L. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

**Answer: D**

**Explanation:**

Explanation

bottlenecks can be avoided with queues (SQS).

**NEW QUESTION 13**

A company is hosting a static website on Amazon S3 and is using Amazon Route 53 for DNS. The website is experiencing increased demand from around the world. The company must decrease latency for users who access the website.

Which solution meets these requirements MOST cost-effectively?

- A. Replicate the S3 bucket that contains the website to all AWS Region
- B. Add Route 53 geolocation routing entries.
- C. Provision accelerators in AWS Global Accelerator
- D. Associate the supplied IP addresses with the S3 bucket
- E. Edit the Route 53 entries to point to the IP addresses of the accelerators.
- F. Add an Amazon CloudFront distribution in front of the S3 bucket
- G. Edit the Route 53 entries to point to the CloudFront distribution.
- H. Enable S3 Transfer Acceleration on the bucket
- I. Edit the Route 53 entries to point to the new endpoint.

**Answer: C**

#### NEW QUESTION 18

A company has thousands of edge devices that collectively generate 1 TB of status alerts each day. Each alert is approximately 2 KB in size. A solutions architect needs to implement a solution to ingest and store the alerts for future analysis. The company wants a highly available solution. However, the company needs to minimize costs and does not want to manage additional infrastructure. Additionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days. What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts. Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket. Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days.
- B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts. Create a script on the EC2 instances that will store the alerts in an Amazon S3 bucket. Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days.
- C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts. Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) cluster. Set up the Amazon ES cluster to take manual snapshots every day and delete data from the cluster that is older than 14 days.
- D. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to ingest the alerts and set the message retention period to 14 days. Configure consumers to poll the SQS queue, check the age of the message, and analyze the message data as needed. If the message is 14 days old, the consumer should copy the message to an Amazon S3 bucket and delete the message from the SQS queue.

**Answer: A**

#### Explanation:

Explanation

<https://aws.amazon.com/kinesis/datafirehose/features/?nc=sn&loc=2#:~:text=into%20Amazon%20S3%2C%20Amazon%20Redshift%2C%20Amazon%20OpenSearch%20Service%2C%20Kinesis,Delivery%20streams>

#### NEW QUESTION 22

A company runs a highly available image-processing application on Amazon EC2 instances in a single VPC. The EC2 instances run inside several subnets across multiple Availability Zones. The EC2 instances do not communicate with each other. However, the EC2 instances download images from Amazon S3 and upload images to Amazon S3 through a single NAT gateway. The company is concerned about data transfer charges. What is the MOST cost-effective way for the company to avoid Regional data transfer charges?

- A. Launch the NAT gateway in each Availability Zone.
- B. Replace the NAT gateway with a NAT instance.
- C. Deploy a gateway VPC endpoint for Amazon S3.
- D. Provision an EC2 Dedicated Host to run the EC2 instances.

**Answer: C**

#### NEW QUESTION 24

A company has an on-premises application that generates a large amount of time-sensitive data that is backed up to Amazon S3. The application has grown and there are user complaints about internet bandwidth limitations. A solutions architect needs to design a long-term solution that allows for both timely backups to Amazon S3 and with minimal impact on internet connectivity for internal users. Which solution meets these requirements?

- A. Establish AWS VPN connections and proxy all traffic through a VPC gateway endpoint.
- B. Establish a new AWS Direct Connect connection and direct backup traffic through this new connection.
- C. Order daily AWS Snowball devices. Load the data onto the Snowball devices and return the devices to AWS each day.
- D. Submit a support ticket through the AWS Management Console. Request the removal of S3 service limits from the account.

**Answer: B**

#### NEW QUESTION 26

A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location. What should a solutions architect do to meet these requirements?

- A. Move the catalog to Amazon ElastiCache for Redis.
- B. Deploy a larger EC2 instance with a larger instance store.
- C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.
- D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

**Answer: A**

#### NEW QUESTION 30

A company needs to store its accounting records in Amazon S3. The records must be immediately accessible for 1 year and then must be archived for an additional 9 years. No one at the company, including administrative users and root users, can be able to delete the records during the entire 10-year period. The

records must be stored with maximum resiliency.  
Which solution will meet these requirements?

- A. Store the records in S3 Glacier for the entire 10-year period
- B. Use an access control policy to deny deletion of the records for a period of 10 years.
- C. Store the records by using S3 Intelligent-Tiering
- D. Use an IAM policy to deny deletion of the records. After 10 years, change the IAM policy to allow deletion.
- E. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 Glacier Deep Archive after 1 year
- F. Use S3 Object Lock in compliance mode for a period of 10 years.
- G. Use an S3 Lifecycle policy to transition the records from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 1 year
- H. Use S3 Object Lock in governance mode for a period of 10 years.

**Answer: C**

#### NEW QUESTION 34

A company has more than 5 TB of file data on Windows file servers that run on premises. Users and applications interact with the data each day. The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency. The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS. What should a solutions architect do to meet these requirements?

- A. Deploy and configure Amazon FSx for Windows File Server on AWS
- B. Move the on-premises file data to FSx for Windows File Server
- C. Reconfigure the workloads to use FSx for Windows File Server on AWS.
- D. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to the S3 File Gateway. Reconfigure the on-premises workloads and the cloud workloads to use the S3 File Gateway.
- E. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to Amazon S3. Reconfigure the workloads to use either Amazon S3 directly or the S3 File Gateway, depending on each workload's location.
- F. Deploy and configure Amazon FSx for Windows File Server on AWS. Deploy and configure an Amazon FSx File Gateway on premises. Move the on-premises file data to the FSx File Gateway. Configure the cloud workloads to use FSx for Windows File Server on AWS. Configure the on-premises workloads to use the FSx File Gateway.

**Answer: D**

#### NEW QUESTION 38

A company is running a high performance computing (HPC) workload on AWS across many Linux based Amazon EC2 instances. The company needs a shared storage system that is capable of sub-millisecond latencies, hundreds of Gbps of throughput and millions of IOPS. Users will store millions of small files. Which solution meets these requirements?

- A. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the file system on each of the EC2 instances.
- B. Create an Amazon S3 bucket. Mount the S3 bucket on each of the EC2 instances.
- C. Ensure that the EC2 instances use Amazon Elastic Block Store (Amazon EBS) optimized. Mount Provisioned IOPS SSD (io2) EBS volumes with Multi-Attach on each instance.
- D. Create an Amazon FSx for Lustre file system.
- E. Mount the file system on each of the EC2 instances.

**Answer: D**

#### NEW QUESTION 40

A company is expecting rapid growth in the near future. A solutions architect needs to configure existing users and grant permissions to new users on AWS. The solutions architect has decided to create IAM groups. The solutions architect will add the new users to IAM groups based on department. Which additional action is the MOST secure way to grant permissions to the new users?

- A. Apply service control policies (SCPs) to manage access permissions.
- B. Create IAM roles that have least privilege permission. Attach the roles to the IAM groups.
- C. Create an IAM policy that grants least privilege permission. Attach the policy to the IAM groups.
- D. Create IAM roles. Associate the roles with a permissions boundary that defines the maximum permissions.

**Answer: C**

#### NEW QUESTION 44

A company is planning to build a high performance computing (HPC) workload as a service solution that is hosted on AWS. A group of 16 Amazon EC2 Linux instances requires the lowest possible latency for node-to-node communication. The instances also need a shared block device volume for high-performing storage. Which solution will meet these requirements?

- A. Use a cluster placement group.
- B. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach.
- C. Use a partition placement group.
- D. Create shared file systems across the instances by using Amazon Elastic File System (Amazon EFS).
- E. Use a spread placement group.
- F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
- G. Use a spread placement group.
- H. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach.

**Answer: A**

#### NEW QUESTION 47

A company is storing sensitive user information in an Amazon S3 bucket. The company wants to provide secure access to this bucket from the application tier running on Amazon EC2 instances inside a VPC.

Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

- A. Configure a VPC gateway endpoint (or Amazon S3 within the VPC)
- B. Create a bucket policy to make the objects in the S3 bucket public
- C. Create a bucket policy that limits access to only the application tier running in the VPC
- D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance
- E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket

**Answer: BD**

#### NEW QUESTION 52

A company is migrating its on-premises PostgreSQL database to Amazon Aurora PostgreSQL. The on-premises database must remain online and accessible during the migration. The Aurora database must remain synchronized with the on-premises database. Which combination of actions must a solutions architect take to meet these requirements? (Select TWO.)

- A. Create an ongoing replication task.
- B. Create a database backup of the on-premises database
- C. Create an AWS Database Migration Service (AWS DMS) replication server
- D. Convert the database schema by using the AWS Schema Conversion Tool (AWS SCT).
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor the database synchronization

**Answer: CD**

#### NEW QUESTION 56

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

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

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

**Answer: CD**

#### NEW QUESTION 61

A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.

Which migration solution will meet these requirements?

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

**Answer: C**

#### NEW QUESTION 62

A solution architect is creating a new Amazon CloudFront distribution for an application. Some of the information submitted by users is sensitive. The application uses HTTPS but needs another layer of security. The sensitive information should be protected throughout the entire application stack. End access to the information should be restricted to certain applications.

Which action should the solutions architect take?

- A. Configure a CloudFront signed URL.
- B. Configure a CloudFront signed cookie.
- C. Configure a CloudFront field-level encryption profile.
- D. Configure CloudFront and set the Origin Protocol Policy setting to HTTPS Only for the Viewer Protocol Policy.

**Answer: C**

#### NEW QUESTION 64

A company wants to build a scalable key management infrastructure to support developers who need to encrypt data in their applications. What should a solutions architect do to reduce the operational burden?

- A. Use multifactor authentication (MFA) to protect the encryption keys.
- B. Use AWS Key Management Service (AWS KMS) to protect the encryption keys.
- C. Use AWS Certificate Manager (ACM) to create, store, and assign the encryption keys.
- D. Use an IAM policy to limit the scope of users who have access permissions to protect the encryption keys.

**Answer: B**

#### NEW QUESTION 69

A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message. The administrator is using an IAM role that has the following IAM policy attached:

What is the cause of the unsuccessful request?

- A. The EC2 Instance has a resource-based policy with a Deny statement.
- B. The principal has not been specified in the policy statement.
- C. The "Action" field does not grant the actions that are required to terminate the EC2 instance.
- D. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0.113.0/24.

**Answer: B**

#### NEW QUESTION 71

A company's website handles millions of requests each day and the number of requests continues to increase. A solutions architect needs to improve the response time of the web application. The solutions architect determines that the application needs to decrease latency when retrieving product details from the Amazon DynamoDB table. Which solution will meet these requirements with the LEAST amount of operational overhead?

- A. Set up a DynamoDB Accelerator (DAX) cluster. Route all read requests through DAX.
- B. Set up Amazon ElastiCache for Redis between the DynamoDB table and the web application. Route all read requests through Redis.
- C. Set up Amazon ElastiCache for Memcached between the DynamoDB table and the web application. Route all read requests through Memcached.
- D. Set up Amazon DynamoDB streams on the table and have AWS Lambda read from the table and populate Amazon ElastiCache. Route all read requests through ElastiCache.

**Answer: A**

#### NEW QUESTION 74

A company's e-commerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections. What should a solutions architect do to meet these requirements?

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

**Answer: B**

#### NEW QUESTION 77

A company is running a critical business application on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances run in an Auto Scaling group and access an Amazon RDS DB instance. The design did not pass an operational review because the EC2 instances and the DB instance are all located in a single Availability Zone. A solutions architect must update the design to use a second Availability Zone. Which solution will make the application highly available?

- A. Provision a subnet in each Availability Zone. Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones. Configure the DB instance with connections to each network.
- B. Provision two subnets that extend across both Availability Zones. Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones.

Configure the DB instance with connections to each network

C. Provision a subnet in each Availability Zone Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones Configure the DB instance for Multi-AZ deployment

D. Provision a subnet that extends across both Availability Zones Configure the Auto Scaling group to distribute the EC2 instances across both Availability Zones Configure the DB instance for Multi-AZ deployment

**Answer: C**

#### NEW QUESTION 82

A gaming company has a web application that displays scores. The application runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS for MySQL database. Users are starting to experience long delays and interruptions that are caused by database read performance. The company wants to improve the user experience while minimizing changes to the application's architecture.

What should a solutions architect do to meet these requirements?

- A. Use Amazon ElastiCache in front of the database.
- B. Use RDS Proxy between the application and the database.
- C. Migrate the application from EC2 instances to AWS Lambda.
- D. Migrate the database from Amazon RDS for MySQL to Amazon DynamoDB.

**Answer: C**

#### NEW QUESTION 83

A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance. A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours. The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue. Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the DB instance to be a Multi-AZ deployment
- B. Create a read replica of the database. Configure the script to query only the read replica.
- C. Instruct the development team to manually export the entries in the database at the end of each day.
- D. Use Amazon ElastiCache to cache the common queries that the script runs against the database.

**Answer: B**

#### NEW QUESTION 86

A company has a web-based map application that provides status information about ongoing repairs. The application sometimes has millions of users. Repair teams have a mobile app that sends current location and status in a JSON message to a REST-based endpoint.

Few repairs occur on most days. The company wants the application to be highly available and to scale when large numbers of repairs occur after nature disasters. Customer use the application most often during these times. The company does not want to pay for idle capacity.

- A. Create a webpage that is based on Amazon S3 to display information.
- B. Use Amazon API Gateway and AWS Lambda to receive the JSON status data. Store the JSON data in Amazon S3.
- C. Use Amazon EC2 instances as web servers across multiple Availability Zones.
- D. Run the EC2 instances in an Auto Scaling group.
- E. Use Amazon API Gateway and AWS Lambda to receive the JSON status data. Store the JSON data in Amazon S3.
- F. Use Amazon EC2 instances as web servers across multiple Availability Zones.
- G. Run the EC2 instances in an Auto Scaling group.
- H. Use a REST endpoint on the EC2 instances to receive the JSON status data.
- I. Store the JSON data in an Amazon RDS Multi-AZ DB instance.
- J. Use Amazon EC2 instances as web servers across multiple Availability Zones. Run the EC2 instances in an Auto Scaling group. Use a REST endpoint on the EC2 instances to receive the JSON status data. Store the JSON data in an Amazon DynamoDB table.

**Answer: D**

#### NEW QUESTION 90

A company has a web application that is based on Java and PHP. The company wants to move the application from on-premises to AWS. The company needs the ability to test new site features frequently. The company also needs a highly available and managed solution that requires minimum operational overhead.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket. Enable static web hosting on the S3 bucket. Upload the static content to the S3 bucket. Use AWS Lambda to process all dynamic content.
- B. Deploy the web application to an AWS Elastic Beanstalk environment. Use URL swapping to switch between multiple Elastic Beanstalk environments for feature testing.
- C. Deploy the web application to Amazon EC2 instances that are configured with Java and PHP. Use Auto Scaling groups and an Application Load Balancer to manage the website's availability.
- D. Containerize the web application. Deploy the web application to Amazon EC2 instances. Use the AWS Load Balancer Controller to dynamically route traffic between containers that contain the new site features for testing.

**Answer: D**

#### NEW QUESTION 91

A company has a stateless asynchronous application that runs in an Apache Hadoop cluster. The application is invoked on demand to run extract, transform, and load (ETL) jobs several times a day.

A solutions architect needs to migrate this application to the AWS Cloud by designing an Amazon EMR cluster for the workload. The cluster must be available immediately to process jobs.

Which implementation meets these requirements MOST cost-effectively?

- A. Use zonal Reserved Instances for the master nodes and the worker nodes. Use a Spot Fleet for the task nodes.

- B. Use zonal Reserved Instances for the master nodes Use Spot instances for the core nodes and the task nodes
- C. Use regional Reserved Instances for the master nodes Use a Spot Fleet for the core nodes and the task nodes
- D. Use regional Reserved Instances for the master node
- E. Use On-Demand Capacity Reservations for the core nodes and the task nodes.

**Answer:** A

#### NEW QUESTION 92

A company is building an ecommerce application and needs to store sensitive customer information. The company needs to give customers the ability to complete purchase transactions on the website. The company also needs to ensure that sensitive customer data is protected, even from database administrators. Which solution meets these requirements?

- A. Store sensitive data in an Amazon Elastic Block Store (Amazon EBS) volume
- B. Use EBS encryption to encrypt the data
- C. Use an IAM instance role to restrict access.
- D. Store sensitive data in Amazon RDS for MySQL
- E. Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.
- F. Store sensitive data in Amazon S3. Use AWS Key Management Service (AWS KMS) service-side encryption the data
- G. Use S3 bucket policies to restrict access.
- H. Store sensitive data in Amazon FSx for Windows Server
- I. Mount the file share on application servers. Use Windows file permissions to restrict access.

**Answer:** C

#### NEW QUESTION 94

A company uses an Amazon Aurora PostgreSQL DB cluster in the us-east-1 Region. The company wants to develop a disaster recovery plan to recover the database in the us-west-1 Region. The company has a recovery time objective (RTO) of 5 minutes and has a recovery point objective (RPO) of 1 minute.

What should a solutions architect do to meet these requirements?

- A. Create a read replica in us-west-1. Set the DB cluster to automatically fail over to the read replica if the primary instance is not responding.
- B. Create an Aurora global database. Set us-west-1 as the secondary Region. Update connections to use the writer and reader endpoints as appropriate.
- C. Set up a second Aurora DB cluster in us-west-1. Use logical replication to keep the databases synchronized. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to change the database endpoint if the primary DB cluster does not respond.
- D. Use Aurora automated snapshots to store data in an Amazon S3 bucket. Enable S3 Versioning.
- E. Configure S3 Cross-Region Replication to us-west-1. Create a second Aurora DB cluster in us-west-1. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to restore the snapshot if the primary DB cluster does not respond.

**Answer:** B

#### NEW QUESTION 95

A company has two AWS accounts in the same AWS Region. One account is a publisher account, and the other account is a subscriber account. Each account has its own Amazon S3 bucket.

An application puts media objects into the publisher account's S3 bucket. The objects are encrypted with server-side encryption with customer-provided encryption keys (SSE-C). The company needs a solution that will automatically copy the objects to the subscriber's account's S3 bucket.

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

- A. Enable S3 Versioning on the publisher account's S3 bucket. Configure S3 Same-Region Replication of the objects to the subscriber account's S3 bucket.
- B. Create an AWS Lambda function that is invoked when objects are published in the publisher account's S3 bucket.
- C. Configure the Lambda function to copy the objects to the subscriber account's S3 bucket.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function when objects are published in the publisher account's S3 bucket. Configure the Lambda function to copy the objects to the subscriber account's S3 bucket.
- E. Configure Amazon EventBridge (Amazon CloudWatch Events) to publish Amazon Simple Notification Service (Amazon SNS) notifications when objects are published in the publisher account's S3 bucket. When notifications are received, use the S3 console to copy the objects to the subscriber account's S3 bucket.

**Answer:** B

#### NEW QUESTION 98

A hospital wants to create digital copies for its large collection of historical written records. The hospital will continue to add hundreds of new documents each day. The hospital's data team will scan the documents and will upload the documents to the AWS Cloud.

A solutions architect must implement a solution to analyze the documents: extract the medical information, and store the documents so that an application can run SQL queries on the data. The solution must maximize scalability and operational efficiency.

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

- A. Write the document information to an Amazon EC2 instance that runs a MySQL database.
- B. Write the document information to an Amazon S3 bucket. Use Amazon Athena to query the data.
- C. Create an Auto Scaling group of Amazon EC2 instances to run a custom application that processes the scanned files and extracts the medical information.
- D. Create an AWS Lambda function that runs when new documents are uploaded. Use Amazon Rekognition to convert the documents to raw text. Use Amazon Transcribe Medical to detect and extract relevant medical information from the text.
- E. Create an AWS Lambda function that runs when new documents are uploaded. Use Amazon Textract to convert the documents to raw text. Use Amazon Comprehend Medical to detect and extract relevant medical information from the text.

**Answer:** AE

#### NEW QUESTION 102

A company has a business system that generates hundreds of reports each day. The business system saves the reports to a network share in CSV format. The company needs to store this data in the AWS Cloud in near-real time for analysis. Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use AWS DataSync to transfer the files to Amazon S3 Create a scheduled task that runs at the end of each day.
- B. Create an Amazon S3 File Gateway Update the business system to use a new network share from the S3 File Gateway.
- C. Use AWS DataSync to transfer the files to Amazon S3 Create an application that uses the DataSync API in the automation workflow.
- D. Deploy an AWS Transfer for SFTP endpoint Create a script that checks for new files on the network share and uploads the new files by using SFTP.

**Answer: B**

#### NEW QUESTION 104

A company needs to ingest and handle large amounts of streaming data that its application generates. The application runs on Amazon EC2 instances and sends data to Amazon Kinesis Data Streams, which is configured with default settings. Every other day the application consumes the data and writes the data to an Amazon S3 bucket for business intelligence (BI) processing. The company observes that Amazon S3 is not receiving all the data that the application sends to Kinesis Data Streams.

What should a solutions architect do to resolve this issue?

- A. Update the Kinesis Data Streams default settings by modifying the data retention period.
- B. Update the application to use the Kinesis Producer Library (KPL) to send the data to Kinesis Data Streams.
- C. Update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams.
- D. Turn on S3 Versioning within the S3 bucket to preserve every version of every object that is ingested in the S3 bucket.

**Answer: A**

#### NEW QUESTION 109

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

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

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

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

**Answer: C**

#### NEW QUESTION 113

A company has enabled AWS CloudTrail logs to deliver log files to an Amazon S3 bucket for each of its developer accounts. The company has created a central AWS account for streamlining management and audit reviews. An internal auditor needs to access the CloudTrail logs, yet access needs to be restricted for all developer account users. The solution must be secure and optimized.

How should a solutions architect meet these requirements?

- A. Configure an AWS Lambda function in each developer account to copy the log files to the central account. Create an IAM role in the central account for the auditor. Attach an IAM policy providing read-only permissions to the bucket.
- B. Configure CloudTrail from each developer account to deliver the log files to an S3 bucket in the central account. Create an IAM user in the central account for the auditor. Attach an IAM policy providing full permissions to the bucket.
- C. Configure CloudTrail from each developer account to deliver the log files to an S3 bucket in the central account. Create an IAM role in the central account for the auditor. Attach an IAM policy providing read-only permissions to the bucket.
- D. Configure an AWS Lambda function in the central account to copy the log files from the S3 bucket in each developer account. Create an IAM user in the central account for the auditor. Attach an IAM policy providing full permissions to the bucket.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-sharing-logs.html>

#### NEW QUESTION 116

A company hosts an application on AWS. The application uses AWS Lambda functions and stores data in Amazon DynamoDB tables. The Lambda functions are connected to a VPC that does not have internet access.

The traffic to access DynamoDB must not travel across the internet. The application must have write access to only specific DynamoDB tables.

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

- A. Attach a VPC endpoint policy for DynamoDB to allow write access to only the specific DynamoDB tables.
- B. Attach a security group to the interface VPC endpoint to allow write access to only the specific DynamoDB tables.
- C. Create a resource-based IAM policy to grant write access to only the specific DynamoDB table.
- D. Attach the policy to the DynamoDB tables.
- E. Create a gateway VPC endpoint for DynamoDB that is associated with the Lambda VPC.
- F. Ensure that the Lambda execution role can access the gateway VPC endpoint.
- G. Create an interface VPC endpoint for DynamoDB that is associated with the Lambda VPC.
- H. Ensure that the Lambda execution role can access the interface VPC endpoint.

**Answer: AD**

#### NEW QUESTION 121

A company has chosen to rehost its application on Amazon EC2 instances. The application occasionally experiences errors that affect parts of its functionality. The company was unaware of this issue until users reported the errors. The company wants to address this problem during the migration and reduce the time it takes to detect issues with the application. Log files for the application are stored on the local disk.

A solutions architect needs to design a solution that will alert staff if there are errors in the application after the application is migrated to AWS. The solution must not require additional changes to the application code.

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

- A. Configure the application to generate custom metrics for the errors. Send these metric data points to Amazon CloudWatch by using the PutMetricData API call. Create a CloudWatch alarm that is based on the custom metrics.
- B. Create an hourly cron job on the instances to copy the application log data to an Amazon S3 bucket. Configure an AWS Lambda function to scan the log file and publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to alert staff if errors are detected.
- C. Install the Amazon CloudWatch agent on the instances. Configure the CloudWatch agent to stream the application log file to Amazon CloudWatch Logs. Run a CloudWatch Logs insights query to search for the relevant pattern in the log file. Create a CloudWatch alarm that is based on the query output.
- D. Install the Amazon CloudWatch agent on the instances. Configure the CloudWatch agent to stream the application log file to Amazon CloudWatch Log.
- E. Create a metric filter for the relevant log group.
- F. Define the filter pattern that is required to determine that there are errors in the application. Create a CloudWatch alarm that is based on the resulting metric.

**Answer: B**

#### NEW QUESTION 125

A company runs an application that receives data from thousands of geographically dispersed remote devices that use UDP. The application processes the data immediately and sends a message back to the device if necessary. No data is stored.

The company needs a solution that minimizes latency for the data transmission from the devices. The solution also must provide rapid failover to another AWS Region.

Which solution will meet these requirements?

- A. Configure an Amazon Route 53 failover routing policy. Create a Network Load Balancer (NLB) in each of the two Regions. Configure the NLB to invoke an AWS Lambda function to process the data.
- B. Use AWS Global Accelerator. Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the NLB. Process the data in Amazon ECS.
- D. Use AWS Global Accelerator. Create an Application Load Balancer (ALB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster.
- E. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.
- F. Configure an Amazon Route 53 failover routing policy. Create an Application Load Balancer (ALB) in each of the two Regions. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.

**Answer: C**

#### NEW QUESTION 126

A company is planning on deploying a newly built application on AWS in a default VPC. The application will consist of a web layer and database layer. The web server was created in public subnets, and the MySQL database was created in private subnet. All subnets are created with the default network ACL settings, and the default security group in the VPC will be replaced with new custom security groups.

- A. Create a database server security group with inbound and outbound rules for MySQL port 3306 traffic to and from anywhere (0.0.0.0/0).
- B. Create a database server security group with an inbound rule for MySQL port 3306 and specify the source as a web server security group.
- C. Create a web server security group within an inbound allow rule for HTTPS port 443 traffic from anywhere (0.0.0.0/0) and an inbound deny rule for IP range 182. 20.0.0/16.
- D. Create a web server security group with an inbound rule for HTTPS port 443 traffic from anywhere (0.0.0.0/0). Create network ACL inbound and outbound deny rules for IP range 182. 20.0.0/16.
- E. Create a web server security group with an inbound and outbound rules for HTTPS port 443 traffic to and from anywhere (0.0.0.0/0). Create a network ACL inbound deny rule for IP range 182. 20.0.0/16.

**Answer: BD**

#### NEW QUESTION 128

A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours. Which solution meets these requirements?

- A. Use Amazon DynamoDB with auto scaling. Use on-demand backups and Amazon DynamoDB Streams.
- B. Use Amazon Redshift.
- C. Configure concurrency scaling.
- D. Activate audit logging.
- E. Perform database snapshots every 4 hours.
- F. Use Amazon RDS with Provisioned IOPS. Activate the database auditing parameter. Perform database snapshots every 5 hours.
- G. Use Amazon Aurora MySQL with auto scaling.
- H. Activate the database auditing parameter.

**Answer: B**

#### NEW QUESTION 129

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

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

- A. Create internal Network Load Balancers in front of the application in each Region.
- B. Create external Application Load Balancers in front of the application in each Region.

- C. Create an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region
- D. Configure Amazon Route 53 to use a geolocation routing policy to distribute the traffic
- E. Configure Amazon CloudFront to handle the traffic and route requests to the application in each Region

**Answer:** AC

#### NEW QUESTION 134

A company wants to use Amazon S3 for the secondary copy of its 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:** C

#### NEW QUESTION 138

A company runs its two-tier e-commerce 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 D8 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? (Select TWO.)

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

**Answer:** AE

#### NEW QUESTION 142

A company is running an application in a private subnet in a VPC with an attached internet gateway. The company needs to provide the application access to the internet while restricting public access to the application. The company does not want to manage additional infrastructure and wants a solution that is highly available and scalable.

Which solution meets these requirements?

- A. Create a NAT gateway in the private subnet.
- B. Create a route table entry from the private subnet to the internet gateway.
- C. Create a NAT gateway in a public subnet. Create a route table entry from the private subnet to the NAT gateway.
- D. Launch a NAT instance in the private subnet. Create a route table entry from the private subnet to the internet gateway.
- E. Launch a NAT instance in a public subnet. Create a route table entry from the private subnet to the NAT instance.

**Answer:** A

#### NEW QUESTION 147

A company wants to run applications in containers in the AWS Cloud. Those applications are stateless and can tolerate disruptions. What should a solutions architect do to meet those requirements?

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

#### NEW QUESTION 150

A company needs to develop a repeatable solution to process time-ordered information from websites around the world. The company collects the data from the websites by using Amazon Kinesis Data Streams and stores the data in Amazon S3. The processing logic needs to collect events and handle data from the last 5 years.

The processing logic also must generate results in an S3 bucket so that a business intelligence application can analyze and compare the results. The processing must be repeated multiple times.

What should a solutions architect do to meet these requirements?

- A. Use Amazon S3 to collect events.
- B. Create an AWS Lambda function to process the event.
- C. Create different Lambda functions to handle repeated processing.
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to collect events. Set AWS Lambda as an event target. Use EventBridge (CloudWatch Events) to create an archive for the events and to replay the events.
- E. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to collect events.
- F. Process the events by using Amazon EC2. Use AWS Step Function to create an archive for the events and to replay the events.
- G. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to collect events.
- H. Process the events by using Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon MSK to create an archive for the events and to replay the events.

**Answer: C**

**NEW QUESTION 153**

A company uses a legacy application to produce data in CSV format. The legacy application stores the output data in Amazon S3. The company is deploying a new commercial off-the-shelf (COTS) application that can perform complex SQL queries to analyze data that is stored in Amazon Redshift and Amazon S3 only. However, the COTS application cannot process the CSV files that the legacy application produces. The company cannot update the legacy application to produce data in another format. The company needs to implement a solution so that the COTS application can use the data that the legacy application produces. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Glue extract, transform, and load (ETL) job that runs on a schedule.
- B. Configure the ETL job to process the .CSV files and store the processed data in Amazon Redshift.
- C. Develop a Python script that runs on Amazon EC2 instances to convert the CSV files to SQL files. Invoke the Python script on a cron schedule to store the output files in Amazon S3.
- D. Create an AWS Lambda function and an Amazon DynamoDB table.
- E. Use an S3 event to invoke the Lambda function.
- F. Configure the Lambda function to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in the DynamoDB table.
- G. Use Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule.
- H. Configure the EMR cluster to perform an extract, transform, and load (ETL) job to process the .CSV files and store the processed data in an Amazon Redshift table.

**Answer: C**

**NEW QUESTION 158**

A company has an application that processes customer orders. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high, the workload does not process orders fast enough. What should a solutions architect do to write the orders reliably to the database as quickly as possible?

- A. Increase the instance size of the EC2 instance when traffic is high.
- B. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
- C. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.
- D. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
- E. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- F. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limit.
- G. Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.

**Answer: B**

**NEW QUESTION 160**

A startup company is hosting a website for its customers on an Amazon EC2 instance. The website consists of a stateless Python application and a MySQL database. The website serves only a small amount of traffic. The company is concerned about the reliability of the instance and needs to migrate to a highly available architecture. The company cannot modify the application code. Which combination of actions should a solution architect take to achieve high availability for the website? (Select TWO.)

- A. Provision an internet gateway in each Availability Zone in use.
- B. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB instance.
- C. Migrate the database to Amazon DynamoDB, and enable DynamoDB auto scaling.
- D. Use AWS DataSync to synchronize the database data across multiple EC2 instances.
- E. Create an Application Load Balancer to distribute traffic to an Auto Scaling group of EC2 instances that are distributed across two Availability Zones.

**Answer: BE**

**NEW QUESTION 163**

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

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

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption>.

**NEW QUESTION 164**

A company has an application with a REST-based interface that allows data to be received in near-real time from a third-party vendor. Once received, the application processes and stores the data for further analysis. The application is running on Amazon EC2 instances. The third-party vendor has received many 503 Service Unavailable errors when sending data to the application. When the data volume spikes, the compute capacity reaches its maximum limit and the application is unable to process all requests. Which design should a solutions architect recommend to provide a more scalable solution?

- A. Use Amazon Kinesis Data Streams to ingest the data. Process the data using an AWS Lambda function.

- B. Use Amazon API Gateway on top of the existing applicatio
- C. Create a usage plan with a quota limit for the third-party vendor
- D. Use Amazon Simple Notification Service (Amazon SNS) to ingest the data Put the EC2 instances in an Auto Scaling group behind an Application Load Balancer
- E. Repackage the application as a container Deploy the application using Amazon Elastic Container Service (Amazon ECS) using the EC2 launch type with an Auto Scaling group

**Answer:** A

#### NEW QUESTION 167

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 polic
- 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 configuratio
- 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 multivalued answer routing polic
- H. Create a health chec
- I. Direct traffic to the website if the health check passe
- J. Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

**Answer:** B

#### NEW QUESTION 170

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

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

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

**Answer:** C

#### Explanation:

[https://docs.aws.amazon.com/vpn/latest/s2svpn/images/Multiple\\_Gateways\\_diagram.png](https://docs.aws.amazon.com/vpn/latest/s2svpn/images/Multiple_Gateways_diagram.png)

"To protect against a loss of connectivity in case your customer gateway device becomes unavailable, you can set up a second Site-to-Site VPN connection to your VPC and virtual private gateway by using a second customer gateway device." <https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html>

#### NEW QUESTION 174

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

What should the solutions architect do to maximize reliability of the application Infrastructure?

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

**Answer:** B

#### Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html>

#### NEW QUESTION 175

A company is developing an Internal application that uses a PostgreSQL database. The company has decided to host the database on Amazon Aurora The application does not need to be highly available but data must be stored in multiple Availability Zones to maximize durability.

Which database configuration meets these requirements MOST cost-effectively?

- A. An Aurora PostgreSQL DB cluster with a single DB Instance
- B. An Aurora PostgreSQL DB cluster with a primary DB instance and a read replica
- C. An Aurora PostgreSQL DB cluster with Multi-AZ deployment enabled
- D. An Aurora PostgreSQL global database cluster

**Answer:** B

#### NEW QUESTION 177

A solutions architect is designing a two-tier web application. The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets. The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet. Security is a high priority for the company. How should security groups be configured in this situation? (Select TWO.)

- A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
- B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
- C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the security group for the web tier.
- D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
- E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

**Answer:** AC

#### Explanation:

"Security groups create an outbound rule for every inbound rule." Not completely right. Stateful does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.

[https://docs.aws.amazon.com/vpc/latest/userguide/VPC\\_SecurityGroups.html#SecurityGroupRules](https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules)

#### NEW QUESTION 182

A company is building a solution that will report Amazon EC2 Auto Scaling events across all the applications in an AWS account. The company needs to use a serverless solution to store the EC2 Auto Scaling status data in Amazon S3. The company then will use the data in Amazon S3 to provide near-real time updates in a dashboard. The solution must not affect the speed of EC2 instance launches.

How should the company move the data to Amazon S3 to meet these requirements?

- A. Use an Amazon CloudWatch metric stream to send the EC2 Auto Scaling status data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.
- B. Launch an Amazon EMR cluster to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a schedule. Configure the Lambda function to send the EC2 Auto Scaling status data directly to Amazon S3.
- D. Use a bootstrap script during the launch of an EC2 instance to install Amazon Kinesis Agent. Configure Kinesis Agent to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.

**Answer:** B

#### NEW QUESTION 184

A solutions architect is using Amazon S3 to design the storage architecture of a new digital media application. The media files must be resilient to the loss of an Availability Zone. Some files are accessed frequently while other files are rarely accessed in an unpredictable pattern. The solutions architect must minimize the costs of storing and retrieving the media files.

Which storage option 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:** B

#### NEW QUESTION 187

A solutions architect is tasked with transferring 750 TB of data from a network-attached file system located at a branch office to Amazon S3 Glacier. The solution must avoid saturating the branch office's low-bandwidth internet connection.

What is the MOST cost-effective solution?

- A. Create a site-to-site VPN tunnel to an Amazon S3 bucket and transfer the files directly.
- B. Create a bucket policy to enforce a VPC endpoint.
- C. Order 10 AWS Snowball appliances and select an S3 Glacier vault as the destination.
- D. Create a bucket policy to enforce a VPC endpoint.
- E. Mount the network-attached file system to Amazon S3 and copy the files directly.
- F. Create a lifecycle policy to transition the S3 objects to Amazon S3 Glacier.
- G. Order 10 AWS Snowball appliances and select an Amazon S3 bucket as the destination.
- H. Create a lifecycle policy to transition the S3 objects to Amazon S3 Glacier.

**Answer:** D

#### NEW QUESTION 192

A company stores data in an Amazon Aurora PostgreSQL DB cluster. The company must store all the data for 5 years and must delete all the data after 5 years. The company also must indefinitely keep audit logs of actions that are performed within the database. Currently, the company has automated backups configured for Aurora.

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

- A. Take a manual snapshot of the DB cluster.
- B. Create a lifecycle policy for the automated backups.
- C. Configure automated backup retention for 5 years.
- D. Configure an Amazon CloudWatch Logs export for the DB cluster.
- E. Use AWS Backup to take the backups and to keep the backups for 5 years.

**Answer:** AD

#### NEW QUESTION 196

A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image. The container needs 50 GB of storage.

available for temporary files. The infrastructure must be serverless.

Which solution meets these requirements with the LEAST operational overhead?

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

**Answer: C**

#### NEW QUESTION 197

A company is creating a new application that will store a large amount of data. The data will be analyzed hourly and will be modified by several Amazon EC2 Linux instances that are deployed across multiple Availability Zones. The needed amount of storage space will continue to grow for the next 6 months.

Which storage solution should a solutions architect recommend to meet these requirements?

- A. Store the data in Amazon S3 Glacier Update the S3 Glacier vault policy to allow access to the application instances.
- B. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume Mount the EBS volume on the application instances.
- C. Store the data in an Amazon Elastic File System (Amazon EFS) file system Mount the file system on the application instances.
- D. Store the data in an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume shared between the application instances.

**Answer: C**

#### NEW QUESTION 198

A company's web application resizes uploaded images for users. The application stores the original images and the resized images in Amazon S3. The company needs to minimize the storage costs for all the images. Original images are viewed frequently, and resized images are viewed infrequently after they are created. Both types of images need to be immediately available.

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

- A. Store the resized images in S3 Standard.
- B. Store the original images in S3 Glacier.
- C. Store the resized images in S3 Glacier.
- D. Store the resized images in S3 One Zone-Infrequent Access (S3 One Zone-IA).

**Answer: AD**

#### NEW QUESTION 201

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

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

**Answer: B**

#### NEW QUESTION 204

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: A**

#### NEW QUESTION 208

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions and Answers in PDF Format

### **SAA-C03 Practice Exam Features:**

- \* SAA-C03 Questions and Answers Updated Frequently
- \* SAA-C03 Practice Questions Verified by Expert Senior Certified Staff
- \* SAA-C03 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* SAA-C03 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The SAA-C03 Practice Test Here](#)**