

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

<https://www.2passeasy.com/dumps/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 company is developing a file-sharing application that will use an Amazon S3 bucket for storage. The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL.

What should a solutions architect do to meet these requirements?

- A. Write individual policies for each S3 bucket to grant read permission for only CloudFront access.
- B. Create an IAM user
- C. Grant the user read permission to objects in the S3 bucket
- D. Assign the user to CloudFront.
- E. Write an S3 bucket policy that assigns the CloudFront distribution ID as the Principal and assigns the target S3 bucket as the Amazon Resource Name (ARN).
- F. Create an origin access identity (OAI). Assign the OAI to the CloudFront distribution
- G. Configure the S3 bucket permissions so that only the OAI has read permission.

Answer: D

Explanation:

Explanation

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-to-amazon-s3/>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3>

NEW QUESTION 3

A company runs an on-premises application that is powered by a MySQL database. The company is migrating the application to AWS to increase the application's elasticity and availability.

The current architecture shows heavy read activity on the database during times of normal operation. Every 4 hours the company's development team pulls a full export of the production database to populate a database in the staging environment. During this period, users experience unacceptable application latency. The development team is unable to use the staging environment until the procedure completes.

A solutions architect must recommend replacement architecture that alleviates the application latency issue. The replacement architecture also must give the development team the ability to continue using the staging environment without delay.

Which solution meets these requirements?

- A. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production.
- B. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.
- C. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production. Use database cloning to create the staging database on-demand.
- D. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production. Use the standby instance for the staging database.
- E. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production.
- F. Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.

Answer: B

NEW QUESTION 4

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: A

NEW QUESTION 5

A development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.

Which solution meets these requirements MOST cost-effectively?

- A. Stop the DB instance when tests are complete
- B. Restart the DB instance when required.
- C. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.
- D. Create a snapshot when tests are complete
- E. Terminate the DB instance and restore the snapshot when required.
- F. Modify the DB instance to a low-capacity instance when tests are complete
- G. Modify the DB instance again when required.

Answer: C

NEW QUESTION 6

A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database table that contains more than 10 million rows. The database has 2 TB of General Purpose SSD storage. There are millions of updates against this data every day through the company's website. The company has noticed that some insert operations are taking 10 seconds or longer. The company has determined that the database storage performance is the problem.

Which solution addresses this performance issue?

- A. Change the storage type to Provisioned IOPS SSD
- B. Change the DB instance to a memory optimized instance class
- C. Change the DB instance to a burstable performance instance class
- D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Answer: A

Explanation:

Explanation

<https://aws.amazon.com/ebs/features/>

"Provisioned IOPS volumes are backed by solid-state drives (SSDs) and are the highest performance EBS volumes designed for your critical, I/O intensive database applications. These volumes are ideal for both IOPS-intensive and throughput-intensive workloads that require extremely low latency."

NEW QUESTION 7

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

Answer: D

NEW QUESTION 8

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 9

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 to 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 10

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 highspeed 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 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 event
- 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

NEW QUESTION 15

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 DynamoDB
- C. Provision a DynamoDB Accelerator (DAX) cluster
- 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 information
- G. Configure the other function to load the information into the database
- 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 information
- K. Configure the other function to load the information into the database
- 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 17

A company needs to review its AWS Cloud deployment to ensure that its Amazon S3 buckets do not have unauthorized configuration changes.

What should a solutions architect do to accomplish this goal?

- A. Turn on AWS Config with the appropriate rules.
- B. Turn on AWS Trusted Advisor with the appropriate checks.
- C. Turn on Amazon Inspector with the appropriate assessment template.
- D. Turn on Amazon S3 server access logging
- E. Configure Amazon EventBridge (Amazon Cloud Watch Events).

Answer: A

NEW QUESTION 18

A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege.

Which solution will meet these requirements?

- A. Share the dashboard from the CloudWatch console
- B. Enter the product manager's email address, and complete the sharing step
- C. Provide a shareable link for the dashboard to the product manager.
- D. Create an IAM user specifically for the product manager
- E. Attach the CloudWatch Read Only Access managed policy to the user
- F. Share the new login credential with the product manager
- G. Share the browser URL of the correct dashboard with the product manager.
- H. Create an IAM user for the company's employees, Attach the View Only Access AWS managed policy to the IAM user
- I. Share the new login credentials with the product manager
- J. Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.
- K. Deploy a bastion server in a public subnet
- L. When the product manager requires access to the dashboard, start the server and share the RDP credential
- M. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Answer: A

NEW QUESTION 22

A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications. Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.

What should a solutions architect recommend to meet these requirements?

- A. Store the transactions data into Amazon DynamoDB. Set up a rule in DynamoDB to remove sensitive data from every transaction upon write. Use DynamoDB Streams to share the transactions data with other applications.
- B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3. Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive data.
- C. Other applications can consume the data stored in Amazon S3.
- D. Stream the transactions data into Amazon Kinesis Data Streams. Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactions data in Amazon DynamoDB. Other applications can consume the transactions data off the Kinesis data stream.
- E. Store the batched transactions data in Amazon S3 as files.
- F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3. The Lambda function then stores the data in Amazon

DynamoDB Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:

Explanation

The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:

- * Amazon OpenSearch Service
- * Amazon S3
- * Datadog
- * Dynatrace
- * Honeycomb
- * HTTP Endpoint
- * Logic Monitor
- * MongoDB Cloud
- * New Relic
- * Splunk
- * Sumo Logic

<https://docs.aws.amazon.com/firehose/latest/dev/create-name.html>

<https://aws.amazon.com/kinesis/data-streams/>

Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 23

A company's application integrates with multiple software-as-a-service (SaaS) sources for data collection. The company runs Amazon EC2 instances to receive the data and to upload the data to an Amazon S3 bucket for analysis. The same EC2 instance that receives and uploads the data also sends a notification to the user when an upload is complete. The company has noticed slow application performance and wants to improve the performance as much as possible. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Auto Scaling group so that EC2 instances can scale out
- B. Configure an S3 event notification to send events to an Amazon Simple Notification Service (Amazon SNS) topic when the upload to the S3 bucket is complete.
- C. Create an Amazon AppFlow flow to transfer data between each SaaS source and the S3 bucket. Configure an S3 event notification to send events to an Amazon Simple Notification Service (Amazon SNS) topic when the upload to the S3 bucket is complete.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for each SaaS source to send output data
- E. Configure the S3 bucket as the rule's target
- F. Create a second EventBridge (CloudWatch Events) rule to send events when the upload to the S3 bucket is complete
- G. Configure an Amazon Simple Notification Service (Amazon SNS) topic as the second rule's target.
- H. Create a Docker container to use instead of an EC2 instance
- I. Host the containerized application on Amazon Elastic Container Service (Amazon ECS). Configure Amazon CloudWatch Container Insights to send events to an Amazon Simple Notification Service (Amazon SNS) topic when the upload to the S3 bucket is complete.

Answer: B

NEW QUESTION 28

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 29

A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability. What should a solutions architect do to meet these requirements?

- A. Create an AWS Lambda function to apply the patch to all EC2 instances.
- B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
- C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
- D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: D

NEW QUESTION 32

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

Answer: DE

NEW QUESTION 37

A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years. What is the MOST operationally efficient solution that meets these requirements?

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Use AWS Backup to create backup schedules and retention policies for the table.
- C. Create an on-demand backup of the table by using the DynamoDB console.
- D. Store the backup in an Amazon S3 bucket.
- E. Set an S3 Lifecycle configuration for the S3 bucket.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function.
- G. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket.
- H. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

NEW QUESTION 40

A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored. Which design should the solutions architect use?

- A. Create an Amazon SNS topic to send the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch configuration that uses the AMI. Create an Auto Scaling group using the launch configuration. Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage.
- B. Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch configuration that uses the AMI. Create an Auto Scaling group using the launch configuration. Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage.
- C. Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch template that uses the AMI. Create an Auto Scaling group using the launch template. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue.
- D. Create an Amazon SNS topic to send the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch template that uses the AMI. Create an Auto Scaling group using the launch template. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic.

Answer: C

Explanation:

"Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue."

In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue. To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the `ApproximateNumberOfMessages` queue attribute to determine the length of the SQS queue.

NEW QUESTION 43

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 partition 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 45

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 46

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 47

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 52

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 57

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 62

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 67

A company's order system sends requests from clients to Amazon EC2 instances. The EC2 instances process the orders and then store the orders in a database on Amazon RDS. Users report that they must reprocess orders when the system fails. The company wants a resilient solution that can process orders

automatically if a system outage occurs.

What should a solutions architect do to meet these requirements?

- A. Move the EC2 instances into an Auto Scaling group
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to target an Amazon Elastic Container Service (Amazon ECS) task
- C. Move the EC2 instances into an Auto Scaling group behind an Application Load Balancer (ALB). Update the order system to send messages to the ALB endpoint
- D. Move the EC2 instances into an Auto Scaling group
- E. Configure the order system to send messages to an Amazon Simple Queue Service (Amazon SQS) queue
- F. Configure the EC2 instances to consume messages from the queue.
- G. Create an Amazon Simple Notification Service (Amazon SNS) topic
- H. Create an AWS Lambda function, and subscribe the function to the SNS topic. Configure the order system to send messages to the SNS topic
- I. Send a command to the EC2 instances to process the messages by using AWS Systems Manager Run Command

Answer: C

NEW QUESTION 68

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 it is deployed. The company is 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) repository
- 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) repository
- 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 instance
- H. Run the containers on EC2 instances that are spread across multiple Availability Zones
- I. Monitor the average CPU utilization in Amazon CloudWatch
- 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 Zones
- L. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Answer: A

NEW QUESTION 70

A company has five organizational units (OUs) as part of its organization in AWS Organization. Each OU correlates to the five businesses that the company owns. The company's research and development (R&D) business is separating from the company and will need its own organization. A solutions architect creates a separate new management account for this purpose.

- A. Have the R&D AWS account be part of both organizations during the transition.
- B. Invite the R&D AWS account to be part of the new organization after the R&D AWS account has left the prior organization.
- C. Create a new R&D AWS account in the new organization
- D. Migrate resources from the prior R&D AWS account to the new R&D AWS account
- E. Have the R&D AWS account join the new organization
- F. Make the new management account a member of the prior organization

Answer: B

NEW QUESTION 74

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 natural disasters. Customers 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 76

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 78

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
- F. Migrate the credential file to the S3 bucket
- G. Point the application to the S3 bucket.
- H. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (or each EC2 instance)
- I. Attach the new EBS volume to each EC2 instance
- J. Migrate the credential file to the new EBS volume
- L. Point the application to the new EBS volume.

Answer: C

NEW QUESTION 82

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 83

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

NEW QUESTION 86

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 87

A company uses Amazon S3 as its data lake. The company has a new partner that must use SFTP to upload data files. A solutions architect needs to implement a highly available SFTP solution that minimizes operational overhead. Which solution will meet these requirements?

- A. Use AWS Transfer Family to configure an SFTP-enabled server with a publicly accessible endpoint. Choose the S3 data lake as the destination.

- B. Use Amazon S3 File Gateway as an SFTP server Expose the S3 File Gateway endpoint URL to the new partner Share the S3 File Gateway endpoint with the new partner
- C. Launch an Amazon EC2 instance in a private subnet in a VPC Instruct the new partner to upload files to the EC2 instance by using a VPN Run a cron job script on the EC2 instance to upload files to the S3 data lake
- D. Launch Amazon EC2 instances in a private subnet in a VPC Place a Network Load Balancer (NLB) in front of the EC2 instances Create an SFTP listener port for the NLB Share the NLB hostname with the new partne
- E. Run a cron job script on the EC2 instances to upload files to the S3 data lake

Answer: A

NEW QUESTION 88

A company is running an ASP.NET MVC application on a single Amazon EC2 instance. A recent increase in application traffic is causing slow response times for users during lunch hours. The company needs to resolve this concern with the least amount of configuration. What should a solutions architect recommend to meet these requirements?

- A. Move the application to AWS Elastic Beanstalk
- B. Configure load-based auto scaling and time-based scaling to handle scaling during lunch hours
- C. Move the application to Amazon Elastic Container Service (Amazon ECS) Create an AWS Lambda function to handle scaling during lunch hours.
- D. Move the application to Amazon Elastic Container Service (Amazon ECS). Configure scheduled scaling for AWS Application Auto Scaling during lunch hours.
- E. Move the application to AWS Elastic Beanstalk
- F. Configure load-based auto scaling, and create an AWS Lambda function to handle scaling during lunch hours.

Answer: A

Explanation:

- Scheduled scaling is the solution here, while "using the least amount of settings possible" - Beanstalk vs moving to ECS - ECS requires MORE CONFIGURATION / SETTINGS (task and service definitions, configuring ECS container agent) than Beanstalk (upload application code)
<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-cfg-autoscaling-scheduledactions.html> Elastic Beanstalk supports time based scaling, since we are aware that the application performance slows down during the lunch hours.
<https://aws.amazon.com/about-aws/whats-new/2015/05/aws-elastic-beanstalk-supports-time-based-scaling/>

NEW QUESTION 89

An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function. The application stores data in an Amazon Aurora PostgreSQL database. During a recent sales event, a sudden surge in customer orders occurred. Some customers experienced timeouts and the application did not process the orders of those customers A solutions architect determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections The solutions architect needs to prevent the timeout errors while making the least possible changes to the application. Which solution will meet these requirements?

- A. Configure provisioned concurrency for the Lambda function Modify the database to be a global database in multiple AWS Regions
- B. Use Amazon RDS Proxy to create a proxy for the database Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint
- C. Create a read replica for the database in a different AWS Region Use query string parameters in API Gateway to route traffic to the read replica
- D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS) Modify the Lambda function to use the OynamoDB table

Answer: C

NEW QUESTION 91

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 connectivit
- 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

NEW QUESTION 95

A company is running a publicly accessible serverless application that uses Amazon API Gateway and AWS Lambda. The application's traffic recently spiked due to fraudulent requests from botnets. Which steps should a solutions architect take to block requests from unauthorized users? (Select TWO.)

- A. Create a usage plan with an API key that it shared with genuine users only.
- B. Integrate logic within the Lambda function to ignore the requests lion- fraudulent IP addresses
- C. Implement an AWS WAF rule to target malicious requests and trigger actions to filler them out
- D. Convert the existing public API to a private API Update the DNS records to redirect users to the new API endpoint
- E. Create an IAM role for each user attempting to access the API A user will assume the role when making the API call

Answer: CD

NEW QUESTION 99

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 100

A company is deploying a web portal. The company wants to ensure that only the web portion of the application is publicly accessible. To accomplish this, the VPC was designed with two public subnets and two private subnets. The application will run on several Amazon EC2 instances in an Auto Scaling group. SSL termination must be offloaded from the EC2 instances.

What should a solutions architect do to ensure these requirements are met? Configure a Network Load Balancer in the public subnets. Configure the Auto Scaling

- A. group in the private subnets and associate it with an Application Load Balancer Configure a Network Load Balancer in the public subnet
- B. Configure the Auto Scaling
- C. group in the public subnets and associate it with an Application Load Balancer.
- D. Configure an Application Load Balancer in the public subnet
- E. Configure the Auto Scaling group in the private subnets and associate it with the Application Load
- F. Balancer, Configure an Application Load Balancer in the private subnet
- G. Configure the Auto Scaling group in the private subnets and associate it with the Application Load Balancer.

Answer: C

NEW QUESTION 104

A company has developed a new content-sharing application that runs on Amazon Elastic Container Service (Amazon ECS). The application runs on Amazon Linux Docker tasks that use the Amazon EC2 launch type. The application requires a storage solution that has the following characteristics:

- Accessibility (or multiple ECS tasks through bind mounts)
- Resiliency across Availability Zones
- Burstable throughput of up to 3 Gbps
- Ability to be scaled up over time

Which storage solution meets these requirements?

- A. Launch an Amazon FSx for Windows File Server Multi-AZ instance
- B. Configure the ECS task definitions to mount the Amazon FSx instance volume at launch.
- C. Launch an Amazon Elastic File System (Amazon EFS) instance
- D. Configure the ECS task definitions to mount the EFS Instance volume at launch.
- E. Create a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach set to enable
- F. Attach the EBS volume to the ECS EC2 instance Configure ECS task definitions to mount the EBS instance volume at launch.
- G. Launch an EC2 instance with several Provisioned IOPS SSD (k>2) Amazon Elastic Block Store (Amazon EBS) volumes attached in a RAID 0 configuration
- H. Configure the EC2 instance as an NFS storage server
- I. Configure ECS task definitions to mount the volumes at launch.

Answer: B

NEW QUESTION 106

An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket. The company has set up S3 event notifications to publish the object creation events to an Amazon SQS queue. The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email.

Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages.

What should the solutions architect do to resolve this issue with the LEAST operational overhead?

- A. Set up long polling in the SQS queue by increasing the ReceiveMessage wait time to 30 seconds.
- B. Change the SQS standard queue to an SQS FIFO queue
- C. Use the message deduplication ID to discard duplicate messages.
- D. Increase the visibility timeout in the SQS queue to a value that is greater than the total of the function timeout and the batch window timeout.
- E. Modify the Lambda function to delete each message from the SQS queue immediately after the message is read before processing.

Answer: B

NEW QUESTION 109

A company wants to establish connectivity between its on-premises data center and AWS (or an existing workload). The workload runs on Amazon EC2 Instances in two VPCs in different AWS Regions. The VPCs need to communicate with each other. The company needs to provide connectivity from its data center to both VPCs. The solution must support a bandwidth of 600 Mbps to the data center.

Which solution will meet these requirements?

- A. Set up an AWS Site-to-Site VPN connection between the data center and one VPC
- B. Create a VPC peering connection between the VPCs.

- C. Set up an AWS Site-to-Site VPN connection between the data center and each VP
- D. Create a VPC peering connection between the VPCs.
- E. Set up an AWS Direct Connect connection between the data center and one VP
- F. Create a VPC peering connection between the VPCs.
- G. Create a transit gatewa
- H. Attach both VPCs to the transit gatewa
- I. Create an AWS Slte-to-Site VPN tunnel to the transit gateway.

Answer: B

NEW QUESTION 114

A company needs to move data from an Amazon EC2 instance to an Amazon S3 bucket. The company mutt ensure that no API calls and no data aim routed through public internet routes Only the EC2 instance can have access to upload data to the S3 bucket. Which solution will meet these requirements?

- A. Create an interlace VPC endpointl for Amazon S3 in the subnet where the EC2 instance is located Attach a resource policy to the S3 bucket to only allow the EC2 instance's 1AM rote for access
- B. Create a gateway VPC endpointl for Amazon S3 in the Availability Zone where the EC2 instance is located Attach appropriate security groups to the endpoint Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM tote for access
- C. Run the nslookup toot from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoint Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket Attach a resource policy to the S3 bucket to only allow the EC2 instance's AM role for access
- D. Use the AWS provided publicly available ip-ranges |son file to obtam the private IP address of the S3 bucket's service API endpoint Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket Attach a resource policy to the S3 bucket to only allow the EC2 instance's 1AM role for access

Answer: B

NEW QUESTION 118

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 Amazo
- B. CloudWatch by using the PutMetricData API call Create a CloudWatch alarm that is based on the custom metrics
- C. 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 rf errors are detected.
- D. 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 lor the relevant pattern in the log file Create a CloudWatch alarm that is based on the query output
- E. Install the Amazon CloudWatch agent on the instances Configure the CloudWatch agent to stream the application log file to Amazon CloudWatch Log
- F. Create a metric fitter for the relevant log grou
- G. 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 120

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 cluste
- 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 122

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 3300 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 125

A company wants to build a data lake on AWS from data that is stored in an on-premises Oracle relational database. The data lake must receive ongoing updates from the on-premises database.

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

- A. Use AWS DataSync to transfer the data to Amazon S3. Use AWS Glue to transform the data and integrate the data into a data lake.
- B. Use AWS Snowball to transfer the data to Amazon S3. Use AWS Batch to transform the data and integrate the data into a data lake.
- C. Use AWS Database Migration Service (AWS DMS) to transfer the data to Amazon S3 Use AWS Glue to transform the data and integrate the data into a data lake.
- D. Use an Amazon EC2 instance to transfer the data to Amazon S3. Configure the EC2 instance to transform the data and integrate the data into a data lake.

Answer: C

NEW QUESTION 130

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: C

NEW QUESTION 133

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

An online photo application lets users upload photos and perform image editing operations The application offers two classes of service free and paid Photos submitted by paid users are processed before those submitted by free users Photos are uploaded to Amazon S3 and the job information is sent to Amazon SQS.

Which configuration should a solutions architect recommend?

- A. Use one SQS FIFO queue Assign a higher priority to the paid photos so they are processed first
- B. Use two SQS FIFO queues: one for paid and one for free Set the free queue to use short polling and the paid queue to use long polling
- C. Use two SQS standard queues one for paid and one for free Configure Amazon EC2 instances to prioritize polling for the paid queue over the free queue.
- D. Use one SQS standard queu
- E. Set the visibility timeout of the paid photos to zero Configure Amazon EC2 instances to prioritize visibility settings so paid photos are processed first

Answer: C

Explanation:

<https://acloud.guru/forums/guru-of-the-week/discussion/-L7Be8rOao3InQxdQcXj/> <https://aws.amazon.com/sqs/features/>

Priority: Use separate queues to provide prioritization of work. <https://aws.amazon.com/sqs/features/>

<https://aws.amazon.com/sqs/features/#:~:text=Priority%3A%20Use%20separate%20queues%20to%20provide%20>

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-polling>.

NEW QUESTION 139

A company runs a 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 D8 instance The EC2 instances and the ROS 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 140

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 143

A business's backup data totals 700 terabytes (TB) and is kept in network-attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.

What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

- A. Order AWS Snowball devices to transfer the data.
- B. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- C. Deploy a VPN connection between the data center and Amazon VPC.
- D. Use the AWS CLI to copy the data from on-premises to Amazon S3 Glacier.
- E. Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3. Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- F. Use AWS DataSync to transfer the data and deploy a DataSync agent on-premise.
- G. Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

Answer: A

NEW QUESTION 145

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 solution 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 146

A company that primarily runs its application servers on-premises has decided to migrate to AWS. The company wants to minimize its need to scale its Internet Small Computer Systems Interface (iSCSI) storage on-premises. The company wants only its recently accessed data to remain stored locally.

Which AWS solution should the company use to meet these requirements?

- A. Amazon S3 File Gateway
- B. AWS Storage Gateway Tape Gateway
- C. AWS Storage Gateway Volume Gateway stored volumes
- D. AWS Storage Gateway Volume Gateway cached volumes

Answer: D

NEW QUESTION 148

A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources. A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege.

Which steps should the solutions architect do in conjunction to reach this goal? (Select two.)

- A. Have the deployment engineer use AWS account root user credentials for performing AWS CloudFormation stack operations.
- B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
- C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the Administrator/Access IAM policy attached.
- D. Create a new IAM user for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
- E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using the IAM role.

Answer: DE

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users.html

NEW QUESTION 152

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for

customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS. The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead. Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RD
- B. Use RDS access controls to limit access.
- C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawle
- D. Use Amazon Athena to query the dat
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formatio
- G. Create an AWS Glue JOBC connection to Amazon RD
- H. Register the S3 bucket in Lake Formatio
- I. Use Lake
- J. Formation access controls to limit acces
- K. Create an Amazon Redshift cluster Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshif
- L. Use Amazon Redshift access controls to limit access.

Answer: C

NEW QUESTION 155

A company has an application that processes customer of tiers. 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 baffle is hig
- 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 158

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 on 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 159

A company's application is running on Amazon EC2 instances within an Auto Scaling group behind an Elastic Load Balancer. Based on the application's history, the company anticipates a spike in traffic during a holiday each year. A solutions architect must design a strategy to ensure that the Auto Scaling group proactively increases capacity to minimize any performance impact on application users. Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm to scale up the EC2 instances when CPU utilization exceeds 90%.
- B. Create a recurring scheduled action to scale up the Auto Scaling group before the expected period of peak demand.
- C. Increase the minimum and maximum number of EC2 instances in the Auto Scaling group during the peak demand period
- D. Configure an Amazon Simple Notification Service (Amazon SNS) notification to send alerts when there are autoscaling EC2_INSTANCE_LAUNCH events

Answer: B

NEW QUESTION 160

The DNS provider that hosts a company's domain name records is experiencing outages that cause service disruption for a website running on AWS. The company needs to migrate to a more resilient managed DNS service and wants the service to run on AWS. What should a solutions architect do to rapidly migrate the DNS hosting service?

- A. Create an Amazon Route 53 public hosted zone for the domain name
- B. Import the zone file containing the domain records hosted by the previous provider.
- C. Create an Amazon Route 53 private hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider
- D. Create a Simple AD directory in AWS
- E. Enable zone transfer between the DNS provider and AWS Directory Service for Microsoft Active Directory for the domain records.
- F. Create an Amazon Route 53 Resolver inbound endpoint in the VPC. Specify the IP addresses that the provider's DNS will forward DNS queries to. Configure the provider's DNS to forward DNS queries for the domain to the IP addresses that are specified in the inbound endpoint.

Answer: B

NEW QUESTION 161

A company that recently started using AWS establishes a Site-to-Site VPN between its on-premises data center and AWS. The company's security mandate states that traffic originating from on premises should stay within the company's private IP space when communicating with an Amazon Elastic Container Service (Amazon ECS) cluster that is hosting a sample web application.

Which solution meets this requirement?

- A. Configure a gateway endpoint for Amazon EC
- B. Modify the route table to include an entry pointing to the ECS cluster.
- C. Create a Network Load Balancer and AWS PrivateLink endpoint for Amazon ECS in the same VPC that is hosting the ECS cluster.
- D. Create a Network Load Balancer in one VPC and an AWS PrivateLink endpoint for Amazon ECS in another VP
- E. Connect the two by using VPC peering.
- F. Configure an Amazon Route record with Amazon ECS as the target
- G. Apply a server certificate to Route 53 from AWS Certificate Manager (ACM) for SSL offloading.

Answer: A

NEW QUESTION 165

A company has deployed a server less application that invokes an AWS Lambda function when new documents are uploaded to an Amazon S3 bucket The application uses the Lambda function to process the documents After a recent marketing campaign the company noticed that the application did not process many of The documents

What should a solutions architect do to improve the architecture of this application?

- A. Set the Lambda function's runtime timeout value to 15 minutes
- B. Configure an S3 bucket replication policy Stage the documents m the S3 bucket for later processing
- C. Deploy an additional Lambda function Load balance the processing of the documents across the two Lambda functions
- D. Create an Amazon Simple Queue Service (Amazon SOS) queue Send the requests to the queue Configure the queue as an event source for Lambda.

Answer: B

NEW QUESTION 169

An online retail company needs to run near-real-time analytics on website traffic to analyze top-selling products across different locations. The product purchase data and the user location details are sent to a third-party application that runs on premises The application processes the data and moves the data into the company's analytics engine

The company needs to implement a cloud-based solution to make the data available for near-real-time analytics.

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

- A. Use Amazon Kinesis Data Streams to ingest the data Use AWS Lambda to transform the data Configure Lambda to write the data to Amazon Amazon OpenSearch Service (Amazon Elasticsearch Service)
- B. Configure Amazon Kinesis Data Streams to write the data to an Amazon S3 bucket Schedule an AWS Glue crawler job to enrich the data and update the AWS Glue Data Catalog Use Amazon Athena for analytics
- C. Configure Amazon Kinesis Data Streams to write the data to an Amazon S3 bucket Add an Apache Spark job on Amazon EMR to enrich the data in the S3 bucket and write the data to Amazon OpenSearch Service (Amazon Elasticsearch Service)
- D. Use Amazon Kinesis Data Firehose to ingest the data Enable Kinesis Data Firehose data transformation with AWS Lambda Configure Kinesis Data Firehose to write the data to Amazon OpenSearch Service (Amazon Elasticsearch Service).

Answer: C

NEW QUESTION 174

A company wants to reduce the cost of its existing three-tier web architect. The web, application, and database servers are running on Amazon EC2 instance EC2 instance for the development, test and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours.

The production EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

- A. Use Spot Instances for the production EC2 instance
- B. Use Reserved Instances for the development and test EC2 instances
- C. Use Reserved Instances for the production EC2 instance
- D. Use On-Demand Instances for the development and test EC2 instances
- E. Use blocks for the production FC2 ins ranges Use Reserved instances for the development and lest EC2 instances
- F. Use On-Demand Instances for the production EC2 instance
- G. Use Spot blocks for the development and test EC2 instances

Answer: B

NEW QUESTION 179

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth.

Which solution will meet these requirements?

- A. Create an S3 bucket Create an IAM role that has permissions to write to the S3 bucke
- B. Use the AWS CLI to copy all files locally to the S3 bucket.
- C. Create an AWS Snowball Edge jo
- D. Receive a Snowball Edge device on premise
- E. Use the Snowball Edge client to transfer data to the devic
- F. Return the device so that AWS can import the data intoAmazon S3.
- G. Deploy an S3 File Gateway on premise
- H. Create a public service endpoint to connect to the S3 File Gateway Create an S3 bucket Create a new NFS file share on the S3 File Gateway Point the new file share to the S3 bucke

- I. Transfer the data from the existing NFS file share to the S3 File Gateway.
- J. Set up an AWS Direct Connect connection between the on-premises network and AW
- K. Deploy an S3 File Gateway on premise
- L. Create a public virtual interlace (VIF) to connect to the S3 File Gatewa
- M. Create an S3 bucke
- N. Create a new NFS file share on the S3 File Gatewa
- O. Point the new file share to the S3 bucke
- P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: C

NEW QUESTION 180

A company runs a photo processing application that needs to frequently upload and download pictures from Amazon S3 buckets that are located in the same AWS Region. A solutions architect has noticed an increased cost in data transfer fees and needs to implement a solution to reduce these costs. How can the solutions architect meet this requirement?

- A. Deploy Amazon API Gateway into a public subnet and adjust the route table to route S3 calls through it
- B. Deploy a NAT gateway into a public subnet and attach an endpoint policy that allows access to the S3 buckets
- C. Deploy the application into a public subnet and allow it to route through an internet gateway to access the S3 buckets
- D. Deploy an S3 VPC gateway endpoint into the VPC and attach an endpoint policy that allows access to the S3 buckets

Answer: D

NEW QUESTION 181

A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer. The application stores data in Amazon Aurora. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy. What should a solutions architect do to meet these requirements?

- A. Deploy the application with the required infrastructure elements in place. Use Amazon Route 53 to configure active-passive failover. Create an Aurora Replica in a second AWS Region.
- B. Host a scaled-down deployment of the application in a second AWS Region. Use Amazon Route 53 to configure active-active failover. Create an Aurora Replica in the second Region.
- C. Replicate the primary infrastructure in a second AWS Region. Use Amazon Route 53 to configure active-active failover. Create an Aurora database that is restored from the latest snapshot.
- D. Back up data with AWS Backup. Use the backup to create the required infrastructure in a second AWS Region. Use Amazon Route 53 to configure active-passive failover. Create an Aurora second primary instance in the second Region.

Answer: C

NEW QUESTION 183

An e-commerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale (or a period of 24 hours). The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon S3 to host the full website in different S3 buckets. Add Amazon CloudFront distributions. Set the S3 buckets as origins for the distributions. Store the order data in Amazon S3.
- B. Deploy the full website on Amazon EC2 instances that run in Auto Scaling groups across multiple Availability Zones. Add an Application Load Balancer (ALB) to distribute the website traffic. Add another ALB for the backend APIs. Store the data in Amazon RDS for MySQL.
- C. Migrate the full application to run in containers. Host the containers on Amazon Elastic Kubernetes Service (Amazon EKS). Use the Kubernetes Cluster Autoscaler to increase and decrease the number of pods to process bursts in traffic. Store the data in Amazon RDS for MySQL.
- D. Use an Amazon S3 bucket to host the website's static content. Deploy an Amazon CloudFront distributio
- E. Set the S3 bucket as the origin. Use Amazon API Gateway and AWS Lambda functions for the backend APIs. Store the data in Amazon DynamoDB.

Answer: D

NEW QUESTION 187

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 191

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAA-C03 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAA-C03 Product From:

<https://www.2passeasy.com/dumps/SAA-C03/>

Money Back Guarantee

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