

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)

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



NEW QUESTION 1

A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants a new solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security.

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

- A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality
- B. Create and deploy an AWS Lambda function to manage and serve the website content
- C. Create the new website and an Amazon S3 bucket. Deploy the website on the S3 bucket with static website hosting enabled
- D. Create the new website
- E. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Answer: AD

NEW QUESTION 2

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 (on each EC2 instance)
- I. Attach the new EBS volume to each EC2 instance
- J. Migrate the credential file to the new EBS volume
- K. Point the application to the new EBS volume.

Answer: B

NEW QUESTION 3

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 4

A company is migrating applications to AWS. The applications are deployed in different accounts. The company manages the accounts centrally by using AWS Organizations. The company's security team needs a single sign-on (SSO) solution across all the company's accounts. The company must continue managing the users and groups in its on-premises self-managed Microsoft Active Directory.

Which solution will meet these requirements?

- A. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- B. Create a one-way forest trust or a one-way domain trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- C. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console
- D. Create a two-way forest trust to connect the company's self-managed Microsoft Active Directory with AWS SSO by using AWS Directory Service for Microsoft Active Directory.
- E. Use AWS Directory Service
- F. Create a two-way trust relationship with the company's self-managed Microsoft Active Directory.
- G. Deploy an identity provider (IdP) on-premise
- H. Enable AWS Single Sign-On (AWS SSO) from the AWS SSO console.

Answer: A

NEW QUESTION 5

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 DB instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available.

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

- A. Use an Auto Scaling group to launch the EC2 instances in private subnet
- B. Deploy an RDS Multi-AZ DB instance in private subnets.

- C. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones. Deploy an Application Load Balancer in the private subnets.
- D. 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.
- E. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zone
- F. Deploy an Application Load Balancer in the public subnet.
- G. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zone
- H. Deploy an Application Load Balancer in the public subnets.

Answer: AE

Explanation:

Explanation

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

NEW QUESTION 6

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 7

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 8

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 9

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 best suite & other Options ruled out [Option A - You can't intruduce 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 10

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 AmazonCloudWatch 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 theSQL queries as needed

Answer: C

Explanation:

Explanation

Amazon Athena can be used to query JSON in S3

NEW QUESTION 10

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores useruploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated 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 13

A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions.

The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) and an associated target group
- B. Associate the target group with the Auto Scaling group
- C. Use the NLB as an AWS Global Accelerator endpoint in each Region.
- D. Deploy an Application Load Balancer (ALB) and an associated target group
- E. Associate the target group with the Auto Scaling group
- F. Use the ALB as an AWS Global Accelerator endpoint in each Region.
- G. Deploy a Network Load Balancer (NLB) and an associated target group
- H. Associate the target group with the Auto Scaling group
- I. Create an Amazon Route 53 latency record that points to aliases for each NL
- J. Create an AmazonCloudFront distribution that uses the latency record as an origin.
- K. Deploy an Application Load Balancer (ALB) and an associated target group
- L. Associate the target group with the Auto Scaling group
- M. Create an Amazon Route 53 weighted record that points to aliases for each AL
- N. Deploy an AmazonCloudFront distribution that uses the weighted record as an origin.

Answer: C

NEW QUESTION 17

A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check.

What should a solutions architect do to accomplish this?

- A. Use AWS Config rules to define and detect resources that are not properly tagged.
- B. Use Cost Explorer to display resources that are not properly tagged
- C. Tag those resources manually.
- D. Write API calls to check all resources for proper tag allocation
- E. Periodically run the code on an EC2 instance.
- F. Write API calls to check all resources for proper tag allocation
- G. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.

Answer: A

NEW QUESTION 20

A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images Which method is the MOST cost-effective for hosting the website?

- A. Containerize the website and host it in AWS Fargate.
- B. Create an Amazon S3 bucket and host the website there
- C. Deploy a web server on an Amazon EC2 instance to host the website.
- D. Configure an Application Load Balancer with an AWS Lambda target that uses the Express.js framework.

Answer: B

Explanation:

Explanation

In Static Websites, Web pages are returned by the server which are prebuilt.

They use simple languages such as HTML, CSS, or JavaScript.

There is no processing of content on the server (according to the user) in Static Websites. Web pages are returned by the server with no change therefore, static Websites are fast.

There is no interaction with databases.

Also, they are less costly as the host does not need to support server-side processing with different languages.

=====

In Dynamic Websites, Web pages are returned by the server which are processed during runtime means they are not prebuilt web pages but they are built during runtime according to the user's demand.

These use server-side scripting languages such as PHP, Node.js, ASP.NET and many more supported by the server.

So, they are slower than static websites but updates and interaction with databases are possible.

NEW QUESTION 21

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 hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources. What should a solutions architect do to meet these requirements?

- A. Use AWS CloudTrail to track configuration changes and AWS Config to record API calls.
- B. Use AWS Config to track configuration changes and AWS CloudTrail to record API calls.
- C. Use AWS Config to track configuration changes and Amazon CloudWatch to record API calls.
- D. Use AWS CloudTrail to track configuration changes and Amazon CloudWatch to record API calls.

Answer: B

NEW QUESTION 27

A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions. The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets. The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions.

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

- A. Create an S3 bucket in each Region. Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Configure replication between the S3 buckets.
- B. Create a customer managed multi-Region KMS key.
- C. Create an S3 bucket in each Region.
- D. Configure replication between the S3 buckets.
- E. Configure the application to use the KMS key with client-side encryption.
- F. Create a customer managed KMS key and an S3 bucket in each Region. Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Configure replication between the S3 buckets.
- G. Create a customer managed KMS key and an S3 bucket in each Region. Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS). Configure replication between the S3 buckets.

Answer: C

Explanation:

Explanation

From <https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html> For most users, the default AWS KMS key store, which is protected by FIPS 140-2 validated cryptographic modules, fulfills their security requirements. There is no need to add an extra layer of maintenance responsibility or a dependency on an additional service. However, you might consider creating a custom key store if your organization has any of the following requirements: Key material cannot be stored in a shared environment. Key material must be subject to a secondary, independent audit path. The HSMs that generate and store key material must be certified at FIPS 140-2 Level 3.

<https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html>

NEW QUESTION 30

A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account. The company needs to create a strategy to access and administer the instances remotely and securely. The company needs to implement a repeatable process that works with native AWS services and follows the AWS Well-Architected Framework.

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

- A. Use the EC2 serial console to directly access the terminal interface of each instance for administration.
- B. Attach the appropriate IAM role to each existing instance and new instance.
- C. Use AWS Systems Manager Session Manager to establish a remote SSH session.
- D. Create an administrative SSH key pair.
- E. Load the public key into each EC2 instance.
- F. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
- G. Establish an AWS Site-to-Site VPN connection.
- H. Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Answer: B

Explanation:

Explanation

<https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-launch-managedinstance.html>

NEW QUESTION 34

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 36

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

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

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

Answer: AB

NEW QUESTION 41

A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size.

Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again.

The company also wants to automate remediation.

What should a solutions architect do to meet these requirements with the LEAST development effort?

- A. Use an Amazon S3 bucket as a secure transfer point.
- B. Use Amazon Inspector to scan the objects in the bucket.
- C. If objects contain PII.
- D. Trigger an S3 Lifecycle policy to remove the objects that contain PII.
- E. Use an Amazon S3 bucket as a secure transfer point.
- F. Use Amazon Macie to scan the objects in the bucket.
- G. If objects contain PII.
- H. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- I. Implement custom scanning algorithms in an AWS Lambda function.
- J. Trigger the function when objects are loaded into the bucket.
- K. If objects contain PII.
- L. Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- M. Implement custom scanning algorithms in an AWS Lambda function.
- N. Trigger the function when objects are loaded into the bucket.
- O. If objects contain PII.
- P. Use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger an S3 Lifecycle policy to remove the objects that contain PII.

Answer: B

NEW QUESTION 46

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 48

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 51

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 53

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 56

A company is preparing to store confidential data in Amazon S3. For compliance reasons the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys must be rotated every year.

Which solution meets these requirements and is the MOST operationally efficient?

- A. Server-side encryption with customer-provided keys (SSE-C)
- B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
- D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automatic rotation

Answer: D

Explanation:

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

When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.

Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use.

AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by

default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 58

A company collects data from thousands of remote devices by using a RESTful web services application that runs on an Amazon EC2 instance. The EC2 instance receives the raw data, transforms the raw data, and stores all the data in an Amazon S3 bucket. The number of remote devices will increase into the millions soon. The company needs a highly scalable solution that minimizes operational overhead.

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

- A. Use AWS Glue to process the raw data in Amazon S3.
- B. Use Amazon Route 53 to route traffic to different EC2 instances.
- C. Add more EC2 instances to accommodate the increasing amount of incoming data.
- D. Send the raw data to Amazon Simple Queue Service (Amazon SQS). Use EC2 instances to process the data.
- E. Use Amazon API Gateway to send the raw data to an Amazon Kinesis data stream.
- F. Configure Amazon Kinesis Data Firehose to use the data stream as a source to deliver the data to Amazon S3.

Answer: BE

NEW QUESTION 62

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 65

A company hosts a serverless application on AWS. The application uses Amazon API Gateway, AWS Lambda, and an Amazon RDS for PostgreSQL database. The company notices an increase in application errors that result from database connection timeouts during times of peak traffic or unpredictable traffic. The company needs a solution that reduces the application failures with the least amount of change to the code.

What should a solutions architect do to meet these requirements?

- A. Reduce the Lambda concurrency rate.
- B. Enable RDS Proxy on the RDS DB instance.
- C. Resize the RDS DB instance class to accept more connections.
- D. Migrate the database to Amazon DynamoDB with on-demand scaling.

Answer: B

NEW QUESTION 66

A company hosts its product information webpages on AWS. The existing solution uses multiple Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group. The website also uses a custom DNS name and communicates with HTTPS only using a dedicated SSL certificate. The company is planning a new product launch and wants to be sure that users from around the world have the best possible experience on the new website.

What should a solutions architect do to meet these requirements?

- A. Redesign the application to use Amazon CloudFront.
- B. Redesign the application to use AWS Elastic Beanstalk.
- C. Redesign the application to use a Network Load Balancer.
- D. Redesign the application to use Amazon S3 static website hosting.

Answer: A

Explanation:

Amazon CloudFront can help provide the best experience for global users. CloudFront integrates seamlessly with ALB and provides an option to use custom DNS and SSL certs.

NEW QUESTION 70

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 72

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. 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 73

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 76

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 80

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 81

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 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 87

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 88

A company runs multiple Windows workloads on AWS. The company's employees use Windows file shares that are hosted on two Amazon EC2 instances. The file shares synchronize data between themselves and maintain duplicate copies. The company wants a highly available and durable storage solution that preserves how users currently access the files.

- A. Migrate all the data to Amazon S3. Set up IAM authentication for users to access files.
- B. Set up an Amazon S3 File Gateway.
- C. Mount the S3 File Gateway on the existing EC2 instances.
- D. Extend the file share environment to Amazon FSx for Windows File Server with a Multi-AZ configuration.
- E. Migrate all the data to FSx for Windows File Server.
- F. Extend the file share environment to Amazon Elastic File System (Amazon EFS) with a Multi-AZ configuration.
- G. Migrate all the data to Amazon EFS.

Answer: C

NEW QUESTION 91

A company is building an e-commerce 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 of 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 runs a high performance computing (HPC) workload on AWS. The workload required low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.

What should a solutions architect propose to improve the performance of the workload?

- A. Choose a cluster placement group while launching Amazon EC2 instances.
- B. Choose dedicated instance tenancy while launching Amazon EC2 instances.
- C. Choose an Elastic Inference accelerator while launching Amazon EC2 instances.
- D. Choose the required capacity reservation while launching Amazon EC2 instances.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-placementgroup.html> "A cluster placement group is a logical grouping of instances within a single Availability Zone that benefit from low network latency, high network throughput"

NEW QUESTION 95

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

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

Which solution meets these requirements MOST cost-effectively?

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

Answer: B

Explanation:

CloudFront uses Edge Locations to cache content while Global Accelerator uses Edge Locations to find an optimal pathway to the nearest regional endpoint.

NEW QUESTION 100

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 103

A company's reporting system delivers hundreds of CSV files to an Amazon S3 bucket each day. The company must convert these files to Apache Parquet format and must store the files in a transformed data bucket.

Which solution will meet these requirements with the LEAST development effort?

- A. Create an Amazon EMR cluster with Apache Spark installed. Write a Spark application to transform the data. Use EMR File System (EMRFS) to write files to the transformed data bucket.
- B. Create an AWS Glue crawler to discover the data. Create an AWS Glue extract, transform, and load (ETL) job to transform the data. Specify the transformed data bucket in the output step.
- C. Use AWS Batch to create a job definition with Bash syntax to transform the data and output the data to the transformed data bucket. Use the job definition to submit a job. Specify an array job as the job type.
- D. Create an AWS Lambda function to transform the data and output the data to the transformed data bucket.
- E. Configure an event notification for the S3 bucket.
- F. Specify the Lambda function as the destination for the event notification.

Answer: D

NEW QUESTION 108

A company is designing a new web application that the company will deploy into a single AWS Region. The application requires a two-tier architecture that will include Amazon EC2 instances and an Amazon RDS DB instance. A solutions architect needs to design the application so that all components are highly available.

- A. Deploy EC2 instances in an additional Region. Create a DB instance with the Multi-AZ option activated.
- B. Deploy all EC2 instances in the same Region and the same Availability Zone.
- C. Create a DB instance with the Multi-AZ option activated.
- D. Deploy the EC2 instances across at least two Availability Zones within the same Region.
- E. Create a DB instance in a single Availability Zone.
- F. Deploy the EC2 instances across at least two Availability Zones within the same Region.

G. Create a DB instance with the Multi-AZ option activated

Answer: D

NEW QUESTION 113

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

NEW QUESTION 114

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 119

A company is running several business applications in three separate VPCs within the us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds of gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.

A solutions architect needs to design a network connectivity solution that maximizes cost-effectiveness. Which solution meets those requirements?

- A. Configure three AWS Site-to-Site VPN connections from the data center to AWS. Establish connectivity by configuring one VPN connection for each VPC.
- B. Launch a third-party virtual network appliance in each VPC. Establish an IPsec VPN tunnel between the Data center and each virtual appliance.
- C. Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway in us-east-1. Establish connectivity by configuring each VPC to use one of the Direct Connect connections.
- D. Set up one AWS Direct Connect connection from the data center to AWS.
- E. Create a transit gateway, and attach each VPC to the transit gateway.
- F. Establish connectivity between the Direct Connect connection and the transit gateway.

Answer: C

NEW QUESTION 121

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 123

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 group in the private subnets and associate it with an Application Load Balancer.
- 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 125

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 127

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
- 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 if 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 for 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 filter for the relevant log group
- 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 131

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 132

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 136

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 137

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 scalin
- D. Activate audit loggin
- 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 scalin
- H. Activate the database auditing parameter

Answer: B

NEW QUESTION 142

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 147

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 150

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

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 153

A company wants to run applications in container 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 154

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 cache volumes

Answer: D

NEW QUESTION 159

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing, and the company is concerned about a potential increase in cost.

What should a solutions architect do to reduce the cost of the website?

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

Answer: A

NEW QUESTION 160

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 RDS
- 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 crawler
- D. Use Amazon Athena to query the data
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formation
- G. Create an AWS Glue JDBC connection to Amazon RDS
- H. Register the S3 bucket in Lake Formation
- I. Use Lake
- J. Formation access controls to limit access
- K. Create an Amazon Redshift cluster. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshift
- L. Use Amazon Redshift access controls to limit access.

Answer: C

NEW QUESTION 162

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 165

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 or EC2 instances that are distributed across two Availability Zones.

Answer: BE

NEW QUESTION 168

A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone. The company needs to redesign its architecture to provide the highest availability with the least operational overhead. What should a solutions architect do to meet these requirements?

- A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
- B. Create a Multi-AZ Auto Scaling group (or EC2 instances that host the applicatio
- C. Create another Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.
- D. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
- E. Create a Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
- F. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- G. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
- H. Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
- I. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- J. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
- K. Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
- L. Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Answer: C

NEW QUESTION 170

A solution architect is using an AWS CloudFormation template to deploy a three-tier web application. The web application consist of a web tier and an application that stores and retrieves user data in Amazon DynamoDB tables. The web and application tiers are hosted on Amazon EC2 instances, and the database tier is not publicly accessible. The application EC2 instances need to access the Dynamo tables Without exposing API credentials in the template. What should the solution architect do to meet the requirements?

- A. Create an IAM role to read the DynamoDB table
- B. Associate the role with the application instances by referencing an instance profile.
- C. Create an IAM role that has the required permissions to read and write from the DynamoDB table
- D. Add the role to the EC2 instance profile, and associate the instances profile with the application instances.
- E. Use the parameter section in the AWS CloudFormation template to have the user input access and secret keys from an already-created IAM user that has the required permissions to read and write from the DynamoDB tables.
- F. Create an IAM user in the AWS CloudFormation template that has the required permissions to read and write from the DynamoDB table
- G. Use the GetAtt function to retrieve the access secret keys, and pass them to the application instances through the user data.

Answer: B

NEW QUESTION 171

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 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 175

A company is migrating a distributed application to AWS The application serves variable workloads The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes The company wants to modernize the application with a solution that maximizes resiliency and scalability How should a solutions architect design the architecture to meet these requirements?

- A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group
- B. Configure EC2 Auto Scaling to use scheduled scaling
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs Implement the compute nodes with Amazon EC2 Instances that are managed in an Auto Scaling group Configure EC2 Auto Scaling based on the size of the queue
- D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group
- E. Configure AWS CloudTrail as a destination for the jobs Configure EC2 Auto Scaling based on the load on the primary server
- F. implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group Configure Amazon EventBridge

(Amazon CloudWatch Events) as a destination for the jobs Configure EC2 Auto Scaling based on the load on the compute nodes

Answer: C

NEW QUESTION 179

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 180

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 EC2 instances Use Reserved instances for the development and test 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 184

A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificate that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.

What should a solutions architect recommend to meet the requirement?

- A. Add a rule in ACM to publish a custom message to an Amazon Simple Notification Service (Amazon SNS) topic every day beginning 30 days before any certificate will expire.
- B. Create an AWS Config rule that checks for certificates that will expire within 30 day
- C. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke a custom alert by way of Amazon Simple Notification Service (Amazon SNS) when AWS Config reports a noncompliant resource
- D. Use AWS trusted Advisor to check for certificates that will expire within to day
- E. Create an Amazon CloudWatch alarm that is based on Trusted Advisor metrics for check status changes Configure the alarm to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS)
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 day
- G. Configure the rule to invoke an AWS Lambda function
- H. Configure the Lambda function to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).

Answer: B

NEW QUESTION 186

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 189

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 bucket
- B. Use the AWS CLI to copy all files locally to the S3 bucket.
- C. Create an AWS Snowball Edge job
- D. Receive a Snowball Edge device on premise
- E. Use the Snowball Edge client to transfer data to the device

- F. Return the device so that AWS can import the data into Amazon 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 bucket
- 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 AWS
- K. Deploy an S3 File Gateway on premise
- L. Create a public virtual interface (VIF) to connect to the S3 File Gateway
- M. Create an S3 bucket
- N. Create a new NFS file share on the S3 File Gateway
- O. Point the new file share to the S3 bucket
- P. Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: C

NEW QUESTION 194

A company wants to create a mobile app that allows users to stream slow-motion video clips on their mobile devices. Currently, the app captures video clips and uploads the video clips in raw format into an Amazon S3 bucket. The app retrieves these video clips directly from the S3 bucket. However, the videos are large in their raw format.

Users are experiencing issues with buffering and playback on mobile devices. The company wants to implement solutions to maximize the performance and scalability of the app while minimizing operational overhead.

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

- A. Deploy Amazon CloudFront for content delivery and caching
- B. Use AWS DataSync to replicate the video files across AWS Regions in other S3 buckets
- C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats
- D. Deploy an Auto Scaling group of Amazon EC2 instances in Local Zones for content delivery and caching
- E. Deploy an Auto Scaling group of Amazon EC2 instances to convert the video files to more appropriate formats

Answer: CD

NEW QUESTION 199

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

NEW QUESTION 201

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 204

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 208

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 210

A company wants to migrate a Windows-based application from on-premises to the AWS Cloud. The application has three tiers, a business tier, and a database tier with Microsoft SQL Server. The company wants to use specific features of SQL Server such as native backups and Data Quality Services. The company also needs to share files for process between the tiers.

How should a solution architect design the architecture to meet these requirements?

- A. Host all three on Amazon instance.
- B. Use Amazon FSx File Gateway for file sharing between tiers.
- C. Host all three on Amazon EC2 instance.
- D. Use Amazon FSx for Windows file sharing between the tiers.
- E. Host the application tier and the business tier on Amazon EC2 instance.
- F. Host the database tier on Amazon RD.
- G. Use Amazon Elastic File System (Amazon EFS) for file sharing between the tiers.
- H. Host the application tier and the business tier on Amazon EC2 instance.
- I. Host the database tier on Amazon RD.
- J. Use a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume for file sharing between the tiers.

Answer: B

NEW QUESTION 215

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 216

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 alias.
- 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 221

.....

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