

Amazon Web Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins Configure Route 53 to route traffic to the CloudFront distribution.
- B. Create an Amazon CloudFront distribution that has the ALB as an origin Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint
- C. Configure Route 53 to route traffic to the CloudFront distribution.
- D. Create an Amazon CloudFront distribution that has the S3 bucket as an origin Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints Create a custom domain name that points to the accelerator DNS name Use the custom domain name as an endpoint for the web application.
- E. Create an Amazon CloudFront distribution that has the ALB as an origin
- F. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint Create two domain name
- G. Point one domain name to the CloudFront DNS name for dynamic content, Point the other domain name to the accelerator DNS name for static content Use the domain names as endpoints for the web application.

Answer: C

Explanation:

Static content can be cached at Cloud front Edge locations from S3 and dynamic content EC2 behind the ALB whose performance can be improved by Global Accelerator whose one endpoint is ALB and other Cloud front. So with regards to custom domain name endpoint is web application is R53 alias records for the custom domain point to web application <https://aws.amazon.com/blogs/networking-and-content-delivery/improving-availability-and-performance-for-application-load-balancers-using-one-click-integration-with-aws-global-accelerator/>

NEW QUESTION 2

- (Topic 1)

A company recently migrated a message processing system to AWS. The system receives messages into an ActiveMQ queue running on an Amazon EC2 instance. Messages are processed by a consumer application running on Amazon EC2. The consumer application processes the messages and writes results to a MySQL database running on Amazon EC2. The company wants this application to be highly available with low operational complexity

Which architecture offers the HIGHEST availability?

- A. Add a second ActiveMQ server to another Availability Zone Add an additional consumer EC2 instance in another Availability Zone
- B. Replicate the MySQL database to another Availability Zone.
- C. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an additional consumer EC2 instance in another Availability Zone
- D. Replicate the MySQL database to another Availability Zone.
- E. Use Amazon MQ with active/standby brokers configured across two Availability Zone
- F. Add an additional consumer EC2 instance in another Availability Zone
- G. Use Amazon RDS for MySQL with Multi-AZ enabled.
- H. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an Auto Scaling group for the consumer EC2 instances across two Availability Zone
- I. Use Amazon RDS for MySQL with Multi-AZ enabled.

Answer: D

Explanation:

Amazon S3 is a highly scalable and durable object storage service that can store and retrieve any amount of data from anywhere on the web¹. Users can configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL. A presigned URL is a URL that gives access to an object in an S3 bucket for a limited time and with a specific action, such as uploading an object². Users can generate a presigned URL programmatically using the AWS SDKs or AWS CLI. By using a presigned URL, users can reduce coupling within the application and improve website performance, as they do not need to send the images to the web server first. AWS Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources³. Users can configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded. S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion. Users can configure S3 Event Notifications to invoke a Lambda function that resizes the image and stores it back in the same or a different S3 bucket. This way, users can offload the image resizing task from the web server to Lambda.

NEW QUESTION 3

- (Topic 1)

A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function.

Which solution meets these requirements?

- A. Add an execution role to the function with lambda: InvokeFunction as the action and * as the principal.
- B. Add an execution role to the function with lambda: InvokeFunction as the action and Service:amazonaws.com as the principal.
- C. Add a resource-based policy to the function with lambda:* as the action and Service:events.amazonaws.com as the principal.
- D. Add a resource-based policy to the function with lambda: InvokeFunction as the action and Service:events.amazonaws.com as the principal.

Answer: D

Explanation:

<https://docs.aws.amazon.com/eventbridge/latest/userguide/resource-based-policies-eventbridge.html#lambda-permissions>

NEW QUESTION 4

- (Topic 1)

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

company wants to have the same functionalities in the AWS Cloud.
Which solution will meet these requirements?

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

Answer: C

Explanation:

AWS Network Firewall supports both inspection and filtering as required

NEW QUESTION 5

- (Topic 2)

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

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

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

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#USER_RestoreFromSnapshot.CON
Under "Encrypt unencrypted resources" - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

NEW QUESTION 6

- (Topic 2)

A company wants to run applications in containers in the AWS Cloud. These applications are stateless and can tolerate disruptions within the underlying infrastructure. The company needs a solution that minimizes cost and operational overhead.

What should a solutions architect do to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 7

- (Topic 3)

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

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

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

Answer: AD

Explanation:

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

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

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

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

NEW QUESTION 8

- (Topic 3)

A solutions architect observes that a nightly batch processing job is automatically scaled up for 1 hour before the desired Amazon EC2 capacity is reached. The peak capacity is the 'same every night and the batch jobs always start at 1 AM. The solutions architect needs to find a cost-effective solution that will allow for the desired EC2 capacity to be reached quickly and allow the Auto Scaling group to scale down after the batch jobs are complete. What should the solutions architect do to meet these requirements?

- A. Increase the minimum capacity for the Auto Scaling group.
- B. Increase the maximum capacity for the Auto Scaling group.
- C. Configure scheduled scaling to scale up to the desired compute level.
- D. Change the scaling policy to add more EC2 instances during each scaling operation.

Answer: C

Explanation:

By configuring scheduled scaling, the solutions architect can set the Auto Scaling group to automatically scale up to the desired compute level at a specific time (IAM) when the batch job starts and then automatically scale down after the job is complete. This will allow the desired EC2 capacity to be reached quickly and also help in reducing the cost.

NEW QUESTION 9

- (Topic 3)

A company has a regional subscription-based streaming service that runs in a single AWS Region. The architecture consists of web servers and application servers on Amazon EC2 instances. The EC2 instances are in Auto Scaling groups behind Elastic Load Balancers. The architecture includes an Amazon Aurora database cluster that extends across multiple Availability Zones. The company wants to expand globally and to ensure that its application has minimal downtime.

- A. Extend the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region
- B. Use an Aurora global database to deploy the database in the primary Region and the second Region
- C. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- D. Deploy the web tier and the application tier to a second Region
- E. Add an Aurora PostgreSQL cross-Region Aurora Replica in the second Region
- F. Use Amazon Route 53 health checks with a failovers routing policy to the second Region, Promote the secondary to primary as needed.
- G. Deploy the web tier and the applicatin tier to a second Region
- H. Create an Aurora PostSQL database in the second Region
- I. Use AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region
- J. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- K. Deploy the web tier and the application tier to a second Region
- L. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Region
- M. Use Amazon Route 53 health checks with a failover routing policy to the second Region
- N. Promote the secondary to primary as needed.

Answer: D

Explanation:

This option is the most efficient because it deploys the web tier and the application tier to a second Region, which provides high availability and redundancy for the application. It also uses an Amazon Aurora global database, which is a feature that allows a single Aurora database to span multiple AWS Regions¹. It also deploys the database in the primary Region and the second Region, which provides low latency global reads and fast recovery from a Regional outage. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which provides data protection by routing traffic to healthy endpoints in different Regions². It also promotes the secondary to primary as needed, which provides data consistency by allowing write operations in one of the Regions at a time³. This solution meets the requirement of expanding globally and ensuring that its application has minimal downtime. Option A is less efficient because it extends the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region, which could incur higher costs and complexity than deploying them separately. It also uses an Aurora global database to deploy the database in the primary Region and the second Region, which is correct. However, it does not use Amazon Route 53 health checks with a failover routing policy to the second Region, which could result in traffic being routed to unhealthy endpoints. Option B is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also adds an Aurora PostgreSQL cross-Region Aurora Replica in the second Region, which provides read scalability across Regions. However, it does not use an Aurora global database, which provides faster replication and recovery than cross-Region replicas. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct. However, it does not promote the secondary to primary as needed, which could result in data inconsistency or loss. Option C is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also creates an Aurora PostgreSQL database in the second Region, which provides data redundancy across Regions. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than creating separate databases. It also uses AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region, which provides data migration between different sources and targets. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than using AWS DMS. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct.

NEW QUESTION 10

- (Topic 3)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 11

- (Topic 3)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 12

- (Topic 3)

A solution architect must create a disaster recovery (DR) plan for a high-volume software as a service (SaaS) platform. All data for the platform is stored in an Amazon Aurora MySQL DB cluster.

The DR plan must replicate data to a secondary AWS Region. Which solution will meet these requirements MOST cost-effectively? Use MySQL binary log replication to an Aurora cluster

- A. Use MySQL binary log replication to an Aurora cluster in the secondary Region Provision one DB instance for the Aurora cluster in the secondary Region.
- B. Set up an Aurora global database for the DB cluster
- C. When setup is complete, remove the DB instance from the secondary Region.
- D. Use AWS Database Migration Service (AWS DMS) to continuously replicate data to an Aurora cluster in the secondary Region Remove the DB instance from the secondary Region.
- E. Set up an Aurora global database for the DB cluster Specify a minimum of one DB instance in the secondary Region

Answer: D

Explanation:

"Replication from the primary DB cluster to all secondaries is handled by the Aurora storage layer rather than by the database engine, so lag time for replicating changes is minimal—typically, less than 1 second. Keeping the database engine out of the replication process means that the database engine is dedicated to processing workloads. It also means that you don't need to configure or manage the Aurora MySQL binlog (binary logging) replication."

NEW QUESTION 13

- (Topic 3)

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

The company wants to ensure that end users retain immediate access to all file types from the on-premises systems without experiencing latency. Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

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

Answer: D

Explanation:

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

NEW QUESTION 14

- (Topic 3)

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

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

customer' accounts.

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

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

Answer: AC

Explanation:

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

NEW QUESTION 15

- (Topic 3)

A solutions architect must secure a VPC network that hosts Amazon EC2 instances. The EC2 instances contain highly sensitive data and run on a private subnet. According to company policy, the EC2 instances that run in the VPC can access only approved third-party software repositories on the internet for software product updates that use the third party's URL. Other internet traffic must be blocked.

Which solution meets these requirements?

- A. Update the route table for the private subnet to route the outbound traffic to an AWS Network Firewall
- B. Configure domain list rule groups
- C. Set up an AWS WAF web ACL
- D. Create a custom set of rules that filter traffic requests based on source and destination IP address range sets.
- E. Implement strict inbound security group rules. Configure an outbound rule that allows traffic only to the authorized software repositories on the internet by specifying the URLs
- F. Configure an Application Load Balancer (ALB) in front of the EC2 instance
- G. Direct an outbound traffic to the ALB. Use a URL-based rule listener in the ALB's target group for outbound access to the internet

Answer: A

Explanation:

Send the outbound connection from EC2 to Network Firewall. In Network Firewall, create stateful outbound rules to allow certain domains for software patch download and deny all other domains. <https://docs.aws.amazon.com/network-firewall/latest/developerguide/suricata-examples.html#suricata-example-domain-filtering>

NEW QUESTION 16

- (Topic 3)

A company has a Microsoft .NET application that runs on an on-premises Windows Server. The application stores data by using an Oracle Database Standard Edition server. The company is planning a migration to AWS and wants to minimize development changes while moving the application. The AWS application environment should be highly available.

Which combination of actions should the company take to meet these requirements? (Select TWO.)

- A. Refactor the application as serverless with AWS Lambda functions running .NET Core
- B. Rehost the application in AWS Elastic Beanstalk with the .NET platform in a Multi-AZ deployment
- C. Replatform the application to run on Amazon EC2 with the Amazon Linux Amazon Machine Image (AMI)
- D. Use AWS Database Migration Service (AWS DMS) to migrate from the Oracle database to Amazon DynamoDB in a Multi-AZ deployment
- E. Use AWS Database Migration Service (AWS DMS) to migrate from the Oracle database to Oracle on Amazon RDS in a Multi-AZ deployment

Answer: BE

Explanation:

To minimize development changes while moving the application to AWS and to ensure a high level of availability, the company can rehost the application in AWS Elastic Beanstalk with the .NET platform in a Multi-AZ deployment. This will allow the application to run in a highly available environment without requiring any changes to the application code.

The company can also use AWS Database Migration Service (AWS DMS) to migrate the Oracle database to Oracle on Amazon RDS in a Multi-AZ deployment. This will allow the company to maintain the existing database platform while still achieving a high level of availability.

NEW QUESTION 17

- (Topic 4)

A company uses AWS Organizations with all features enabled and runs multiple Amazon EC2 workloads in the ap-southeast-2 Region. The company has a service control policy (SCP) that prevents any resources from being created in any other Region. A security policy requires the company to encrypt all data at rest. An audit discovers that employees have created Amazon Elastic Block Store (Amazon EBS) volumes for EC2 instances without encrypting the volumes. The company wants any new EC2 instances that any IAM user or root user launches in ap-southeast-2 to use encrypted EBS volumes. The company wants a solution that will have minimal effect on employees who create EBS volumes.

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

- A. In the Amazon EC2 console, select the EBS encryption account attribute and define a default encryption key.
- B. Create an IAM permission boundary
- C. Attach the permission boundary to the root organizational unit (OU). Define the boundary to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
- D. Create an SCP. Attach the SCP to the root organizational unit (OU). Define the SCP to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
- E. Update the IAM policies for each account to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false.
- F. In the Organizations management account, specify the Default EBS volume encryption setting.

Answer: C

Explanation:

A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. You can use an SCP to deny the ec2:CreateVolume action when the ec2:Encrypted condition equals false, which means that any user or role in the accounts under the root OU will not be able to create unencrypted EBS volumes. This solution will have minimal effect on employees who create EBS volumes, as they can still create encrypted volumes as needed. References: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

NEW QUESTION 18

- (Topic 4)

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
- D. 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.
- E. Write orders to Amazon Simple Notification Service (Amazon SNS) Subscribe the database endpoint to the SNS topic 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

Explanation:

Amazon SQS is a fully managed message queuing service that can decouple and scale microservices, distributed systems, and serverless applications. By writing orders to an SQS queue, the application can handle spikes in traffic without losing any orders. The EC2 instances in an Auto Scaling group can read from the SQS queue and process orders into the database at a steady pace. The Application Load Balancer can distribute the load across the EC2 instances and provide health checks. This solution meets all the requirements of the question, while the other options do not. References:

? <https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/welcome.html>

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

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

NEW QUESTION 19

- (Topic 4)

A company has one million users that use its mobile app. The company must analyze the data usage in near-real time. The company also must encrypt the data in near-real time and must store the data in a centralized location in Apache Parquet format for further processing. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon Kinesis data stream to store the data in Amazon S3. Create an Amazon Kinesis Data Analytics application to analyze the data
- B. Invoke an AWS Lambda function to send the data to the Kinesis Data Analytics application.
- C. Create an Amazon Kinesis data stream to store the data in Amazon S3. Create an Amazon EMR cluster to analyze the data
- D. Invoke an AWS Lambda function to send the data to the EMR cluster.
- E. Create an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Create an Amazon EMR cluster to analyze the data.
- F. Create an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Create an Amazon Kinesis Data Analytics application to analyze the data

Answer: D

Explanation:

This solution will meet the requirements with the least operational overhead as it uses Amazon Kinesis Data Firehose, which is a fully managed service that can automatically handle the data collection, data transformation, encryption, and data storage in near-real time. Kinesis Data Firehose can automatically store the data in Amazon S3 in Apache Parquet format for further processing. Additionally, it allows you to create an Amazon Kinesis Data Analytics application to analyze the data in near real-time, with no need to manage any infrastructure or invoke any Lambda function. This way you can process a large amount of data with the least operational overhead.

NEW QUESTION 20

- (Topic 4)

A research company uses on-premises devices to generate data for analysis. The company wants to use the AWS Cloud to analyze the data. The devices generate .csv files and support writing the data to SMB file share. Company analysts must be able to use SQL commands to query the data. The analysts will run queries periodically throughout the day.

Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

- A. Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode.
- B. Deploy an AWS Storage Gateway on premises in Amazon FSx File Gateway mode.
- C. Set up an AWS Glue crawler to create a table based on the data that is in Amazon S3.
- D. Set up an Amazon EMR cluster with EMR Fife System (EMRFS) to query the data that is in Amazon S3. Provide access to analysts.
- E. Set up an Amazon Redshift cluster to query the data that is in Amazon S3. Provide access to analysts.
- F. Set up Amazon Athena to query the data that is in Amazon S3. Provide access to analysts.

Answer: ACF

Explanation:

To meet the requirements of the use case in a cost-effective way, the following steps are recommended:

? Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode.

This will allow the company to write the .csv files generated by the devices to an SMB file share, which will be stored as objects in Amazon S3 buckets. AWS Storage Gateway is a hybrid cloud storage service that integrates on-premises environments with AWS storage. Amazon S3 File Gateway mode provides a seamless way to connect to Amazon S3 and access a virtually unlimited amount of cloud storage1.

? Set up an AWS Glue crawler to create a table based on the data that is in Amazon

S3. This will enable the company to use standard SQL to query the data stored in Amazon S3 buckets. AWS Glue is a serverless data integration service that simplifies data preparation and analysis. AWS Glue crawlers can automatically discover and classify data from various sources, and create metadata tables in the AWS Glue Data Catalog². The Data Catalog is a central repository that stores information about data sources and how to access them³.

? Set up Amazon Athena to query the data that is in Amazon S3. This will provide

the company analysts with a serverless and interactive query service that can analyze data directly in Amazon S3 using standard SQL. Amazon Athena is integrated with the AWS Glue Data Catalog, so users can easily point Athena at the data source tables defined by the crawlers. Amazon Athena charges only for the queries that are run, and offers a pay-per-query pricing model, which makes it a cost-effective option for periodic queries⁴.

The other options are not correct because they are either not cost-effective or not suitable for the use case. Deploying an AWS Storage Gateway on premises in Amazon FSx File Gateway mode is not correct because this mode provides low-latency access to fully managed Windows file shares in AWS, which is not required for the use case. Setting up an Amazon EMR cluster with EMR File System (EMRFS) to query the data that is in Amazon S3 is not correct because this option involves setting up and managing a cluster of EC2 instances, which adds complexity and cost to the solution. Setting up an Amazon Redshift cluster to query the data that is in Amazon S3 is not correct because this option also involves provisioning and managing a cluster of nodes, which adds overhead and cost to the solution.

References:

? What is AWS Storage Gateway?

? What is AWS Glue?

? AWS Glue Data Catalog

? What is Amazon Athena?

NEW QUESTION 21

- (Topic 4)

A company runs an application using Amazon ECS. The application creates resized versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3.

How can a solutions architect ensure that the application has permission to access Amazon S3?

- A. Update the S3 role in AWS IAM to allow read/write access from Amazon ECS, and then relaunch the container.
- B. Create an IAM role with S3 permissions, and then specify that role as the taskRoleArn in the task definition.
- C. Create a security group that allows access from Amazon ECS to Amazon S3, and update the launch configuration used by the ECS cluster.
- D. Create an IAM user with S3 permissions, and then relaunch the Amazon EC2 instances for the ECS cluster while logged in as this account.

Answer: B

Explanation:

This answer is correct because it allows the application to access Amazon S3 by using an IAM role that is associated with the ECS task. The task role grants permissions to the containers running in the task, and can be used to make AWS API calls from the application code. The taskRoleArn is a parameter in the task definition that specifies the IAM role to use for the task.

References:

? <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-iam-roles.html>

? https://docs.aws.amazon.com/AmazonECS/latest/APIReference/API_TaskDefinition.html

NEW QUESTION 22

- (Topic 4)

A company has a new mobile app. Anywhere in the world, users can see local news on topics they choose. Users also can post photos and videos from inside the app.

Users access content often in the first minutes after the content is posted. New content quickly replaces older content, and then the older content disappears. The local nature of the news means that users consume 90% of the content within the AWS Region where it is uploaded.

Which solution will optimize the user experience by providing the LOWEST latency for content uploads?

- A. Upload and store content in Amazon S3. Use Amazon CloudFront for the uploads.
- B. Upload and store content in Amazon S3. Use S3 Transfer Acceleration for the uploads.
- C. Upload content to Amazon EC2 instances in the Region that is closest to the user.
- D. Copy the data to Amazon S3.
- E. Upload and store content in Amazon S3 in the Region that is closest to the user.
- F. Use multiple distributions of Amazon CloudFront.

Answer: B

Explanation:

The most suitable solution for optimizing the user experience by providing the lowest latency for content uploads is to upload and store content in Amazon S3 and use S3 Transfer Acceleration for the uploads. This solution will enable the company to leverage the AWS global network and edge locations to speed up the data transfer between the users and the S3 buckets.

Amazon S3 is a storage service that provides scalable, durable, and highly available object storage for any type of data. Amazon S3 allows users to store and retrieve data from anywhere on the web, and offers various features such as encryption, versioning, lifecycle management, and replication¹.

S3 Transfer Acceleration is a feature of Amazon S3 that helps users transfer data to and from S3 buckets more quickly. S3 Transfer Acceleration works by using optimized network paths and Amazon's backbone network to accelerate data transfer speeds. Users can enable S3 Transfer Acceleration for their buckets and use a distinct URL to access them, such as <bucket>.s3-accelerate.amazonaws.com².

The other options are not correct because they either do not provide the lowest latency or are not suitable for the use case. Uploading and storing content in Amazon S3 and using

Amazon CloudFront for the uploads is not correct because this solution is not designed for optimizing uploads, but rather for optimizing downloads. Amazon CloudFront is a content delivery network (CDN) that helps users distribute their content globally with low latency and high transfer speeds. CloudFront works by caching the content at edge locations around the world, so that users can access it quickly and easily from anywhere³. Uploading content to Amazon EC2 instances in the Region that is closest to the user and copying the data to Amazon S3 is not correct because this solution adds unnecessary complexity and cost to the process. Amazon EC2 is a computing service that provides scalable and secure virtual servers in the cloud. Users can launch, stop, or terminate EC2 instances as needed, and choose from various instance types, operating systems, and configurations⁴. Uploading and storing content in Amazon S3 in the Region that is closest to the user and using multiple distributions of Amazon CloudFront is not correct because this solution is not cost-effective or efficient for the use case. As mentioned above, Amazon CloudFront is a CDN that helps users distribute their content globally with low latency and high transfer speeds. However, creating multiple CloudFront distributions for each Region would incur additional charges and management overhead, and would not be necessary since 90% of the content is consumed within the same Region where it is uploaded³.

References:

? What Is Amazon Simple Storage Service? - Amazon Simple Storage Service

? Amazon S3 Transfer Acceleration - Amazon Simple Storage Service

- ? What Is Amazon CloudFront? - Amazon CloudFront
- ? What Is Amazon EC2? - Amazon Elastic Compute Cloud

NEW QUESTION 23

- (Topic 4)

A company wants to deploy its containerized application workloads to a VPC across three Availability Zones. The company needs a solution that is highly available across Availability Zones. The solution must require minimal changes to the application.

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

- A. Use Amazon Elastic Container Service (Amazon ECS). Configure Amazon ECS Service Auto Scaling to use target tracking scaling
- B. Set the minimum capacity to 3. Set the task placement strategy type to spread with an Availability Zone attribute.
- C. Use Amazon Elastic Kubernetes Service (Amazon EKS) self-managed node
- D. Configure Application Auto Scaling to use target tracking scaling
- E. Set the minimum capacity to 3.
- F. Use Amazon EC2 Reserved Instance
- G. Launch three EC2 instances in a spread placement group
- H. Configure an Auto Scaling group to use target tracking scaling
- I. Set the minimum capacity to 3.
- J. Use an AWS Lambda function
- K. Configure the Lambda function to connect to a VPC
- L. Configure Application Auto Scaling to use Lambda as a scalable target
- M. Set the minimum capacity to 3.

Answer: A

Explanation:

The company wants to deploy its containerized application workloads to a VPC across three Availability Zones, with high availability and minimal changes to the application. The solution that will meet these requirements with the least operational overhead is:

? Use Amazon Elastic Container Service (Amazon ECS). Amazon ECS is a fully managed container orchestration service that allows you to run and scale containerized applications on AWS. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. Amazon ECS also integrates with other AWS services, such as VPC, ELB, CloudFormation, CloudWatch, IAM, and more.

? Configure Amazon ECS Service Auto Scaling to use target tracking scaling.

Amazon ECS Service Auto Scaling allows you to automatically adjust the number of tasks in your service based on the demand or custom metrics. Target tracking scaling is a policy type that adjusts the number of tasks in your service to keep a specified metric at a target value. For example, you can use target tracking scaling to maintain a target CPU utilization or request count per task for your service.

? Set the minimum capacity to 3. This ensures that your service always has at least

three tasks running across three Availability Zones, providing high availability and fault tolerance for your application.

? Set the task placement strategy type to spread with an Availability Zone attribute.

This ensures that your tasks are evenly distributed across the Availability Zones in your cluster, maximizing the availability of your service.

This solution will provide high availability across Availability Zones, require minimal changes to the application, and reduce the operational overhead of managing your own cluster infrastructure.

References:

? Amazon Elastic Container Service

? Amazon ECS Service Auto Scaling

? Target Tracking Scaling Policies for Amazon ECS Services

? Amazon ECS Task Placement Strategies

NEW QUESTION 24

- (Topic 4)

A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance. The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.

Which solution will allow the node to join the cluster?

- A. Grant the required permission in AWS Identity and Access Management (IAM) to the AmazonEKSNodeRole IAM role.
- B. Create interface VPC endpoints to allow nodes to access the control plane.
- C. Recreate nodes in the public subnet. Restrict security groups for EC2 nodes.
- D. Allow outbound traffic in the security group of the nodes.

Answer: B

Explanation:

Kubernetes API requests within your cluster's VPC (such as node to control plane communication) use the private VPC endpoint.

<https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html>

NEW QUESTION 25

- (Topic 4)

A company has resources across multiple AWS Regions and accounts. A newly hired solutions architect discovers a previous employee did not provide details about the resources inventory. The solutions architect needs to build and map the relationship details of the various workloads across all accounts.

Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: C

Explanation:

Workload Discovery on AWS (formerly called AWS Perspective) is a tool that visualizes AWS Cloud workloads. It maintains an inventory of the AWS resources across your accounts and Regions, mapping relationships between them, and displaying them in a web UI. It also allows you to query AWS Cost and Usage Reports, search for resources, save and export architecture diagrams, and more¹. By using Workload Discovery on AWS, the solution can build and map the relationship details of the various workloads across all accounts with the least operational effort.

- * A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report. This solution will not meet the requirement of building and mapping the relationship details of the various workloads across all accounts, as AWS Systems Manager Inventory is a feature that collects metadata from your managed instances and stores it in a central Amazon S3 bucket. It does not provide a map view or architecture diagrams of the workloads².
- * B. Use AWS Step Functions to collect workload details Build architecture diagrams of the work-loads manually. This solution will not meet the requirement of the least operational effort, as it involves creating and managing state machines to orchestrate the workload details collection, and building architecture diagrams manually.
- * D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships. This solution will not meet the requirement of the least operational effort, as it involves instrumenting your applications with X-Ray SDKs to collect workload details, and building architecture diagrams manually.

Reference URL: <https://aws.amazon.com/solutions/implementations/workload-discovery-on-aws/>

NEW QUESTION 26

- (Topic 4)

A company needs to integrate with a third-party data feed. The data feed sends a webhook to notify an external service when new data is ready for consumption A developer wrote an AWS Lambda function to retrieve data when the company receives a webhook callback The developer must make the Lambda function available for the third party to call.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Create a function URL for the Lambda function
- B. Provide the Lambda function URL to the third party for the webhook.
- C. Deploy an Application Load Balancer (ALB) in front of the Lambda function
- D. Provide the ALB URL to the third party for the webhook
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic
- F. Attach the topic to the Lambda function
- G. Provide the public hostname of the SNS topic to the third party for the webhook.
- H. Create an Amazon Simple Queue Service (Amazon SQS) queue
- I. Attach the queue to the Lambda function
- J. Provide the public hostname of the SQS queue to the third party for the webhook.

Answer: A

Explanation:

A function URL is a unique identifier for a Lambda function that can be used to invoke the function over HTTPS. It is composed of the API endpoint of the AWS Region where the function is deployed, and the name or ARN of the function¹. By creating a function URL for the Lambda function, the solution can make the Lambda function available for the third party to call with the most operational efficiency.

- * B. Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook. This solution will not meet the requirement of the most operational efficiency, as it involves creating and managing an additional resource (ALB) that is not necessary for invoking a Lambda function over HTTPS².
- * C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook. This solution will not work, as Amazon SNS topics do not have public hostnames that can be used as webhooks. SNS topics are used to publish messages to subscribers, not to receive messages from external sources³.
- * D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lambda function. Provide the public hostname of the SQS queue to the third party for the webhook. This solution will not work, as Amazon SQS queues do not have public hostnames that can be used as webhooks. SQS queues are used to send, store, and receive messages between AWS services, not to receive messages from external sources. Reference URL: <https://docs.aws.amazon.com/lambda/latest/dg/lambda-api-permissions-ref.html>

NEW QUESTION 27

- (Topic 4)

A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account.

The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.

Which authentication option will meet these requirements MOST securely?

- A. Integrate DynamoDB with AWS Secrets Manager in the inventory application account
- B. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB table
- C. Schedule secret rotation for every 30 days.
- D. In every business account, create an IAM user that has programmatic access
- E. Configure the application to use the correct IAM user access key ID and secret access key to authenticate and read the DynamoDB table
- F. Manually rotate IAM access keys every 30 days.
- G. In every business account, create an IAM role named BU_ROLE with a policy that gives the role access to the DynamoDB table and a trust policy to trust a specific role in the inventory application account
- H. In the inventory account, create a role named APP_ROLE that allows access to the STS AssumeRole API operation
- I. Configure the application to use APP_ROLE and assume the cross-account role BU_ROLE to read the DynamoDB table.
- J. Integrate DynamoDB with AWS Certificate Manager (ACM). Generate identity certificates to authenticate DynamoDB
- K. Configure the application to use the correct certificate to authenticate and read the DynamoDB table.

Answer: C

Explanation:

This solution meets the requirements most securely because it uses IAM roles and the STS AssumeRole API operation to authenticate and authorize the inventory application to access the DynamoDB tables in different accounts. IAM roles are more secure than IAM users or certificates because they do not require long-term credentials or passwords. Instead, IAM roles provide temporary security credentials that are automatically rotated and can be configured with a limited duration. The STS AssumeRole API operation enables you to request temporary credentials for a role that you are allowed to assume. By using this operation, you can delegate access to resources that are in different AWS accounts that you own or that are owned by third parties. The trust policy of the role defines which entities can assume the role, and the permissions policy of the role defines which actions can be performed on the resources. By using this solution, you can avoid hard-coding credentials or certificates in the inventory application, and you can also avoid storing them in Secrets Manager or ACM. You can also leverage the

built-in security features of IAM and STS, such as MFA, access logging, and policy conditions.

References:

? IAM Roles

? STS AssumeRole

? Tutorial: Delegate Access Across AWS Accounts Using IAM Roles

NEW QUESTION 28

- (Topic 4)

A company has a mobile chat application with a data store based in Amazon DynamoDB. Users would like new messages to be read with as little latency as possible. A solutions architect needs to design an optimal solution that requires minimal application changes. Which method should the solutions architect select?

- A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages table.
- B. Update the code to use the DAX endpoint.
- C. Add DynamoDB read replicas to handle the increased read load.
- D. Update the application to point to the read endpoint for the read replicas.
- E. Double the number of read capacity units for the new messages table in DynamoDB.
- F. Continue to use the existing DynamoDB endpoint.
- G. Add an Amazon ElastiCache for Redis cache to the application stack.
- H. Update the application to point to the Redis cache endpoint instead of DynamoDB.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-high-latency/>

Amazon DynamoDB Accelerator (DAX) is a fully managed in-memory cache for DynamoDB that improves the performance of DynamoDB tables by up to 10 times and

provides a microsecond level of response time at any scale. It is compatible with DynamoDB API operations and requires minimal code changes to use. By configuring DAX for the

new messages table, the solution can reduce the latency for reading new messages with minimal application changes.

* B. Add DynamoDB read replicas to handle the increased read load. Update the application to point to the read endpoint for the read replicas. This solution will not work, as DynamoDB does not support read replicas as a feature. Read replicas are available for Amazon RDS, not for DynamoDB.

* C. Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint. This solution will not meet the requirement of reading new messages with as little latency as possible, as increasing the read capacity units will only increase the throughput of DynamoDB, not the performance or latency.

* D. Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB. This solution will not meet the requirement of minimal application changes, as adding ElastiCache for Redis will require significant code changes to implement caching logic, such as querying cache first, updating cache after writing to DynamoDB, and invalidating cache when needed. Reference URL:

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

NEW QUESTION 29

- (Topic 4)

A company containerized a Windows job that runs on .NET 6 Framework under a Windows container. The company wants to run this job in the AWS Cloud. The job runs every 10 minutes. The job's runtime varies between 1 minute and 3 minutes.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Lambda function based on the container image of the job.
- B. Configure Amazon EventBridge to invoke the function every 10 minutes.
- C. Use AWS Batch to create a job that uses AWS Fargate resources.
- D. Configure the job scheduling to run every 10 minutes.
- E. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job.
- F. Create a scheduled task based on the container image of the job to run every 10 minutes.
- G. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate to run the job.
- H. Create a standalone task based on the container image of the job.
- I. Use Windows task scheduler to run the job every 10 minutes.

Answer: A

Explanation:

AWS Lambda supports container images as a packaging format for functions. You can use existing container development workflows to package and deploy Lambda functions as container images of up to 10 GB in size. You can also use familiar tools such as Docker CLI to build, test, and push your container images to Amazon Elastic Container Registry (Amazon ECR). You can then create an AWS Lambda function based on the container image of your job and configure Amazon EventBridge to invoke the function every 10 minutes using a cron expression. This solution will be cost-effective as you only pay for the compute time you consume when your function runs. References: <https://docs.aws.amazon.com/lambda/latest/dg/images-create.html>

<https://docs.aws.amazon.com/eventbridge/latest/userguide/run-lambda-schedule.html>

NEW QUESTION 30

- (Topic 4)

A company uses an organization in AWS Organizations to manage AWS accounts that contain applications. The company sets up a dedicated monitoring member account in the organization. The company wants to query and visualize observability data across the accounts by using Amazon CloudWatch.

Which solution will meet these requirements?

- A. Enable CloudWatch cross-account observability for the monitoring account.
- B. Deploy an AWS CloudFormation template provided by the monitoring account in each AWS account to share the data with the monitoring account.
- C. Set up service control policies (SCPs) to provide access to CloudWatch in the monitoring account under the Organizations root organizational unit (OU).
- D. Configure a new IAM user in the monitoring account.
- E. In each AWS account, configure an IAM policy to have access to query and visualize the CloudWatch data in the account.
- F. Attach the new IAM policy to the new IAM user.
- G. Create a new IAM user in the monitoring account.
- H. Create cross-account IAM policies in each AWS account.

I. Attach the 1AM policies to the new IAM user.

Answer: A

Explanation:

This solution meets the requirements because it allows the monitoring account to query and visualize observability data across the accounts by using CloudWatch. CloudWatch cross-account observability is a feature that enables a central monitoring account to view and interact with observability data shared by other accounts. To enable cross-account observability, the monitoring account needs to configure the types of data to be shared (metrics, logs, and traces) and the source accounts to be linked. The source accounts can be specified by account IDs, organization IDs, or organization paths. To share the data with the monitoring account, the source accounts need to deploy an AWS CloudFormation template provided by the monitoring account. This template creates an observability link resource that represents the link between the source account and the monitoring account. The template also creates a sink resource that represents an attachment point in the monitoring account. The source accounts can share their observability data with the sink in the monitoring account. The monitoring account can then use the CloudWatch console, API, or CLI to search, analyze, and correlate the observability data across the accounts. References: CloudWatch cross-account observability, Setting up CloudWatch cross-account observability, [Observability Access Manager API Reference]

NEW QUESTION 31

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAA-C03 Practice Exam Features:

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

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