

Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate



NEW QUESTION 1

- (Topic 4)

A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image. The container needs 50 GB of storage available for temporary files. The infrastructure must be serverless.

Which solution meets these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:

The AWS Fargate launch type is a serverless way to run containers on Amazon ECS, without having to manage any underlying infrastructure. You only pay for the resources required to run your containers, and AWS handles the provisioning, scaling, and security of the cluster. Amazon EFS is a fully managed, elastic, and scalable file system that can be mounted to multiple containers, and provides high availability and durability. By using AWS Fargate and Amazon EFS, you can run your Docker container image with 50 GB of storage available for temporary files, with the least operational overhead. This solution meets the requirements of the question.

References:

? [AWS Fargate](#)

? [Amazon Elastic File System](#)

? [Using Amazon EFS file systems with Amazon ECS](#)

NEW QUESTION 2

- (Topic 4)

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

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

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

Answer: C

Explanation:

This answer is correct because it provides redundancy for the VPN connection between the Management VPC and the data center. If one customer gateway device or one VPN tunnel becomes unavailable, the traffic can still flow over the second customer gateway device and the second VPN tunnel. This way, the single point of failure in the VPN connection is mitigated.

References:

? <https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-redundant-connection.html>

? <https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/VPC/vpn-tunnel-redundancy.html>

NEW QUESTION 3

- (Topic 4)

A company wants to use an AWS CloudFormation stack for its application in a test environment. The company stores the CloudFormation template in an Amazon S3 bucket that blocks public access. The company wants to grant CloudFormation access to the template in the S3 bucket based on specific user requests to create the test environment. The solution must follow security best practices.

Which solution will meet these requirements?

- A. Create a gateway VPC endpoint for Amazon S3. Configure the CloudFormation stack to use the S3 object URL.
- B. Create an Amazon API Gateway REST API that has the S3 bucket as the target.
- C. Configure the CloudFormation stack to use the API Gateway URL.
- D. Create a presigned URL for the template object. Configure the CloudFormation stack to use the presigned URL.
- E. Allow public access to the template object in the S3 bucket.
- F. Block the public access after the test environment is created.

Answer: C

Explanation:

It allows CloudFormation to access the template in the S3 bucket without granting public access or creating additional resources. A presigned URL is a URL that is signed with the access key of an IAM user or role that has permission to access the object. The presigned URL can be used by anyone who receives it, but it expires after a specified time. By creating a presigned URL for the template object and configuring the CloudFormation stack to use it, the company can grant CloudFormation access to the template based on specific user requests and follow security best practices. References:

? [Using Amazon S3 Presigned URLs](#)

? [Using Amazon S3 Buckets](#)

NEW QUESTION 4

- (Topic 4)

A company has an application that uses Docker containers in its local data center. The application runs on a container host that stores persistent data in a volume

on the host. The container instances use the stored persistent data.

The company wants to move the application to a fully managed service because the company does not want to manage any servers or storage infrastructure. Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS) with self-managed node
- B. Create an Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instance
- C. Use the EBS volume as a persistent volume mounted in the containers.
- D. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch type
- E. Create an Amazon Elastic File System (Amazon EFS) volume
- F. Add the EFS volume as a persistent storage volume mounted in the containers.
- G. Use Amazon Elastic Container Service (Amazon ECS) with an AWS Fargate launch type
- H. Create an Amazon S3 bucket
- I. Map the S3 bucket as a persistent storage volume mounted in the containers.
- J. Use Amazon Elastic Container Service (Amazon ECS) with an Amazon EC2 launch type
- K. Create an Amazon Elastic File System (Amazon EFS) volume
- L. Add the EFS volume as a persistent storage volume mounted in the containers.

Answer: B

Explanation:

This solution meets the requirements because it allows the company to move the application to a fully managed service without managing any servers or storage infrastructure. AWS Fargate is a serverless compute engine for containers that runs the Amazon ECS tasks. With Fargate, the company does not need to provision, configure, or scale clusters of virtual machines to run containers. Amazon EFS is a fully managed file system that can be accessed by multiple containers concurrently. With EFS, the company does not need to provision and manage storage capacity. EFS provides a simple interface to create and configure file systems quickly and easily. The company can use the EFS volume as a persistent storage volume mounted in the containers to store the persistent data. The company can also use the EFS mount helper to simplify the mounting process. References: Amazon ECS on AWS Fargate, Using Amazon EFS file systems with Amazon ECS, Amazon EFS mount helper.

NEW QUESTION 5

- (Topic 4)

A company wants to analyze and generate reports to track the usage of its mobile app. The app is popular and has a global user base. The company uses a custom report building program to analyze application usage.

The program generates multiple reports during the last week of each month. The program takes less than 10 minutes to produce each report. The company rarely uses the program to generate reports outside of the last week of each month. The company wants to generate reports in the least amount of time when the reports are requested.

Which solution will meet these requirements MOST cost-effectively?

- A. Run the program by using Amazon EC2 On-Demand Instance
- B. Create an Amazon EventBridge rule to start the EC2 instances when reports are requested
- C. Run the EC2 instances continuously during the last week of each month.
- D. Run the program in AWS Lambda
- E. Create an Amazon EventBridge rule to run a Lambda function when reports are requested.
- F. Run the program in Amazon Elastic Container Service (Amazon ECS). Schedule Amazon ECS to run the program when reports are requested.
- G. Run the program by using Amazon EC2 Spot Instance
- H. Create an Amazon EventBridge rule to start the EC2 instances when reports are requested
- I. Run the EC2 instances continuously during the last week of each month.

Answer: B

Explanation:

This solution meets the requirements most cost-effectively because it leverages the serverless and event-driven capabilities of AWS Lambda and Amazon EventBridge. AWS Lambda allows you to run code without provisioning or managing servers, and you pay only for the compute time you consume. Amazon EventBridge is a serverless event bus service that lets you connect your applications with data from various sources and routes that data to targets such as AWS Lambda. By using Amazon EventBridge, you can create a rule that triggers a Lambda function to run the program when reports are requested, and you can also schedule the rule to run during the last week of each month. This way, you can generate reports in the least amount of time and pay only for the resources you use.

References:

? AWS Lambda

? Amazon EventBridge

NEW QUESTION 6

- (Topic 4)

A company stores raw collected data in an Amazon S3 bucket. The data is used for several types of analytics on behalf of the company's customers. The type of analytics requested determines the access pattern on the S3 objects.

The company cannot predict or control the access pattern. The company wants to reduce its S3 costs.

Which solution will meet these requirements?

- A. Use S3 replication to transition infrequently accessed objects to S3 Standard-Infrequent Access (S3 Standard-1A)
- B. Use S3 Lifecycle rules to transition objects from S3 Standard to Standard-Infrequent Access (S3 Standard-1A).
- C. Use S3 Lifecycle rules to transition objects from S3 Standard to S3 Intelligent-Tiering.
- D. Use S3 Inventory to identify and transition objects that have not been accessed from S3 Standard to S3 Intelligent-Tiering.

Answer: C

Explanation:

S3 Intelligent-Tiering is a storage class that automatically reduces storage costs by moving data to the most cost-effective access tier based on access frequency. It has two access tiers: frequent access and infrequent access. Data is stored in the frequent access tier by default, and moved to the infrequent access tier after 30 consecutive days of no access. If the data is accessed again, it is moved back to the frequent access tier. By using S3 Lifecycle rules to transition objects from S3 Standard to S3 Intelligent-Tiering, the solution can reduce S3 costs for data with unknown or changing access patterns.

* A. Use S3 replication to transition infrequently accessed objects to S3 Standard-Infrequent Access (S3 Standard-1A). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 replication is a feature that copies objects across buckets or Regions for

redundancy or compliance purposes. It does not automatically move objects to a different storage class based on access frequency².

* B. Use S3 Lifecycle rules to transition objects from S3 Standard to Standard-Infrequent Access (S3 Standard-IA). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 Standard-IA is a storage class that offers lower storage costs than S3 Standard, but charges a retrieval fee for accessing the data. It is suitable for long-lived and infrequently accessed data, not for data with changing access patterns¹.

* D. Use S3 Inventory to identify and transition objects that have not been accessed from S3 Standard to S3 Intelligent-Tiering. This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 Inventory is a feature that provides a report of the objects in a bucket and their metadata on a daily or weekly basis. It does not automatically move objects to a different storage class based on access frequency³.

Reference URL: <https://aws.amazon.com/s3/storage-classes/intelligent-tiering/>

S3 Intelligent-Tiering is the best solution for reducing S3 costs when the access pattern is unpredictable or changing. S3 Intelligent-Tiering automatically moves objects between two access tiers (frequent and infrequent) based on the access frequency, without any performance impact or retrieval fees. S3 Intelligent-Tiering also has an optional archive tier for objects that are rarely accessed. S3 Lifecycle rules can be used to transition objects from S3 Standard to S3 Intelligent-Tiering.

Reference URLs:

1 <https://aws.amazon.com/s3/storage-classes/intelligent-tiering/>

2 <https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-intelligent-tiering.html>

3 <https://docs.aws.amazon.com/AmazonS3/latest/userguide/intelligent-tiering-overview.html>

NEW QUESTION 7

- (Topic 4)

A company is building a shopping application on AWS. The application offers a catalog that changes once each month and needs to scale with traffic volume. The company wants the lowest possible latency from the application. Data from each user's shopping cart needs to be highly available. User session data must be available even if the user is disconnected and reconnects.

What should a solutions architect do to ensure that the shopping cart data is preserved at all times?

- A. Configure an Application Load Balancer to enable the sticky sessions feature (session affinity) for access to the catalog in Amazon Aurora.
- B. Configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- C. Configure Amazon OpenSearch Service to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- D. Configure an Amazon EC2 instance with Amazon Elastic Block Store (Amazon EBS) storage for the catalog and shopping cart.
- E. Configure automated snapshots.

Answer: B

Explanation:

To ensure that the shopping cart data is preserved at all times, a solutions architect should configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session. This solution has the following benefits:

? It offers the lowest possible latency from the application, as ElastiCache for Redis

is a blazing fast in-memory data store that provides sub-millisecond latency to power internet-scale real-time applications¹.

? It scales with traffic volume, as ElastiCache for Redis supports horizontal scaling

by adding more nodes or shards to the cluster, and vertical scaling by changing the node type².

? It is highly available, as ElastiCache for Redis supports replication across multiple

Availability Zones and automatic failover in case of a primary node failure³.

? It preserves user session data even if the user is disconnected and reconnects, as ElastiCache for Redis can store session data, such as user login information and shopping cart contents, in a persistent and durable manner using snapshots or AOF (append-only file) persistence⁴.

References:

? 1: <https://aws.amazon.com/elasticache/redis/>

? 2: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Scaling.html>

? 3: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.html>

? 4: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/backups.html>

NEW QUESTION 8

- (Topic 4)

A company maintains an Amazon RDS database that maps users to cost centers. The company has accounts in an organization in AWS Organizations. The company needs a solution that will tag all resources that are created in a specific AWS account in the organization. The solution must tag each resource with the cost center ID of the user who created the resource.

Which solution will meet these requirements?

- A. Move the specific AWS account to a new organizational unit (OU) in Organizations from the management account.
- B. Create a service control policy (SCP) that requires all existing resources to have the correct cost center tag before the resources are created.
- C. Apply the SCP to the new OU.
- D. Create an AWS Lambda function to tag the resources after the Lambda function looks up the appropriate cost center from the RDS database.
- E. Configure an Amazon EventBridge rule that reacts to AWS CloudTrail events to invoke the Lambda function.
- F. Create an AWS CloudFormation stack to deploy an AWS Lambda function.
- G. Configure the Lambda function to look up the appropriate cost center from the RDS database and to tag resource.
- H. Create an Amazon EventBridge scheduled rule to invoke the CloudFormation stack.
- I. Create an AWS Lambda function to tag the resources with a default value.
- J. Configure an Amazon EventBridge rule that reacts to AWS CloudTrail events to invoke the Lambda function when a resource is missing the cost center tag.

Answer: B

Explanation:

AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be used to tag resources with the cost center ID of the user who created the resource, by querying the RDS database that maps users to cost centers. Amazon EventBridge is a serverless event bus service that enables event-driven architectures. EventBridge can be configured to react to AWS CloudTrail events, which are recorded API calls made by or on behalf of the AWS account. EventBridge can invoke the Lambda function when a resource is created in the specific AWS account, passing the user identity and resource information as parameters. This solution will meet the requirements, as it enables automatic tagging of resources based on the user and cost center mapping.

References:

? 1 provides an overview of AWS Lambda and its benefits.

? 2 provides an overview of Amazon EventBridge and its benefits.

? 3 explains the concept and benefits of AWS CloudTrail events.

NEW QUESTION 9

- (Topic 4)

A company has an organization in AWS Organizations that has all features enabled. The company requires that all API calls and logins in any existing or new AWS account must be audited. The company needs a managed solution to prevent additional work and to minimize costs. The company also needs to know when any AWS account is not compliant with the AWS Foundational Security Best Practices (FSBP) standard.

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

- A. Deploy an AWS Control Tower environment in the Organizations management account. Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.
- B. Deploy an AWS Control Tower environment in a dedicated Organizations member account. Enable AWS Security Hub and AWS Control Tower Account Factory in the environment.
- C. Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ). Submit an RFC to self-service provision Amazon GuardDuty in the MALZ.
- D. Use AWS Managed Services (AMS) Accelerate to build a multi-account landing zone (MALZ). Submit an RFC to self-service provision AWS Security Hub in the MALZ.

Answer: A

Explanation:

AWS Control Tower is a fully managed service that simplifies the setup and governance of a secure, compliant, multi-account AWS environment. It establishes a landing zone that is based on best-practices blueprints, and it enables governance using controls you can choose from a pre-packaged list. The landing zone is a well-architected, multi-account baseline that follows AWS best practices. Controls implement governance rules for security, compliance, and operations. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts. It aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer, as well as from AWS Partner solutions. AWS Security Hub continuously monitors your environment using automated compliance checks based on the AWS best practices and industry standards, such as the AWS Foundational Security Best Practices (FSBP) standard. AWS Control Tower Account Factory is a feature that automates the provisioning of new AWS accounts that are preconfigured to meet your business, security, and compliance requirements. By deploying an AWS Control Tower environment in the Organizations management account, you can leverage the existing organization structure and policies, and enable AWS Security Hub and AWS Control Tower Account Factory in the environment. This way, you can audit all API calls and logins in any existing or new AWS account, monitor the compliance status of each account with the FSBP standard, and provision new accounts with ease and consistency. This solution meets the requirements with the least operational overhead, as you do not need to manage any infrastructure, perform any data migration, or submit any requests for changes. References:

? AWS Control Tower

? [AWS Security Hub]

? [AWS Control Tower Account Factory]

NEW QUESTION 10

- (Topic 4)

A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL-compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory.

The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.

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

- A. Provision an Amazon DynamoDB database with default read and write capacity settings.
- B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacity unit (ACU).
- C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
- D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

Answer: C

Explanation:

It allows the company to migrate the on-premises database to a managed AWS service that supports auto scaling capabilities and has the least administrative overhead. Amazon Aurora Serverless v2 is a configuration of Amazon Aurora that automatically scales compute capacity based on workload demand. It can scale from hundreds to hundreds of thousands of transactions in a fraction of a second. Amazon Aurora Serverless v2 also supports MySQL-compatible databases and AWS Direct Connect connectivity. References:

? Amazon Aurora Serverless v2

? Connecting to an Amazon Aurora DB Cluster

NEW QUESTION 10

- (Topic 4)

A company offers a food delivery service that is growing rapidly. Because of the growth, the company's order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:

- A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
- Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders

The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.

A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize

utilization of the company's AWS resources. Which solution meets these requirements?

- A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
- B. Configure each Auto Scaling group's minimum capacity according to peak workload values.
- C. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
- D. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
- E. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment
- F. Configure the EC2 instances to poll their respective queue
- G. Scale the Auto Scaling groups based on notifications that the queues send.
- H. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillment
- I. Configure the EC2 instances to poll their respective queue
- J. Create a metric based on a backlog per instance calculation
- K. Scale the Auto Scaling groups based on this metric.

Answer: D

Explanation:

The number of instances in your Auto Scaling group can be driven by how long it takes to process a message and the acceptable amount of latency (queue delay). The solution is to use a backlog per instance metric with the target value being the acceptable backlog per instance to maintain.

NEW QUESTION 14

- (Topic 4)

A company has an on-premises server that uses an Oracle database to process and store customer information. The company wants to use an AWS database service to achieve higher availability and to improve application performance. The company also wants to offload reporting from its primary database system. Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions. Point the reporting functions toward a separate DB instance from the primary DB instance.
- B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database. Create a read replica in the same zone as the primary DB instance.
- C. Direct the reporting functions to the read replica.
- D. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database. Direct the reporting functions to use the reader instance in the cluster deployment.
- E. Use Amazon RDS deployed in a Multi-AZ instance deployment to create an Amazon Aurora database.
- F. Direct the reporting functions to the reader instances.

Answer: D

Explanation:

Amazon Aurora is a fully managed relational database that is compatible with MySQL and PostgreSQL. It provides up to five times better performance than MySQL and

up to three times better performance than PostgreSQL. It also provides high availability and durability by replicating data across multiple Availability Zones and continuously backing up data to Amazon S3. By using Amazon RDS deployed in a Multi-AZ instance deployment

to create an Amazon Aurora database, the solution can achieve higher availability and improve application performance. Amazon Aurora supports read replicas, which are separate instances that share the same underlying storage as the primary instance. Read replicas can be used to offload read-only queries from the primary instance and improve performance. Read replicas can also be used for reporting functions. By directing the reporting functions to the reader instances, the solution can offload reporting from its primary database system.

* A. Use AWS Database Migration Service (AWS DMS) to create an Amazon RDS DB instance in multiple AWS Regions. Point the reporting functions toward a separate DB instance from the primary DB instance. This solution will not meet the requirement of using an AWS database service, as AWS DMS is a service that helps users migrate databases to AWS, not a database service itself. It also involves creating multiple DB instances in different Regions, which may increase complexity and cost.

* B. Use Amazon RDS in a Single-AZ deployment to create an Oracle database. Create a read replica in the same zone as the primary DB instance. Direct the reporting functions to the read replica. This solution will not meet the requirement of achieving higher availability, as a Single-AZ deployment does not provide failover protection in case of an Availability Zone outage. It also involves using Oracle as the database engine, which may not provide better performance than Aurora.

* C. Use Amazon RDS deployed in a Multi-AZ cluster deployment to create an Oracle database. Direct the reporting functions to use the reader instance in the cluster deployment. This solution will not meet the requirement of improving application performance, as Oracle may not provide better performance than Aurora. It also involves using a cluster deployment, which is only supported for Aurora, not for Oracle. Reference URL: <https://aws.amazon.com/rds/aurora/>

NEW QUESTION 17

- (Topic 4)

A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture. The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service. Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO.)

- A. Create a new organization in AWS Organizations with all features turned on.
- B. Create the new AWS accounts in the organization.
- C. Set up an Amazon Cognito identity pool.
- D. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication.
- E. Configure a service control policy (SCP) to manage the AWS account.
- F. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service.
- G. Create a new organization in AWS Organization.
- H. Configure the organization's authentication mechanism to use AWS Directory Service directly.
- I. Set up AWS IAM Identity Center (AWS Single Sign-On) in the organization.
- J. Configure IAM Identity Center, and integrate it with the company's corporate directory service.

Answer: AE

Explanation:

AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts. By creating a new organization in AWS Organizations with all features turned on, the solution can consolidate and manage the new AWS accounts for different business units.

AWS IAM Identity Center (formerly known as AWS Single Sign-On) is a service that provides single sign-on access for all of your AWS accounts and cloud applications. It connects with Microsoft Active Directory through AWS Directory Service to allow users in that directory to sign in to a personalized AWS access portal using their existing Active Directory user names and passwords. From the AWS access portal, users have access to all the AWS accounts and cloud applications that they have permissions for. By setting up IAM Identity Center in the organization and integrating it with the company's corporate directory service, the solution can authenticate access to these AWS accounts using a centralized corporate directory service.

* B. Set up an Amazon Cognito identity pool. Configure AWS IAM Identity Center (AWS Single Sign-On) to accept Amazon Cognito authentication. This solution will not meet the requirement of authenticating access to these AWS accounts by using a centralized corporate directory service, as Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications, not for corporate directory services.

* C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS IAM Identity Center (AWS Single Sign-On) to AWS Directory Service. This solution will not work, as SCPs are used to restrict the actions that can be performed by the accounts in an organization, not to manage the accounts themselves. Also, IAM Identity Center cannot be added to AWS Directory Service, as it is a separate service that connects with Microsoft Active Directory through AWS Directory Service.

* D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly. This solution

will not work, as AWS Organizations does not have an authentication mechanism that can use AWS Directory Service directly. AWS Organizations relies on IAM Identity Center to provide single sign-on access for the accounts in an organization.

Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_integrate_services.html

NEW QUESTION 20

- (Topic 4)

An online video game company must maintain ultra-low latency for its game servers. The game servers run on Amazon EC2 instances. The company needs a solution that can

handle millions of UDP internet traffic requests each second.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an Application Load Balancer with the required protocol and ports for the internet traffic
- B. Specify the EC2 instances as the targets.
- C. Configure a Gateway Load Balancer for the internet traffic
- D. Specify the EC2 instances as the targets.
- E. Configure a Network Load Balancer with the required protocol and ports for the internet traffic
- F. Specify the EC2 instances as the targets.
- G. Launch an identical set of game servers on EC2 instances in separate AWS Region
- H. Route internet traffic to both sets of EC2 instances.

Answer: C

Explanation:

The most cost-effective solution for the online video game company is to configure a Network Load Balancer with the required protocol and ports for the internet traffic and specify the EC2 instances as the targets. This solution will enable the company to handle millions of UDP requests per second with ultra-low latency and high performance. A Network Load Balancer is a type of Elastic Load Balancing that operates at the connection level (Layer 4) and routes traffic to targets (EC2 instances, microservices, or containers) within Amazon VPC based on IP protocol data. A Network Load Balancer is ideal for load balancing of both TCP and UDP traffic, as it is capable of handling millions of requests per second while maintaining high throughput at ultra-low latency. A Network Load Balancer also preserves the source IP address of the clients to the back-end applications, which can be useful for logging or security purposes¹.

NEW QUESTION 24

- (Topic 4)

A company runs a highly available SFTP service. The SFTP service uses two Amazon EC2

Linux instances that run with elastic IP addresses to accept traffic from trusted IP sources on the internet. The SFTP service is backed by shared storage that is attached to the instances. User accounts are created and managed as Linux users in the SFTP servers.

The company wants a serverless option that provides high IOPS performance and highly configurable security. The company also wants to maintain control over user permissions.

Which solution will meet these requirements?

- A. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume
- B. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses
- C. Attach the EBS volume to the SFTP service endpoint
- D. Grant users access to the SFTP service.
- E. Create an encrypted Amazon Elastic File System (Amazon EFS) volume
- F. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing access
- G. Attach a security group to the endpoint that allows only trusted IP addresses
- H. Attach the EFS volume to the SFTP service endpoint
- I. Grant users access to the SFTP service.
- J. Create an Amazon S3 bucket with default encryption enabled
- K. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresses
- L. Attach the S3 bucket to the SFTP service endpoint
- M. Grant users access to the SFTP service.
- N. Create an Amazon S3 bucket with default encryption enabled
- O. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subnet
- P. Attach a security group that allows only trusted IP addresses
- Q. Attach the S3 bucket to the SFTP service endpoint
- R. Grant users access to the SFTP service.

Answer: C

Explanation:

AWS Transfer Family is a secure transfer service that enables you to transfer files into and out of AWS storage services using SFTP, FTPS, FTP, and AS2 protocols. You can use AWS Transfer Family to create an SFTP-enabled server with a public endpoint that allows only trusted IP addresses. You can also attach an Amazon S3 bucket with default encryption enabled to the SFTP service endpoint, which will provide high IOPS performance and highly configurable security for your data at rest. You can also maintain control over user permissions by granting users access to the SFTP service using IAM roles or service-managed identities. References: <https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html>

<https://docs.aws.amazon.com/transfer/latest/userguide/create-server-s3.html>

NEW QUESTION 26

- (Topic 4)

A company is using an Application Load Balancer (ALB) to present its application to the internet. The company finds abnormal traffic access patterns across the application. A solutions architect needs to improve visibility into the infrastructure to help the company understand these abnormalities better.

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

- A. Create a table in Amazon Athena for AWS CloudTrail log
- B. Create a query for the relevant information.
- C. Enable ALB access logging to Amazon S3. Create a table in Amazon Athena, and query the logs.
- D. Enable ALB access logging to Amazon S3. Open each file in a text editor, and search each line for the relevant information
- E. Use Amazon EMR on a dedicated Amazon EC2 instance to directly query the ALB to acquire traffic access log information.

Answer: B

Explanation:

This solution meets the requirements because it allows the company to improve visibility into the infrastructure by using ALB access logging and Amazon Athena. ALB access logging is a feature that captures detailed information about requests sent to the load balancer, such as the client's IP address, request path, response code, and latency. By enabling ALB access logging to Amazon S3, the company can store the access logs in an S3 bucket as compressed files. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. By creating a table in Amazon Athena for the access logs, the company can query the logs and get results in seconds. This way, the company can better understand the abnormal traffic access patterns across the application.

References:

- ? Access logs for your Application Load Balancer
- ? Querying Application Load Balancer Logs

NEW QUESTION 30

- (Topic 4)

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:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Condition": {
        "NotIpAddress": {
          "aws:SourceIp": [
            "192.0.2.0/24",
            "203.0.113.0/24"
          ]
        }
      }
    }
  ],
  "Resource": ["*"]
}
```

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: D

NEW QUESTION 34

- (Topic 4)

A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL.

The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time.

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

- A. Migrate the web tier and the application tier to Amazon EC2 instances in private subnet
- B. Migrate the database tier to Amazon RDS for MySQL in private subnets.
- C. Migrate the web tier to Amazon EC2 instances in public subnet
- D. Migrate the application tier to EC2 instances in private subnet
- E. Migrate the database tier to Amazon Aurora MySQL in private subnets.
- F. Migrate the web tier to Amazon EC2 instances in public subnet
- G. Migrate the application tier to EC2 instances in private subnet
- H. Migrate the database tier to Amazon RDS for MySQL in private subnets.
- I. Migrate the web tier and the application tier to Amazon EC2 instances in public subnet
- J. Migrate the database tier to Amazon Aurora MySQL in public subnets.

Answer: C

Explanation:

The solution that meets the requirements with the least operational overhead is to migrate the web tier to Amazon EC2 instances in public subnets, migrate the application tier to EC2 instances in private subnets, and migrate the database tier to Amazon RDS for MySQL in private subnets. This solution allows the company to migrate its three-tier application to AWS by making minimal changes to the architecture, as it preserves the same web, application, and database tiers and uses the same MySQL database engine. The solution also provides a database solution that can restore data to a specific point in time, as Amazon RDS for MySQL supports automated backups and point-in-time recovery. This solution also reduces the operational overhead by using managed services such as Amazon EC2 and Amazon RDS, which handle tasks such as provisioning, patching, scaling, and monitoring.

The other solutions do not meet the requirements as well as the first one because they either involve more changes to the architecture, do not provide point-in-time recovery, or do not follow best practices for security and availability. Migrating the database tier to Amazon Aurora MySQL would require changing the database engine and potentially modifying the application code to ensure compatibility. Migrating the web tier and the application tier to public subnets would expose them to more security risks and reduce their availability in case of a subnet failure. Migrating the database tier to public subnets would also compromise its security and performance. References:

- ? Migrate Your Application Database to Amazon RDS
- ? Amazon RDS for MySQL
- ? Amazon Aurora MySQL
- ? Amazon VPC

NEW QUESTION 38

- (Topic 4)

A company maintains about 300 TB in Amazon S3 Standard storage month after month. The S3 objects are each typically around 50 GB in size and are frequently replaced with multipart uploads by their global application. The number and size of S3 objects remain constant but the company's S3 storage costs are increasing each month.

How should a solutions architect reduce costs in this situation?

- A. Switch from multipart uploads to Amazon S3 Transfer Acceleration.
- B. Enable an S3 Lifecycle policy that deletes incomplete multipart uploads.
- C. Configure S3 inventory to prevent objects from being archived too quickly.
- D. Configure Amazon CloudFront to reduce the number of objects stored in Amazon S3.

Answer: B

Explanation:

This option is the most cost-effective way to reduce the S3 storage costs in this situation. Incomplete multipart uploads are parts of objects that are not completed or aborted by the application. They consume storage space and incur charges until they are deleted. By enabling an S3 Lifecycle policy that deletes incomplete multipart uploads, you can automatically remove them after a specified period of time (such as one day) and free up the storage space. This will reduce the S3 storage costs and also improve the performance of the application by avoiding unnecessary retries or errors.

Option A is not correct because switching from multipart uploads to Amazon S3 Transfer Acceleration will not reduce the S3 storage costs. Amazon S3 Transfer Acceleration is a feature that enables faster data transfers to and from S3 by using the AWS edge network. It is useful for improving the upload speed of large objects over long distances, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the feature.

Option C is not correct because configuring S3 inventory to prevent objects from being archived too quickly will not reduce the S3 storage costs. Amazon S3 Inventory is a feature that provides a report of the objects and their metadata in an S3 bucket. It is useful for managing and auditing the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by generating additional S3 objects for the inventory reports.

Option D is not correct because configuring Amazon CloudFront to reduce the number of objects stored in Amazon S3 will not reduce the S3 storage costs. Amazon CloudFront is a content delivery network (CDN) that distributes the S3 objects to edge locations for faster and lower latency access. It is useful for improving the download speed and availability of the S3 objects, but it does not affect the storage space or charges. In fact, it may increase the costs by adding a data transfer fee for using the service. References:

- ? Managing your storage lifecycle
- ? Using multipart upload
- ? Amazon S3 Transfer Acceleration
- ? Amazon S3 Inventory
- ? What Is Amazon CloudFront?

NEW QUESTION 40

- (Topic 4)

A company runs an SMB file server in its data center. The file server stores large files that the company frequently accesses for up to 7 days after the file creation date. After 7 days, the company needs to be able to access the files with a maximum retrieval time of 24 hours.

Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to increase the company's storage space.
- C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
- D. Create an Amazon FSx File Gateway to increase the company's storage space.
- E. Create an Amazon S3 Lifecycle policy to transition the data after 7 days.
- F. Configure access to Amazon S3 for each use.
- G. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:

Amazon S3 File Gateway is a service that provides a file-based interface to Amazon S3, which appears as a network file share. It enables you to store and retrieve Amazon S3 objects through standard file storage protocols such as SMB. S3 File Gateway can also cache frequently accessed data locally for low-latency access. S3 Lifecycle policy is a feature that allows you to define rules that automate the management of your objects throughout their lifecycle. You can use S3 Lifecycle policy to transition objects to different storage classes based on their age and access patterns. S3 Glacier Deep Archive is a storage class that offers the lowest cost for long-term data archiving, with a retrieval time of 12 hours or 48 hours. This solution will meet the requirements, as it allows the company to store large files in S3 with SMB file access, and to move the files to S3 Glacier Deep Archive after 7 days for cost savings and compliance.

References:

- ? 1 provides an overview of Amazon S3 File Gateway and its benefits.
- ? 2 explains how to use S3 Lifecycle policy to manage object storage lifecycle.

? 3 describes the features and use cases of S3 Glacier Deep Archive storage class.

NEW QUESTION 42

- (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 47

- (Topic 4)

A company is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store.

Which solution will meet these requirements?

- A. Create a new AWS Key Management Service (AWS KMS) key. Use AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS.
- B. Create a new AWS Key Management Service (AWS KMS) key. Enable Amazon EKS KMS secrets encryption on the Amazon EKS cluster.
- C. Create the Amazon EKS cluster with default options. Use the Amazon Elastic Block Store (Amazon EBS) Container Storage Interface (CSI) driver as an add-on.
- D. Create a new AWS Key Management Service (AWS KMS) key with the alias `aws/ebs` and enable default Amazon EBS volume encryption for the account.

Answer: B

Explanation:

This option is the most secure and simple way to encrypt the secrets that are stored in Amazon EKS. AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys that can be used to encrypt your data. Amazon EKS KMS secrets encryption is a feature that enables you to use a KMS key to encrypt the secrets that are stored in the Kubernetes etcd key-value store. This provides an additional layer of protection for your sensitive data, such as passwords, tokens, and keys. You can create a new KMS key or use an existing one, and then enable the Amazon EKS KMS secrets encryption on the Amazon EKS cluster. You can also use IAM policies to control who can access or use the KMS key.

Option A is not correct because using AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS is not necessary or efficient. AWS Secrets Manager is a service that helps you securely store, retrieve, and rotate your secrets, such as database credentials, API keys, and passwords. You can use it to manage secrets that are used by your applications or services outside of Amazon EKS, but it is not designed to encrypt the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using AWS Secrets Manager would incur additional costs and complexity, and it would not leverage the native Kubernetes secrets management capabilities.

Option C is not correct because using the Amazon EBS Container Storage Interface (CSI) driver as an add-on does not encrypt the secrets that are stored in Amazon EKS. The Amazon EBS CSI driver is a plugin that allows you to use Amazon EBS volumes as persistent storage for your Kubernetes pods. It is useful for providing durable and scalable storage for your applications, but it does not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the Amazon EBS CSI driver would require additional configuration and resources, and it would not provide the same level of security as using a KMS key.

Option D is not correct because creating a new AWS KMS key with the alias `aws/ebs` and enabling default Amazon EBS volume encryption for the account does not encrypt the secrets that are stored in Amazon EKS. The alias `aws/ebs` is a reserved alias that is used by AWS to create a default KMS key for your account. This key is used to encrypt the Amazon EBS volumes that are created in your account, unless you specify a different KMS key. Enabling default Amazon EBS volume encryption for the account is a setting that ensures that all new Amazon EBS volumes are encrypted by default. However, these features do not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the default KMS key or the default encryption setting would not provide the same level of control and security as using a custom KMS key and enabling the Amazon EKS KMS secrets encryption feature. References:

? Encrypting secrets used in Amazon EKS

? What Is AWS Key Management Service?

? What Is AWS Secrets Manager?

? Amazon EBS CSI driver

? Encryption at rest

NEW QUESTION 49

- (Topic 4)

A company used an Amazon RDS for MySQL DB instance during application testing. Before terminating the DB instance at the end of the test cycle, a solutions architect created two backups. The solutions architect created the first backup by using the mysqldump utility to create a database dump. The solutions architect created the second backup by enabling the final DB snapshot option on RDS termination. The company is now planning for a new test cycle and wants to create a new DB instance from the most recent backup. The company has chosen a MySQL-compatible edition of Amazon Aurora to host the DB instance. Which solutions will create the new DB instance? (Select TWO.)

- A. Import the RDS snapshot directly into Aurora.
- B. Upload the RDS snapshot to Amazon S3. Then import the RDS snapshot into Aurora.
- C. Upload the database dump to Amazon S3. Then import the database dump into Aurora.
- D. Use AWS Database Migration Service (AWS DMS) to import the RDS snapshot into Aurora.
- E. Upload the database dump to Amazon S3. Then use AWS Database Migration Service (AWS DMS) to import the database dump into Aurora.

Answer: AC

Explanation:

These answers are correct because they meet the requirements of creating a new DB instance from the most recent backup and using a MySQL-compatible edition of Amazon Aurora to host the DB instance. You can import the RDS snapshot directly into Aurora if the MySQL DB instance and the Aurora DB cluster are running the same version of MySQL. For example, you can restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.6, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is simple and requires the fewest number of steps. You can upload the database dump to Amazon S3 and then import the database dump into Aurora if the MySQL DB instance and the Aurora DB cluster are running different versions of MySQL. For example, you can import a MySQL version 5.6 database dump into Aurora MySQL version 5.7, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is more flexible and allows you to migrate across different versions of MySQL.

References:

- ? <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.RDSMySQL.Import.html>
- ? <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.RDSMySQL.Dump.html>

NEW QUESTION 50

- (Topic 4)

A company stores multiple Amazon Machine Images (AMIs) in an AWS account to launch its Amazon EC2 instances. The AMIs contain critical data and configurations that are necessary for the company's operations. The company wants to implement a solution that will recover accidentally deleted AMIs quickly and efficiently.

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

- A. Create Amazon Elastic Block Store (Amazon EBS) snapshots of the AMI
- B. Store the snapshots in a separate AWS account.
- C. Copy all AMIs to another AWS account periodically.
- D. Create a retention rule in Recycle Bin.
- E. Upload the AMIs to an Amazon S3 bucket that has Cross-Region Replication.

Answer: C

Explanation:

Recycle Bin is a data recovery feature that enables you to restore accidentally deleted Amazon EBS snapshots and EBS-backed AMIs. When using Recycle Bin, if your resources are deleted, they are retained in the Recycle Bin for a time period that you specify before being permanently deleted. You can restore a resource from the Recycle Bin at any time before its retention period expires. This solution has the least operational overhead, as you do not need to create, copy, or upload any additional resources. You can also manage tags and permissions for AMIs in the Recycle Bin. AMIs in the Recycle Bin do not incur any additional charges.

References:

- ? [Recover AMIs from the Recycle Bin](#)
- ? [Recover an accidentally deleted Linux AMI](#)

NEW QUESTION 51

- (Topic 4)

A company runs an infrastructure monitoring service. The company is building a new feature that will enable the service to monitor data in customer AWS accounts. The new feature will call AWS APIs in customer accounts to describe Amazon EC2 instances and read Amazon CloudWatch metrics.

What should the company do to obtain access to customer accounts in the MOST secure way?

- A. Ensure that the customers create an IAM role in their account with read-only EC2 and CloudWatch permissions and a trust policy to the company's account.
- B. Create a serverless API that implements a token vending machine to provide temporary AWS credentials for a role with read-only EC2 and CloudWatch permissions.
- C. Ensure that the customers create an IAM user in their account with read-only EC2 and CloudWatch permission
- D. Encrypt and store customer access and secret keys in a secrets management system.
- E. Ensure that the customers create an Amazon Cognito user in their account to use an IAM role with read-only EC2 and CloudWatch permission
- F. Encrypt and store the Amazon Cognito user and password in a secrets management system.

Answer: A

Explanation:

By having customers create an IAM role with the necessary permissions in their own accounts, the company can use AWS Identity and Access Management (IAM) to establish cross-account access. The trust policy allows the company's AWS account to assume the customer's IAM role temporarily, granting access to the specified resources (EC2 instances and CloudWatch metrics) within the customer's account. This approach follows the principle of least privilege, as the company only requests the necessary permissions and does not require long-term access keys or user credentials from the customers.

NEW QUESTION 53

- (Topic 4)

A company has deployed a multiplayer game for mobile devices. The game requires live location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations. The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to

maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly. What should a solutions architect do to improve the performance of the data tier?

- A. Take a snapshot of the existing DB instance
- B. Restore the snapshot with Multi-AZ enabled.
- C. Migrate from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards.
- D. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance
- E. Modify the game to use DAX.
- F. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance
- G. Modify the game to use Redis.

Answer: D

Explanation:

The solution that will improve the performance of the data tier is to deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance and modify the game to use Redis. This solution will enable the game to store and retrieve the location data of the players in a fast and scalable way, as Redis is an in-memory data store that supports geospatial data types and commands. By using ElastiCache for Redis, the game can reduce the load on the RDS for PostgreSQL DB instance, which is not optimized for high-frequency updates and queries of location data. ElastiCache for Redis also supports replication, sharding, and auto scaling to handle the increasing user base of the game. The other solutions are not as effective as the first one because they either do not improve the performance, do not support geospatial data, or do not leverage caching. Taking a snapshot of the existing DB instance and restoring it with Multi-AZ enabled will not improve the performance of the data tier, as it only provides high availability and durability, but not scalability or low latency. Migrating from Amazon RDS to Amazon OpenSearch Service with OpenSearch Dashboards will not improve the performance of the data tier, as OpenSearch Service is mainly designed for full-text search and analytics, not for real-time location tracking. OpenSearch Service also does not support geospatial data types and commands natively, unlike Redis. Deploying Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance and modifying the game to use DAX will not improve the performance of the data tier, as DAX is only compatible with DynamoDB, not with RDS for PostgreSQL. DAX also does not support geospatial data types and commands.

References:

- ? Amazon ElastiCache for Redis
- ? Geospatial Data Support - Amazon ElastiCache for Redis
- ? Amazon RDS for PostgreSQL
- ? Amazon OpenSearch Service
- ? Amazon DynamoDB Accelerator (DAX)

NEW QUESTION 57

- (Topic 4)

A company has a web application hosted over 10 Amazon EC2 instances with traffic directed by Amazon Route 53. The company occasionally experiences a timeout error when attempting to browse the application. The networking team finds that some DNS queries return IP addresses of unhealthy instances, resulting in the timeout error.

What should a solutions architect implement to overcome these timeout errors?

- A. Create a Route 53 simple routing policy record for each EC2 instance
- B. Associate a health check with each record.
- C. Create a Route 53 failover routing policy record for each EC2 instance
- D. Associate a health check with each record.
- E. Create an Amazon CloudFront distribution with EC2 instances as its origin
- F. Associate a health check with the EC2 instances.
- G. Create an Application Load Balancer (ALB) with a health check in front of the EC2 instance
- H. Route to the ALB from Route 53.

Answer: D

Explanation:

An Application Load Balancer (ALB) allows you to distribute incoming traffic across multiple backend instances, and can automatically route traffic to healthy instances while removing traffic from unhealthy instances. By using an ALB in front of the EC2 instances and routing traffic to it from Route 53, the load balancer can perform health checks on the instances and only route traffic to healthy instances, which should help to reduce or eliminate timeout errors caused by unhealthy instances.

NEW QUESTION 59

- (Topic 4)

A company is making a prototype of the infrastructure for its new website by manually provisioning the necessary infrastructure. This infrastructure includes an Auto Scaling group, an Application Load Balancer, and an Amazon RDS database. After the configuration has been thoroughly validated, the company wants the capability to immediately deploy the infrastructure for development and production use in two Availability Zones in an automated fashion.

What should a solutions architect recommend to meet these requirements?

- A. Use AWS Systems Manager to replicate and provision the prototype infrastructure in two Availability Zones.
- B. Define the infrastructure as a template by using the prototype infrastructure as a guide
- C. Deploy the infrastructure with AWS CloudFormation
- D. Use AWS Config to record the inventory of resources that are used in the prototype infrastructure
- E. Use AWS Config to deploy the prototype infrastructure into two Availability Zones.
- F. Use AWS Elastic Beanstalk and configure it to use an automated reference to the prototype infrastructure to automatically deploy new environments in two Availability Zones

Answer: B

Explanation:

AWS CloudFormation is a service that helps you model and set up your AWS resources by using templates that describe all the resources that you want, such as Auto Scaling groups, load balancers, and databases. You can use AWS CloudFormation to deploy your infrastructure in an automated and consistent way across multiple environments and regions. You can also use AWS CloudFormation to update or delete your infrastructure as a single unit.

Reference URLs:

- 1 <https://aws.amazon.com/cloudformation/>
- 2 <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>

3 <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-concepts.html>

NEW QUESTION 63

- (Topic 4)

A company runs a three-tier application in two AWS Regions. The web tier, the application tier, and the database tier run on Amazon EC2 instances. The company uses Amazon RDS for Microsoft SQL Server Enterprise for the database tier. The database tier is experiencing high load when weekly and monthly reports are run. The company wants to reduce the load on the database tier.

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

- A. Create read replica
- B. Configure the reports to use the new read replicas.
- C. Convert the RDS database to Amazon DynamoDB. Configure the reports to use DynamoDB
- D. Modify the existing RDS DB instances by selecting a larger instance size.
- E. Modify the existing RDS DB instances and put the instances into an Auto Scaling group.

Answer: A

Explanation:

It allows the company to create read replicas of its RDS database and reduce the load on the database tier. By creating read replicas, the company can offload read traffic from the primary database instance to one or more replicas. By configuring the reports to use the new read replicas, the company can improve performance and availability of its database tier. References:

? Working with Read Replicas

? Read Replicas for Amazon RDS for SQL Server

NEW QUESTION 65

- (Topic 4)

A company hosts a data lake on Amazon S3. The data lake ingests data in Apache Parquet format from various data sources. The company uses multiple transformation steps to prepare the ingested data. The steps include filtering of anomalies, normalizing of data to standard date and time values, and generation of aggregates for analyses.

The company must store the transformed data in S3 buckets that data analysts access. The company needs a prebuilt solution for data transformation that does not require code. The solution must provide data lineage and data profiling. The company needs to share the data transformation steps with employees throughout the company.

Which solution will meet these requirements?

- A. Configure an AWS Glue Studio visual canvas to transform the data
- B. Share the transformation steps with employees by using AWS Glue jobs.
- C. Configure Amazon EMR Serverless to transform the data
- D. Share the transformation steps with employees by using EMR Serverless jobs.
- E. Configure AWS Glue DataBrew to transform the data
- F. Share the transformation steps with employees by using DataBrew recipes.
- G. Create Amazon Athena tables for the data
- H. Write Athena SQL queries to transform the data
- I. Share the Athena SQL queries with employees.

Answer: C

Explanation:

The most suitable solution for the company's requirements is to configure AWS Glue DataBrew to transform the data and share the transformation steps with employees by using DataBrew recipes. This solution will provide a prebuilt solution for data transformation that does not require code, and will also provide data lineage and data profiling. The company can easily share the data transformation steps with employees throughout the company by using DataBrew recipes. AWS Glue DataBrew is a visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data for analytics or machine learning by up to 80% faster. Users can upload their data from various sources, such as Amazon S3, Amazon RDS, Amazon Redshift, Amazon Aurora, or Glue Data Catalog, and use a point-and-click interface to apply over 250 built-in transformations. Users can also preview the results of each transformation step and see how it affects the quality and distribution of the data¹.

A DataBrew recipe is a reusable set of transformation steps that can be applied to one or more datasets. Users can create recipes from scratch or use existing ones from the DataBrew recipe library. Users can also export, import, or share recipes with other users or groups within their AWS account or organization². DataBrew also provides data lineage and data profiling features that help users understand and improve their data quality. Data lineage shows the source and destination of each dataset and how it is transformed by each recipe step. Data profiling shows various statistics and metrics about each dataset, such as column

NEW QUESTION 70

- (Topic 4)

A company wants to use high-performance computing and artificial intelligence to improve its fraud prevention and detection technology. The company requires distributed processing to complete a single workload as quickly as possible.

Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS) and multiple containers.
- B. Use AWS ParallelCluster and the Message Passing Interface (MPI) libraries.
- C. Use an Application Load Balancer and Amazon EC2 instances.
- D. Use AWS Lambda functions.

Answer: B

Explanation:

AWS ParallelCluster is a service that allows you to create and manage high-performance computing (HPC) clusters on AWS. It supports multiple schedulers, including AWS Batch, which can run distributed workloads across multiple EC2 instances¹.

MPI is a standard for message passing between processes in parallel computing. It provides functions for sending and receiving data, synchronizing processes, and managing communication groups².

By using AWS ParallelCluster and MPI libraries, you can take advantage of the following benefits:

? You can easily create and configure HPC clusters that meet your specific requirements, such as instance type, number of nodes, network configuration, and storage options¹.

- ? You can leverage the scalability and elasticity of AWS to run large-scale parallel workloads without worrying about provisioning or managing servers¹.
- ? You can use MPI libraries to optimize the performance and efficiency of your parallel applications by enabling inter-process communication and data exchange².
- ? You can choose from a variety of MPI implementations that are compatible with AWS ParallelCluster, such as Open MPI, Intel MPI, and MPICH3.

NEW QUESTION 72

- (Topic 4)

A company runs a container application by using Amazon Elastic Kubernetes Service (Amazon EKS). The application includes microservices that manage customers and place orders. The company needs to route incoming requests to the appropriate microservices. Which solution will meet this requirement MOST cost-effectively?

- A. Use the AWS Load Balancer Controller to provision a Network Load Balancer.
- B. Use the AWS Load Balancer Controller to provision an Application Load Balancer.
- C. Use an AWS Lambda function to connect the requests to Amazon EKS.
- D. Use Amazon API Gateway to connect the requests to Amazon EKS.

Answer: B

Explanation:

An Application Load Balancer is a type of Elastic Load Balancer that operates at the application layer (layer 7) of the OSI model. It can distribute incoming traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. It can also route requests based on the content of the request, such as the host name, path, or query parameters¹.

The AWS Load Balancer Controller is a controller that helps you manage Elastic Load Balancers for your Kubernetes cluster. It can provision Application Load Balancers or Network Load Balancers when you create Kubernetes Ingress or Service resources².

By using the AWS Load Balancer Controller to provision an Application Load Balancer for your Amazon EKS cluster, you can achieve the following benefits:

- ? You can route incoming requests to the appropriate microservices based on the rules you define in your Ingress resource. For example, you can route requests with different host names or paths to different microservices that handle customers and orders².
- ? You can improve the performance and availability of your container applications by distributing the load across multiple targets and enabling health checks and automatic scaling¹.
- ? You can reduce the cost and complexity of managing your load balancers by using a single controller that integrates with Amazon EKS and Kubernetes. You do not need to manually create or configure load balancers or update them when your cluster changes².

NEW QUESTION 77

- (Topic 4)

A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance. New company management wants to ensure the application is highly available. What should a solutions architect do to meet this requirement?

- A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer.
- B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region.
- C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application.
- D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer.

Answer: A

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability-zone.html>

NEW QUESTION 82

- (Topic 4)

A company has a three-tier application for image sharing. The application uses an Amazon EC2 instance for the front-end layer, another EC2 instance for the application layer, and a third EC2 instance for a MySQL database. A solutions architect must design a scalable and highly available solution that requires the least amount of change to the application.

Which solution meets these requirements?

- A. Use Amazon S3 to host the front-end layer.
- B. Use AWS Lambda functions for the application layer.
- C. Move the database to an Amazon DynamoDB table.
- D. Use Amazon S3 to store and serve users' images.
- E. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application layer.
- F. Move the database to an Amazon RDS DB instance with multiple read replicas to serve users' images.
- G. Use Amazon S3 to host the front-end layer.
- H. Use a fleet of EC2 instances in an Auto Scaling group for the application layer.
- I. Move the database to a memory optimized instance type to store and serve users' images.
- J. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application layer.
- K. Move the database to an Amazon RDS Multi-AZ DB instance.
- L. Use Amazon S3 to store and serve users' images.

Answer: D

Explanation:

for "Highly available": Multi-AZ & for "least amount of changes to the application": Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring.

NEW QUESTION 83

- (Topic 4)

A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will

continue to grow

and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL.

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

- A. Send activity data to an Amazon Kinesis data stream
- B. Configure the stream to deliver the data to an Amazon S3 bucket.
- C. Send activity data to an Amazon Kinesis Data Firehose delivery stream
- D. Configure the stream to deliver the data to an Amazon Redshift cluster.
- E. Place activity data in an Amazon S3 bucket
- F. Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
- G. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zones
- H. Configure the service to forward data to an Amazon RDS Multi-AZ database.

Answer: B

Explanation:

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This allows you to use your data to gain new insights for your business and customers. The first step to create a data warehouse is to launch a set of nodes, called an Amazon Redshift cluster. After you provision your cluster, you can upload your data set and then perform data analysis queries. Regardless of the size of the data set, Amazon Redshift offers fast query performance using the same SQL-based tools and business intelligence applications that you use today.

NEW QUESTION 87

- (Topic 4)

A company is deploying an application that processes streaming data in near-real time. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to provide the lowest possible latency between nodes.

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

- A. Enable and configure enhanced networking on each EC2 instance
- B. Group the EC2 instances in separate accounts
- C. Run the EC2 instances in a cluster placement group
- D. Attach multiple elastic network interfaces to each EC2 instance
- E. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

Answer: AC

Explanation:

These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth, lower latency, and lower jitter. Enhanced networking uses single root I/O virtualization (SR-IOV) or Elastic Fabric Adapter (EFA) to provide direct access to the network hardware. You can enable and configure enhanced networking by choosing a supported instance type and a compatible operating system, and installing the required drivers. Option C runs the EC2 instances in a cluster placement group, which is a logical grouping of instances within a single Availability Zone that are placed close together on the same underlying hardware. Cluster placement groups provide the lowest network latency and the highest network throughput among the placement group options. You can run the EC2 instances in a cluster placement group by creating a placement group and launching the instances into it. Option B is not suitable because grouping the EC2 instances in separate accounts does not provide the lowest possible latency between nodes. Separate accounts are used to isolate and organize resources for different purposes, such as security, billing, or compliance. However, they do not affect the network performance or proximity of the instances. Moreover, grouping the EC2 instances in separate accounts would incur additional costs and complexity, and it would require setting up cross-account networking and permissions.

Option D is not suitable because attaching multiple elastic network interfaces to each EC2 instance does not provide the lowest possible latency between nodes. Elastic network interfaces are virtual network interfaces that can be attached to EC2 instances to provide additional network capabilities, such as multiple IP addresses, multiple subnets, or enhanced security. However, they do not affect the network performance or proximity of the instances. Moreover, attaching multiple elastic network interfaces to each EC2 instance would consume additional resources and limit the instance type choices.

Option E is not suitable because using Amazon EBS optimized instance types does not provide the lowest possible latency between nodes. Amazon EBS optimized instance types are instances that provide dedicated bandwidth for Amazon EBS volumes, which are block storage volumes that can be attached to EC2 instances. EBS optimized instance types improve the performance and consistency of the EBS volumes, but they do not affect the network performance or proximity of the instances. Moreover, using EBS optimized instance types would incur additional costs and may not be necessary for the streaming data workload.

References:

- ? Enhanced networking on Linux
- ? Placement groups
- ? Elastic network interfaces
- ? Amazon EBS-optimized instances

NEW QUESTION 88

- (Topic 4)

A company previously migrated its data warehouse solution to AWS. The company also has an AWS Direct Connect connection. Corporate office users query the data warehouse using a visualization tool. The average size of a query returned by the data warehouse is 50 MB and each webpage sent by the visualization tool is approximately 500 KB. Result sets returned by the data warehouse are not cached.

Which solution provides the LOWEST data transfer egress cost for the company?

- A. Host the visualization tool on premises and query the data warehouse directly over the internet.
- B. Host the visualization tool in the same AWS Region as the data warehouse
- C. Access it over the internet.
- D. Host the visualization tool on premises and query the data warehouse directly over a Direct Connect connection at a location in the same AWS Region.
- E. Host the visualization tool in the same AWS Region as the data warehouse and access it over a Direct Connect connection at a location in the same Region.

Answer: D

Explanation:

<https://aws.amazon.com/directconnect/pricing/> <https://aws.amazon.com/blogs/aws/aws-data-transfer-prices-reduced/>

NEW QUESTION 90

- (Topic 4)

A company runs a container application on a Kubernetes cluster in the company's data center. The application uses Advanced Message Queuing Protocol (AMQP) to communicate with a message queue. The data center cannot scale fast enough to meet the company's expanding business needs. The company wants to migrate the workloads to AWS.

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

- A. Migrate the container application to Amazon Elastic Container Service (Amazon ECS). Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages.
- B. Migrate the container application to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon MQ to retrieve the messages.
- C. Use highly available Amazon EC2 instances to run the application. Use Amazon MQ to retrieve the messages.
- D. Use AWS Lambda functions to run the application. Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages.

Answer: B

Explanation:

This option is the best solution because it allows the company to migrate the container application to AWS with minimal changes and leverage a managed service to run the Kubernetes cluster and the message queue. By using Amazon EKS, the company can run the container application on a fully managed Kubernetes control plane that is compatible with the existing Kubernetes tools and plugins. Amazon EKS handles the provisioning, scaling, patching, and security of the Kubernetes cluster, reducing the operational overhead and complexity. By using Amazon MQ, the company can use a fully managed message broker service that supports AMQP and other popular messaging protocols. Amazon MQ handles the administration, maintenance, and scaling of the message broker, ensuring high availability, durability, and security of the messages.

* A. Migrate the container application to Amazon Elastic Container Service (Amazon ECS). Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages. This option is not optimal because it requires the company to change the container orchestration platform from Kubernetes to ECS, which can introduce additional complexity and risk. Moreover, it requires the company to change the messaging protocol from AMQP to SQS, which can also affect the application logic and performance. Amazon ECS and Amazon SQS are both fully managed services that simplify the deployment and management of containers and messages, but they may not be compatible with the existing application architecture and requirements.

* C. Use highly available Amazon EC2 instances to run the application. Use Amazon MQ to retrieve the messages. This option is not ideal because it requires the company to manage the EC2 instances that host the container application. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, the company would need to install and maintain the Kubernetes software on the EC2 instances, which can also add complexity and risk. Amazon MQ is a fully managed message broker service that supports AMQP and other popular messaging protocols, but it cannot compensate for the lack of a managed Kubernetes service.

* D. Use AWS Lambda functions to run the application. Use Amazon Simple Queue Service (Amazon SQS) to retrieve the messages. This option is not feasible because AWS Lambda does not support running container applications directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run container applications on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the container API, which can introduce additional complexity and overhead. Moreover, Lambda functions have limitations in terms of available CPU, memory, and runtime, which may not suit the application needs. Amazon SQS is a fully managed message queue service that supports asynchronous communication, but it does not support AMQP or other messaging protocols.

References:

- ? 1 Amazon Elastic Kubernetes Service - Amazon Web Services
- ? 2 Amazon MQ - Amazon Web Services
- ? 3 Amazon Elastic Container Service - Amazon Web Services
- ? 4 AWS Lambda FAQs - Amazon Web Services

NEW QUESTION 93

- (Topic 4)

A company wants to create an application to store employee data in a hierarchical structured relationship. The company needs a minimum-latency response to high-traffic queries for the employee data and must protect any sensitive data. The company also needs to receive monthly email messages if any financial information is present in the employee data.

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

- A. Use Amazon Redshift to store the employee data in hierarchy.
- B. Unload the data to Amazon S3 every month.
- C. Use Amazon DynamoDB to store the employee data in hierarchy.
- D. Export the data to Amazon S3 every month.
- E. Configure Amazon IAM for the AWS account.
- F. Integrate Macie with Amazon EventBridge to send monthly events to AWS Lambda.
- G. Use Amazon Athena to analyze the employee data in Amazon S3. Integrate Athena with Amazon QuickSight to publish analysis dashboards and share the dashboards with users.
- H. Configure Amazon Macie for the AWS account. Integrate Macie with Amazon EventBridge to send monthly notifications through an Amazon Simple Notification Service (Amazon SNS) subscription.

Answer: BE

Explanation:

Generally, for building a hierarchical relationship model, a graph database such as Amazon Neptune is a better choice. In some cases, however, DynamoDB is a better choice for hierarchical data modeling because of its flexibility, security, performance, and scale. <https://docs.aws.amazon.com/prescriptive-guidance/latest/dynamodb-hierarchical-data-model/introduction.html>

NEW QUESTION 96

- (Topic 4)

An e-commerce company stores terabytes of customer data in the AWS Cloud. The data contains personally identifiable information (PII). The company wants to use the data in three applications. Only one of the applications needs to process the PII. The PII must be removed before the other two applications process the data.

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

- A. Store the data in an Amazon DynamoDB table.
- B. Create a proxy application layer to intercept and process the data that each application requests.
- C. Store the data in an Amazon S3 bucket.
- D. Process and transform the data by using S3 Object Lambda before returning the data to the requesting application.
- E. Process the data and store the transformed data in three separate Amazon S3 buckets so that each application has its own custom dataset.

- F. Point each application to its respective S3 bucket.
- G. Process the data and store the transformed data in three separate Amazon DynamoDB tables so that each application has its own custom database.
- H. Point each application to its respective DynamoDB table.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/introducing-amazon-s3-object-lambda-use-your-code-to-process-data-as-it-is-being-retrieved-from-s3/>
S3 Object Lambda is a new feature of Amazon S3 that enables customers to add their own code to process data retrieved from S3 before returning it to the application. By using S3 Object Lambda, the data can be processed and transformed in real-time, without the need to store multiple copies of the data in separate S3 buckets or DynamoDB tables.
In this case, the PII can be removed from the data by the code added to S3 Object Lambda before returning the data to the two applications that do not need to process PII. The one application that requires PII can be pointed to the original S3 bucket where the PII is still stored.
Using S3 Object Lambda is the simplest and most cost-effective solution, as it eliminates the need to maintain multiple copies of the same data in different buckets or tables, which can result in additional storage costs and operational overhead.

NEW QUESTION 101

- (Topic 4)

A company stores critical data in Amazon DynamoDB tables in the company's AWS account. An IT administrator accidentally deleted a DynamoDB table. The deletion caused a significant loss of data and disrupted the company's operations. The company wants to prevent this type of disruption in the future. Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure a trail in AWS CloudTrail
- B. Create an Amazon EventBridge rule for delete action
- C. Create an AWS Lambda function to automatically restore deleted DynamoDB tables.
- D. Create a backup and restore plan for the DynamoDB table
- E. Recover the DynamoDB tables manually.
- F. Configure deletion protection on the DynamoDB tables.
- G. Enable point-in-time recovery on the DynamoDB tables.

Answer: C

Explanation:

Deletion protection is a feature of DynamoDB that prevents accidental deletion of tables. When deletion protection is enabled, you cannot delete a table unless you explicitly disable it first. This adds an extra layer of security and reduces the risk of data loss and operational disruption. Deletion protection is easy to enable and disable using the AWS Management Console, the AWS CLI, or the DynamoDB API. This solution has the least operational overhead, as you do not need to create, manage, or invoke any additional resources or services. References:

- ? [Using deletion protection to protect your table](#)
- ? [Preventing Accidental Table Deletion in DynamoDB](#)
- ? [Amazon DynamoDB now supports table deletion protection](#)

NEW QUESTION 106

- (Topic 4)

A company has an online gaming application that has TCP and UDP multiplayer gaming capabilities. The company uses Amazon Route 53 to point the application traffic to multiple Network Load Balancers (NLBs) in different AWS Regions. The company needs to improve application performance and decrease latency for the online game in preparation for user growth. Which solution will meet these requirements?

- A. Add an Amazon CloudFront distribution in front of the NLB
- B. Increase the Cache-Control: max-age parameter.
- C. Replace the NLBs with Application Load Balancers (ALBs). Configure Route 53 to use latency-based routing.
- D. Add AWS Global Accelerator in front of the NLB
- E. Configure a Global Accelerator endpoint to use the correct listener ports.
- F. Add an Amazon API Gateway endpoint behind the NLB
- G. Enable API caching
- H. Override method caching for the different stages.

Answer: C

Explanation:

This answer is correct because it improves the application performance and decreases latency for the online game by using AWS Global Accelerator. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your public applications. Global Accelerator provides two global static public IPs that act as a fixed entry point to your application endpoints, such as NLBs, in different AWS Regions. Global Accelerator uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies that you configure. Global Accelerator also terminates TCP and UDP traffic at the edge locations, which reduces the number of hops and improves the network performance. By adding AWS Global Accelerator in front of the NLBs, you can achieve up to 60% improvement in latency for your online game.

References:

- ? <https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html>
- ? <https://aws.amazon.com/global-accelerator/>

NEW QUESTION 111

- (Topic 4)

A company has NFS servers in an on-premises data center that need to periodically back up small amounts of data to Amazon S3. Which solution meets these requirements and is MOST cost-effective?

- A. Set up AWS Glue to copy the data from the on-premises servers to Amazon S3.
- B. Set up an AWS DataSync agent on the on-premises servers, and sync the data to Amazon S3.
- C. Set up an SFTP sync using AWS Transfer for SFTP to sync data from on-premises to Amazon S3.
- D. Set up an AWS Direct Connect connection between the on-premises data center and a VPC, and copy the data to Amazon S3.

Answer: B

Explanation:

AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and AWS storage services. AWS DataSync can transfer data at speeds up to 10 times faster than open-source tools by using a purpose-built network protocol and parallelizing data transfers. AWS DataSync also handles encryption, data integrity verification, and bandwidth optimization. To use AWS DataSync, users need to deploy a DataSync agent on their on-premises servers, which connects to the NFS servers and syncs the data to Amazon S3. Users can schedule periodic or one-time sync tasks and monitor the progress and status of the transfers.

The other options are not correct because they are either not cost-effective or not suitable for the use case. Setting up AWS Glue to copy the data from the on-premises servers to Amazon S3 is not cost-effective because AWS Glue is a serverless data integration service that is mainly used for extract, transform, and load (ETL) operations, not for simple data backup. Setting up an SFTP sync using AWS Transfer for SFTP to sync data from on-premises to Amazon S3 is not cost-effective because AWS Transfer for SFTP is a fully managed service that provides secure file transfer using the SFTP protocol, which is more suitable for exchanging data with third parties than for backing up data. Setting up an AWS Direct Connect connection between the on-premises data center and a VPC, and copying the data to Amazon S3 is not cost-effective because AWS Direct Connect is a dedicated network connection between AWS and the on-premises location, which has high upfront costs and requires additional configuration.

References:

- ? AWS DataSync
- ? How AWS DataSync works
- ? AWS DataSync FAQs

NEW QUESTION 113

- (Topic 4)

A company uses AWS Organizations. The company wants to operate some of its AWS accounts with different budgets. The company wants to receive alerts and automatically prevent provisioning of additional resources on AWS accounts when the allocated budget threshold is met during a specific period.

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

- A. Use AWS Budgets to create a budget
- B. Set the budget amount under the Cost and Usage Reports section of the required AWS accounts.
- C. Use AWS Budgets to create a budget
- D. Set the budget amount under the Billing dashboards of the required AWS accounts.
- E. Create an IAM user for AWS Budgets to run budget actions with the required permissions.
- F. Create an IAM role for AWS Budgets to run budget actions with the required permissions.
- G. Add an alert to notify the company when each account meets its budget threshold
- H. Add a budget action that selects the IAM identity created with the appropriate config rule to prevent provisioning of additional resources.
- I. Add an alert to notify the company when each account meets its budget threshold
- J. Add a budget action that selects the IAM identity created with the appropriate service control policy (SCP) to prevent provisioning of additional resources.

Answer: BDF

Explanation:

To use AWS Budgets to create and manage budgets for different AWS accounts, the company needs to do the following steps:

? Use AWS Budgets to create a budget for each AWS account that needs a different budget amount. The budget can be based on cost or usage metrics, and can have different time periods, filters, and thresholds. The company can set the budget amount under the Billing dashboards of the required AWS accounts¹.

? Create an IAM role for AWS Budgets to run budget actions with the required permissions. A budget action is a response that AWS Budgets initiates when a budget exceeds a specified threshold. The IAM role allows AWS Budgets to perform actions on behalf of the company, such as applying an IAM policy or a service control policy (SCP) to restrict the provisioning of additional resources².

? Add an alert to notify the company when each account meets its budget threshold.

The alert can be sent via email or Amazon SNS. The company can also add a budget action that selects the IAM role created and the appropriate SCP to prevent provisioning of additional resources. An SCP is a type of policy that can be applied to an AWS account or an organizational unit (OU) within AWS Organizations. An SCP can limit the actions that users and roles can perform in the account or OU³.

References:

- ? 4: <https://aws.amazon.com/budgets/>
- ? 1: <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html>
- ? 2: <https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-controls.html>
- ? 3: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

NEW QUESTION 115

- (Topic 4)

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.
- D. Add a rule to the web ACL to cache static files.
- E. Create a second ALB in an alternative AWS Region. Route user traffic to the closest Region to minimize data transfer costs.

Answer: A

Explanation:

Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website. CloudFront

can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not. References:

- ? <https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/>
- ? <https://nodeployfriday.com/posts/static-website-hosting/>

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

NEW QUESTION 120

- (Topic 4)

A company is running its production and nonproduction environment workloads in multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to design a solution that will prevent the modification of cost usage tags. Which solution will meet these requirements?

- A. Create a custom AWS Config rule to prevent tag modification except by authorized principals.
- B. Create a custom trail in AWS CloudTrail to prevent tag modification
- C. Create a service control policy (SCP) to prevent tag modification except by authorized principals.
- D. Create custom Amazon CloudWatch logs to prevent tag modification.

Answer: C

Explanation:

This solution meets the requirements because it uses SCPs to restrict the actions that can be performed on cost usage tags in the organization. SCPs are a type of organization policy that you can use to manage permissions in your organization. SCPs specify the maximum permissions for an organization, organizational unit (OU), or account. You can use SCPs to enforce consistent tag policies across your organization and prevent unauthorized or accidental changes to your tags. You can also create exceptions for authorized principals, such as administrators or auditors, who need to modify tags for legitimate purposes.

References:

? [Service control policies \(SCPs\) - AWS Organizations](#)

? [Tag policies - AWS Organizations](#)

NEW QUESTION 122

- (Topic 4)

A company is storing 700 terabytes of data on a large network-attached storage (NAS) system in its corporate data center. The company has a hybrid environment with a 10 Gbps AWS Direct Connect connection.

After an audit from a regulator, the company has 90 days to move the data to the cloud. The company needs to move the data efficiently and without disruption. The company still needs to be able to access and update the data during the transfer window.

Which solution will meet these requirements?

- A. Create an AWS DataSync agent in the corporate data center
- B. Create a data transfer task
- C. Start the transfer to an Amazon S3 bucket.
- D. Back up the data to AWS Snowball Edge Storage Optimized device
- E. Ship the devices to an AWS data center
- F. Mount a target Amazon S3 bucket on the on-premises file system.
- G. Use rsync to copy the data directly from local storage to a designated Amazon S3 bucket over the Direct Connect connection.
- H. Back up the data on tape
- I. Ship the tapes to an AWS data center
- J. Mount a target Amazon S3 bucket on the on-premises file system.

Answer: A

Explanation:

This answer is correct because it meets the requirements of moving the data efficiently and without disruption, and still being able to access and update the data during the transfer window. AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS and helps you move data quickly and securely between on-premises storage, edge locations, other clouds, and AWS Storage. You can create an AWS DataSync agent in the corporate data center to connect your NAS system to AWS over the Direct Connect connection. You can create a data transfer task to specify the source location, destination location, and options for transferring the data. You can start the transfer to an Amazon S3 bucket and monitor the progress of the task. DataSync automatically encrypts data in transit and verifies data integrity during transfer. DataSync also supports incremental transfers, which means that only files that have changed since the last transfer are copied. This way, you can ensure that your data is synchronized between your NAS system and S3 bucket, and you can access and update the data during the transfer window.

References:

? <https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>

? <https://docs.aws.amazon.com/datasync/latest/userguide/how-datasync-works.html>

NEW QUESTION 127

- (Topic 4)

A company needs a solution to prevent photos with unwanted content from being uploaded to the company's web application. The solution must not involve training a machine learning (ML) model. Which solution will meet these requirements?

- A. Create and deploy a model by using Amazon SageMaker Autopilot
- B. Create a real-time endpoint that the web application invokes when new photos are uploaded.
- C. Create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content
- D. Create a Lambda function URL that the web application invokes when new photos are uploaded.
- E. Create an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content
- F. Associate the function with the web application.
- G. Create an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content
- H. Create a Lambda function URL that the web application invokes when new photos are uploaded.

Answer: B

Explanation:

The solution that will meet the requirements is to create an AWS Lambda function that uses Amazon Rekognition to detect unwanted content, and create a Lambda function URL that the web application invokes when new photos are uploaded. This solution does not involve training a machine learning model, as Amazon Rekognition is a fully managed service that provides pre-trained computer vision models for image and video analysis. Amazon Rekognition can detect unwanted content such as explicit or suggestive adult content, violence, weapons, drugs, and more. By using AWS Lambda, the company can create a serverless function that can be triggered by an HTTP request from the web application. The Lambda function can use the Amazon Rekognition API to analyze the uploaded

photos and return a response indicating whether they contain unwanted content or not.

The other solutions are not as effective as the first one because they either involve training a machine learning model, do not support image analysis, or do not work with photos. Creating and deploying a model by using Amazon SageMaker Autopilot involves training a machine learning model, which is not required for the scenario. Amazon SageMaker Autopilot is a service that automatically creates, trains, and tunes the best machine learning models for classification or regression based on the data provided by the user. Creating an Amazon CloudFront function that uses Amazon Comprehend to detect unwanted content does not support image analysis, as Amazon Comprehend is a natural language processing service that analyzes text, not images. Amazon Comprehend can extract insights and relationships from text such as language, sentiment, entities, topics, and more. Creating an AWS Lambda function that uses Amazon Rekognition Video to detect unwanted content does not work with photos, as Amazon Rekognition Video is designed for analyzing video streams, not static images. Amazon Rekognition Video can detect activities, objects, faces, celebrities, text, and more in video streams.

References:

- ? Amazon Rekognition
- ? AWS Lambda
- ? Detecting unsafe content - Amazon Rekognition
- ? Amazon SageMaker Autopilot
- ? Amazon Comprehend

NEW QUESTION 131

- (Topic 4)

A social media company runs its application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. The application has more than a billion images stored in an Amazon S3 bucket and processes thousands of images each second. The company wants to resize the images dynamically and serve appropriate formats to clients.

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

- A. Install an external image management library on an EC2 instance
- B. Use the image management library to process the images.
- C. Create a CloudFront origin request policy
- D. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.
- E. Use a Lambda@Edge function with an external image management library
- F. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.
- G. Create a CloudFront response headers policy
- H. Use the policy to automatically resize images and to serve the appropriate format based on the User-Agent HTTP header in the request.

Answer: C

Explanation:

To resize images dynamically and serve appropriate formats to clients, a Lambda@Edge function with an external image management library can be used. Lambda@Edge allows running custom code at the edge locations of CloudFront, which can process the images on the fly and optimize them for different devices and browsers. An external image management library can provide various image manipulation and optimization features. References:

- ? Lambda@Edge
- ? Resizing Images with Amazon CloudFront & Lambda@Edge

NEW QUESTION 136

- (Topic 4)

A company manages AWS accounts in AWS Organizations. AWS IAM Identity Center (AWS Single Sign-On) and AWS Control Tower are configured for the accounts. The company wants to manage multiple user permissions across all the accounts.

The permissions will be used by multiple IAM users and must be split between the developer and administrator teams. Each team requires different permissions. The company wants a solution that includes new users that are hired on both teams.

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

- A. Create individual users in IAM Identity Center (or each account)
- B. Create separate developer and administrator groups in IAM Identity Center
- C. Assign the users to the appropriate groups Create a custom IAM policy for each group to set fine-grained permissions.
- D. Create individual users in IAM Identity Center for each account
- E. Create separate developer and administrator groups in IAM Identity Center
- F. Assign the users to the appropriate group
- G. Attach AWS managed IAM policies to each user as needed for fine-grained permissions.
- H. Create individual users in IAM Identity Center Create new developer and administrator groups in IAM Identity Center
- I. Create new permission sets that include the appropriate IAM policies for each group
- J. Assign the new groups to the appropriate accounts Assign the new permission sets to the new groups When new users are hired, add them to the appropriate group.
- K. Create individual users in IAM Identity Center
- L. Create new permission sets that include the appropriate IAM policies for each user
- M. Assign the users to the appropriate account
- N. Grant additional IAM permissions to the users from within specific account
- O. When new users are hired, add them to IAM Identity Center and assign them to the accounts.

Answer: C

Explanation:

This solution meets the requirements with the least operational overhead because it leverages the features of IAM Identity Center and AWS Control Tower to centrally manage multiple user permissions across all the accounts. By creating new groups and permission sets, the company can assign fine-grained permissions to the developer and administrator teams based on their roles and responsibilities. The permission sets are applied to the groups at the organization level, so they are automatically inherited by all the accounts in the organization. When new users are hired, the company only needs to add them to the appropriate group in IAM Identity Center, and they will automatically get the permissions assigned to that group. This simplifies the user management and reduces the manual effort of assigning permissions to each user individually.

References:

- ? Managing access to AWS accounts and applications
- ? Managing permission sets
- ? Managing groups

NEW QUESTION 137

- (Topic 4)

A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS). The company's workload is not consistent throughout the day. The company wants Amazon EKS to scale in and out according to the workload.

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

- A. Use an AWS Lambda function to resize the EKS cluster
- B. Use the Kubernetes Metrics Server to activate horizontal pod autoscaling.
- C. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
- D. Use Amazon API Gateway and connect it to Amazon EKS
- E. Use AWS App Mesh to observe network activity.

Answer: BC

Explanation:

<https://docs.aws.amazon.com/eks/latest/userguide/horizontal-pod-autoscaler.html> <https://docs.aws.amazon.com/eks/latest/userguide/autoscaling.html>
Horizontal pod autoscaling is a feature of Kubernetes that automatically scales the number of pods in a deployment, replication controller, or replica set based on that resource's CPU utilization. It requires a metrics source such as the Kubernetes Metrics Server to provide CPU usage data¹. Cluster autoscaling is a feature of Kubernetes that automatically adjusts the number of nodes in a cluster when pods fail or are rescheduled onto other nodes. It requires an integration with AWS Auto Scaling groups to manage the EC2 instances that join the cluster². By using both horizontal pod autoscaling and cluster autoscaling, the solution can ensure that Amazon EKS scales in and out according to the workload.

NEW QUESTION 141

- (Topic 4)

A company uses high concurrency AWS Lambda functions to process a constantly increasing number of messages in a message queue during marketing events. The Lambda functions use CPU intensive code to process the messages. The company wants to reduce the compute costs and to maintain service latency for its customers.

Which solution will meet these requirements?

- A. Configure reserved concurrency for the Lambda function
- B. Decrease the memory allocated to the Lambda functions.
- C. Configure reserved concurrency for the Lambda function
- D. Increase the memory according to AWS Compute Optimizer recommendations.
- E. Configure provisioned concurrency for the Lambda function
- F. Decrease the memory allocated to the Lambda functions.
- G. Configure provisioned concurrency for the Lambda function
- H. Increase the memory according to AWS Compute Optimizer recommendations.

Answer: D

Explanation:

The company wants to reduce the compute costs and maintain service latency for its Lambda functions that process a constantly increasing number of messages in a message queue. The Lambda functions use CPU intensive code to process the messages. To meet these requirements, a solutions architect should recommend the following solution:

? Configure provisioned concurrency for the Lambda functions. Provisioned concurrency is the number of pre-initialized execution environments that are allocated to the Lambda functions. These execution environments are prepared to respond immediately to incoming function requests, reducing the cold start latency. Configuring provisioned concurrency also helps to avoid throttling errors due to reaching the concurrency limit of the Lambda service.

? Increase the memory according to AWS Compute Optimizer recommendations.

AWS Compute Optimizer is a service that provides recommendations for optimal AWS resource configurations based on your utilization data. By increasing the memory allocated to the Lambda functions, you can also increase the CPU power and improve the performance of your CPU intensive code. AWS Compute Optimizer can help you find the optimal memory size for your Lambda functions based on your workload characteristics and performance goals.

This solution will reduce the compute costs by avoiding unnecessary over-provisioning of memory and CPU resources, and maintain service latency by using provisioned concurrency and optimal memory size for the Lambda functions.

References:

? Provisioned Concurrency

? AWS Compute Optimizer

NEW QUESTION 144

- (Topic 4)

A company wants to run its payment application on AWS. The application receives payment notifications from mobile devices. Payment notifications require a basic validation before they are sent for further processing.

The backend processing application is long running and requires compute and memory to be adjusted. The company does not want to manage the infrastructure. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices. Configure the rule to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS). Anywhere. Create a standalone cluster.
- B. Create an Amazon API Gateway API. Integrate the API with an AWS Step Functions state machine to receive payment notifications from mobile devices. Invoke the state machine to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure an EKS cluster with self-managed nodes.
- C. Create an Amazon Simple Queue Service (Amazon SQS) queue. Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices. Configure the rule to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon EC2 Spot Instances. Configure a Spot Fleet with a default allocation strategy.
- D. Create an Amazon API Gateway API. Integrate the API with AWS Lambda to receive payment notifications from mobile devices. Invoke a Lambda function to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Container Service (Amazon ECS). Configure Amazon ECS with an AWS Fargate launch type.

Answer: D

Explanation:

This option is the best solution because it allows the company to run its payment application on AWS with minimal operational overhead and infrastructure management. By using Amazon API Gateway, the company can create a secure and scalable API to receive payment notifications from mobile devices. By using AWS Lambda, the company can run a serverless function to validate the payment notifications and send them to the backend application. Lambda handles the provisioning, scaling, and security of the function, reducing the operational complexity and cost. By using Amazon ECS with AWS Fargate, the company can run the backend application on a fully managed container service that scales the compute resources automatically and does not require any EC2 instances to manage. Fargate allocates the right amount of CPU and memory for each container and adjusts them as needed.

* A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices. Configure the rule to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS) Anywhere. Create a standalone cluster. This option is not optimal because it requires the company to manage the Kubernetes cluster that runs the backend application. Amazon EKS Anywhere is a deployment option that allows the company to create and operate Kubernetes clusters on-premises or in other environments outside AWS. The company would need to provision, configure, scale, patch, and monitor the cluster nodes, which can increase the operational overhead and complexity. Moreover, the company would need to ensure the connectivity and security between the AWS services and the EKS Anywhere cluster, which can also add challenges and risks.

* B. Create an Amazon API Gateway API. Integrate the API with an AWS Step Functions state machine to receive payment notifications from mobile devices. Invoke the state machine to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure an EKS cluster with self-managed nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the Kubernetes cluster that runs the backend application. Amazon EKS is a fully managed service that runs Kubernetes on AWS, but it still requires the company to manage the worker nodes that run the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using AWS Step Functions to validate the payment notifications may be unnecessary and complex, as the validation logic can be implemented in a simpler way with Lambda or other services.

* C. Create an Amazon Simple Queue Service (Amazon SQS) queue. Integrate the queue with an Amazon EventBridge rule to receive payment notifications from mobile devices. Configure the rule to validate payment notifications and send the notifications to the backend application. Deploy the backend application on Amazon EC2 Spot Instances. Configure a Spot Fleet with a default allocation strategy. This option is not cost-effective because it requires the company to manage the EC2 instances that run the backend application. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs. Moreover, using Spot Instances can introduce the risk of interruptions, as Spot Instances are reclaimed by AWS when the demand for On-Demand Instances increases. The company would need to handle the interruptions gracefully and ensure the availability and reliability of the backend application.

References:

- ? 1 Amazon API Gateway - Amazon Web Services
- ? 2 AWS Lambda - Amazon Web Services
- ? 3 Amazon Elastic Container Service - Amazon Web Services
- ? 4 AWS Fargate - Amazon Web Services

NEW QUESTION 146

- (Topic 4)

A company has multiple AWS accounts with applications deployed in the us-west-2 Region. Application logs are stored within Amazon S3 buckets in each account. The company wants

to build a centralized log analysis solution that uses a single S3 bucket. Logs must not leave us-west-2, and the company wants to incur minimal operational overhead.

Which solution meets these requirements and is MOST cost-effective?

- A. Create an S3 Lifecycle policy that copies the objects from one of the application S3 buckets to the centralized S3 bucket.
- B. Use S3 Same-Region Replication to replicate logs from the S3 buckets to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.
- C. Write a script that uses the PutObject API operation every day to copy the entire contents of the buckets to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.
- D. Write AWS Lambda functions in these accounts that are triggered every time logs are delivered to the S3 buckets (s3:ObjectCreated:* event). Copy the logs to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.

Answer: B

Explanation:

This solution meets the following requirements:

? It is cost-effective, as it only charges for the storage and data transfer of the replicated objects, and does not require any additional AWS services or custom scripts. S3 Same-Region Replication (SRR) is a feature that automatically replicates objects across S3 buckets within the same AWS Region. SRR can help you aggregate logs from multiple sources to a single destination for analysis and auditing. SRR also preserves the metadata, encryption, and access control of the source objects.

? It is operationally efficient, as it does not require any manual intervention or scheduling. SRR replicates objects as soon as they are uploaded to the source bucket, ensuring that the destination bucket always has the latest log data. SRR also handles any updates or deletions of the source objects, keeping the destination bucket in sync. SRR can be enabled with a few clicks in the S3 console or with a simple API call.

? It is secure, as it does not allow the logs to leave the us-west-2 Region. SRR only replicates objects within the same AWS Region, ensuring that the data sovereignty and compliance requirements are met. SRR also supports encryption of the source and destination objects, using either server-side encryption with AWS KMS or S3-managed keys, or client-side encryption.

References:

- ? Same-Region Replication - Amazon Simple Storage Service
- ? How do I replicate objects across S3 buckets in the same AWS Region?
- ? Centralized Logging on AWS | AWS Solutions | AWS Solutions Library

NEW QUESTION 148

.....

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

AWS-Solution-Architect-Associate Practice Exam Features:

- * AWS-Solution-Architect-Associate Questions and Answers Updated Frequently
- * AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Solution-Architect-Associate Practice Test Here](#)