

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 is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket. During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible.

Which solution will meet these requirements?

- A. Enable S3 Intelligent-Tiering for the S3 bucket.
- B. Enable S3 Transfer Acceleration for the S3 bucket.
- C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC.
- D. Create an interface endpoint for Amazon S3 in the VPC.
- E. Associate this endpoint with all route tables in the VPC.

Answer: C

Explanation:

A gateway VPC endpoint for Amazon S3 enables private connections between the VPC and Amazon S3 that do not require an internet gateway or NAT device. This minimizes costs and prevents traffic from traversing the internet. A gateway VPC endpoint uses a prefix list as the route target in a VPC route table to route traffic privately to Amazon S3. Associating the endpoint with all route tables in the VPC ensures that all subnets can access Amazon S3 through the endpoint. Option A is incorrect because S3 Intelligent-Tiering is a storage class that optimizes storage costs by automatically moving objects between two access tiers based on changing access patterns. It does not affect the network traffic between the VPC and Amazon S3.

Option B is incorrect because S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between clients and an S3 bucket. It does not prevent traffic from traversing the internet.

Option D is incorrect because an interface VPC endpoint for Amazon S3 is powered by AWS PrivateLink, which requires an elastic network interface (ENI) with a private IP address in each subnet. This adds complexity and cost to the solution. Moreover, an interface VPC endpoint does not support cross-Region access to Amazon S3. Reference URL: 1: <https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html> 2:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-dynamic-data-access> 3:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html> : <https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/>

NEW QUESTION 3

- (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 4

- (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 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 solutions architect wants to use the following JSON text as an identity-based policy to grant specific permissions:

```
{
  "Statement": [
    {
      "Action": [
        "ssm:ListDocuments",
        "ssm:GetDocument"
      ],
      "Effect": "Allow",
      "Resource": "*",
      "Sid": ""
    }
  ],
  "Version": "2012-10-17"
}
```

Which IAM principals can the solutions architect attach this policy to? (Select TWO.)

- A. Role
- B. Group
- C. Organization
- D. Amazon Elastic Container Service (Amazon ECS) resource
- E. Amazon EC2 resource

Answer: AB

Explanation:

This JSON text is an identity-based policy that grants specific permissions. The IAM principals that the solutions architect can attach this policy to are Role and Group. This is because the policy is written in JSON and is an identity-based policy, which can be attached to IAM principals such as users, groups, and roles. Identity-based policies are permissions policies that you attach to IAM identities (users, groups, or roles) and explicitly state what that identity is allowed (or denied) to do¹. Identity-based policies are different from resource-based policies, which define the permissions around the specific resource¹. Resource-based policies are attached to a resource, such as an Amazon S3 bucket or an Amazon EC2 instance¹. Resource-based policies can also specify a principal, which is the entity that is allowed or denied access to the resource¹. Organization is not an IAM principal, but a feature of AWS Organizations that allows you to manage multiple AWS accounts centrally². Amazon ECS resource and Amazon EC2 resource are not IAM principals, but AWS resources that can have resource-based policies attached to them^{3,4}. References:

- ? Identity-based policies and resource-based policies
- ? AWS Organizations
- ? Amazon ECS task role
- ? Amazon EC2 instance profile

NEW QUESTION 7

- (Topic 4)

A company has a multi-tier payment processing application that is based on virtual machines (VMs). The communication between the tiers occurs asynchronously through a third-party middleware solution that guarantees exactly-once delivery.

The company needs a solution that requires the least amount of infrastructure management. The solution must guarantee exactly-once delivery for application messaging

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

- A. Use AWS Lambda for the compute layers in the architecture.
- B. Use Amazon EC2 instances for the compute layers in the architecture.
- C. Use Amazon Simple Notification Service (Amazon SNS) as the messaging component between the compute layers.
- D. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues as the messaging component between the compute layers.
- E. Use containers that are based on Amazon Elastic Kubernetes Service (Amazon EKS) for the compute layers in the architecture.

Answer: AD

Explanation:

This solution meets the requirements because it requires the least amount of infrastructure management and guarantees exactly-once delivery for application messaging. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You only pay for the compute time you consume. Lambda scales automatically with the size of your workload. Amazon SQS FIFO queues are designed to ensure that messages are processed exactly once, in the exact order that they are sent. FIFO queues have high availability and deliver messages in a strict first-in, first-out order. You can use Amazon SQS to decouple and scale microservices, distributed systems, and serverless applications. References: AWS Lambda, Amazon SQS FIFO queues

NEW QUESTION 8

- (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 to 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 for 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-1A). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 Standard-1A 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 9

- (Topic 4)

A company needs to migrate a MySQL database from its on-premises data center to AWS within 2 weeks. The database is 20 TB in size. The company wants to complete the migration with minimal downtime.

Which solution will migrate the database MOST cost-effectively?

- A. Order an AWS Snowball Edge Storage Optimized device
- B. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing change
- C. Send the Snowball Edge device to AWS to finish the migration and continue the ongoing replication.
- D. Order an AWS Snowmobile vehicle
- E. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
- F. Send the Snowmobile vehicle back to AWS to finish the migration and continue the ongoing replication.
- G. Order an AWS Snowball Edge Compute Optimized with GPU device
- H. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
- I. Send the Snowball device to AWS to finish the migration and continue the ongoing replication.
- J. Order a 1 GB dedicated AWS Direct Connect connection to establish a connection with the data center
- K. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing changes.

Answer: A

Explanation:

This answer is correct because it meets the requirements of migrating a 20 TB MySQL database within 2 weeks with minimal downtime and cost-effectively. The AWS Snowball Edge Storage Optimized device has up to 80 TB of usable storage space, which is enough to fit the database. The AWS Database Migration Service (AWS DMS) can migrate data from MySQL to Amazon Aurora, Amazon RDS for MySQL, or MySQL on Amazon EC2 with minimal downtime by continuously replicating changes from the source to the target. The AWS Schema Conversion Tool (AWS SCT) can convert the source schema and code to a format compatible with the target database. By using these services together, the company can migrate the database to AWS with minimal downtime and cost. The Snowball Edge device can be shipped back to AWS to finish the migration and continue the ongoing replication until the database is fully migrated.

References:

- ? <https://docs.aws.amazon.com/snowball/latest/developer-guide/device-differences.html>
- ? https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.MySQL.html
- ? https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/CHAP_Source.MySQL.htm

NEW QUESTION 10

- (Topic 4)

A financial company needs to handle highly sensitive data. The company will store the data in an Amazon S3 bucket. The company needs to ensure that the data is encrypted in transit and at rest. The company must manage the encryption keys outside the AWS Cloud.

Which solution will meet these requirements?

- A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key
- B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key
- C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE)
- D. Encrypt the data at the company's data center before storing the data in the S3 bucket

Answer: D

Explanation:

This option is the only solution that meets the requirements because it allows the company to encrypt the data with its own encryption keys and tools outside the AWS Cloud. By encrypting the data at the company's data center before storing the data in the S3 bucket, the company can ensure that the data is encrypted in transit and at rest, and that the company has full control over the encryption keys and processes. This option also avoids the need to use any AWS encryption services or features, which may not be compatible with the company's security policies or compliance standards.

- * A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.
- * B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.
- * C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards. References:
 - ? 1 Protecting data with encryption - Amazon Simple Storage Service
 - ? 2 Protecting data with server-side encryption - Amazon Simple Storage Service
 - ? 3 Protecting data by using client-side encryption - Amazon Simple Storage Service
 - ? 4 AWS Key Management Service Concepts - AWS Key Management Service

NEW QUESTION 10

- (Topic 4)

A company is designing a new web service that will run on Amazon EC2 instances behind an Elastic Load Balancing (ELB) load balancer. However, many of the web service clients can only reach IP addresses authorized on their firewalls. What should a solutions architect recommend to meet the clients' needs?

- A. A Network Load Balancer with an associated Elastic IP address.
- B. An Application Load Balancer with an associated Elastic IP address.
- C. An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address.
- D. An EC2 instance with a public IP address running as a proxy in front of the load balancer.

Answer: A

Explanation:

A Network Load Balancer can be assigned one Elastic IP address for each Availability Zone it uses¹. This allows the clients to reach the load balancer using a static IP address that can be authorized on their firewalls. An Application Load Balancer cannot be assigned an Elastic IP address². An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address would not work because the load balancer would still use its own IP address as the source of the forwarded requests to the web service. An EC2 instance with a public IP address running as a proxy in front of the load balancer would add unnecessary complexity and cost, and would not provide the same scalability and availability as a Network Load Balancer. References: 1: Network Load Balancers - Elastic Load Balancing³, IP address type section²: How to assign Elastic IP to Application Load Balancer in AWS⁴, answer section.

NEW QUESTION 12

- (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 accounts¹. 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 for². 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 services³.
- * 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 themselves¹. 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 Service².
- * 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 13

- (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 15

- (Topic 4)

A company is implementing new data retention policies for all databases that run on Amazon RDS DB instances. The company must retain daily backups for a minimum period of 2 years. The backups must be consistent and restorable.

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

- A. Create a backup vault in AWS Backup to retain RDS backup
- B. Create a new backup plan with a daily schedule and an expiration period of 2 years after creation
- C. Assign the RDS DB instances to the backup plan.
- D. Configure a backup window for the RDS DB instances for daily snapshots
- E. Assign a snapshot retention policy of 2 years to each RDS DB instance
- F. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule snapshot deletions.
- G. Configure database transaction logs to be automatically backed up to Amazon CloudWatch Logs with an expiration period of 2 years.
- H. Configure an AWS Database Migration Service (AWS DMS) replication task
- I. Deploy a replication instance, and configure a change data capture (CDC) task to stream database changes to Amazon S3 as the target
- J. Configure S3 Lifecycle policies to delete the snapshots after 2 years.

Answer: A

Explanation:

AWS Backup is a fully managed service that enables users to centralize and automate the backup of data across AWS services. It can create and manage backup plans that specify the frequency and retention period of backups. It can also assign backup resources to backup vaults, which are containers that store backup data. By using AWS Backup, the solution can ensure that the RDS backups are consistent, restorable, and retained for a minimum period of 2 years.

* B. Configure a backup window for the RDS DB instances for daily snapshots. Assign a snapshot retention policy of 2 years to each RDS DB instance. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule snapshot deletions. This solution will not meet the requirement of ensuring that the backups are consistent and restorable, as Amazon DLM is not compatible with RDS snapshots and cannot be used to schedule snapshot deletions.

* C. Configure database transaction logs to be automatically backed up to Amazon CloudWatch Logs with an expiration period of 2 years. This solution will not meet the requirement of ensuring that the backups are consistent and restorable, as database transaction logs are not sufficient to restore a database to a point in time. They only capture the changes made to the database, not the full state of the database.

* D. Configure an AWS Database Migration Service (AWS DMS) replication task. Deploy a replication instance, and configure a change data capture (CDC) task to stream database changes to Amazon S3 as the target. Configure S3 Lifecycle policies to delete the snapshots after 2 years. This solution will not meet the requirement of ensuring that the backups are consistent and restorable, as AWS DMS is a service that helps users migrate databases to AWS, not back up databases. It also requires additional resources and configuration, such as replication instances and CDC tasks.

Reference URL: <https://docs.aws.amazon.com/aws-backup/latest/devguide/what-is-backup.html>

NEW QUESTION 16

- (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 19

- (Topic 4)

A company is deploying an application that processes large quantities of data in parallel. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to prevent groups of nodes from sharing the same underlying hardware.

Which networking solution meets these requirements?

- A. Run the EC2 instances in a spread placement group.
- B. Group the EC2 instances in separate accounts.
- C. Configure the EC2 instances with dedicated tenancy.
- D. Configure the EC2 instances with shared tenancy.

Answer: A

Explanation:

it allows the company to deploy an application that processes large quantities of data in parallel and prevent groups of nodes from sharing the same underlying hardware. By running the EC2 instances in a spread placement group, the company can launch a small number of instances across distinct underlying hardware to reduce correlated failures. A spread placement group ensures that each instance is isolated from each other at the rack level. References:

- ? Placement Groups
- ? Spread Placement Groups

NEW QUESTION 21

- (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 23

- (Topic 4)

A solutions architect is designing a new API using Amazon API Gateway that will receive requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made.

Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

- A. An AWS Glue job
- B. An AWS Lambda function
- C. A containerized service hosted in Amazon Elastic Kubernetes Service (Amazon EKS)
- D. A containerized service hosted in Amazon ECS with Amazon EC2

Answer: B

Explanation:

API Gateway + Lambda is the perfect solution for modern applications with serverless architecture.

NEW QUESTION 24

- (Topic 4)

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

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

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

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

Answer: D

Explanation:

This answer is correct because it meets the requirements of running the ML models as independent microservices that can handle irregular and unpredictable usage patterns. By directing the requests from the API into an Amazon SQS queue, the company can decouple the request processing from the model execution, and ensure that no requests are lost due to spikes in demand. By deploying the models as Amazon ECS services that read from the queue, the company can leverage containers to isolate and package each model as a microservice, and fetch the model data from S3 at startup. By enabling AWS Auto Scaling on Amazon ECS for both the cluster and copies of the service based on the queue size, the company can automatically scale up or down the number of EC2 instances in the cluster and the number of tasks in each service to match the demand and optimize performance.

References:

- ? <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html>
- ? <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>
- ? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/autoscaling-ecs.html>

NEW QUESTION 26

- (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 30

- (Topic 4)

A company is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company's application. A solutions architect wants to implement a solution that is highly available, fault tolerant, and automatically scalable. What should the solutions architect recommend?

- A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
- B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in different Availability Zones.
- C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
- D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

Answer: C

Explanation:

If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway's Availability Zone is down, resources in the other Availability Zones lose internet access. To create an Availability Zone-independent architecture, create a NAT gateway in each Availability Zone and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone. <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html#nat-gateway-basics>

NEW QUESTION 32

- (Topic 4)

A solutions architect is designing a REST API in Amazon API Gateway for a cash payback service. The application requires 1 GB of memory and 2 GB of storage for its computation resources. The application will require that the data is in a relational format.

Which additional combination of AWS services will meet these requirements with the LEAST administrative effort? (Select TWO.)

- A. Amazon EC2
- B. AWS Lambda
- C. Amazon RDS
- D. Amazon DynamoDB
- E. Amazon Elastic Kubernetes Services (Amazon EKS)

Answer: BC

Explanation:

AWS Lambda is a service that lets users run code without provisioning or managing servers. It automatically scales and manages the underlying compute resources for the code. It supports multiple languages, such as Java, Python, Node.js, and G10. By using AWS Lambda for the REST API, the solution can meet the requirements of 1 GB of memory and minimal administrative effort.

Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups. It supports multiple database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server2. By using Amazon RDS for the data store, the solution can meet the requirements of 2 GB of storage and a relational format.

* A. Amazon EC2. This solution will not meet the requirement of minimal administrative effort, as Amazon EC2 is a service that provides virtual servers in the cloud that users have to configure and manage themselves. It requires users to choose an instance type, an operating system, a security group, and other options3.

* D. Amazon DynamoDB. This solution will not meet the requirement of a relational format, as Amazon DynamoDB is a service that provides a key-value and document database that delivers single-digit millisecond performance at any scale. It is a non-relational or NoSQL database that does not support joins or transactions.

* E. Amazon Elastic Kubernetes Services (Amazon EKS). This solution will not meet the requirement of minimal administrative effort, as Amazon EKS is a service that provides a fully managed Kubernetes service that users have to configure and manage themselves. It requires users to create clusters, nodes groups, pods, services, and other Kubernetes resources.

Reference URL: <https://aws.amazon.com/lambda/>

NEW QUESTION 34

- (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 37

- (Topic 4)

A retail company has several businesses. The IT team for each business manages its own AWS account. Each team account is part of an organization in AWS Organizations. Each team monitors its product inventory levels in an Amazon DynamoDB table in the team's own AWS account. The company is deploying a central inventory reporting application into a shared AWS account. The application must be able to read items from all the teams' DynamoDB tables.

Which authentication option will meet these requirements MOST securely?

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

Answer: C

Explanation:

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

References:

- ? IAM Roles
- ? STS AssumeRole
- ? Tutorial: Delegate Access Across AWS Accounts Using IAM Roles

NEW QUESTION 38

- (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 41

- (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 45

- (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 46

- (Topic 4)

A company runs its applications on Amazon EC2 instances. The company performs periodic financial assessments of its AWS costs. The company recently identified unusual spending.

The company needs a solution to prevent unusual spending. The solution must monitor costs and notify responsible stakeholders in the event of unusual spending. Which solution will meet these requirements?

- A. Use an AWS Budgets template to create a zero spend budget
- B. Create an AWS Cost Anomaly Detection monitor in the AWS Billing and Cost Management console.
- C. Create AWS Pricing Calculator estimates for the current running workload pricing details_
- D. Use Amazon CloudWatch to monitor costs and to identify unusual spending

Answer: B

Explanation:

it allows the company to monitor costs and notify responsible stakeholders in the event of unusual spending. By creating an AWS Cost Anomaly Detection monitor in the AWS Billing and Cost Management console, the company can use a machine learning service that automatically detects and alerts on anomalous spend. By configuring alert thresholds, notification preferences, and root cause analysis, the company can prevent unusual spending and identify its source. References:

? AWS Cost Anomaly Detection

? Creating a Cost Anomaly Monitor

NEW QUESTION 47

- (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 48

- (Topic 4)

A company wants to use an event-driven programming model with AWS Lambda. The company wants to reduce startup latency for Lambda functions that run on Java 11. The company does not have strict latency requirements for the applications. The company wants to reduce cold starts and outlier latencies when a function scales up.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure Lambda provisioned concurrency.
- B. Increase the timeout of the Lambda functions.
- C. Increase the memory of the Lambda functions.
- D. Configure Lambda SnapStart.

Answer: D

Explanation:

To reduce startup latency for Lambda functions that run on Java 11, Lambda SnapStart is a suitable solution. Lambda SnapStart is a feature that enables faster cold starts and lower outlier latencies for Java 11 functions. Lambda SnapStart uses a pre-initialized Java Virtual Machine (JVM) to run the functions, which reduces the initialization time and memory footprint. Lambda SnapStart does not incur any additional charges. References:

? Lambda SnapStart for Java 11 Functions

? Lambda SnapStart FAQs

NEW QUESTION 53

- (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 56

- (Topic 4)

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

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

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

Answer: C

Explanation:

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

* A. Use AWS Systems Manager Inventory to generate a map view from the detailed view report. This solution will not meet the requirement of building and mapping the relationship details of the various workloads across all accounts, as AWS Systems Manager Inventory is a feature that collects metadata from your managed instances and stores it in a central Amazon S3 bucket. It does not provide a map view or architecture diagrams of the workloads².

* B. Use AWS Step Functions to collect workload details Build architecture diagrams of the workloads manually. This solution will not meet the requirement of the

least operational effort, as it involves creating and managing state machines to orchestrate the workload details collection, and building architecture diagrams manually.

* D. Use AWS X-Ray to view the workload details Build architecture diagrams with relationships. This solution will not meet the requirement of the least operational effort, as it involves instrumenting your applications with X-Ray SDKs to collect workload details, and building architecture diagrams manually.

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

NEW QUESTION 60

- (Topic 4)

A company sends AWS CloudTrail logs from multiple AWS accounts to an Amazon S3 bucket in a centralized account. The company must keep the CloudTrail logs. The company must also be able to query the CloudTrail logs at any time

Which solution will meet these requirements?

- A. Use the CloudTrail event history in the centralized account to create an Amazon Athena table
- B. Query the CloudTrail logs from Athena.
- C. Configure an Amazon Neptune instance to manage the CloudTrail log
- D. Query the CloudTrail logs from Neptune.
- E. Configure CloudTrail to send the logs to an Amazon DynamoDB table
- F. Create a dashboard in Amazon QuickSight to query the logs in the table.
- G. use Amazon Athena to create an Athena notebook
- H. Configure CloudTrail to send the logs to the notebook
- I. Run queries from Athena.

Answer: A

Explanation:

it allows the company to keep the CloudTrail logs and query them at any time. By using the CloudTrail event history in the centralized account, the company can view, filter, and download recent API activity across multiple AWS accounts. By creating an Amazon Athena table from the CloudTrail event history, the company can use a serverless interactive query service that makes it easy to analyze data in S3 using standard SQL. By querying the CloudTrail logs from Athena, the company can gain insights into user activity and resource changes. References:

? Viewing Events with CloudTrail Event History

? Querying AWS CloudTrail Logs

? Amazon Athena

NEW QUESTION 62

- (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 66

- (Topic 4)

A company runs analytics software on Amazon EC2 instances The software accepts job requests from users to process data that has been uploaded to Amazon S3 Users report that some submitted data is not being processed Amazon CloudWatch reveals that the EC2 instances have a consistent CPU utilization at or near 100% The company wants to improve system performance and scale the system based on user load.

What should a solutions architect do to meet these requirements?

- A. Create a copy of the instance Place all instances behind an Application Load Balancer
- B. Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint
- C. Stop the EC2 instance
- D. Modify the instance type to one with a more powerful CPU and more memor
- E. Restart the instances.
- F. Route incoming requests to Amazon Simple Queue Service (Amazon SQS) Configure an EC2 Auto Scaling group based on queue size Update the software to read from thequeue.

Answer: D

Explanation:

This option is the best solution because it allows the company to decouple the analytics software from the user requests and scale the EC2 instances dynamically based on the demand. By using Amazon SQS, the company can create a queue that stores the user requests and acts as a buffer between the users and the analytics software. This way, the software can process the requests at its own pace without losing any data or overloading the EC2 instances. By using EC2 Auto Scaling, the company can create an Auto Scaling group that launches or terminates EC2 instances automatically based on the size of the queue. This way, the company can ensure that there are enough instances to handle the load and optimize the cost and performance of the system. By updating the software to read from the queue, the company can enable the analytics software to consume the requests from the queue and process the data from Amazon S3.

* A. Create a copy of the instance Place all instances behind an Application Load Balancer. This option is not optimal because it does not address the root cause of the problem, which is the high CPU utilization of the EC2 instances. An Application Load Balancer can distribute the incoming traffic across multiple instances, but it cannot scale the instances based on the load or reduce the processing time of the analytics software. Moreover, this option can incur additional costs for the load balancer and the extra instances.

* B. Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint. This option is not effective because it does not solve the issue of the high CPU utilization of the EC2 instances. An S3 VPC endpoint can enable the EC2 instances to access Amazon S3 without going through the internet, which can improve the network performance and security. However, it cannot reduce the processing time of the analytics software or scale the instances based on the load.

* C. Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances. This option is not scalable because it does not account for the variability of the user load. Changing the instance type to a more powerful one can improve the performance of the analytics software, but it cannot adjust the number of instances based on the demand. Moreover, this option can increase the cost of the system and cause downtime during the instance modification.

References:

- ? 1 Using Amazon SQS queues with Amazon EC2 Auto Scaling - Amazon EC2 Auto Scaling
- ? 2 Tutorial: Set up a scaled and load-balanced application - Amazon EC2 Auto Scaling
- ? 3 Amazon EC2 Auto Scaling FAQs

NEW QUESTION 69

- (Topic 4)

A company has a production workload that is spread across different AWS accounts in various AWS Regions. The company uses AWS Cost Explorer to continuously monitor costs and usage. The company wants to receive notifications when the cost and usage spending of the workload is unusual. Which combination of steps will meet these requirements? (Select TWO.)

- A. In the AWS accounts where the production workload is running, create a linked account budget by using Cost Explorer in the AWS Cost Management console
- B. In ys AWS accounts where the production workload is running, create a linked account monitor by using AWS Cost Anomaly Detection in the AWS Cost Management console
- C. In the AWS accounts where the production workload is running, create a Cost and Usage Report by using Cost Anomaly Detection in the AWS Cost Management console.
- D. Create a report and send email messages to notify the company on a weekly basis.
- E. Create a subscription with the required threshold and notify the company by using weekly summaries.

Answer: BE

Explanation:

AWS Cost Anomaly Detection allows you to create monitors that track the cost and usage of your AWS resources and alert you when there is an unusual spending pattern. You can create monitors based on different dimensions, such as AWS services, accounts, tags, or cost categories. You can also create alert subscriptions that notify you by email or Amazon SNS when an anomaly is detected. You can specify the threshold and frequency of the alerts, and choose to receive weekly summaries of your anomalies. Reference URLs:

- 1 <https://aws.amazon.com/aws-cost-management/aws-cost-anomaly-detection/>
- 2 <https://docs.aws.amazon.com/cost-management/latest/userguide/getting-started-ad.html>
- 3 <https://docs.aws.amazon.com/cost-management/latest/userguide/manage-ad.html>

NEW QUESTION 72

- (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 74

- (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 hierarchie
- B. Unload the data to Amazon S3 every month.
- C. Use Amazon DynamoDB to store the employee data in hierarchie
- D. Export the data to Amazon S3 every month.
- E. Configure Amazon fvlacie for the AWS accoun
- 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 75

- (Topic 4)

A company runs a web application on Amazon EC2 instances in an Auto Scaling group that has a target group. The company desgned the application to work with session affinity (sticky sessions) for a better user experience.

The application must be available publicly over the internet as an endpoint_ A WAF must be applied to the endpoint for additional security. Session affinity (sticky sessions) must be configured on the endpoint

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

- A. Create a public Network Load Balancer Specify the application target group.
- B. Create a Gateway Load Balancer Specify the application target group.
- C. Create a public Application Load Balancer Specify the application target group.
- D. Create a second target grou
- E. Add Elastic IP addresses to the EC2 instances
- F. Create a web ACL in AWS WAF Associate the web ACL with the endpoint

Answer: CE

Explanation:

C and E are the correct answers because they allow the company to create a public endpoint for its web application that supports session affinity (sticky sessions) and has a WAF applied for additional security. By creating a public Application Load Balancer, the company can distribute incoming traffic across multiple EC2 instances in an Auto Scaling group and specify the application target group. By creating a web ACL in AWS WAF and associating it with the Application Load Balancer, the company can protect its web application from common web exploits. By enabling session stickiness on the Application Load Balancer, the company can ensure that subsequent requests from a user during a session are routed to the same target. References:

- ? Application Load Balancers
- ? AWS WAF
- ? Target Groups for Your Application Load Balancers
- ? How Application Load Balancer Works with Sticky Sessions

NEW QUESTION 78

- (Topic 4)

A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming data. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting.

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

- A. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.
- B. Use AWS Glue to deliver streaming data to Amazon S3.
- C. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.
- D. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.

Answer: A

Explanation:

To ingest and process high-volume streaming data with the least operational overhead, Amazon Kinesis Data Firehose is a suitable solution. Amazon Kinesis Data Firehose can capture, transform, and deliver streaming data to Amazon S3 or other destinations. Amazon Kinesis Data Firehose can scale automatically to match the throughput of the data and handle any amount of data. Amazon Kinesis Data Firehose is also a fully managed service that does not require any servers to provision or manage. References:

- ? [What Is Amazon Kinesis Data Firehose?](#)
- ? [Amazon Kinesis Data Firehose Pricing](#)

NEW QUESTION 83

- (Topic 4)

A company needs to store data from its healthcare application. The application's data frequently changes. A new regulation requires audit z access at all levels of the stored data.

The company hosts the application on an on-premises infrastructure that is running out of storage capacity. A solutions architect must securely migrate the existing data to AWS while satisfying the new regulation.

Which solution will meet these requirements?

- A. Use AWS DataSync to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- B. Use AWS Snowcone to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.
- C. Use Amazon S3 Transfer Acceleration to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- D. Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.

Answer: A

Explanation:

This answer is correct because it meets the requirements of securely migrating the existing data to AWS and satisfying the new regulation. AWS DataSync is a service that makes it easy to move large amounts of data online between on-premises storage and Amazon S3. DataSync automatically encrypts data in transit and verifies data integrity during transfer. AWS CloudTrail is a service that records AWS API calls for your account and delivers log files to Amazon S3. CloudTrail can log data events, which show the resource operations performed on or within a resource in your AWS account, such as S3 object-level API activity. By using CloudTrail to log data events, you can audit access at all levels of the stored data.

References:

- ? <https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>
- ? <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/logging-data-events-with-cloudtrail.html>

NEW QUESTION 86

- (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 87

- (Topic 4)

A company has a large workload that runs every Friday evening. The workload runs on Amazon EC2 instances that are in two Availability Zones in the us-east-1 Region. Normally, the company must run no more than two instances at all times. However, the company wants to scale up to six instances each Friday to handle a regularly repeating increased workload.

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

- A. Create a reminder in Amazon EventBridge to scale the instances.
- B. Create an Auto Scaling group that has a scheduled action.
- C. Create an Auto Scaling group that uses manual scaling.
- D. Create an Auto Scaling group that uses automatic scaling.

Answer: B

Explanation:

An Auto Scaling group is a collection of EC2 instances that share similar characteristics and can be scaled in or out automatically based on demand. An Auto Scaling group can have a scheduled action, which is a configuration that tells the group to scale to a specific size at a specific time. This way, the company can scale up to six instances each Friday evening to handle the increased workload, and scale down to two instances at other times to save costs. This solution meets the requirements with the least operational overhead, as it does not require manual intervention or custom scripts. References:

- ? 1 explains how to create a scheduled action for an Auto Scaling group.
- ? 2 describes the concept and benefits of an Auto Scaling group.

NEW QUESTION 91

- (Topic 4)

A company is designing a new web application that will run on Amazon EC2 Instances. The application will use Amazon DynamoDB for backend data storage. The application traffic will be unpredictable. The company expects that the application read and write throughput to the database will be moderate to high. The company needs to scale in response to application traffic.

Which DynamoDB table configuration will meet these requirements MOST cost-effectively?

- A. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard table class
- B. Set DynamoDB auto scaling to a maximum defined capacity.
- C. Configure DynamoDB in on-demand mode by using the DynamoDB Standard table class.
- D. Configure DynamoDB with provisioned read and write by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class
- E. Set DynamoDB auto scaling to a maximum defined capacity.
- F. Configure DynamoDB in on-demand mode by using the DynamoDB Standard Infrequent Access (DynamoDB Standard-IA) table class.

Answer: B

Explanation:

The most cost-effective DynamoDB table configuration for the web application is to configure DynamoDB in on-demand mode by using the DynamoDB Standard table class. This configuration will allow the company to scale in response to application traffic and pay only for the read and write requests that the application performs on the table.

On-demand mode is a flexible billing option that can handle thousands of requests per second without capacity planning. On-demand mode automatically adjusts the table's capacity based on the incoming traffic, and charges only for the read and write requests that are actually performed. On-demand mode is suitable for applications with unpredictable or variable workloads, or applications that prefer the ease of paying for only what they use¹.

The DynamoDB Standard table class is the default and recommended table class for most workloads. The DynamoDB Standard table class offers lower throughput costs than the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and is more cost-effective for tables where throughput is the dominant cost. The DynamoDB Standard table class also offers the same performance, durability, and availability as the DynamoDB Standard-IA table class². The other options are not correct because they are either not cost-effective or not suitable for the use case. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration requires manual estimation and management of the table's capacity, which adds complexity and cost to the solution. Provisioned mode is a billing option that requires users to specify the amount of read and write capacity units for their tables, and charges for the reserved capacity regardless of usage. Provisioned mode is suitable for applications with predictable or stable workloads, or applications that require finer-grained control over their capacity settings¹. Configuring DynamoDB with provisioned read and write by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class, and setting DynamoDB auto scaling to a maximum defined capacity is not correct because this configuration is not cost-effective for tables with moderate to high throughput. The DynamoDB Standard-IA table class offers lower storage costs than the DynamoDB Standard table class, but higher throughput costs. The DynamoDB Standard-IA table class is optimized for tables where storage is the dominant cost, such as tables that store infrequently accessed data². Configuring DynamoDB in on-demand mode by using the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class is not correct because this configuration is not cost-effective for tables with moderate to high throughput. As mentioned above, the DynamoDB Standard-IA table class has higher throughput costs than the DynamoDB Standard table class, which can offset the savings from lower storage costs.

References:

- ? Table classes - Amazon DynamoDB
- ? Read/write capacity mode - Amazon DynamoDB

NEW QUESTION 95

- (Topic 4)

The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After successful scheduling, this application stores the meeting information in an Amazon DynamoDB database.

As the company expands, customers report that their meeting invitations are taking longer to arrive.

What should a solutions architect recommend to resolve this issue?

- A. Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.
- B. Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.
- C. Add an Amazon CloudFront distributio
- D. Set the origin as the web application that accepts the appointment requests.
- E. Add an Auto Scaling group for the application that sends meeting invitation
- F. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Answer: D

Explanation:

To resolve the issue of longer delivery times for meeting invitations, the solutions architect can recommend adding an Auto Scaling group for the application that sends meeting invitations and configuring the Auto Scaling group to scale based on the depth of the SQS queue. This will allow the application to scale up as the number of appointment requests increases, improving the performance and delivery times of the meeting invitations.

NEW QUESTION 99

- (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 100

- (Topic 4)

An IoT company is releasing a mattress that has sensors to collect data about a user's sleep. The sensors will send data to an Amazon S3 bucket. The sensors collect approximately 2 MB of data every night for each mattress. The company must process and summarize the data for each mattress. The results need to be available as soon as possible. Data processing will require 1 GB of memory and will finish within 30 seconds.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Glue with a Scalajob.
- B. Use Amazon EMR with an Apache Spark script.
- C. Use AWS Lambda with a Python script.
- D. Use AWS Glue with a PySpark job.

Answer: C

Explanation:

AWS Lambda charges you based on the number of invocations and the execution time of your function. Since the data processing job is relatively small (2 MB of data), Lambda is a cost-effective choice. You only pay for the actual usage without the need to provision and maintain infrastructure.

NEW QUESTION 101

- (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 103

- (Topic 4)

A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month.

What is the MOST cost-effective solution to connect these VPCs?

- A. Implement AWS Transit Gateway to connect the VPC
- B. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
- C. Implement an AWS Site-to-Site VPN tunnel between the VPC
- D. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
- E. Set up a VPC peering connection between the VPC
- F. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
- G. Set up a 1 GB AWS Direct Connect connection between the VPC
- H. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Answer: C

Explanation:

To connect two VPCs in the same Region within the same AWS account, VPC peering is the most cost-effective solution. VPC peering allows direct network traffic between the VPCs without requiring a gateway, VPN connection, or AWS Transit Gateway. VPC peering also does not incur any additional charges for data transfer between the VPCs.

References:

? What Is VPC Peering?

? VPC Peering Pricing

NEW QUESTION 105

- (Topic 4)

A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.

However, the company's security policy states that any external service cannot initiate a connection to the EC2 instances.

What should a solutions architect recommend to resolve this issue?

- A. Create a NAT gateway and make it the destination of the subnet's route table.
- B. Create an internet gateway and make it the destination of the subnet's route table
- C. Create a virtual private gateway and make it the destination of the subnet's route table.
- D. Create an egress-only internet gateway and make it the destination of the subnet's route table.

Answer: D

Explanation:

An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in your VPC to the internet, and prevents the internet from initiating an IPv6 connection with your instances. This meets the company's security policy and requirements. To use an egress-only internet gateway, you need to add a route in the subnet's route table that routes IPv6 internet traffic (::/0) to the egress-only internet gateway.

Reference URLs:

1 <https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

2 <https://dev.to/aws-builders/what-is-an-egress-only-internet-gateways-in-aws-7gp>

3 <https://docs.aws.amazon.com/vpc/latest/userguide/route-table-options.html>

NEW QUESTION 108

- (Topic 4)

A company has a nightly batch processing routine that analyzes report files that an on-premises file system receives daily through SFTP. The company wants to move the solution to the AWS Cloud. The solution must be highly available and resilient. The solution also must minimize operational effort.

Which solution meets these requirements?

- A. Deploy AWS Transfer for SFTP and an Amazon Elastic File System (Amazon EFS) file system for storage
- B. Use an Amazon EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.
- C. Deploy an Amazon EC2 instance that runs Linux and an SFTP service
- D. Use an Amazon Elastic Block Store (Amazon EBS) volume for storage
- E. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
- F. Deploy an Amazon EC2 instance that runs Linux and an SFTP service
- G. Use an Amazon Elastic File System (Amazon EFS) file system for storage
- H. Use an Auto Scaling group with the minimum number of instances and desired number of instances set to 1.
- I. Deploy AWS Transfer for SFTP and an Amazon S3 bucket for storage
- J. Modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processing
- K. Use an EC2 instance in an Auto Scaling group with a scheduled scaling policy to run the batch operation.

Answer: D

Explanation:

The solution that meets the requirements of high availability, performance, security, and static IP addresses is to use Amazon CloudFront, Application Load Balancers (ALBs), Amazon Route 53, and AWS WAF. This solution allows the company to distribute its HTTP-based application globally using CloudFront, which is a content delivery network (CDN) service that caches content at edge locations and provides static IP addresses for each edge location. The company can also use Route 53 latency-based routing to route requests to the closest ALB in each Region, which balances the load across the EC2 instances. The company can also deploy AWS WAF on the CloudFront distribution to protect the application against common web exploits by creating rules that allow, block, or count web requests based on conditions that are defined. The other solutions do not meet all the requirements because they either use Network Load Balancers (NLBs), which do not support HTTP-based applications, or they do not use CloudFront, which provides better performance and security than AWS Global Accelerator.

References :=

? Amazon CloudFront

? Application Load Balancer

? Amazon Route 53

? AWS WAF

NEW QUESTION 109

- (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 112

- (Topic 4)

A company has a serverless website with millions of objects in an Amazon S3 bucket. The company uses the S3 bucket as the origin for an Amazon CloudFront distribution. The company did not set encryption on the S3 bucket before the objects were loaded. A solutions architect needs to enable encryption for all existing objects and for all objects that are added to the S3 bucket in the future.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Create a new S3 bucket
- B. Turn on the default encryption settings for the new S3 bucket
- C. Download all existing objects to temporary local storage
- D. Upload the objects to the new S3 bucket.
- E. Turn on the default encryption settings for the S3 bucket
- F. Use the S3 Inventory feature to create a .csv file that lists the unencrypted objects
- G. Run an S3 Batch Operations job that uses the copy command to encrypt those objects.
- H. Create a new encryption key by using AWS Key Management Service (AWS KMS). Change the settings on the S3 bucket to use server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Turn on versioning for the S3 bucket.
- I. Navigate to Amazon S3 in the AWS Management Console
- J. Browse the S3 bucket's objects
- K. Sort by the encryption field
- L. Select each unencrypted object
- M. Use the Modify button to apply default encryption settings to every unencrypted object in the S3 bucket.

Answer: B

Explanation:

<https://spin.atomicobject.com/2020/09/15/aws-s3-encrypt-existing-objects/>

NEW QUESTION 113

- (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 115

- (Topic 4)

A company uses multiple vendors to distribute digital assets that are stored in Amazon S3 buckets The company wants to ensure that its vendor AWS accounts have the minimum access that is needed to download objects in these S3 buckets Which solution will meet these requirements with the LEAST operational overhead?

- A. Design a bucket policy that has anonymous read permissions and permissions to list all buckets.
- B. Design a bucket policy that gives read-only access to user
- C. Specify IAM entities as principals
- D. Create a cross-account IAM role that has a read-only access policy specified for the IAM role.
- E. Create a user policy and vendor user groups that give read-only access to vendor users

Answer: C

Explanation:

A cross-account IAM role is a way to grant users from one AWS account access to resources in another AWS account. The cross-account IAM role can have a read-only access policy attached to it, which allows the users to download objects from the S3 buckets without modifying or deleting them. The cross-account IAM role also reduces the operational overhead of managing multiple IAM users and policies in each account. The cross-account IAM role meets all the requirements of the question, while the other options do not. References:

- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-walkthroughs-managing-access-example2.html>
- ? <https://aws.amazon.com/blogs/storage/setting-up-cross-account-amazon-s3-access-with-s3-access-points/>
- ? https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html

NEW QUESTION 116

- (Topic 4)

A company wants to run its experimental workloads in the AWS Cloud. The company has a budget for cloud spending. The company's CFO is concerned about cloud spending accountability for each department. The CFO wants to receive notification when the spending threshold reaches 60% of the budget. Which solution will meet these requirements?

- A. Use cost allocation tags on AWS resources to label owner
- B. Create usage budgets in AWS Budget
- C. Add an alert threshold to receive notification when spending exceeds 60% of the budget.
- D. Use AWS Cost Explorer forecasts to determine resource owner
- E. Use AWS Cost Anomaly Detection to create alert threshold notifications when spending exceeds 60% of the budget.
- F. Use cost allocation tags on AWS resources to label owner
- G. Use AWS Support API on AWS Trusted Advisor to create alert threshold notifications when spending exceeds 60% of the budget
- H. Use AWS Cost Explorer forecasts to determine resource owner
- I. Create usage budgets in AWS Budget
- J. Add an alert threshold to receive notification when spending exceeds 60% of the budget.

Answer: A

Explanation:

This solution meets the requirements because it allows the company to track and manage its cloud spending by using cost allocation tags to assign costs to different departments, creating usage budgets to set spending limits, and adding alert thresholds to receive notifications when the spending reaches a certain percentage of the budget. This way, the company can monitor its experimental workloads and avoid overspending on the cloud.

References:

- ? Using Cost Allocation Tags
- ? Creating an AWS Budget
- ? Creating an Alert for an AWS Budget

NEW QUESTION 121

.....

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](#)