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) volum
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 imag
F. Create a service with that task definition.

Answer: C

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

NEW QUESTION 2
- (Topic 4)
A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.
Which migration solution will meet these requirements?

A. Use native MySQL tools to migrate the database to Amazon RDS for MySQ
B. Configureelastic storage scaling.
C. Migrate the database to Amazon Redshift by using the mysqldump utilit
D. Turn on Auto Scaling for the Amazon Redshift cluster.
E. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Auror
F. Turn on Aurora Auto Scaling.
G. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon DynamoD
H. Configure an Auto Scaling policy.

Answer: C

Explanation:
To migrate a MySQL database to AWS with compatibility and scalability, Amazon Aurora is a suitable option. Aurora is compatible with MySQL and can scale automatically with Aurora Auto Scaling. AWS Database Migration Service (AWS DMS) can be used to migrate the database from on-premises to Aurora with minimal downtime. References:
? What Is Amazon Aurora?
? Using Amazon Aurora Auto Scaling with Aurora Replicas
? What Is AWS Database Migration Service?

NEW QUESTION 3
- (Topic 4)
A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours.
Which solution meets these requirements with the LEAST operational overhead?

A. Configure point-in-time recovery for the table.
B. Use AWS Backup for the table.
C. Use an AWS Lambda function to make an on-demand backup of the table every hour.
D. Turn on streams on the table to capture a log of all changes to the table in the last 24 hours Store a copy of the stream in an Amazon S3 bucket.

Answer: A

Explanation:
Point-in-time recovery (PITR) for DynamoDB is a feature that enables you to restore your table data to any point in time during the last 35 days. PITR helps protect your table from accidental write or delete operations, such as a test script writing to a production table or a user issuing a wrong command. PITR is easy to use, fully managed, fast, and scalable. You can enable PITR with a single click in the DynamoDB console or with a simple API call. You can restore a table to a new table using the console, the AWS CLI, or the DynamoDB API. PITR does not consume any provisioned table capacity and has no impact on the performance or availability of your production applications. PITR meets the requirements of the company with the least operational overhead, as it does not require any manual backup creation, scheduling, or maintenance. It also provides per-second granularity for restoring the table to any point within the last 24 hours.
References:
? Point-in-time recovery for DynamoDB - Amazon DynamoDB
? Amazon DynamoDB point-in-time recovery (PITR)
? Enable Point-in-Time Recovery (PITR) for Dynamodb global tables

? Restoring a DynamoDB table to a point in time - Amazon DynamoDB
? Point-in-time recovery: How it works - Amazon DynamoDB

NEW QUESTION 4
- (Topic 4)
A company migrated a MySQL database from the company's on-premises data center to an Amazon RDS for MySQL DB instance. The company sized the RDS DB instance to meet the company's average daily workload. Once a month, the database performs slowly when the company runs queries for a report. The company wants to have the ability to run reports and maintain the performance of the daily workloads.
Which solution will meet these requirements?

A. Create a read replica of the databas
B. Direct the queries to the read replica.
C. Create a backup of the databas
D. Restore the backup to another DB instanc
E. Direct the queries to the new database.
F. Export the data to Amazon S3. Use Amazon Athena to query the S3 bucket.
G. Resize the DB instance to accommodate the additional workload.

Answer: C

Explanation:
Amazon Athena is a service that allows you to run SQL queries on data stored in Amazon S3. It is serverless, meaning you do not need to provision or manage any infrastructure. You only pay for the queries you run and the amount of data scanned1.
By using Amazon Athena to query your data in Amazon S3, you can achieve the following benefits:
? You can run queries for your report without affecting the performance of your
Amazon RDS for MySQL DB instance. You can export your data from your DB instance to an S3 bucket and use Athena to query the data in the bucket. This way, you can avoid the overhead and contention of running queries on your DB instance.
? You can reduce the cost and complexity of running queries for your report. You do
not need to create a read replica or a backup of your DB instance, which would incur additional charges and require maintenance. You also do not need to resize your DB instance to accommodate the additional workload, which would increase your operational overhead.
? You can leverage the scalability and flexibility of Amazon S3 and Athena. You can
store large amounts of data in S3 and query them with Athena without worrying about capacity or performance limitations. You can also use different formats, compression methods, and partitioning schemes to optimize your data storage and query performance1.

NEW QUESTION 5
- (Topic 4)
A solutions architect wants to use the following JSON text as an identity-based policy to grant specific permissions:

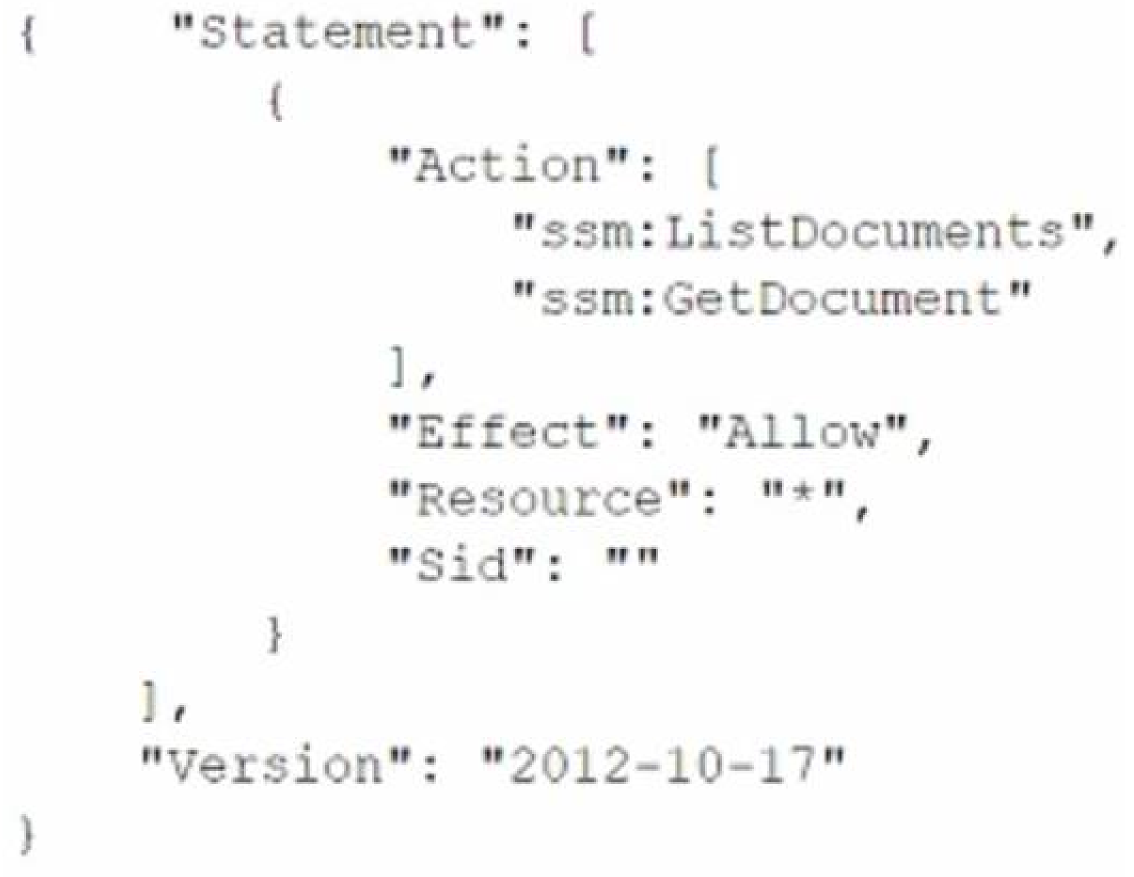

```
{     "Statement": [
          {
               "Action": [
                    "ssm:ListDocuments",
                    "ssm:GetDocument"
               ],
               "Effect": "Allow",
               "Resource": "*",
               "Sid": ""
          }
     ],
     "Version": "2012-10-17"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

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 them34. References:
? Identity-based policies and resource-based policies
? AWS Organizations
? Amazon ECS task role
? Amazon EC2 instance profile

NEW QUESTION 6
- (Topic 4)
A company needs to configure a real-time data ingestion architecture for its application. The company needs an API. a process that transforms data as the data is streamed, and a storage solution for the data.
Which solution will meet these requirements with the LEAST operational overhead?

A. Deploy an Amazon EC2 instance to host an API that sends data to an Amazon Kinesis data strea
B. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
C. Use AWS Lambda functions to transform the dat
D. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
E. Deploy an Amazon EC2 instance to host an API that sends data to AWS Glu
F. Stop source/destination checking on the EC2 instanc
G. Use AWS Glue to transform the data and to send the data to Amazon S3.
H. Configure an Amazon API Gateway API to send data to an Amazon Kinesis datastrea
I. Create an Amazon Kinesis Data Firehose delivery stream that uses the Kinesis data stream as a data sourc
J. Use AWS Lambda functions to transform the dat
K. Use the Kinesis Data Firehose delivery stream to send the data to Amazon S3.
L. Configure an Amazon API Gateway API to send data to AWS Glu
M. Use AWS Lambda functions to transform the dat
N. Use AWS Glue to send the data to Amazon S3.

Answer: C

Explanation:
It uses Amazon Kinesis Data Firehose which is a fully managed service for delivering real-time streaming data to destinations such as Amazon S3. This service requires less operational overhead as compared to option A, B, and D. Additionally, it also uses Amazon API Gateway which is a fully managed service for creating, deploying, and managing APIs. These services help in reducing the operational overhead and automating the data ingestion process.

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 Kubemetes 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 tie1r. 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-IA). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 replication is a feature that copies objects across buckets or Regions for redundancy or compliance purposes. It does not automatically move objects to a different storage class based on access frequency2.

* B. Use S3 Lifecycle rules to transition objects from S3 Standard to Standard-Infrequent Access (S3 Standard-IA). This solution will not meet the requirement of reducing S3 costs for data with unknown or changing access patterns, as S3 Standard-IA is a storage class that offers lower storage costs than S3 Standard, but charges a retrieval fee for accessing the data. It is suitable for long-lived and infrequently accessed data, not for data with changing access patterns1.
* D. Use S3 Inventory to identify and transition objects that have not been accessed from S3 Stand-ard 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 devic
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 vehicl
E. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database wjgh 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 devic
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 cente
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_So urce.MySQL.htm

NEW QUESTION 10
- (Topic 4)
A solutions architect is designing a highly available Amazon ElastiCache for Redis based solution. The solutions architect needs to ensure that failures do not result in performance degradation or loss of data locally and within an AWS Region. The solution needs to provide high availability at the node level and at the Region level.
Which solution will meet these requirements?

A. Use Multi-AZ Redis replication groups with shards that contain multiple nodes.
B. Use Redis shards that contain multiple nodes with Redis append only files (AOF) tured on.
C. Use a Multi-AZ Redis cluster with more than one read replica in the replication group.
D. Use Redis shards that contain multiple nodes with Auto Scaling turned on.

Answer: A

Explanation:
This answer is correct because it provides high availability at the node level and at the Region level for the ElastiCache for Redis solution. A Multi-AZ Redis replication group consists of a primary cluster and up to five read replica clusters, each in a different Availability Zone. If the primary cluster fails, one of the read replicas is automatically promoted to be the new primary cluster. A Redis replication group with shards enables partitioning of the data across multiple nodes, which increases the scalability and performance of the solution. Each shard can have one or more replicas to provide redundancy and read scaling.
References:
? https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html
? https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Shards.html

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 15
- (Topic 4)
A company maintains an Amazon RDS database that maps users to cost centers. The company has accounts in an organization in AWS Organizations. The company needs a solution that will tag all resources that are created in a specific AWS account in the organization. The solution must tag each resource with the cost center ID of the user who created the resource.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 20
- (Topic 4)
A company has an organization in AWS Organizations that has all features enabled The company requires that all API calls and logins in any existing or new AWS account must be audited The company needs a managed solution to prevent additional work and to minimize costs The company also needs to know when any AWS account is not compliant with the AWS Foundational Security Best Practices (FSBP) standard.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

Explanation:
AWS Control Tower is a fully managed service that simplifies the setup and governance of a secure, compliant, multi-account AWS environment. It establishes a

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

NEW QUESTION 25
- (Topic 4)
A company needs to use its on-premises LDAP directory service to authenticate its users to the AWS Management Console. The directory service is not compatible with Security Assertion Markup Language (SAML).
Which solution meets these requirements?

A. Enable AWS 1AM Identity Center (AWS Single Sign-On) between AWS and the on- premises LDAP.
B. Create an 1AM policy that uses AWS credentials, and integrate the policy into LDAP.
C. Set up a process that rotates the I AM credentials whenever LDAP credentials are updated.
D. Develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials.

Answer: D

Explanation:
The solution that meets the requirements is to develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials. This solution allows the company to use its existing LDAP directory service to authenticate its users to the AWS Management Console, without requiring SAML compatibility. The custom identity broker application or process can act as a proxy between the LDAP directory service and AWS STS, and can request temporary security credentials for the users based on their LDAP attributes and roles. The users can then use these credentials to access the AWS Management Console via a sign-in URL generated by the identity broker. This solution also enhances security by using short-lived credentials that expire after a specified duration.
The other solutions do not meet the requirements because they either require SAML compatibility or do not provide access to the AWS Management Console. Enabling AWS IAM Identity Center (AWS Single Sign-On) between AWS and the on-premises LDAP would require the LDAP directory service to support SAML 2.0, which is not the case for this scenario. Creating an IAM policy that uses AWS credentials and integrating the policy into LDAP would not provide access to the AWS Management Console, but only to the AWS APIs. Setting up a process that rotates the IAM credentials whenever LDAP credentials are updated would also not provide access to the AWS Management Console, but only to the AWS CLI. Therefore, these solutions are not suitable for the given requirements.

NEW QUESTION 28
- (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 AWS?4, answer section.

NEW QUESTION 31
- (Topic 4)
A company offers a food delivery service that is growing rapidly. Because of the growth, the company’s order processing system is experiencing scaling problems during peak traffic hours. The current architecture includes the following:
• A group of Amazon EC2 instances that run in an Amazon EC2 Auto Scaling group to collect orders from the application
• Another group of EC2 instances that run in an Amazon EC2 Auto Scaling group to fulfill orders
The order collection process occurs quickly, but the order fulfillment process can take longer. Data must not be lost because of a scaling event.
A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize
utilization of the company’s AWS resources. Which solution meets these requirements?

A. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
B. Configure each Auto Scaling group’s minimum capacity according to peak workload values.
C. Use Amazon CloudWatch metrics to monitor the CPU of each instance in the Auto Scaling group
D. Configure a CloudWatch alarm to invoke an Amazon Simple Notification Service (Amazon SNS) topic that creates additional Auto Scaling groups on demand.
E. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillmen

F. Configure the EC2 instances to poll their respective queu
G. Scale the Auto Scaling groups based on notifications that the queues send.
H. Provision two Amazon Simple Queue Service (Amazon SQS) queues: one for order collection and another for order fulfillmen
I. Configure the EC2 instances to poll their respective queu
J. Create a metric based on a backlog per instance calculatio
K. Scale the Auto Scaling groups based on this metric.

Answer: D

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

NEW QUESTION 35
- (Topic 4)
A manufacturing company runs its report generation application on AWS. The application generates each report in about 20 minutes. The application is built as a monolith that runs on a single Amazon EC2 instance. The application requires frequent updates to its tightly coupled modules. The application becomes complex to maintain as the company adds new features.
Each time the company patches a software module, the application experiences downtime. Report generation must restart from the beginning after any interruptions. The company wants to redesign the application so that the application can be flexible, scalable, and gradually improved. The company wants to minimize application downtime.
Which solution will meet these requirements?

A. Run the application on AWS Lambda as a single function with maximum provisioned concurrency.
B. Run the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy.
C. Run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling.
D. Run the application on AWS Elastic Beanstalk as a single application environment with an all-at-once deployment strategy.

Answer: C

Explanation:
The solution that will meet the requirements is to run the application on Amazon Elastic Container Service (Amazon ECS) as microservices with service auto scaling. This solution will allow the application to be flexible, scalable, and gradually improved, as well as minimize application downtime. By breaking down the monolithic application into microservices, the company can decouple the modules and update them independently, without affecting the whole application. By running the microservices on Amazon ECS, the company can leverage the benefits of containerization, such as portability, efficiency, and isolation. By enabling service auto scaling, the company can adjust the number of containers running for each microservice based on demand, ensuring optimal performance and cost. Amazon ECS also supports various deployment strategies, such as rolling update or blue/green deployment, that can reduce or eliminate downtime during updates.
The other solutions are not as effective as the first one because they either do not meet the requirements or introduce new challenges. Running the application on AWS Lambda as a single function with maximum provisioned concurrency will not meet the requirements, as it will not break down the monolith into microservices, nor will it reduce the complexity of maintenance. Lambda functions are also limited by execution time (15 minutes), memory size (10 GB), and concurrency quotas, which may not be sufficient for the report generation application. Running the application on Amazon EC2 Spot Instances as microservices with a Spot Fleet default allocation strategy will not meet the requirements, as it will introduce the risk of interruptions due to spot price fluctuations. Spot Instances are not guaranteed to be available or stable, and may be reclaimed by AWS at any time with a two-minute warning. This may cause report generation to fail or restart from scratch. Running the application on AWS Elastic Beanstalk as a single application environment with an all-at- once deployment strategy will not meet the requirements, as it will not break down the monolith into microservices, nor will it minimize application downtime. The all-at-once deployment strategy will deploy updates to all instances simultaneously, causing a brief outage for the application.
References:
? Amazon Elastic Container Service
? Microservices on AWS
? Service Auto Scaling - Amazon Elastic Container Service
? AWS Lambda
? Amazon EC2 Spot Instances
? [AWS Elastic Beanstalk]

NEW QUESTION 38
- (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 o
B. Create the new AWS accounts in the organization.
C. Set up an Amazon Cognito identity poo
D. Configure AWS 1AM 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 1AM 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 1AM Identity Center (AWS Single Sign-On) in the organizatio
J. Configure 1AM 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 1AM 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 1AM Identi-ty 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 43
- (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) volum
B. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresse
C. Attach the EBS volume to the SFTP service endpoin
D. Grant users access to the SFTP service.
E. Create an encrypted Amazon Elastic File System (Amazon EFS) volum
F. Create an AWS Transfer Family SFTP service with elastic IP addresses and a VPC endpoint that has internet-facing acces
G. Attach a security group to the endpoint that allows only trusted IP addresse
H. Attach the EFS volume to the SFTP service endpoin
I. Grant users access to the SFTP service.
J. Create an Amazon S3 bucket with default encryption enable
K. Create an AWS Transfer Family SFTP service with a public endpoint that allows only trusted IP addresse
L. Attach the S3 bucket to the SFTP service endpoin
M. Grant users access to the SFTP service.
N. Create an Amazon S3 bucket with default encryption enable
O. Create an AWS Transfer Family SFTP service with a VPC endpoint that has internal access in a private subne
P. Attach a security group that allows only trusted IP addresse
Q. Attach the S3 bucket to the SFTP service endpoin
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 48
- (Topic 4)
A company uses Amazon EC2 instances to host its internal systems. As part of a deployment operation, an administrator tries to use the AWS CLI to terminate an EC2 instance. However, the administrator receives a 403 (Access Denied) error message.
The administrator is using an IAM role that has the following IAM policy attached:

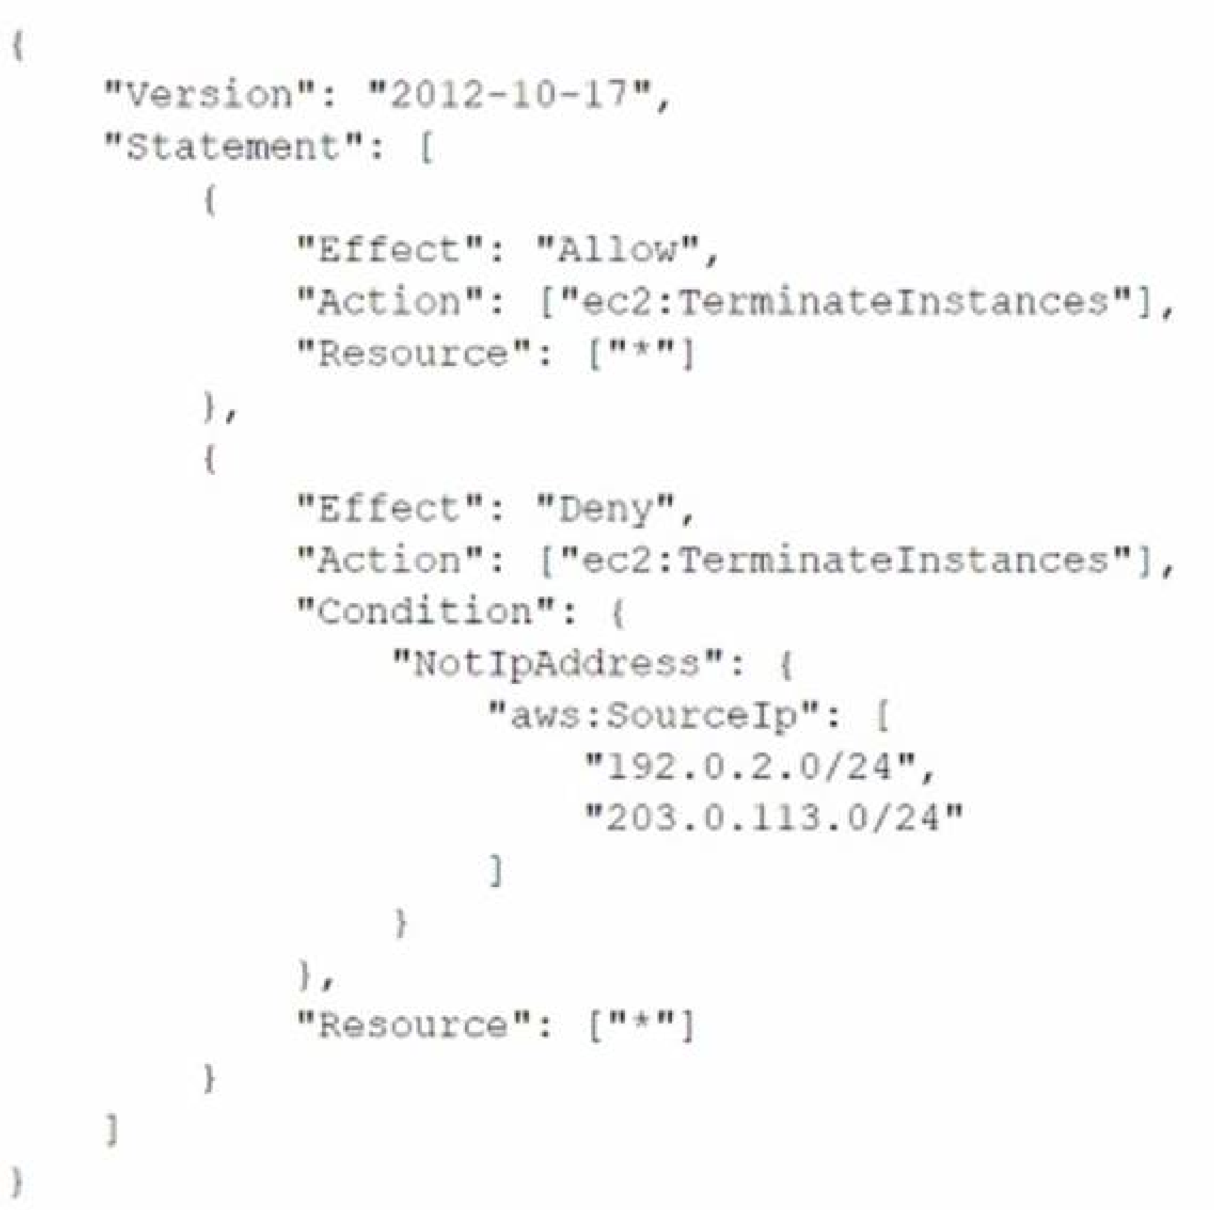

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": ["ec2:TerminateInstances"],
            "Resource": ["*"]
        },
        {
            "Effect": "Deny",
            "Action": ["ec2:TerminateInstances"],
            "Condition":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

What is the cause of the unsuccessful request?

A. The EC2 instance has a resource-based policy with a Deny statement.
B. The principal has not been specified in the policy statement
C. The "Action" field does not grant the actions that are required to terminate the EC2 instance.
D. The request to terminate the EC2 instance does not originate from the CIDR blocks 192.0.2.0/24 or 203.0 113.0/24

Answer: D

NEW QUESTION 53
- (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 creatio
C. Assign the RDS DB instances to the backup plan.
D. Configure a backup window for the RDS DB instances for daily snapshot
E. Assign a snapshot retention policy of 2 years to each RDS DB instanc
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 tas
I. Deploy a replication instance, and configure a change data capture (CDC) task to stream database changes to Amazon S3 as the targe
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 data1. 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 deletions2.
* 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 database3.
* 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/whatisbackup.html

NEW QUESTION 57
- (Topic 4)

A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent accidental deletion of the documents and ensure that all versions of the documents are available. Users must be able to download, modify, and upload documents.
Which combination of actions should be taken to meet these requirements? (Choose two.)

A. Enable a read-only bucket ACL.
B. Enable versioning on the bucket.
C. Attach an IAM policy to the bucket.
D. Enable MFA Delete on the bucket.
E. Encrypt the bucket using AWS KMS.

Answer: BD

Explanation:
Versioning is a feature of Amazon S3 that allows users to keep multiple versions of the same object in a bucket. It can help prevent accidental deletion of the documents and ensure that all versions of the documents are available1. MFA Delete is a feature of Amazon S3 that adds an extra layer of security by requiring two forms of authentication to delete a version or change the versioning state of a bucket. It can help prevent unauthorized or accidental deletion of the documents2. By enabling both versioning and MFA Delete on the bucket, the solution can meet the requirements.
* A. Enable a read-only bucket ACL. This solution will not meet the requirement of allowing
users to download, modify, and upload documents, as a read-only bucket ACL will prevent write access to the bucket3.
* C. Attach an IAM policy to the bucket. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as an IAM policy is used to grant or deny permissions to users or roles, not to enable versioning or MFA Delete4.
* E. Encrypt the bucket using AWS KMS. This solution will not meet the requirement of preventing accidental deletion of the documents and ensuring that all versions of the documents are available, as encrypting the bucket using AWS KMS is a method of protecting data at rest, not enabling versioning or MFA Delete.
Reference URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html

NEW QUESTION 58
- (Topic 4)
A company runs an SMB file server in its data center. The file server stores large files that the company frequently accesses for up to 7 days after the file creation date. After 7 days, the company needs to be able to access the files with a maximum retrieval time of 24 hours.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 63
- (Topic 4)
A company wants to migrate an on-premises legacy application to AWS. The application ingests customer order files from an on-premises enterprise resource planning (ERP) system. The application then uploads the files to an SFTP server. The application uses a scheduled job that checks for order files every hour.
The company already has an AWS account that has connectivity to the on-premises network. The new application on AWS must support integration with the existing ERP system. The new application must be secure and resilient and must use the SFTP protocol to process orders from the ERP system immediately.
Which solution will meet these requirements?

A. Create an AWS Transfer Family SFTP internet-facing server in two Availability Zone
B. Use Amazon S3 storag
C. Create an AWS Lambda function to process order file
D. Use S3 Event Notifications to send s3: ObjectCreated: * events to the Lambda function.
E. Create an AWS Transfer Family SFTP internet-facing server in one Availability Zon
F. Use Amazon Elastic File System (Amazon EFS) storag
G. Create an AWS Lambda function to process order file
H. Use a Transfer Family managed workflow to invoke the Lambda function.
I. Create an AWS Transfer Family SFTP internal server in two Availability Zone
J. Use Amazon Elastic File System (Amazon EFS) storag
K. Create an AWS Step Functions state machine to process order file
L. Use Amazon EventBridge Scheduler to invoke the state machine to periodically check Amazon EFS for order files.
M. Create an AWS Transfer Family SFTP internal server in two Availability Zone
N. Use Amazon S3 storag
O. Create an AWS Lambda function to process order file
P. Use a Transfer Family managed workflow to invoke the Lambda function.

Answer: D

Explanation:

This solution meets the requirements because it uses the following components and features:
? AWS Transfer Family SFTP internal server: This allows the application to securely
transfer order files from the on-premises ERP system to AWS using the SFTP protocol over a private connection. The internal server is deployed in two Availability Zones for high availability and fault tolerance.
? Amazon S3 storage: This provides scalable, durable, and cost-effective object
storage for the order files. Amazon S3 also supports encryption at rest and in transit, as well as lifecycle policies and versioning for data protection and compliance.
? AWS Lambda function: This enables the application to process the order files in a
serverless manner, without provisioning or managing servers. The Lambda function can perform any custom logic or transformation on the order files, such as validating, parsing, or enriching the data.
? Transfer Family managed workflow: This simplifies the orchestration of the file
processing tasks by triggering the Lambda function as soon as a file is uploaded to the SFTP server. The managed workflow also provides error handling, retry policies, and logging capabilities.

NEW QUESTION 65

- (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 67

- (Topic 4)
A company is building an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for its workloads. All secrets that are stored in Amazon EKS must be encrypted in the Kubernetes etcd key-value store.
Which solution will meet these requirements?

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

Answer: B

Explanation:

This option is the most secure and simple way to encrypt the secrets that are stored in Amazon EKS. AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys that can be used to encrypt your data. Amazon EKS KMS secrets encryption is a feature that enables you to use a KMS key to encrypt the secrets that are stored in the Kubernetes etcd key-value store. This provides an additional layer of protection for your sensitive data, such as passwords, tokens, and keys. You can create a new KMS key or use an existing one, and then enable the Amazon EKS KMS secrets encryption on the Amazon EKS cluster. You can also use IAM policies to control who can access or use the KMS key.
Option A is not correct because using AWS Secrets Manager to manage, rotate, and store all secrets in Amazon EKS is not necessary or efficient. AWS Secrets Manager is a service that helps you securely store, retrieve, and rotate your secrets, such as database credentials, API keys, and passwords. You can use it to manage secrets that are used by your applications or services outside of Amazon EKS, but it is not designed to encrypt the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using AWS Secrets Manager would incur additional costs and complexity, and it would not leverage the native Kubernetes secrets management capabilities.
Option C is not correct because using the Amazon EBS Container Storage Interface (CSI) driver as an add-on does not encrypt the secrets that are stored in Amazon EKS. The Amazon EBS CSI driver is a plugin that allows you to use Amazon EBS volumes as persistent storage for your Kubernetes pods. It is useful for providing durable and scalable storage for your applications, but it does not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the Amazon EBS CSI driver would require additional configuration and resources, and it would not provide the same level of security as using a KMS key.
Option D is not correct because creating a new AWS KMS key with the alias aws/ebs and enabling default Amazon EBS volume encryption for the account does not encrypt the secrets that are stored in Amazon EKS. The alias aws/ebs is a reserved alias that is used by AWS to create a default KMS key for your account. This key is used to encrypt the Amazon EBS volumes that are created in your account, unless you specify a different KMS key. Enabling default Amazon EBS volume encryption for the account is a setting that ensures that all new Amazon EBS volumes are encrypted by default. However, these features do not affect the encryption of the secrets that are stored in the Kubernetes etcd key-value store. Moreover, using the default KMS key or the default encryption setting would not provide the same level of control and security as using a custom KMS key and enabling the Amazon EKS KMS secrets encryption feature. References:
? Encrypting secrets used in Amazon EKS
? What Is AWS Key Management Service?
? What Is AWS Secrets Manager?
? Amazon EBS CSI driver
? Encryption at rest

NEW QUESTION 70

- (Topic 4)
A company used an Amazon RDS for MySQL DB instance during application testing. Before terminating the DB instance at the end of the test cycle, a solutions architect created two backups. The solutions architect created the first backup by using the mysqldump utility to create a database dump. The solutions architect

created the second backup by enabling the final DB snapshot option on RDS termination.
The company is now planning for a new test cycle and wants to create a new DB instance from the most recent backup. The company has chosen a MySQL-compatible edition of Amazon Aurora to host the DB instance.
Which solutions will create the new DB instance? (Select TWO.)

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

Answer: AC

Explanation:
These answers are correct because they meet the requirements of creating a new DB instance from the most recent backup and using a MySQL-compatible edition of Amazon Aurora to host the DB instance. You can import the RDS snapshot directly into Aurora if the MySQL DB instance and the Aurora DB cluster are running the same version of MySQL. For example, you can restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.6, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is simple and requires the fewest number of steps. You can upload the database dump to Amazon S3 and then import the database dump into Aurora if the MySQL DB instance and the Aurora DB cluster are running different versions of MySQL. For example, you can import a MySQL version 5.6 database dump into Aurora MySQL version 5.7, but you can't restore a MySQL version 5.6 snapshot directly to Aurora MySQL version 5.7. This method is more flexible and allows you to migrate across different versions of MySQL.
References:
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Import.html
? https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL
.Migrating.RDSMySQL.Dump.html

NEW QUESTION 71
- (Topic 4)
A company needs to create an AWS Lambda function that will run in a VPC in the
company's primary AWS account. The Lambda function needs to access files that the company stores
in an Amazon Elastic File System (Amazon EFS) file system. The EFS file system is located in a secondary AWS account. As the company adds files to the file system the solution must scale to meet the demand.
Which solution will meet these requirements MOST cost-effectively?

A. Create a new EPS file system in the primary account Use AWS DataSync to copy the contents of the original EPS file system to the new EPS file system
B. Create a VPC peering connection between the VPCs that are in the primary account and the secondary account
C. Create a second Lambda function In the secondary account that has a mount that is configured for the file syste
D. Use the primary account's Lambda function to invoke the secondary account's Lambda function
E. Move the contents of the file system to a Lambda Layer's Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

Answer: B

Explanation:
This option is the most cost-effective and scalable way to allow the Lambda function in the primary account to access the EFS file system in the secondary account. VPC peering enables private connectivity between two VPCs without requiring gateways, VPN connections, or dedicated network connections. The Lambda function can use the VPC peering connection to mount the EFS file system as a local file system and access the files as needed. The solution does not incur additional data transfer or storage costs, and it leverages the existing EFS file system without duplicating or moving the data.
Option A is not cost-effective because it requires creating a new EFS file system and using AWS DataSync to copy the data from the original EFS file system. This would incur additional storage and data transfer costs, and it would not provide real-time access to the files.
Option C is not scalable because it requires creating a second Lambda function in the secondary account and configuring cross-account permissions to invoke it from the primary account. This would add complexity and latency to the solution, and it would increase the Lambda invocation costs.
Option D is not feasible because Lambda layers are not designed to store large amounts of data or provide file system access. Lambda layers are used to share common code or libraries across multiple Lambda functions. Moving the contents of the EFS file system to a Lambda layer would exceed the size limit of 250 MB for a layer, and it would not allow the Lambda function to read or write files to the layer. References:
? What Is VPC Peering?
? Using Amazon EFS file systems with AWS Lambda
? What Are Lambda Layers?

NEW QUESTION 76
- (Topic 4)
A company runs an infrastructure monitoring service. The company is building a new feature that will enable the service to monitor data in customer AWS accounts. The new feature will call AWS APIs in customer accounts to describe Amazon EC2 instances and read Amazon CloudWatch metrics.
What should the company do to obtain access to customer accounts in the MOST secure way?

A. Ensure that the customers create an 1AM role in their account with read-only EC2 and CloudWatch permissions and a trust policy to the company's account.
B. Create a serverless API that implements a token vending machine to provide temporary AWS credentials for a role with read-only EC2 and CloudWatch permissions.
C. Ensure that the customers create an 1AM user in their account with read-only EC2 and CloudWatch permission
D. Encrypt and store customer access and secret keys in a secrets management system.
E. Ensure that the customers create an Amazon Cognito user in their account to use an 1AM role with read-only EC2 and CloudWatch permission
F. Encrypt and store the Amazon Cognito user and password in a secrets management system.

Answer: A

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

NEW QUESTION 77
- (Topic 4)
A company is conducting an internal audit. The company wants to ensure that the data in an Amazon S3 bucket that is associated with the company's AWS Lake Formation data lake does not contain sensitive customer or employee data. The company wants to discover personally identifiable information (Pll) or financial information, including passport numbers and credit card numbers.
Which solution will meet these requirements?

A. Configure AWS Audit Manager on the accoun
B. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
C. Configure Amazon S3 Inventory on the S3 bucke
D. Configure Amazon Athena to query the inventory.
E. Configure Amazon Macie to run a data discovery job that uses managed identifiers for the required data types.
F. Use Amazon S3 Select to run a report across the S3 bucket.

Answer: C

Explanation:
Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS. Macie can run data discovery jobs that use managed identifiers for various types of PII or financial information, such as passport numbers and credit card numbers. Macie can also generate findings that alert you to potential issues or risks with your data. References:
https://docs.aws.amazon.com/macie/latest/userguide/macie-identifiers.html

NEW QUESTION 79
- (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 G1o. 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 81
- (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 instanc
B. Associate a health check with each record.
C. Create a Route 53 failover routing policy record for each EC2 instanc
D. Associate a health check with each record.
E. Create an Amazon CloudFront distribution with EC2 instances as its origi
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 85

- (Topic 4)
A solutions architect has created two IAM policies: Policy1 and Policy2. Both policies are attached to an IAM group.

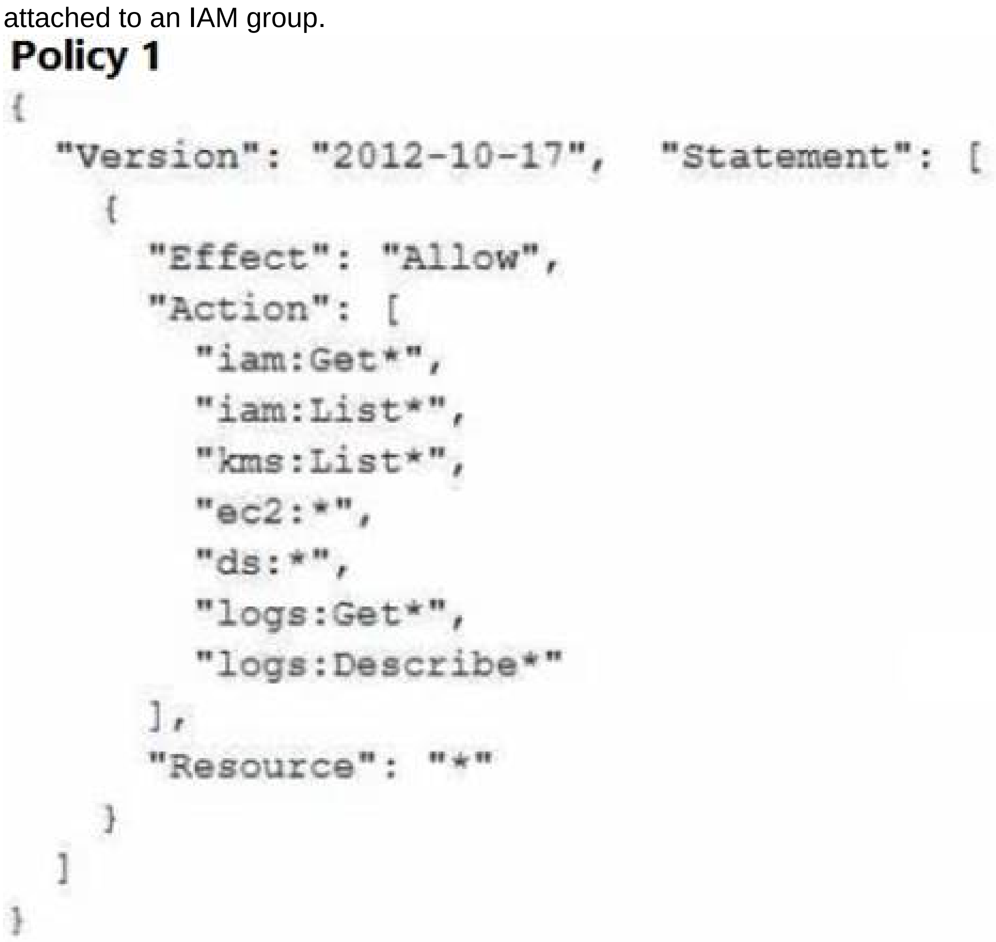

Policy 1

```
{
  "Version": "2012-10-17",  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:Get*",
        "iam:List*",
        "kms:List*",
        "ec2:*",
        "ds:*",
        "logs:Get*",
        "logs:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

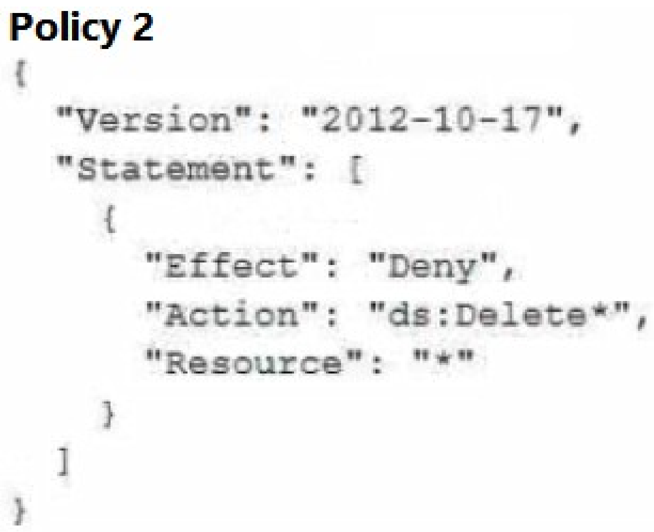

Policy 2

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "ds:Delete*",
      "Resource": "*"
    }
  ]
}
```

A cloud engineer is added as an IAM user to the IAM group. Which action will the cloud engineer be able to perform?

A. Deleting IAM users
B. Deleting directories
C. Deleting Amazon EC2 instances
D. Deleting logs from Amazon CloudWatch Logs

Answer: C

Explanation:
https://awscli.amazonaws.com/v2/documentation/api/latest/reference/ds/index.html

NEW QUESTION 90
- (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 accoun
B. Configure the application to use the correct secret from Secrets Manager to authenticate and read the DynamoDB tabl
C. Schedule secret rotation for every 30 days.
D. In every business account, create an 1AM user that has programmatic acces
E. Configure the application to use the correct 1AM user access key ID and secret access key to authenticate and read the DynamoDB tabl
F. Manually rotate 1AM access keys every 30 days.
G. In every business account, create an 1AM 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 accoun
H. In the inventory account, create a role named APP_ROLE that allows access to the STS AssumeRole API operatio
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 DynamoD
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 security 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 91
- (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 guid
C. Deploy the infrastructure with AWS CloudFormation
D. Use AWS Config to record the inventory of resources that are used in the prototype infrastructur
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-whatis- concepts.html

NEW QUESTION 94
- (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 dat
B. Share the transformation steps with employees by using AWS Glue jobs.
C. Configure Amazon EMR Serverless to transform the dat
D. Share the transformation steps with employees by using EMR Serveriess jobs.
E. Configure AWS Glue DataBrew to transform the dat
F. Share the transformation steps with employees by using DataBrew recipes.
G. Create Amazon Athena tables for the dat
H. Write Athena SQL queries to transform the dat
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 98
- (Topic 4)
A company needs to integrate with a third-party data feed. The data feed sends a webhook to notify an external service when new data is ready for consumption A developer wrote an AWS Lambfe function to retrieve data when the company receives a webhook callback The developer must make the Lambda function available for the third party to call.
Which solution will meet these requirements with the MOST operational efficiency?

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

Answer: A

Explanation:
A function URL is a unique identifier for a Lambda function that can be used to invoke the function over HTTPS. It is composed of the API endpoint of the AWS Region where the function is deployed, and the name or ARN of the function1. By creating a function URL for the Lambda function, the solution can make the Lambda function available for the third party to call with the most operational efficiency.
* B. Deploy an Application Load Balancer (ALB) in front of the Lambda function. Provide the ALB URL to the third party for the webhook. This solution will not meet the requirement of the most operational efficiency, as it involves creating and managing an additional resource (ALB) that is not necessary for invoking a Lambda function over HTTPS2.
* C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Attach the topic to the Lambda function. Provide the public hostname of the SNS topic to the third party for the webhook. This solution will not work, as Amazon SNS topics do not have public hostnames that can be used as webhooks. SNS topics are used to publish messages to subscribers, not to receive messages from external sources3.
* D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Attach the queue to the Lamb-da function. Provide the public hostname of the SQS queue to the third party for the webhook. This solution will not work, as Amazon SQS queues do not have public hostnames that can be used as webhooks. SQS queues are used to send, store, and receive messages between AWS services, not to receive messages from external sources. Reference URL:
https://docs.aws.amazon.com/lambda/latest/dg/lambda-api-permissions- ref.html

NEW QUESTION 102
- (Topic 4)
A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance New company management wants to ensure the application is highly available.
What should a solutions architect do to meet this requirement?

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

Answer: A

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

NEW QUESTION 107
- (Topic 4)
A gaming company wants to launch a new internet-facing application in multiple AWS Regions The application will use the TCP and UDP protocols for communication. The company needs to provide high availability and minimum latency for global users.
Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

A. Create internal Network Load Balancers in front of the application in each Region.
B. Create external Application Load Balancers in front of the application in each Region.
C. Create an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region.
D. Configure Amazon Route 53 to use a geolocation routing policy to distribute the traffic.
E. Configure Amazon CloudFront to handle the traffic and route requests to the application in each Region.

Answer: BC

Explanation:
This combination of actions will provide high availability and minimum latency for global users by using AWS Global Accelerator and Application Load Balancers. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your internet-facing applications by using the AWS global network. It provides two global static public IPs that act as a fixed entry point to your application endpoints, such as Application Load Balancers, in multiple Regions1. Global Accelerator uses the AWS backbone network to route traffic to the optimal regional
endpoint based on health, client location, and policies that you configure. It also offers TCP and UDP support, traffic encryption, and DDoS protection2. Application Load Balancers are external load balancers that distribute incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. They support both HTTP and HTTPS (SSL/TLS) protocols, and offer advanced features such as content-based routing, health checks, and integration with other AWS services3. By creating external Application Load Balancers in front of the application in each Region, you can ensure that the application can handle varying load patterns and scale on demand. By creating an AWS Global Accelerator accelerator to route traffic to the load balancers in each Region, you can leverage the performance, security, and availability of the AWS global network to deliver the best possible user experience.
References: 1: What is AWS Global Accelerator? - AWS Global Accelerator4, Overview section2: Network Acceleration Service - AWS Global Accelerator - AWS5, Why AWS Global Accelerator? section. 3: What is an Application Load Balancer? - Elastic Load Balancing6, Overview section.

NEW QUESTION 110
- (Topic 4)
A company runs multiple workloads in its on-premises data center. The company's data center cannot scale fast enough to meet the company's expanding business needs. The company wants to collect usage and configuration data about the on-premises servers and workloads to plan a migration to AWS.
Which solution will meet these requirements?

A. Set the home AWS Region in AWS Migration Hu
B. Use AWS Systems Manager to collect data about the on-premises servers.

C. Set the home AWS Region in AWS Migration Hu
D. Use AWS Application Discovery Service to collect data about the on-premises servers.
E. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant template
F. Use AWS Trusted Advisor to collect data about the on-premises servers.
G. Use the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates.Use AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers.

Answer: B

Explanation:
The most suitable solution for the company's requirements is to set the home AWS Region in AWS Migration Hub and use AWS Application Discovery Service to collect data about the on-premises servers. This solution will enable the company to gather usage and configuration data of its on-premises servers and workloads, and plan a migration to AWS.
AWS Migration Hub is a service that simplifies and accelerates migration tracking by aggregating migration status information into a single console. Users can view the discovered servers, group them into applications, and track the migration status of each application from the Migration Hub console in their home Region. The home Region is the AWS Region where users store their migration data, regardless of which Regions they migrate into1.
AWS Application Discovery Service is a service that helps users plan their migration to AWS by collecting usage and configuration data about their on-premises servers and databases. Application Discovery Service is integrated with AWS Migration Hub and supports two methods of performing discovery: agentless discovery and agent-based discovery. Agentless discovery can be performed by deploying the Application Discovery Service Agentless Collector through VMware vCenter, which collects static configuration data and utilization data for virtual machines (VMs) and databases. Agent-based discovery can be performed by deploying the AWS Application Discovery Agent on each of the VMs and physical servers, which collects static configuration data, detailed time-series system-performance information, inbound and outbound network connections, and processes that are running2.
The other options are not correct because they do not meet the requirements or are not relevant for the use case. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Trusted Advisor to collect data about the on- premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another, such as from Oracle to PostgreSQL3. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits4. Using the AWS Schema Conversion Tool (AWS SCT) to create the relevant templates and using AWS Database Migration Service (AWS DMS) to collect data about the on-premises servers is not correct because this solution is not suitable for collecting usage and configuration data of on-premises servers and workloads. As mentioned above, AWS SCT is a tool that helps users convert database schemas and code objects from one database engine to another. AWS DMS is a service that helps users migrate relational databases, non-relational databases, and other types of data stores to
AWS with minimal downtime5. References:
? Home Region - AWS Migration Hub
? What is AWS Application Discovery Service? - AWS Application Discovery Service
? AWS Schema Conversion Tool - Amazon Web Services
? What Is Trusted Advisor? - Trusted Advisor
? What Is AWS Database Migration Service? - AWS Database Migration Service

NEW QUESTION 112
- (Topic 4)
A company hosts an application used to upload files to an Amazon S3 bucket Once uploaded, the files are processed to extract metadata which takes less than 5 seconds The volume and frequency of the uploads varies from a few files each hour to hundreds of concurrent uploads The company has asked a solutions architect to design a cost-effective architecture that will meet these requirements.
What should the solutions architect recommend?

A. Configure AWS CloudTrail trails to tog S3 API calls Use AWS AppSync to process the files.
B. Configure an object-created event notification within the S3 bucket to invoke an AWS Lambda function to process the files.
C. Configure Amazon Kinesis Data Streams to process and send data to Amazon S3. Invoke an AWS Lambda function to process the files.
D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to process the files uploaded to Amazon S3 Invoke an AWS Lambda function to process the files.

Answer: B

Explanation:
This option is the most cost-effective and scalable way to process the files uploaded to S3. AWS CloudTrail is used to log API calls, not to trigger actions based on them. AWS AppSync is a service for building GraphQL APIs, not for processing files. Amazon Kinesis Data Streams is used to ingest and process streaming data, not to send data to S3. Amazon SNS is a pub/sub service that can be used to notify subscribers of events, not to process files. References:
? Using AWS Lambda with Amazon S3
? AWS CloudTrail FAQs
? What Is AWS AppSync?
? [What Is Amazon Kinesis Data Streams?]
? [What Is Amazon Simple Notification Service?]

NEW QUESTION 116
- (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 117
- (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 122
- (Topic 4)
A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a third-party provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead. What should the solutions architect do to meet these requirements?

A. Deploy a NAT instance in the VP
B. Route all the internet-based traffic through the NAT instance.
C. Deploy a NAT gateway in the public subnet
D. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
E. Configure an internet gateway and attach it to the VP
F. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
G. Configure a virtual private gateway and attach it to the VP
H. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

Answer: B

Explanation:
To allow the MySQL database in the private subnets to access the internet without exposing it to the public, a NAT gateway is a suitable solution. A NAT gateway enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. A NAT gateway resides in the public subnets and can handle high throughput of traffic with low latency. A NAT gateway is also a managed service that does not require any operational overhead. References:
? NAT Gateways
? NAT Gateway Pricing

NEW QUESTION 126
- (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 127
- (Topic 4)
A company runs a website that uses a content management system (CMS) on Amazon EC2. The CMS runs on a single EC2 instance and uses an Amazon Aurora MySQL Multi- AZ DB instance for the data tier. Website images are stored on an Amazon Elastic Block Store (Amazon EBS) volume that is mounted inside the EC2 instance.
Which combination of actions should a solutions architect take to improve the performance and resilience of the website? (Select TWO.)

A. Move the website images into an Amazon S3 bucket that is mounted on every EC2 instance.
B. Share the website images by using an NFS share from the primary EC2 instanc
C. Mountthis share on the other EC2 instances.
D. Move the website images onto an Amazon Elastic File System (Amazon EFS) file system that is mounted on every EC2 instance.
E. Create an Amazon Machine Image (AMI) from the existing EC2 instance Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling grou
F. Configure the Auto Scaling group to maintain a minimum of two instance
G. Configure an accelerator in AWS Global Accelerator for the website.
H. Create an Amazon Machine Image (AMI) from the existing EC2 instanc
I. Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling grou
J. Configure the Auto Scaling group to maintain a minimum of two instance
K. Configure an Amazon CloudFront distribution for the website.

Answer: CE

Explanation:
Option C provides moving the website images onto an Amazon EFS file system that is mounted on every EC2 instance. Amazon EFS provides a scalable and fully managed file storage solution that can be accessed concurrently from multiple EC2 instances. This ensures that the website images can be accessed efficiently and consistently by all instances, improving performance In Option E The Auto Scaling group maintains a minimum of two instances, ensuring resilience by automatically replacing any unhealthy instances. Additionally, configuring an Amazon CloudFront distribution for the website further improves performance by caching content at edge locations closer to the end-users, reducing latency and improving content delivery. Hence combining these actions, the website's performance is improved through efficient image storage and content delivery

NEW QUESTION 130
- (Topic 4)
A company is running a microservices application on Amazon EC2 instances. The company wants to migrate the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for scalability. The company must configure the Amazon EKS control plane with endpoint private access set to true and endpoint public access set to false to maintain security compliance The company must also put the data plane in private subnets. However, the company has received error notifications because the node cannot join the cluster.
Which solution will allow the node to join the cluster?

A. Grant the required permission in AWS Identity and Access Management (1AM) to the AmazonEKSNodeRole 1AM role.
B. Create interface VPC endpoints to allow nodes to access the control plane.
C. Recreate nodes in the public subnet Restrict security groups for EC2 nodes
D. Allow outbound traffic in the security group of the nodes.

Answer: B

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

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

NEW QUESTION 132
- (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 136
- (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 cachin
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 141
- (Topic 4)
A company is moving its on-premises Oracle database to Amazon Aurora PostgreSQL. The database has several applications that write to the same tables. The applications need to be migrated one by one with a month in between each migration. Management has expressed concerns that the database has a high number of reads and writes. The data must be kept in sync across both databases throughout the migration.
What should a solutions architect recommend?

A. Use AWS DataSync for the initial migratio
B. Use AWS Database Migration Service (AWS DMS) to create a change data capture (CDC) replication task and a table mapping to select all tables.
C. Use AWS DataSync for the initial migratio
D. Use AWS Database Migration Service (AWS DMS) to create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
E. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a memory optimized replication instanc
F. Create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
G. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a compute optimized replication instanc
H. Create a full load plus change data capture (CDC) replication task and a table mapping to select the largest tables.

Answer: C

Explanation:
https://aws.amazon.com/ko/premiumsupport/knowledge-center/dms- memory-optimization/

NEW QUESTION 144
- (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 budge
B. Set the budget amount under the Cost and Usage Reports section of the required AWS accounts.
C. Use AWS Budgets to create a budge
D. Set the budget amount under the Billing dashboards of the required AWS accounts.
E. Create an 1AM user for AWS Budgets to run budget actions with the required permissions.
F. Create an 1AM 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 threshol
H. Add a budget action that selects the 1AM 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 threshol
J. Add a budget action that selects the 1AM 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_policie s_scps.html

NEW QUESTION 145
- (Topic 4)
A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB) The website serves static content Website traffic is increasing and the company is concerned about a potential increase in cost.
What should a solutions architect do to reduce the cost of the website?

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

Answer: A

Explanation:
Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website.
CloudFront
can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not. References:
? https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/
? https://nodeployfriday.com/posts/static-website-hosting/
? https://aws.amazon.com/cloudfront/

NEW QUESTION 146
- (Topic 4)
A company runs multiple Amazon EC2 Linux instances in a VPC across two Availability Zones. The instances host applications that use a hierarchical directory structure. The applications need to read and write rapidly and concurrently to shared storage.
What should a solutions architect do to meet these requirements?

A. Create an Amazon S3 bucke
B. Allow access from all the EC2 instances in the VPC.
C. Create an Amazon Elastic File System (Amazon EFS) file syste
D. Mount the EFS file system from each EC2 instance.
E. Create a file system on a Provisioned IOPS SSD (102) Amazon Elastic Block Store (Amazon EBS) volum
F. Attach the EBS volume to all the EC2 instances.
G. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instanc
H. Synchromze the EBS volumes across the different EC2 instances.

Answer: B

Explanation:
it allows the EC2 instances to read and write rapidly and concurrently to shared storage across two Availability Zones. Amazon EFS provides a scalable, elastic, and highly available file system that can be mounted from multiple EC2 instances. Amazon EFS supports high levels of throughput and IOPS, and consistent low latencies. Amazon EFS also supports NFSv4 lock upgrading and downgrading, which enables high levels of concurrency. References:

? Amazon EFS Features
? Using Amazon EFS with Amazon EC2

NEW QUESTION 148
- (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 152
- (Topic 4)
A solutions architect must provide an automated solution for a company's compliance policy that states security groups cannot include a rule that allows SSH from 0.0.0.0/0. The company needs to be notified if there is any breach in the policy. A solution is needed as soon as possible.
What should the solutions architect do to meet these requirements with the LEAST operational overhead?

A. Write an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one.
B. Enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created.
C. Create an 1AM role with permissions to globally open security groups and network ACL
D. Create an Amazon Simple Notification Service (Amazon SNS) topic to generate a notification every time the role is assumed by a user.
E. Configure a service control policy (SCP) that prevents non-administrative users from creating or editing security group
F. Create a notification in the ticketing system when a user requests a rule that needs administrator permissions.

Answer: B

Explanation:
The most suitable solution for the company's compliance policy is to enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created. This solution has the least operational overhead because it uses a predefined rule that is already available in AWS Config, which is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. The restricted-ssh rule checks whether security groups that are in use have inbound rules that allow SSH from 0.0.0.0/0 addresses, and reports them as noncompliant1. Users can configure the rule to send notifications to an Amazon SNS topic when a noncompliant change occurs, and subscribe to the topic to receive alerts via email, SMS, or other methods2.
The other options are not correct because they either have more operational overhead or do not meet the requirements. Writing an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one is not correct because it requires custom code development and maintenance, which adds complexity and cost to the solution. Creating an IAM role with permissions to globally open security groups and network ACLs, and creating an Amazon SNS topic to generate a notification every time the role is assumed by a user is not correct because it does not prevent or detect the creation of noncompliant rules by other users or roles, and it does not address the existing rules that may violate the policy. Configuring a service control policy (SCP) that prevents non-administrative users from creating or editing security groups, and creating a notification in the ticketing system when a user requests a rule that needs administrator permissions is not correct because it does not provide an automated solution for the policy enforcement and notification, and it may limit the flexibility and productivity of the users.
References:
? restricted-ssh - AWS Config
? Getting Notifications When Your Resources Change - AWS Config

NEW QUESTION 156
......