

Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate



NEW QUESTION 1

- (Topic 4)

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

Which solution meets these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:

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

References:

- ? AWS Fargate
- ? Amazon Elastic File System
- ? Using Amazon EFS file systems with Amazon ECS

NEW QUESTION 2

- (Topic 4)

A company has 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 3

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

Option D is incorrect because an interface VPC endpoint for Amazon S3 is powered by AWS PrivateLink, which requires an elastic network interface (ENI) with a private IP address in each subnet. This adds complexity and cost to the solution. Moreover, an interface VPC endpoint does not support cross-Region access to Amazon S3. Reference URL: 1: <https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html> 2: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html#sc-dynamic-data-access> 3: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html> : <https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/>

NEW QUESTION 4

- (Topic 4)

An ecommerce company is running a seasonal online sale. The company hosts its website on Amazon EC2 instances spanning multiple Availability Zones. The company wants its website to manage sudden traffic increases during the sale.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an Auto Scaling group that is large enough to handle peak traffic load
- B. Stop half of the Amazon EC2 instance
- C. Configure the Auto Scaling group to use the stopped instances to scale out when traffic increases.
- D. Create an Auto Scaling group for the website
- E. Set the minimum size of the Auto Scaling group so that it can handle high traffic volumes without the need to scale out.
- F. Use Amazon CloudFront and Amazon ElastiCache to cache dynamic content with an Auto Scaling group set as the origin
- G. Configure the Auto Scaling group with the instances necessary to populate CloudFront and ElastiCache
- H. Scale in after the cache is fully populated.
- I. Configure an Auto Scaling group to scale out as traffic increases
- J. Create a launch template to start new instances from a preconfigured Amazon Machine Image (AMI).

Answer: D

Explanation:

The solution that meets the requirements of high availability, resiliency, and minimal operational effort is to use AWS Transfer for SFTP and an Amazon S3 bucket for storage. This solution allows the company to securely transfer files over SFTP to Amazon S3, which is a durable and scalable object storage service. The company can then modify the application to pull the batch files from Amazon S3 to an Amazon EC2 instance for processing. The EC2 instance can be part of an Auto Scaling group with a scheduled scaling policy to run the batch operation only at night. This way, the company can save costs by scaling down the EC2 instances when they are not needed. The other solutions do not meet all the requirements because they either use Amazon EFS or Amazon EBS for storage, which are more expensive and less scalable than Amazon S3, or they do not use a scheduled scaling policy to optimize the EC2 instances usage. References :=

? AWS Transfer for SFTP

? Amazon S3

? Amazon EC2 Auto Scaling

NEW QUESTION 5

- (Topic 4)

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: A

Explanation:

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

NEW QUESTION 6

- (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) turned 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 7

- (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 8

- (Topic 4)

An image hosting company uploads its large assets to Amazon S3 Standard buckets. The company uses multipart upload in parallel by using S3 APIs and overwrites if the same object is uploaded again. For the first 30 days after upload, the objects will be accessed frequently. The objects will be used less frequently after 30 days, but the access patterns for each object will be inconsistent. The company must optimize its S3 storage costs while maintaining high availability and resiliency of stored assets.

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

- A. Move assets to S3 Intelligent-Tiering after 30 days.
- B. Configure an S3 Lifecycle policy to clean up incomplete multipart uploads.
- C. Configure an S3 Lifecycle policy to clean up expired object delete markers.
- D. Move assets to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- E. Move assets to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

Answer: AB

Explanation:

S3 Intelligent-Tiering is a storage class that automatically moves data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead¹. It is ideal for data with unknown or changing access patterns, such as the company's assets. By moving assets to S3 Intelligent-Tiering after 30 days, the company can optimize its storage costs while maintaining high availability and resilience of stored assets.

S3 Lifecycle is a feature that enables you to manage your objects so that they are stored cost effectively throughout their lifecycle². You can create lifecycle rules to define actions that Amazon S3 applies to a group of objects. One of the actions is to abort incomplete multipart uploads that can occur when an upload is interrupted. By configuring an S3 Lifecycle policy to clean up incomplete multipart uploads, the company can reduce its storage costs and avoid paying for parts that are not used.

Option C is incorrect because expired object delete markers are automatically deleted by Amazon S3 and do not incur any storage costs³. Therefore, configuring an S3 Lifecycle policy to clean up expired object delete markers will not have any effect on the company's storage costs.

Option D is incorrect because S3 Standard-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed¹. It has a lower storage cost than S3 Standard, but it has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 Standard-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally.

Option E is incorrect because S3 One Zone-IA is a storage class for data that is accessed less frequently, but requires rapid access when needed¹. It has a lower storage cost than S3 Standard-IA, but it stores data in only one Availability Zone and has less resilience than other storage classes. It also has a higher retrieval cost and a minimum storage duration charge of 30 days. Therefore, moving assets to S3 One Zone-IA after 30 days may not optimize the company's storage costs if the assets are still accessed occasionally or require high availability. Reference URL: 1: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html> 2:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html> 3: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/delete-or-empty-bucket.html#delete-bucket-considerations> : <https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html> :

<https://aws.amazon.com/certification/certified-solutions-architect-associate/>

NEW QUESTION 9

- (Topic 4)

A company has users all around the world accessing its HTTP-based application deployed on Amazon EC2 instances in multiple AWS Regions. The company wants to improve the availability and performance of the application. The company also wants to protect the application against common web exploits that may affect availability, compromise security, or consume excessive resources. Static IP addresses are required.

What should a solutions architect recommend to accomplish this?

- A. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region
- B. Deploy AWS WAF on the NLB
- C. Create an accelerator using AWS Global Accelerator and register the NLBs as endpoints.
- D. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region
- E. Deploy AWS WAF on the ALB
- F. Create an accelerator using AWS Global Accelerator and register the ALBs as endpoints.
- G. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region
- H. Deploy AWS WAF on the NLB
- I. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the NLBs.
- J. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region
- K. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the ALB
- L. Deploy AWS WAF on the CloudFront distribution.

Answer: A

Explanation:

The company wants to improve the availability and performance of the application, as well as protect it against common web exploits. The company also needs static IP addresses for the application. To meet these requirements, a solutions architect should recommend the following solution:

? Put the EC2 instances behind Network Load Balancers (NLBs) in each Region.

NLBs are designed to handle millions of requests per second while maintaining high throughput at ultra-low latency. NLBs also support static IP addresses for each Availability Zone, which can be useful for whitelisting or firewalling purposes.

? Deploy AWS WAF on the NLBs. AWS WAF is a web application firewall that helps

protect web applications from common web exploits that could affect availability, security, or performance. AWS WAF lets you define customizable web security rules that control which traffic to allow or block to your web applications.

? Create an accelerator using AWS Global Accelerator and register the NLBs as

endpoints. AWS Global Accelerator is a service that improves the availability and performance of your applications with local or global users. It provides static IP addresses that act as a fixed entry point to your application endpoints in any AWS Region. It uses the AWS global network to optimize the path from your users to your applications, improving the performance of your TCP and UDP traffic.

This solution will provide high availability across Availability Zones and Regions, improve performance by routing traffic over the AWS global network, protect the application from common web attacks, and provide static IP addresses for the application.

References:

? Network Load Balancer

? AWS WAF

? AWS Global Accelerator

NEW QUESTION 10

- (Topic 4)

A company hosts a multi-tier web application on Amazon Linux Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company observes that the Auto Scaling group launches more On-Demand Instances when the application's end users access high volumes of static web content. The company wants to optimize cost.

What should a solutions architect do to redesign the application MOST cost-effectively?

- A. Update the Auto Scaling group to use Reserved Instances instead of On-Demand Instances.
- B. Update the Auto Scaling group to scale by launching Spot Instances instead of On-Demand Instances.
- C. Create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket.
- D. Create an AWS Lambda function behind an Amazon API Gateway API to host the static website contents.

Answer: C

Explanation:

This answer is correct because it meets the requirements of optimizing cost and reducing the workload on the database. Amazon CloudFront is a content delivery network (CDN) service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the request is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance. You can create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket, which is an origin that you define for CloudFront. This way, you can offload the requests for static web content from your EC2 instances to CloudFront, which can improve the performance and availability of your website, and reduce the cost of running your EC2 instances.

References:

? <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>

NEW QUESTION 10

- (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 IAM Identity Center (AWS Single Sign-On) between AWS and the on-premises LDAP.
- B. Create an IAM policy that uses AWS credentials, and integrate the policy into LDAP.
- C. Set up a process that rotates the IAM 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 14

- (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 targ
- 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 18

- (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 20

- (Topic 4)

A company is reviewing a recent migration of a three-tier application to a VPC. The security team discovers that the principle of least privilege is not being applied to Amazon EC2 security group ingress and egress rules between the application tiers.

What should a solutions architect do to correct this issue?

- A. Create security group rules using the instance ID as the source or destination.
- B. Create security group rules using the security group ID as the source or destination.
- C. Create security group rules using the VPC CIDR blocks as the source or destination.
- D. Create security group rules using the subnet CIDR blocks as the source or destination.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group-rules.html>

NEW QUESTION 21

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

References:

- ? 1 provides an overview of Amazon S3 File Gateway and its benefits.
- ? 2 explains how to use S3 Lifecycle policy to manage object storage lifecycle.
- ? 3 describes the features and use cases of S3 Glacier Deep Archive storage class.

NEW QUESTION 26

- (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 30

- (Topic 4)

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

Which method should the solutions architect select?

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

Answer: A

Explanation:

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

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

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

configuring DAX for the

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

* B. Add DynamoDB read replicas to handle the increased read load. Update the application to point to the read endpoint for the read replicas. This solution will not work, as DynamoDB does not support read replicas as a feature. Read replicas are available for Amazon RDS, not for DynamoDB2.

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

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

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

NEW QUESTION 34

- (Topic 4)

A company operates a two-tier application for image processing. The application uses two Availability Zones, each with one public subnet and one private subnet. An Application Load Balancer (ALB) for the web tier uses the public subnets. Amazon EC2 instances for the application tier use the private subnets.

Users report that the application is running more slowly than expected. A security audit of the web server log files shows that the application is receiving millions of illegitimate requests from a small number of IP addresses. A solutions architect needs to resolve the immediate performance problem while the company investigates a more permanent solution.

What should the solutions architect recommend to meet this requirement?

- A. Modify the inbound security group for the web tier
- B. Add a deny rule for the IP addresses that are consuming resources.
- C. Modify the network ACL for the web tier subnet
- D. Add an inbound deny rule for the IP addresses that are consuming resources
- E. Modify the inbound security group for the application tier
- F. Add a deny rule for the IP addresses that are consuming resources.
- G. Modify the network ACL for the application tier subnet
- H. Add an inbound deny rule for the IP addresses that are consuming resources

Answer: B

Explanation:

Deny the request from the first entry at the public subnet, don't allow it to cross and get to the private subnet.

In this scenario, the security audit reveals that the application is receiving millions of illegitimate requests from a small number of IP addresses. To address this issue, it is recommended to modify the network ACL (Access Control List) for the web tier subnets. By adding an inbound deny rule specifically targeting the IP addresses that are consuming resources, the network ACL can block the illegitimate traffic at the subnet level before it reaches the web servers. This will help alleviate the excessive load on the web tier and improve the application's performance.

NEW QUESTION 38

- (Topic 4)

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

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

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

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

Answer: D

Explanation:

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

References:

? <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html>

? <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>

? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/autoscaling-ecs.html>

NEW QUESTION 42

- (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 (PII) or financial information, including passport numbers and credit card numbers.

Which solution will meet these requirements?

- A. Configure AWS Audit Manager on the account
- B. Select the Payment Card Industry Data Security Standards (PCI DSS) for auditing.
- C. Configure Amazon S3 Inventory on the S3 bucket
- 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 44

- (Topic 4)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 45

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

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

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

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

? You can leverage the scalability and elasticity of AWS to run large-scale parallel workloads without worrying about provisioning or managing servers¹.

? You can use MPI libraries to optimize the performance and efficiency of your parallel applications by enabling inter-process communication and data exchange².

? You can choose from a variety of MPI implementations that are compatible with AWS ParallelCluster, such as Open MPI, Intel MPI, and MPICH3.

NEW QUESTION 46

- (Topic 4)

A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance. The 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 49

- (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 Hub
- B. Use AWS Systems Manager to collect data about the on-premises servers.
- C. Set the home AWS Region in AWS Migration Hub
- 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 50

- (Topic 4)

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

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

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

Answer: BE

Explanation:

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

NEW QUESTION 55

- (Topic 4)

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

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

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

- A. Create a public Network Load Balancer Specify the application target group.

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

Answer: CE

Explanation:

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

NEW QUESTION 60

- (Topic 4)

A company has multiple Windows file servers on premises. The company wants to migrate and consolidate its files into an Amazon FSx for Windows File Server file system. File permissions must be preserved to ensure that access rights do not change. Which solutions will meet these requirements? (Select TWO.)

- A. Deploy AWS DataSync agents on premise
- B. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- C. Copy the shares on each file server into Amazon S3 buckets by using the AWS CLI Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- D. Remove the drives from each file server Ship the drives to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system
- E. Order an AWS Snowcone device
- F. Connect the device to the on-premises network
- G. Launch AWS DataSync agents on the device
- H. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system,
- I. Order an AWS Snowball Edge Storage Optimized device
- J. Connect the device to the on- premises network
- K. Copy data to the device by using the AWS CL
- L. Ship the device back to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.

Answer: AD

Explanation:

A This option involves deploying DataSync agents on your on-premises file servers and using DataSync to transfer the data directly to the FSx for Windows File Server. DataSync ensures that file permissions are preserved during the migration process. D This option involves using an AWS Snowcone device, a portable data transfer device. You would connect the Snowcone device to your on-premises network, launch DataSync agents on the device, and schedule DataSync tasks to transfer the data to FSx for Windows File Server. DataSync handles the migration process while preserving file permissions.

NEW QUESTION 61

- (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 64

- (Topic 4)

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

- A. Add an Amazon CloudFront distribution in front of the NLB
- B. Increase the Cache- Control: max-age parameter.

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 69

- (Topic 4)

A company uses AWS Organizations. The company wants to operate some of its AWS accounts with different budgets. The company wants to receive alerts and automatically prevent provisioning of additional resources on AWS accounts when the allocated budget threshold is met during a specific period. Which combination of solutions will meet these requirements? (Select THREE.)

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

Answer: BDF

Explanation:

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

- ? Use AWS Budgets to create a budget for each AWS account that needs a different budget amount. The budget can be based on cost or usage metrics, and can have different time periods, filters, and thresholds. The company can set the budget amount under the Billing dashboards of the required AWS accounts¹.
- ? Create an IAM role for AWS Budgets to run budget actions with the required permissions. A budget action is a response that AWS Budgets initiates when a budget exceeds a specified threshold. The IAM role allows AWS Budgets to perform actions on behalf of the company, such as applying an IAM policy or a service control policy (SCP) to restrict the provisioning of additional resources².
- ? Add an alert to notify the company when each account meets its budget threshold. The alert can be sent via email or Amazon SNS. The company can also add a budget action that selects the IAM role created and the appropriate SCP to prevent provisioning of additional resources. An SCP is a type of policy that can be applied to an AWS account or an organizational unit (OU) within AWS Organizations. An SCP can limit the actions that users and roles can perform in the account or OU³.

References:

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

NEW QUESTION 74

- (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 bucket
- B. Allow access from all the EC2 instances in the VPC.
- C. Create an Amazon Elastic File System (Amazon EFS) file system
- 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) volume
- 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 instance
- H. Synchronize 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 79

- (Topic 4)

A company is deploying an application in three AWS Regions using an Application Load Balancer. Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST high-performing experience?

- A. Create an A record with a latency policy.
- B. Create an A record with a geolocation policy.
- C. Create a CNAME record with a failover policy.
- D. Create a CNAME record with a geoproximity policy.

Answer: A

Explanation:

To provide the most high-performing experience for the users of the application, a solutions architect should use a latency routing policy for the Route 53 A record. This policy allows Route 53 to route traffic to the AWS Region that provides the lowest possible latency for the users¹. A latency routing policy can also improve the availability of the application, as Route 53 can automatically route traffic to another Region if the primary Region becomes unavailable².

References:

? 1: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-latency>

? 2: https://aws.amazon.com/route53/faqs/#Latency_Based_Routing

NEW QUESTION 81

- (Topic 4)

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

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

What should a solutions architect recommend to resolve this issue?

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

Answer: D

Explanation:

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

Reference URLs:

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

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

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

NEW QUESTION 82

- (Topic 4)

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

After an audit from a regulator, the company has 90 days to move the data to the cloud. The company needs to move the data efficiently and without disruption.

The company still needs to be able to access and update the data during the transfer window.

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

References:

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

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

NEW QUESTION 84

- (Topic 4)

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

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

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 87

- (Topic 4)

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

The permissions will be used by multiple IAM users and must be split between the developer and administrator teams. Each team requires different permissions.

The company wants a solution that includes new users that are hired on both teams.

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

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 88

- (Topic 4)

A company runs applications on AWS that connect to the company's Amazon RDS database. The applications scale on weekends and at peak times of the year. The company wants to scale the database more effectively for its applications that connect to the database.

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

- A. Use Amazon DynamoDB with connection pooling with a target group configuration for the databases
- B. Change the applications to use the DynamoDB endpoint.
- C. Use Amazon RDS Proxy with a target group for the databases
- D. Change the applications to use the RDS Proxy endpoint.
- E. Use a custom proxy that runs on Amazon EC2 as an intermediary to the databases
- F. Change the applications to use the custom proxy endpoint.
- G. Use an AWS Lambda function to provide connection pooling with a target group configuration for the databases
- H. Change the applications to use the Lambda function.

Answer: B

Explanation:

Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure¹. RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability². RDS Proxy also reduces failover times for Aurora and RDS databases by up to 66% and enables IAM authentication and Secrets Manager integration for database access¹. RDS Proxy can be enabled for most applications with no code changes².

NEW QUESTION 91

- (Topic 4)

A company has multiple AWS accounts with applications deployed in the us-west-2 Region. Application logs are stored within Amazon S3 buckets in each account. The company wants to build a centralized log analysis solution that uses a single S3 bucket. Logs must not leave us-west-2, and the company wants to incur minimal operational overhead.

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

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

Answer: B

Explanation:

This solution meets the following requirements:

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

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

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

References:

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

NEW QUESTION 92

- (Topic 4)

The following IAM policy is attached to an IAM group. This is the only policy applied to the group.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "1",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    },
    {
      "Sid": "2",
      "Effect": "Deny",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "BoolIfExists": {"aws:MultiFactorAuthPresent": false}
      }
    }
  ]
}
```

- A. Group members are permitted any Amazon EC2 action within the us-east-1 Region
- B. Statements after the Allow permission are not applied.
- C. Group members are denied any Amazon EC2 permissions in the us-east-1 Region unless they are logged in with multi-factor authentication (MFA).
- D. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for all Regions when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action.
- E. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for the us-east-1 Region only when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action within the us-east-1 Region.

Answer: D

Explanation:

This answer is correct because it reflects the effect of the IAM policy on the group members. The policy has two statements: one with an Allow effect and one with a Deny effect. The Allow statement grants permission to perform any EC2 action on any resource within the us-east-1 Region. The Deny statement overrides the Allow statement and denies permission to perform the ec2:StopInstances and ec2:TerminateInstances actions on any resource within the us-east-1 Region, unless the group member is logged in with MFA. Therefore, the group members can perform any EC2 action except stopping or terminating instances in the us-east-1 Region, unless they use MFA.

NEW QUESTION 94

- (Topic 4)

A company wants to experiment with individual AWS accounts for its engineer team. The company wants to be notified as soon as the Amazon EC2 instance usage for a given month exceeds a specific threshold for each account.

What should a solutions architect do to meet this requirement MOST cost-effectively?

- A. Use Cost Explorer to create a daily report of costs by service
- B. Filter the report by EC2 instance
- C. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
- D. Use Cost Explorer to create a monthly report of costs by service
- E. Filter the report by EC2 instance
- F. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
- G. Use AWS Budgets to create a cost budget for each account
- H. Set the period to month
- I. Set the scope to EC2 instance
- J. Set an alert threshold for the budget
- K. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.
- L. Use AWS Cost and Usage Reports to create a report with hourly granularity
- M. Integrate the report data with Amazon Athena
- N. Use Amazon EventBridge to schedule an Athena query
- O. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.

Answer: C

Explanation:

AWS Budgets allows you to create budgets for your AWS accounts and set alerts when usage exceeds a certain threshold. By creating a budget for each account, specifying the period as monthly and the scope as EC2 instances, you can effectively track the EC2 usage for each account and be notified when a threshold is exceeded. This solution is the most cost-effective option as it does not require additional resources such as Amazon Athena or Amazon EventBridge.

NEW QUESTION 96

- (Topic 4)

A company has a financial application that produces reports. The reports average 50 KB in size and are stored in Amazon S3. The reports are frequently accessed during the first week after production and must be stored for several years. The reports must be retrievable within 6 hours.

Which solution meets these requirements MOST cost-effectively?

- A. Use S3 Standard
- B. Use an S3 Lifecycle rule to transition the reports to S3 Glacier after 7 days.
- C. Use S3 Standard
- D. Use an S3 Lifecycle rule to transition the reports to S3 Standard- Infrequent Access (S3 Standard-IA) after 7 days.
- E. Use S3 Intelligent-Tiering
- F. Configure S3 Intelligent-Tiering to transition the reports to S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Glacier.
- G. Use S3 Standard
- H. Use an S3 Lifecycle rule to transition the reports to S3 Glacier Deep Archive after 7 days.

Answer: A

Explanation:

To store and retrieve reports that are frequently accessed during the first week and must be stored for several years, S3 Standard and S3 Glacier are suitable solutions. S3 Standard offers high durability, availability, and performance for frequently accessed data. S3 Glacier offers secure and durable storage for long-term data archiving at a low cost. S3 Lifecycle rules can be used to transition the reports from S3 Standard to S3 Glacier after 7 days, which can reduce storage costs. S3 Glacier also supports retrieval within 6 hours.

References:

? Storage Classes

? Object Lifecycle Management

? Retrieving Archived Objects from Amazon S3 Glacier

NEW QUESTION 97

- (Topic 4)

A company wants to migrate its on-premises Microsoft SQL Server Enterprise edition database to AWS. The company's online application uses the database to process transactions. The data analysis team uses the same production database to run reports for analytical processing. The company wants to reduce operational overhead by moving to managed services wherever possible.

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

- A. Migrate to Amazon RDS for Microsoft SQL Server
- B. Use read replicas for reporting purposes.
- C. Migrate to Microsoft SQL Server on Amazon EC2. Use Always On read replicas for reporting purposes.
- D. Migrate to Amazon DynamoDB
- E. Use DynamoDB on-demand replicas for reporting purposes.
- F. Migrate to Amazon Aurora MySQL
- G. Use Aurora read replicas for reporting purposes.

Answer: A

Explanation:

Amazon RDS for Microsoft SQL Server is a fully managed service that offers SQL Server 2014, 2016, 2017, and 2019 editions while offloading database administration tasks such as backups, patching, and scaling. Amazon RDS supports read replicas, which are read-only copies of the primary database that can be used for reporting purposes without affecting the performance of the online application. This solution will meet the requirements with the least operational overhead, as it does not require any code changes or manual intervention.

References:

? 1 provides an overview of Amazon RDS for Microsoft SQL Server and its benefits.

? 2 explains how to create and use read replicas with Amazon RDS.

NEW QUESTION 99

- (Topic 4)

A company runs an application on a group of Amazon Linux EC2 instances. For compliance reasons, the company must retain all application log files for 7 years. The log files will be analyzed by a reporting tool that must be able to access all the files concurrently.

Which storage solution meets these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon EC2 instance store
- D. Amazon S3

Answer: D

Explanation:

<https://docs.aws.amazon.com/efs/latest/ug/transfer-data-to-efs.html>

NEW QUESTION 103

- (Topic 4)

A company hosts an internal serverless application on AWS by using Amazon API Gateway and AWS Lambda. The company's employees report issues with high latency when they begin using the application each day. The company wants to reduce latency.

Which solution will meet these requirements?

- A. Increase the API Gateway throttling limit.
- B. Set up a scheduled scaling to increase Lambda provisioned concurrency before employees begin to use the application each day.
- C. Create an Amazon CloudWatch alarm to initiate a Lambda function as a target for the alarm at the beginning of each day.
- D. Increase the Lambda function memory.

Answer: B

Explanation:

AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda scales automatically based on the incoming requests, but it may take some time to initialize new instances of your function if there is a sudden increase in demand. This may result in high latency or cold starts for your application. To avoid this, you can use provisioned concurrency, which ensures that your function is initialized and ready to respond at any time. You can also set up a scheduled scaling policy that increases the provisioned concurrency before employees begin to use the application each day, and decreases it when the demand is low. References: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html>

NEW QUESTION 104

- (Topic 4)

A company has multiple AWS accounts for development work. Some staff consistently use oversized Amazon EC2 instances, which causes the company to exceed the yearly budget for the development accounts. The company wants to centrally restrict the creation of AWS resources in these accounts. Which solution will meet these requirements with the LEAST development effort?

- A. Develop AWS Systems Manager templates that use an approved EC2 creation process.
- B. Use the approved Systems Manager templates to provision EC2 instances.
- C. Use AWS Organizations to organize the accounts into organizational units (OUs). Define and attach a service control policy (SCP) to control the usage of EC2 instance types.
- D. Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is created.
- E. Stop disallowed EC2 instance types.
- F. Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types. Ensure that staff can deploy EC2 instances only by using the Service Catalog products.

Answer: B

Explanation:

AWS Organizations is a service that helps users centrally manage and govern multiple AWS accounts. It allows users to create organizational units (OUs) to group accounts based on business needs or other criteria. It also allows users to define and attach service control policies (SCPs) to OUs or accounts to restrict the actions that can be performed by the accounts. By using AWS Organizations, the solution can centrally restrict the creation of AWS resources in the development accounts.

- * A. Develop AWS Systems Manager templates that use an approved EC2 creation process. Use the approved Systems Manager templates to provision EC2 instances. This solution will not meet the requirement of the least development effort, as it involves developing and maintaining custom templates for EC2 creation, and relying on the staff to use the approved templates instead of enforcing a restriction.
- * C. Configure an Amazon EventBridge rule that invokes an AWS Lambda function when an EC2 instance is created. Stop disallowed EC2 instance types. This solution will not meet the requirement of the least development effort, as it involves writing custom code for Lambda functions, and handling events and errors for EC2 creation.
- * D. Set up AWS Service Catalog products for the staff to create the allowed EC2 instance types. Ensure that staff can deploy EC2 instances only by using the Service Catalog products. This solution will not meet the requirement of the least development effort, as it involves setting up and managing Service Catalog products for EC2 creation, and ensuring that staff can only use Service Catalog products instead of enforcing a restriction. Reference URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

NEW QUESTION 105

- (Topic 4)

An application runs on an Amazon EC2 instance that has an Elastic IP address in VPC A. The application requires access to a database in VPC B. Both VPCs are in the same AWS account. Which solution will provide the required access MOST securely?

- A. Create a DB instance security group that allows all traffic from the public IP address of the application server in VPC A.
- B. Configure a VPC peering connection between VPC A and VPC B.
- C. Make the DB instance publicly accessible.
- D. Assign a public IP address to the DB instance.
- E. Launch an EC2 instance with an Elastic IP address into VPC B.
- F. Proxy all requests through the new EC2 instance.

Answer: B

Explanation:

A VPC peering connection is a networking connection between two VPCs that enables users to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network. A VPC peering connection can be created between VPCs in the same or different AWS accounts and Regions. By configuring a VPC peering connection between VPC A and VPC B, the solution can provide the required access most securely.

- * A. Create a DB instance security group that allows all traffic from the public IP address of the application server in VPC A. This solution will not provide the required access most securely, as it involves exposing the DB instance to the public internet and relying on a single IP address for access control.
 - * C. Make the DB instance publicly accessible. Assign a public IP address to the DB instance. This solution will not provide the required access most securely, as it involves exposing the DB instance to the public internet and allowing any source to connect to it.
 - * D. Launch an EC2 instance with an Elastic IP address into VPC B. Proxy all requests through the new EC2 instance. This solution will not provide the required access most securely, as it involves creating an additional resource and configuring a proxy server that may introduce latency and complexity.
- Reference URL: <https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html>

NEW QUESTION 110

- (Topic 4)

A company needs to extract the names of ingredients from recipe records that are stored as text files in an Amazon S3 bucket. A web application will use the ingredient names to query an Amazon DynamoDB table and determine a nutrition score. The application can handle non-food records and errors. The company does not have any employees who have machine learning knowledge to develop this solution. Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur. Program the Lambda function to analyze the object and extract the ingredient names by using Amazon Comprehend. Store the Amazon Comprehend output in the DynamoDB table.

- B. Use an Amazon EventBridge rule to invoke an AWS Lambda function when PutObject requests occur
- C. Program the Lambda function to analyze the object by using Amazon Forecast to extract the ingredient names. Store the Forecast output in the DynamoDB table.
- D. Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur. Use Amazon Polly to create audio recordings of the recipe record.
- E. Save the audio files in the S3 bucket. Use Amazon Simple Notification Service (Amazon SNS) to send a URL as a message to employees. Instruct the employees to listen to the audio files and calculate the nutrition score. Store the ingredient names in the DynamoDB table.
- F. Use an Amazon EventBridge rule to invoke an AWS Lambda function when a PutObject request occurs. Program the Lambda function to analyze the object and extract the ingredient names by using Amazon SageMaker. Store the inference output from the SageMaker endpoint in the DynamoDB table.

Answer: A

Explanation:

This solution meets the following requirements:

- ? It is cost-effective, as it only uses serverless components that are charged based on usage and do not require any upfront provisioning or maintenance.
- ? It is scalable, as it can handle any number of recipe records that are uploaded to the S3 bucket without any performance degradation or manual intervention.
- ? It is easy to implement, as it does not require any machine learning knowledge or complex data processing logic. Amazon Comprehend is a natural language processing service that can automatically extract entities such as ingredients from text files. The Lambda function can simply invoke the Comprehend API and store the results in the DynamoDB table.
- ? It is reliable, as it can handle non-food records and errors gracefully. Amazon Comprehend can detect the language and domain of the text files and return an appropriate response. The Lambda function can also implement error handling and logging mechanisms to ensure the data quality and integrity.

References:

- ? Using AWS Lambda with Amazon S3 - AWS Lambda
- ? What Is Amazon Comprehend? - Amazon Comprehend
- ? Working with Tables - Amazon DynamoDB

NEW QUESTION 114

- (Topic 4)

A group requires permissions to list an Amazon S3 bucket and delete objects from that bucket. An administrator has created the following IAM policy to provide access to the bucket and applied that policy to the group. The group is not able to delete objects in the bucket. The company follows least-privilege access rules.

A)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:DeleteObject"
      ],
      "Resource": [
        "arn:aws:s3:::bucket-name"
      ],
      "Effect": "Allow"
    }
  ]
}
```

B)

```
"Action": [
  "s3:*Object"
],
"Resource": [
  "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

C)

```
"Action": [
  "s3:DeleteObject"
],
"Resource": [
  "arn:aws:s3:::bucket-name*"
],
"Effect": "Allow"
```

D)

```
"Action": [
  "s3:DeleteObject"
],
"Resource": [
  "arn:aws:s3:::bucket-name/*"
],
"Effect": "Allow"
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation:

```
{ "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [ "s3:ListBucket", "s3:DeleteObject" ],
      "Resource": [ "arn:aws:s3:::<bucket-name>" ],
      "Effect": "Allow",
    },
    {
      "Action": "s3:*DeleteObject", "Resource": [
        "arn:aws:s3:::<bucket-name>/*" # <- The policy clause kludge "added" to match the solution (Q248.1) example
      ],
      "Effect": "Allow"
    }
  ]
}
```

NEW QUESTION 115

- (Topic 4)

A solutions architect needs to review a company's Amazon S3 buckets to discover personally identifiable information (PII). The company stores the PII data in the us-east-1 Region and us-west-2 Region.

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

- A. Configure Amazon Macie in each Region
- B. Create a job to analyze the data that is in Amazon S3_
- C. Configure AWS Security Hub for all Region
- D. Create an AWS Config rule to analyze the data that is in Amazon S3_
- E. Configure Amazon Inspector to analyze the data that IS in Amazon S3.
- F. Configure Amazon GuardDuty to analyze the data that is in Amazon S3.

Answer: A

Explanation:

it allows the solutions architect to review the S3 buckets to discover personally identifiable information (PII) with the least operational overhead. Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. Amazon Macie can analyze data in S3 buckets across multiple regions and provide insights into the type, location, and level of sensitivity of the data. References:

? Amazon Macie

? Analyzing data with Amazon Macie

NEW QUESTION 118

- (Topic 4)

A company that uses AWS needs a solution to predict the resources needed for manufacturing processes each month. The solution must use historical values that are currently stored in an Amazon S3 bucket The company has no machine learning (ML) experience and wants to use a managed service for the training and predictions.

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

- A. Deploy an Amazon SageMaker model
- B. Create a SageMaker endpoint for inference.
- C. Use Amazon SageMaker to train a model by using the historical data in the S3 bucket.
- D. Configure an AWS Lambda function with a function URL that uses Amazon SageMaker endpoints to create predictions based on the inputs.
- E. Configure an AWS Lambda function with a function URL that uses an Amazon Forecast predictor to create a prediction based on the inputs.
- F. Train an Amazon Forecast predictor by using the historical data in the S3 bucket.

Answer: BE

Explanation:

To predict the resources needed for manufacturing processes each month using historical values that are currently stored in an Amazon S3 bucket, a solutions architect should use Amazon SageMaker to train a model by using the historical data in the S3 bucket, and deploy an Amazon SageMaker model and create a SageMaker endpoint for inference. Amazon SageMaker is a fully managed service that provides an easy way to build, train, and deploy machine learning (ML) models. The solutions architect can use the built-in algorithms or frameworks provided by SageMaker, or bring their own custom code, to train a model using the

historical data in the S3 bucket as input. The trained model can then be deployed to a SageMaker endpoint, which is a scalable and secure web service that can handle requests for predictions from the application. The solutions architect does not need to have any ML experience or manage any infrastructure to use SageMaker.

NEW QUESTION 120

- (Topic 4)

An ecommerce company has noticed performance degradation of its Amazon RDS based web application. The performance degradation is attributed to an increase in the number of read-only SQL queries triggered by business analysts. A solutions architect needs to solve the problem with minimal changes to the existing web application.

What should the solutions architect recommend?

- A. Export the data to Amazon DynamoDB and have the business analysts run their queries.
- B. Load the data into Amazon ElastiCache and have the business analysts run their queries.
- C. Create a read replica of the primary database and have the business analysts run their queries.
- D. Copy the data into an Amazon Redshift cluster and have the business analysts run their queries

Answer: C

Explanation:

Creating a read replica of the primary RDS database will offload the read- only SQL queries from the primary database, which will help to improve the performance of the web application. Read replicas are exact copies of the primary database that can be used to handle read-only traffic, which will reduce the load on the primary database and improve the performance of the web application. This solution can be implemented with minimal changes to the existing web application, as the business analysts can continue to run their queries on the read replica without modifying the code.

NEW QUESTION 125

- (Topic 4)

A company's marketing data is uploaded from multiple sources to an Amazon S3 bucket A series of data preparation jobs aggregate the data for reporting The data preparation jobs need to run at regular intervals in parallel A few jobs need to run in a specific order later

The company wants to remove the operational overhead of job error handling retry logic, and state management

Which solution will meet these requirements?

- A. Use an AWS Lambda function to process the data as soon as the data is uploaded to the S3 bucket Invoke Other Lambda functions at regularly scheduled intervals
- B. Use Amazon Athena to process the data Use Amazon EventBridge Scheduler to invoke Athena on a regular interval
- C. Use AWS Glue DataBrew to process the data Use an AWS Step Functions state machine to run the DataBrew data preparation jobs
- D. Use AWS Data Pipeline to process the data
- E. Schedule Data Pipeline to process the data once at midnight.

Answer: C

Explanation:

AWS Glue DataBrew is a visual data preparation tool that allows you to easily clean, normalize, and transform your data without writing any code. You can create and run data preparation jobs on your data stored in Amazon S3, Amazon Redshift, or other data sources. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. You can use Step Functions to orchestrate your DataBrew jobs, define the order and parallelism of execution, handle errors and retries, and monitor the state of your workflow. By using AWS Glue DataBrew and AWS Step Functions, you can meet the requirements of the company with minimal operational overhead, as you do not need to write any code, manage any servers, or deal with complex dependencies.

References:

? AWS Glue DataBrew

? AWS Step Functions

? Orchestrate AWS Glue DataBrew jobs using AWS Step Functions

NEW QUESTION 128

- (Topic 4)

A company hosts an application on Amazon EC2 instances that run in a single Availability Zone. The application is accessible by using the transport layer of the Open Systems Interconnection (OSI) model. The company needs the application architecture to have high availability

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

- A. Configure new EC2 instances in a different Availability Zone
- B. Use Amazon Route 53 to route traffic to all instances.
- C. Configure a Network Load Balancer in front of the EC2 instances.
- D. Configure a Network Load Balancer for TCP traffic to the instance
- E. Configure an Application Load Balancer for HTTP and HTTPS traffic to the instances.
- F. Create an Auto Scaling group for the EC2 instance
- G. Configure the Auto Scaling group to use multiple Availability Zone
- H. Configure the Auto Scaling group to run application health checks on the instances_
- I. Create an Amazon CloudWatch alarm
- J. Configure the alarm to restart EC2 instances that transition to a stopped state

Answer: AD

Explanation:

To achieve high availability for an application that runs on EC2 instances, the application should be deployed across multiple Availability Zones and use a load balancer to distribute traffic. An Auto Scaling group can be used to launch and manage EC2 instances in multiple Availability Zones and perform health checks on them. A Network Load Balancer can be used to handle transport layer traffic to the EC2 instances. References:

? Auto Scaling Groups

? What Is a Network Load Balancer?

NEW QUESTION 131

- (Topic 4)

A company has multiple AWS accounts that use consolidated billing. The company runs several active high performance Amazon RDS for Oracle On-Demand DB instances for 90 days. The company's finance team has access to AWS Trusted Advisor in the consolidated billing account and all other AWS accounts. The finance team needs to use the appropriate AWS account to access the Trusted Advisor check recommendations for RDS. The finance team must review the appropriate Trusted Advisor check to reduce RDS costs.

Which combination of steps should the finance team take to meet these requirements? (Select TWO.)

- A. Use the Trusted Advisor recommendations from the account where the RDS instances are running.
- B. Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time.
- C. Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization.
- D. Review the Trusted Advisor check for Amazon RDS Idle DB Instances.
- E. Review the Trusted Advisor check for Amazon Redshift Reserved Node Optimization.

Answer: BC

Explanation:

B. Use the Trusted Advisor recommendations from the consolidated billing account to see all RDS instance checks at the same time. The consolidated billing account has access to all the other AWS accounts that use consolidated billing. Using the Trusted Advisor recommendations from the consolidated billing account will allow the finance team to see all RDS instance checks for all accounts at the same time.
 * C. Review the Trusted Advisor check for Amazon RDS Reserved Instance Optimization. The Trusted Advisor check for Amazon RDS Reserved Instance Optimization provides recommendations for purchasing reserved instances to reduce RDS costs. By reviewing this check, the finance team can identify which RDS instances can be converted to reserved instances to save costs.

NEW QUESTION 132

- (Topic 4)

A company has an AWS Direct Connect connection from its corporate data center to its VPC in the us-east-1 Region. The company recently acquired a corporation that has several VPCs and a Direct Connect connection between its on-premises data center and the eu-west-2 Region. The CIDR blocks for the VPCs of the company and the corporation do not overlap. The company requires connectivity between two Regions and the data centers. The company needs a solution that is scalable while reducing operational overhead.

What should a solutions architect do to meet these requirements?

- A. Set up inter-Region VPC peering between the VPC in us-east-1 and the VPCs in eu-west-2.
- B. Create private virtual interfaces from the Direct Connect connection in us-east-1 to the VPCs in eu-west-2.
- C. Establish VPN appliances in a fully meshed VPN network hosted by Amazon EC2. Use AWS VPN CloudHub to send and receive data between the data centers and each VPC.
- D. Connect the existing Direct Connect connection to a Direct Connect gateway
- E. Route traffic from the virtual private gateways of the VPCs in each Region to the Direct Connect gateway.

Answer: D

Explanation:

This solution meets the requirements because it allows the company to use a single Direct Connect connection to connect to multiple VPCs in different Regions using a Direct Connect gateway. A Direct Connect gateway is a globally available resource that enables you to connect your on-premises network to VPCs in any AWS Region, except the AWS China Regions. You can associate a Direct Connect gateway with a transit gateway or a virtual private gateway in each Region. By routing traffic from the virtual private gateways of the VPCs to the Direct Connect gateway, you can enable inter-Region and on-premises connectivity for your VPCs. This solution is scalable because you can add more VPCs in different Regions to the Direct Connect gateway without creating additional connections. This solution also reduces operational overhead because you do not need to manage multiple VPN appliances, VPN connections, or VPC peering connections.

References:

- ? Direct Connect gateways
- ? Inter-Region VPC peering

NEW QUESTION 137

- (Topic 4)

A company has an application that processes customer orders. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high: the workload does not process orders fast enough.

What should a solutions architect do to write the orders reliably to the database as quickly as possible?

- A. Increase the instance size of the EC2 instance when traffic is high
- B. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
- C. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue
- D. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database.
- E. Write orders to Amazon Simple Notification Service (Amazon SNS) Subscribe the database endpoint to the SNS topic Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- F. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limit
- G. Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database

Answer: B

Explanation:

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

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

NEW QUESTION 139

- (Topic 4)

A company moved its on-premises PostgreSQL database to an Amazon RDS for PostgreSQL DB instance. The company successfully launched a new product. The workload on the database has increased.

The company wants to accommodate the larger workload without adding infrastructure. Which solution will meet these requirements MOST cost-effectively?

- A. Buy reserved DB instances for the total workload
- B. Make the Amazon RDS for PostgreSQL DB instance larger.
- C. Make the Amazon RDS for PostgreSQL DB instance a Multi-AZ DB instance.
- D. Buy reserved DB instances for the total workload
- E. Add another Amazon RDS for PostgreSQL DB instance.
- F. Make the Amazon RDS for PostgreSQL DB instance an on-demand DB instance.

Answer: A

Explanation:

This answer is correct because it meets the requirements of accommodating the larger workload without adding infrastructure and minimizing the cost. Reserved DB instances are a billing discount applied to the use of certain on-demand DB instances in your account. Reserved DB instances provide you with a significant discount compared to on-demand DB instance pricing. You can buy reserved DB instances for the total workload and choose between three payment options: No Upfront, Partial Upfront, or All Upfront. You can make the Amazon RDS for PostgreSQL DB instance larger by modifying its instance type to a higher performance class. This way, you can increase the CPU, memory, and network capacity of your DB instance and handle the increased workload. References:

? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithReservedDBInstances.html

ReservedDBInstances.html

? <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.DBInstanceClass.html>

NEW QUESTION 144

- (Topic 4)

A solutions architect needs to optimize storage costs. The solutions architect must identify any Amazon S3 buckets that are no longer being accessed or are rarely accessed.

Which solution will accomplish this goal with the LEAST operational overhead?

- A. Analyze bucket access patterns by using the S3 Storage Lens dashboard for advanced activity metrics.
- B. Analyze bucket access patterns by using the S3 dashboard in the AWS Management Console.
- C. Turn on the Amazon CloudWatch BucketSizeBytes metric for bucket
- D. Analyze bucket access patterns by using the metrics data with Amazon Athena.
- E. Turn on AWS CloudTrail for S3 object monitoring
- F. Analyze bucket access patterns by using CloudTrail logs that are integrated with Amazon CloudWatch Logs.

Answer: A

Explanation:

S3 Storage Lens is a fully managed S3 storage analytics solution that provides a comprehensive view of object storage usage, activity trends, and recommendations to optimize costs. Storage Lens allows you to analyze object access patterns across all of your S3 buckets and generate detailed metrics and reports.

NEW QUESTION 145

- (Topic 4)

A company is building an ecommerce application and needs to store sensitive customer information. The company needs to give customers the ability to complete purchase transactions on the website. The company also needs to ensure that sensitive customer data is protected, even from database administrators.

Which solution meets these requirements?

- A. Store sensitive data in an Amazon Elastic Block Store (Amazon EBS) volume
- B. Use EBS encryption to encrypt the data
- C. Use an IAM instance role to restrict access.
- D. Store sensitive data in Amazon RDS for MySQL
- E. Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.
- F. Store sensitive data in Amazon S3. Use AWS Key Management Service (AWS KMS) server-side encryption to encrypt the data
- G. Use S3 bucket policies to restrict access.
- H. Store sensitive data in Amazon FSx for Windows Server
- I. Mount the file share on application server
- J. Use Windows file permissions to restrict access.

Answer: B

Explanation:

It allows the company to store sensitive customer information in a managed AWS service and give customers the ability to complete purchase transactions on the website. By using AWS Key Management Service (AWS KMS) client-side encryption, the company can encrypt the data before sending it to Amazon RDS for MySQL. This ensures that sensitive customer data is protected, even from database administrators, as only the application has access to the encryption keys.

References:

? [Using Encryption with Amazon RDS for MySQL](#)

? [Encrypting Amazon RDS Resources](#)

NEW QUESTION 150

- (Topic 4)

A company hosts its application in the AWS Cloud. The application runs on Amazon EC2 instances behind an Elastic Load Balancer in an Auto Scaling group and with an Amazon DynamoDB table. The company wants to ensure the application can be made available in another AWS Region with minimal downtime.

What should a solutions architect do to meet these requirements with the LEAST amount of downtime?

- A. Create an Auto Scaling group and a load balancer in the disaster recovery Region
- B. Configure the DynamoDB table as a global table
- C. Configure DNS failover to point to the new disaster recovery Region's load balancer.

- D. Create an AWS CloudFormation template to create EC2 instances, load balancers, and DynamoDB tables to be launched when needed
- E. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- F. Create an AWS CloudFormation template to create EC2 instances and a load balancer to be launched when needed
- G. Configure the DynamoDB table as a global table
- H. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- I. Create an Auto Scaling group and load balancer in the disaster recovery Region
- J. Configure the DynamoDB table as a global table
- K. Create an Amazon CloudWatch alarm to trigger an AWS Lambda function that updates Amazon Route 53 pointing to the disaster recovery load balancer.

Answer: A

Explanation:

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

References:

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

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

NEW QUESTION 152

- (Topic 4)

A company has stored 10 TB of log files in Apache Parquet format in an Amazon S3 bucket. The company occasionally needs to use SQL to analyze the log files. Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon Aurora MySQL database. Migrate the data from the S3 bucket into Aurora by using AWS Database Migration Service (AWS DMS). Issue SQL statements to the Aurora database.
- B. Create an Amazon Redshift cluster. Use Redshift Spectrum to run SQL statements directly on the data in the S3 bucket.
- C. Create an AWS Glue crawler to store and retrieve table metadata from the S3 bucket. Use Amazon Athena to run SQL statements directly on the data in the S3 bucket.
- D. Create an Amazon EMR cluster. Use Apache Spark SQL to run SQL statements directly on the data in the S3 bucket.

Answer: C

Explanation:

AWS Glue is a serverless data integration service that can crawl, catalog, and prepare data for analysis. AWS Glue can automatically discover the schema and partitioning of the data stored in Apache Parquet format in S3, and create a table in the AWS Glue Data Catalog. Amazon Athena is a serverless interactive query service that can run SQL queries directly on data in S3, without requiring any data loading or transformation. Athena can use the table metadata from the AWS Glue Data Catalog to query the data in S3. By using AWS Glue and Athena, you can analyze the log files in S3 most cost-effectively, as you only pay for the resources consumed by the crawler and the queries, and you do not need to provision or manage any servers or clusters.

References:

? AWS Glue

? Amazon Athena

? Analyzing Data in S3 using Amazon Athena

NEW QUESTION 156

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AWS-Solution-Architect-Associate Practice Exam Features:

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

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