

Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 4)

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

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

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

Answer: C

Explanation:

Lambda@Edge is a service that allows you to run Lambda functions at CloudFront edge locations. It can be used to modify requests and responses that flow through CloudFront. CloudFront origin request policy is a policy that controls the values (URL query strings, HTTP headers, and cookies) that are included in requests that CloudFront sends to the origin. It can be used to collect additional information at the origin or to customize the origin response. CloudFront response headers policy is a policy that specifies the HTTP headers that CloudFront removes or adds in responses that it sends to viewers. It can be used to add security or custom headers to responses.

Based on these definitions, the solution that will meet the requirements with the least operational overhead is:

* C. Use a Lambda@Edge function with an external image management library. Associate the Lambda@Edge function with the CloudFront behaviors that serve the images.

This solution would allow the application to use a Lambda@Edge function to resize the images dynamically and serve appropriate formats to clients based on the User-Agent HTTP header in the request. The Lambda@Edge function would run at the edge locations, reducing latency and load on the origin. The application code would only need to include an external image management library that can perform image manipulation tasks.

NEW QUESTION 2

- (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 overhead. 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 lifecycle. 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 costs. 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 needed. 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 needed. 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 3

- (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 4

- (Topic 4)

A company has deployed a multiplayer game for mobile devices. The game requires live location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations. The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly. What should a solutions architect do to improve the performance of the data tier?

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 5

- (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 into. 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 6

- (Topic 4)

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

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

Answer: B

Explanation:

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

- ? NAT Gateways
- ? NAT Gateway Pricing

NEW QUESTION 7

- (Topic 4)

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

Which solution will allow the node to join the cluster?

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

Answer: B

Explanation:

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

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

NEW QUESTION 8

- (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 9

- (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 10

- (Topic 4)

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

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

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

Answer: C

Explanation:

To resize images dynamically and serve appropriate formats to clients, a Lambda@Edge function with an external image management library can be used.

Lambda@Edge allows running custom code at the edge locations of CloudFront, which can process the images on the fly and optimize them for different devices and browsers. An external image management library can provide various image manipulation and optimization features. References:

? Lambda@Edge

? Resizing Images with Amazon CloudFront & Lambda@Edge

NEW QUESTION 11

- (Topic 4)

A gaming company uses Amazon DynamoDB to store user information such as geographic location, player data, and leaderboards. The company needs to configure continuous backups to an Amazon S3 bucket with a minimal amount of coding. The backups must not affect availability of the application and must not affect the read capacity units (RCUs) that are defined for the table. Which solution meets these requirements?

- A. Use an Amazon EMR cluster
- B. Create an Apache Hive job to back up the data to Amazon S3.
- C. Export the data directly from DynamoDB to Amazon S3 with continuous backup
- D. Turn on point-in-time recovery for the table.
- E. Configure Amazon DynamoDB Stream
- F. Create an AWS Lambda function to consume the stream and export the data to an Amazon S3 bucket.
- G. Create an AWS Lambda function to export the data from the database tables to Amazon S3 on a regular basis
- H. Turn on point-in-time recovery for the table.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/>
<https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/>

NEW QUESTION 12

- (Topic 4)

A company wants to use Amazon Elastic Container Service (Amazon ECS) clusters and Amazon RDS DB instances to build and run a payment processing application. The company will run the application in its on-premises data center for compliance purposes.

A solutions architect wants to use AWS Outposts as part of the solution. The solutions architect is working with the company's operational team to build the application.

Which activities are the responsibility of the company's operational team? (Select THREE.)

- A. Providing resilient power and network connectivity to the Outposts racks
- B. Managing the virtualization hypervisor, storage systems, and the AWS services that run on Outposts
- C. Physical security and access controls of the data center environment
- D. Availability of the Outposts infrastructure including the power supplies, servers, and networking equipment within the Outposts racks
- E. Physical maintenance of Outposts components
- F. Providing extra capacity for Amazon ECS clusters to mitigate server failures and maintenance events

Answer: ACF

Explanation:

These answers are correct because they reflect the customer's responsibilities for using AWS Outposts as part of the solution. According to the AWS shared responsibility model, the customer is responsible for providing resilient power and network connectivity to the Outposts racks, ensuring physical security and access controls of the data center environment, and providing extra capacity for Amazon ECS clusters to mitigate server failures and maintenance events. AWS is responsible for managing the virtualization hypervisor, storage systems, and the AWS services that run on Outposts, as well as the availability of the Outposts infrastructure including the power supplies, servers, and networking equipment within the Outposts racks, and the physical maintenance of Outposts components.

References:

? <https://docs.aws.amazon.com/outposts/latest/userguide/what-is-outposts.html>
? <https://www.contino.io/insights/the-sandwich-responsibility-model-aws-outposts/>

NEW QUESTION 13

- (Topic 4)

A company uses an on-premises network-attached storage (NAS) system to provide file shares to its high performance computing (HPC) workloads. The company wants to migrate its latency-sensitive HPC workloads and its storage to the AWS Cloud. The company must be able to provide NFS and SMB multi-protocol access from the file system.

Which solution will meet these requirements with the LEAST latency? (Select TWO.)

- A. Deploy compute optimized EC2 instances into a cluster placement group.
- B. Deploy compute optimized EC2 instances into a partition placement group.
- C. Attach the EC2 instances to an Amazon FSx for Lustre file system.
- D. Attach the EC2 instances to an Amazon FSx for OpenZFS file system.
- E. Attach the EC2 instances to an Amazon FSx for NetApp ONTAP file system.

Answer: AE

Explanation:

A cluster placement group is a logical grouping of EC2 instances within a single Availability Zone that are placed close together to minimize network latency. This is suitable for latency-sensitive HPC workloads that require high network performance. A compute optimized EC2 instance is an instance type that has a high ratio of vCPUs to memory, which is ideal for compute-intensive applications. Amazon FSx for NetApp ONTAP is a fully managed service that provides NFS and SMB multi-protocol access from the file system, as well as features such as data deduplication, compression, thin provisioning, and snapshots. This solution will meet the requirements with the least latency,

as it leverages the low-latency network and storage performance of AWS.

References:

? 1 explains how cluster placement groups work and their benefits.
? 2 describes the characteristics and use cases of compute optimized EC2 instances.
? 3 provides an overview of Amazon FSx for NetApp ONTAP and its features.

NEW QUESTION 14

- (Topic 4)

A pharmaceutical company is developing a new drug. The volume of data that the

company generates has grown exponentially over the past few months. The company's researchers regularly require a subset of the entire dataset to be immediately available with minimal lag. However the entire dataset does not need to be accessed on a daily basis. All the data currently resides in on-premises

storage arrays, and the company wants to reduce ongoing capital expenses.
 Which storage solution should a solutions architect recommend to meet these requirements?

- A. Run AWS DataSync as a scheduled cron job to migrate the data to an Amazon S3 bucket on an ongoing basis.
- B. Deploy an AWS Storage Gateway file gateway with an Amazon S3 bucket as the target storage. Migrate the data to the Storage Gateway appliance.
- C. Deploy an AWS Storage Gateway volume gateway with cached volumes with an Amazon S3 bucket as the target storage.
- D. Migrate the data to the Storage Gateway appliance.
- E. Configure an AWS Site-to-Site VPN connection from the on-premises environment to AWS.
- F. Migrate data to an Amazon Elastic File System (Amazon EFS) file system.

Answer: C

Explanation:

AWS Storage Gateway is a hybrid cloud storage service that allows you to seamlessly integrate your on-premises applications with AWS cloud storage. Volume Gateway is a type of Storage Gateway that presents cloud-backed iSCSI block storage volumes to your on-premises applications. Volume Gateway operates in either cache mode or stored mode. In cache mode, your primary data is stored in Amazon S3, while retaining your frequently accessed data locally in the cache for low latency access. In stored mode, your primary data is stored locally and your entire dataset is available for low latency access on premises while also asynchronously getting backed up to Amazon S3.

For the pharmaceutical company's use case, cache mode is the most suitable option, as it meets the following requirements:

- ? It reduces the need to scale the on-premises storage infrastructure, as most of the data is stored in Amazon S3, which is scalable, durable, and cost-effective.
- ? It provides low latency access to the subset of the data that the researchers regularly require, as it is cached locally in the Storage Gateway appliance.
- ? It does not require the entire dataset to be accessed on a daily basis, as it is stored in Amazon S3 and can be retrieved on demand.
- ? It offers flexible data protection and recovery options, as it allows taking point-in-time copies of the volumes using AWS Backup, which are stored in AWS as Amazon EBS snapshots.

Therefore, the solutions architect should recommend deploying an AWS Storage Gateway volume gateway with cached volumes with an Amazon S3 bucket as the target storage and migrating the data to the Storage Gateway appliance.

References:

- ? Volume Gateway | Amazon Web Services
- ? How Volume Gateway works (architecture) - AWS Storage Gateway
- ? AWS Storage Volume Gateway - Cached volumes - Stack Overflow

NEW QUESTION 15

- (Topic 3)

A company hosts a frontend application that uses an Amazon API Gateway API backend that is integrated with AWS Lambda. When the API receives requests, the Lambda function loads many libraries. Then the Lambda function connects to an Amazon RDS database, processes the data, and returns the data to the frontend application. The company wants to ensure that response latency is as low as possible for all its users with the fewest number of changes to the company's operations.

Which solution will meet these requirements?

- A. Establish a connection between the frontend application and the database to make queries faster by bypassing the API.
- B. Configure provisioned concurrency for the Lambda function that handles the requests.
- C. Cache the results of the queries in Amazon S3 for faster retrieval of similar datasets.
- D. Increase the size of the database to increase the number of connections Lambda can establish at one time.

Answer: B

Explanation:

Configure provisioned concurrency for the Lambda function that handles the requests. Provisioned concurrency allows you to set the amount of compute resources that are available to the Lambda function, so that it can handle more requests at once and reduce latency. Caching the results of the queries in Amazon S3 could also help to reduce latency, but it would not be as effective as setting up provisioned concurrency. Increasing the size of the database would not help to reduce latency, as this would not increase the number of connections the Lambda function could establish, and establishing a direct connection between the frontend application and the database would bypass the API, which would not be the best solution either.

Using AWS Lambda with Amazon API Gateway - AWS Lambda <https://docs.aws.amazon.com/lambda/latest/dg/services-apigateway.html>

AWS Lambda FAQs <https://aws.amazon.com/lambda/faqs/>

NEW QUESTION 16

- (Topic 3)

A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch runs. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application.

What should a solution architect recommend to ensure the application is able to handle the workload and avoid downtime?

- A. Configure an Amazon CloudFront distribution in front of the ALB.
- B. Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
- C. Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
- D. Configure Amazon ElasticCache to remove some of the workload from the EC2 instances.

Answer: C

Explanation:

Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule is the best option because it allows for the proactive scaling of the EC2 instances before the monthly batch run begins. This will ensure that the application is able to handle the increased workload without experiencing downtime. The scheduled scaling policy can be configured to increase the number of instances in the Auto Scaling group a few hours before the batch run and then decrease the number of instances after the batch run is complete. This will ensure that the resources are available when needed and not wasted when not needed. The most appropriate solution to handle the increased workload during the monthly batch run and avoid downtime would be to configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule. <https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-scheduled-scaling.html>

NEW QUESTION 17

- (Topic 3)

A company wants to use Amazon S3 for the secondary copy of its on-premises dataset. The company would rarely need to access this copy. The storage solution's cost should be minimal.
Which storage solution meets these requirements?

- A. S3 Standard
- B. S3 Intelligent-Tiering
- C. S3 Standard-Infrequent Access (S3 Standard-IA)
- D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: D

Explanation:

S3 One Zone-IA is a storage class that is designed for data that is accessed less frequently, but requires rapid access when needed. Unlike other S3 Storage Classes which store data in a minimum of three Availability Zones (AZs), S3 One Zone-IA stores data in a single AZ and costs 20% less than S3 Standard-IA. This storage class meets the requirements of the company because it provides a low-cost solution for the secondary copy of its on-premises dataset that would rarely need to be accessed. The other storage classes are either more expensive or not suitable for infrequently accessed data.
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html>

NEW QUESTION 18

- (Topic 3)

A company is moving its data management application to AWS. The company wants to transition to an event-driven architecture. The architecture needs to be more distributed and to use serverless concepts when performing the different aspects of the workflow. The company also wants to minimize operational overhead. Which solution will meet these requirements?

- A. Build out the workflow in AWS Glue Use AWS Glue to invoke AWS Lambda functions to process the workflow steps
- B. Build out the workflow in AWS Step Functions Deploy the application on Amazon EC2 Instances Use Step Functions to invoke the workflow steps on the EC2 instances
- C. Build out the workflow in Amazon EventBridge
- D. Use EventBridge to invoke AWS Lambda functions on a schedule to process the workflow steps.
- E. Build out the workflow in AWS Step Functions Use Step Functions to create a state machine Use the state machine to invoke AWS Lambda functions to process the workflow steps

Answer: D

Explanation:

This answer is correct because it meets the requirements of transitioning to an event-driven architecture, using serverless concepts, and minimizing operational overhead. AWS Step Functions is a serverless service that lets you coordinate multiple AWS services into workflows using state machines. State machines are composed of tasks and transitions that define the logic and order of execution of the workflow steps. AWS Lambda is a serverless function-as-a-service platform that lets you run code without provisioning or managing servers. Lambda functions can be invoked by Step Functions as tasks in a state machine, and can perform different aspects of the data management workflow, such as data ingestion, transformation, validation, and analysis. By using Step Functions and Lambda, the company can benefit from the following advantages:

? Event-driven: Step Functions can trigger Lambda functions based on events, such as timers, API calls, or other AWS service events. Lambda functions can also emit events to other services or state machines, creating an event-driven architecture.

? Serverless: Step Functions and Lambda are fully managed by AWS, so the company does not need to provision or manage any servers or infrastructure. The company only pays for the resources consumed by the workflows and functions, and can scale up or down automatically based on demand.

? Operational overhead: Step Functions and Lambda simplify the development and deployment of workflows and functions, as they provide built-in features such as monitoring, logging, tracing, error handling, retry logic, and security. The company can focus on the business logic and data processing rather than the operational details.

References:

? What is AWS Step Functions?

? What is AWS Lambda?

NEW QUESTION 19

- (Topic 3)

A company has launched an Amazon RDS for MySQL D6 instance. Most of the connections to the database come from serverless applications. Application traffic to the database changes significantly at random intervals. At times of high demand, users report that their applications experience database connection rejection errors.

Which solution will resolve this issue with the LEAST operational overhead?

- A. Create a proxy in RDS Proxy Configure the users' applications to use the DB instance through RDS Proxy
- B. Deploy Amazon ElastiCache for Memcached between the users' application and the DB instance
- C. Migrate the DB instance to a different instance class that has higher I/O capacity
- D. Configure the users' applications to use the new DB instance.
- E. Configure Multi-AZ for the DB instance Configure the users' application to switch between the DB instances.

Answer: A

Explanation:

Many applications, including those built on modern serverless architectures, can have a large number of open connections to the database server and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. (<https://aws.amazon.com/pt/rds/proxy/>)

NEW QUESTION 20

- (Topic 3)

A solutions architect is designing a two-tiered architecture that includes a public subnet and a database subnet. The web servers in the public subnet must be open to the internet on port 443. The Amazon RDS for MySQL D6 instance in the database subnet must be accessible only to the web servers on port 3306.

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

- A. Create a network ACL for the public subnet Add a rule to deny outbound traffic to 0.0.0.0/0 on port 3306
- B. Create a security group for the DB instance Add a rule to allow traffic from the public subnet CIDR block on port 3306

- C. Create a security group for the web servers in the public subnet Add a rule to allow traffic from 0 0 0 0 on port 443
- D. Create a security group for the DB instance Add a rule to allow traffic from the web servers' security group on port 3306
- E. Create a security group for the DB instance Add a rule to deny all traffic except traffic from the web servers' security group on port 3306

Answer: BC

Explanation:

Security groups are virtual firewalls that protect AWS instances and can be applied to EC2, ELB and RDS1. Security groups have rules for inbound and outbound traffic and are stateful, meaning that responses to allowed inbound traffic are allowed to flow out of the instance2. Network ACLs are different from security groups in several ways. They cover entire subnets, not individual instances, and are stateless, meaning that they require rules for both inbound and outbound traffic2. Network ACLs also support deny rules, while security groups only support allow rules2.

To meet the requirements of the scenario, the solutions architect should create two security groups: one for the DB instance and one for the web servers in the public subnet. The security group for the DB instance should allow traffic from the public subnet CIDR block on port 3306, which is the default port for MySQL3. This way, only the web servers in the public subnet can access the DB instance on that port. The security group for the web servers should allow traffic from 0 0 0 0 on port 443, which is the default port for HTTPS 4. This way, the web servers can accept secure connections from the internet on that port.

NEW QUESTION 21

- (Topic 3)

A company wants to deploy a new public web application on AWS The application includes a web server tier that uses Amazon EC2 instances The application also includes a database tier that uses an Amazon RDS for MySQL DB instance

The application must be secure and accessible for global customers that have dynamic IP addresses

How should a solutions architect configure the security groups to meet these requirements'?

- A. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0. 0/0) Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers
- B. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers Configure the security group for the DB instance to allow inbound traffic on port 3306 from the security group of the web servers
- C. Configure the security group for the web servers to allow inbound traffic on port 443 from the IP addresses of the customers Configure the security group for the DB instance to allow inbound traffic on port 3306 from the IP addresses of the customers
- D. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0 Configure the security group for the DB instance to allow inbound traffic on port 3306 from 0.0.0.0/0)

Answer: A

Explanation:

? Restricting inbound access to the web servers to only port 443, which is used for HTTPS traffic, and allowing access from any IP address (0.0.0.0/0), since the application is public and accessible for global customers.

? Restricting inbound access to the DB instance to only port 3306, which is used for MySQL traffic, and allowing access only from the security group of the web servers, which creates a secure connection between the two tiers and prevents unauthorized access to the database.

? Restricting outbound access to the minimum required for both tiers, which is not specified in the question but can be assumed to be similar to the inbound rules.

References:

? Security groups - Amazon Virtual Private Cloud

? 5 Best Practices for AWS Security Groups - DZone

NEW QUESTION 22

- (Topic 2)

A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer The application stores data in Amazon Aurora. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy

What should a solutions architect do to meet these requirements?

- A. Deploy the application with the required infrastructure elements in place Use Amazon Route 53 to configure active-passive failover Create an Aurora Replica in a second AWS Region
- B. Host a scaled-down deployment of the application in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora Replica in the second Region
- C. Replicate the primary infrastructure in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora database that is restored from the latest snapshot
- D. Back up data with AWS Backup Use the backup to create the required infrastructure in a second AWS Region Use Amazon Route 53 to configure active-passive failover Create an Aurora second primary instance in the second Region

Answer: A

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html>

NEW QUESTION 23

- (Topic 2)

A company uses AWS Organizations to create dedicated AWS accounts for each business unit to manage each business unit's account independently upon request. The root email recipient missed a notification that was sent to the root user email address of one account. The company wants to ensure that all future notifications are not missed. Future notifications must be limited to account administrators.

Which solution will meet these requirements?

- A. Configure the company's email server to forward notification email messages that are sent to the AWS account root user email address to all users in the organization.
- B. Configure all AWS account root user email addresses as distribution lists that go to a few administrators who can respond to alert
- C. Configure AWS account alternate contacts in the AWS Organizations console or programmatically.
- D. Configure all AWS account root user email messages to be sent to one administrator who is responsible for monitoring alerts and forwarding those alerts to the appropriate groups.
- E. Configure all existing AWS accounts and all newly created accounts to use the same root user email address
- F. Configure AWS account alternate contacts in the AWS Organizations console or programmatically.

Answer: B

Explanation:

Use a group email address for the management account's root user https://docs.aws.amazon.com/organizations/latest/userguide/orgs_best-practices_mgmt-acct.html#best-practices_mgmt-acct_email-address

NEW QUESTION 24

- (Topic 2)

A company wants to move its application to a serverless solution. The serverless solution needs to analyze existing and new data by using SL. The company stores the data in an Amazon S3 bucket. The data requires encryption and must be replicated to a different AWS Region. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new S3 bucket
- B. Load the data into the new S3 bucket
- C. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- D. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon Athena to query the data.
- E. Create a new S3 bucket
- F. Load the data into the new S3 bucket
- G. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- H. Use server-side encryption with AWS KMS multi-Region keys (SSE-KMS). Use Amazon RDS to query the data.
- I. Load the data into the existing S3 bucket
- J. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- K. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon Athena to query the data.
- L. Load the data into the existing S3 bucket
- M. Use S3 Cross-Region Replication (CRR) to replicate encrypted objects to an S3 bucket in another Region
- N. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use Amazon RDS to query the data.

Answer: A

Explanation:

This solution meets the requirements of a serverless solution, encryption, replication, and SQL analysis with the least operational overhead. Amazon Athena is a serverless interactive query service that can analyze data in S3 using standard SQL. S3 Cross-Region Replication (CRR) can replicate encrypted objects to an S3 bucket in another Region automatically. Server-side encryption with AWS KMS multi-Region keys (SSE-KMS) can encrypt the data at rest using keys that are replicated across multiple Regions. Creating a new S3 bucket can avoid potential conflicts with existing data or configurations.

Option B is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. Option C is incorrect because server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not use KMS keys and it does not support multi-Region replication. Option D is incorrect because Amazon RDS is not a serverless solution and it cannot query data in S3 directly. It is also incorrect for the same reason as option C. References:

- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-4.html>
- ? <https://aws.amazon.com/blogs/storage/considering-four-different-replication-options-for-data-in-amazon-s3/>
- ? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingEncryption.html>
- ? <https://aws.amazon.com/athena/>

NEW QUESTION 25

- (Topic 2)

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).

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

- A. Use Amazon Athena for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- B. Use Amazon Kinesis Data Analytics for one-time queries Use Amazon QuickSight to create dashboards for KPIs
- C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster
- D. Use an AWS Glue extract transform, and load (ETL) job to convert the data into JSON format Load the data into multiple Amazon OpenSearch Service (Amazon Elasticsearch Service) clusters
- E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake Use AWS Glue to crawl the source extract the data and load the data into Amazon S3 in Apache Parquet format

Answer: AE

Explanation:

Amazon Athena is the best choice for running one-time queries on streaming data. Although Amazon Kinesis Data Analytics provides an easy and familiar standard SQL language to analyze streaming data in real-time, it is designed for continuous queries rather than one-time queries[1]. On the other hand, Amazon Athena is a serverless interactive query service that allows querying data in Amazon S3 using SQL. It is optimized for ad-hoc querying and is ideal for running one-time queries on streaming data[2]. AWS Lake Formation uses as a central place to have all your data for analytics purposes (E). Athena integrates perfectly with S3 and can make queries (A).

NEW QUESTION 26

- (Topic 1)

A solutions architect is designing a two-tier web application. The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets. The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet. Security is a high priority for the company. How should security groups be configured in this situation? (Select TWO)

- A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
- B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
- C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the security group for the web tier.
- D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
- E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

Answer: AC

Explanation:

"Security groups create an outbound rule for every inbound rule." Not completely right. Stateful does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request in the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those incoming requests.

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules

NEW QUESTION 27

- (Topic 1)

A company wants to improve its ability to clone large amounts of production data into a test environment in the same AWS Region. The data is stored in Amazon EC2 instances on Amazon Elastic Block Store (Amazon EBS) volumes. Modifications to the cloned data must not affect the production environment. The software that accesses this data requires consistently high I/O performance.

A solutions architect needs to minimize the time that is required to clone the production data into the test environment.

Which solution will meet these requirements?

- A. Take EBS snapshots of the production EBS volume
- B. Restore the snapshots onto EC2 instance store volumes in the test environment.
- C. Configure the production EBS volumes to use the EBS Multi-Attach feature
- D. Take EBS snapshots of the production EBS volume
- E. Attach the production EBS volumes to the EC2 instances in the test environment.
- F. Take EBS snapshots of the production EBS volume
- G. Create and initialize new EBS volume
- H. Attach the new EBS volumes to EC2 instances in the test environment before restoring the volumes from the production EBS snapshots.
- I. Take EBS snapshots of the production EBS volume
- J. Turn on the EBS fast snapshot restore feature on the EBS snapshot
- K. Restore the snapshots into new EBS volume
- L. Attach the new EBS volumes to EC2 instances in the test environment.

Answer: C

Explanation:

To clone the production data into the test environment with high I/O performance and without affecting the production environment, the best option is to take EBS snapshots of the production EBS volumes and restore them onto new EBS volumes in the test environment. Then, attach the new EBS volumes to EC2 instances in the test environment. This option minimizes the time required to clone the data and ensures that modifications to the cloned data do not affect the production environment. Therefore, option C is the correct answer.

Reference: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-restoring-volume.html>

NEW QUESTION 28

- (Topic 1)

A company is running a popular social media website. The website gives users the ability to upload images to share with other users. The company wants to make sure that the images do not contain inappropriate content. The company needs a solution that minimizes development effort.

What should a solutions architect do to meet these requirements?

- A. Use Amazon Comprehend to detect inappropriate content
- B. Use human review for low-confidence predictions.
- C. Use Amazon Rekognition to detect inappropriate content
- D. Use human review for low-confidence predictions.
- E. Use Amazon SageMaker to detect inappropriate content
- F. Use ground truth to label low-confidence predictions.
- G. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate content
- H. Use ground truth to label low-confidence predictions.

Answer: B

Explanation:

<https://docs.aws.amazon.com/rekognition/latest/dg/moderation.html?pg=ln&sec=ft> <https://docs.aws.amazon.com/rekognition/latest/dg/a2i-rekognition.html>

NEW QUESTION 29

- (Topic 4)

A company is designing a containerized application that will use Amazon Elastic Container Service (Amazon ECS). The application needs to access a shared file system that is highly durable and can recover data to another AWS Region with a recovery point objective (RPO) of 8 hours. The file system needs to provide a mount target in each Availability Zone within a Region.

A solutions architect wants to use AWS Backup to manage the replication to another Region.

Which solution will meet these requirements?

- A. 'Amazon FSx for Windows File Server with a Multi-AZ deployment
- B. Amazon FSx for NetApp ONTAP with a Multi-AZ deployment
- C. 'Amazon Elastic File System (Amazon EFS) with the Standard storage class
- D. Amazon FSx for OpenZFS

Answer: B

Explanation:

This answer is correct because it meets the requirements of accessing a shared file system that is highly durable, can recover data to another AWS Region, and can provide a mount target in each Availability Zone within a Region. Amazon FSx for NetApp ONTAP is a fully managed service that provides enterprise-grade data management and storage for your Windows and Linux applications. You can use Amazon FSx for NetApp ONTAP to create file systems that span multiple Availability Zones within an AWS Region, providing high availability and durability. You can also use AWS Backup to manage the replication of your file systems to

another AWS Region, with a recovery point objective (RPO) of 8 hours or less. AWS Backup is a fully managed backup service that automates and centralizes backup of data over AWS services. You can use AWS Backup to create backup policies and monitor activity for your AWS resources in one place.

References:

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

? <https://docs.aws.amazon.com/aws-backup/latest/devguide/whatisbackup.html>

NEW QUESTION 30

- (Topic 4)

A company is developing a mobile gaming app in a single AWS Region. The app runs on multiple Amazon EC2 instances in an Auto Scaling group. The company stores the app data in Amazon DynamoDB. The app communicates by using TCP traffic and UDP traffic between the users and the servers. The application will be used globally. The company wants to ensure the lowest possible latency for all users.

Which solution will meet these requirements?

- A. Use AWS Global Accelerator to create an accelerator
- B. Create an Application Load Balancer (ALB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP port
- C. Update the Auto Scaling group to register instances on the ALB.
- D. Use AWS Global Accelerator to create an accelerator
- E. Create a Network Load Balancer (NLB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP port
- F. Update the Auto Scaling group to register instances on the NLB
- G. Create an Amazon CloudFront content delivery network (CDN) endpoint
- H. Create a Network Load Balancer (NLB) behind the endpoint and listening on the TCP and UDP port
- I. Update the Auto Scaling group to register instances on the NL
- J. Update CloudFront to use the NLB as the origin.
- K. Create an Amazon Cloudfront content delivery network (CDN) endpoint
- L. Create an Application Load Balancer (ALB) behind the endpoint and listening on the TCP and UDP port
- M. Update the Auto Scaling group to register instances on the AL
- N. Update CloudFront to use the ALB as the origin

Answer: B

Explanation:

AWS Global Accelerator is a networking service that improves the performance and availability of applications for global users. It uses the AWS global network to route user traffic to the optimal endpoint based on performance and health. It also provides static IP addresses that act as a fixed entry point to the applications and support both TCP and UDP protocols¹. By using AWS Global Accelerator, the solution can ensure the lowest possible latency for all users.

* A. Use AWS Global Accelerator to create an accelerator. Create an Application Load Balancer (ALB) behind an accelerator endpoint that uses Global Accelerator integration and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB. This solution will not work, as ALB does not support UDP protocol².

* C. Create an Amazon CloudFront content delivery network (CDN) endpoint. Create a Network Load Balancer (NLB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the NLB. Update CloudFront to use the NLB as the origin. This solution will not work, as CloudFront does not support UDP protocol³.

* D. Create an Amazon Cloudfront content delivery network (CDN) endpoint. Create an Application Load Balancer (ALB) behind the endpoint and listening on the TCP and UDP ports. Update the Auto Scaling group to register instances on the ALB. Update CloudFront to use the ALB as the origin. This solution will not work, as CloudFront and ALB do not support UDP protocol³.

Reference URL: <https://aws.amazon.com/global-accelerator/>

NEW QUESTION 31

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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](#)