



Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

NEW QUESTION 1

- (Topic 4)

A company has an on-premises MySQL database that handles transactional data. The company is migrating the database to the AWS Cloud. The migrated database must maintain compatibility with the company's applications that use the database. The migrated database also must scale automatically during periods of increased demand.

Which migration solution will meet these requirements?

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

Answer: C

Explanation:

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

? [What Is Amazon Aurora?](#)

? [Using Amazon Aurora Auto Scaling with Aurora Replicas](#)

? [What Is AWS Database Migration Service?](#)

NEW QUESTION 2

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 3

- (Topic 4)

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

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

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

Answer: C

Explanation:

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

References:

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

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

NEW QUESTION 4

- (Topic 4)

A company is moving its data and applications to AWS during a multiyear migration project. The company wants to securely access data on Amazon S3 from the company's AWS Region and from the company's on-premises location. The data must not traverse the internet. The company has established an AWS Direct Connect connection between its Region and its on-premises location

Which solution will meet these requirements?

- A. Create gateway endpoints for Amazon S3. Use the gateway endpoints to securely access the data from the Region and the on-premises location.
- B. Create a gateway in AWS Transit Gateway to access Amazon S3 securely from the Region and the on-premises location.
- C. Create interface endpoints for Amazon S3. Use the interface endpoints to securely access the data from the Region and the on-premises location.
- D. Use an AWS Key Management Service (AWS KMS) key to access the data securely from the Region and the on-premises location.

Answer: B

Explanation:

A gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service¹. Amazon S3 does not support gateway endpoints, only interface endpoints². Therefore, option A is incorrect.

An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported service¹. An interface endpoint can provide secure access to Amazon S3 from within the Region, but not from the on-premises location. Therefore, option C is incorrect.

AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data³. AWS KMS does not provide a way to access data on Amazon S3 without traversing the internet. Therefore, option D is incorrect. AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. You can create a gateway in AWS Transit Gateway to access Amazon S3 securely from both the Region and the on-premises location using AWS Direct Connect. Therefore, option B is correct.

NEW QUESTION 5

- (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 instances
- 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 6

- (Topic 4)

A company needs to configure a real-time data ingestion architecture for its application. The company needs an API, a process that transforms data as the data is streamed, and a storage solution for the data.

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

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

Answer: C

Explanation:

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

NEW QUESTION 7

- (Topic 4)

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

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

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

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

Answer: AD

Explanation:

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

NEW QUESTION 8

- (Topic 4)

A company uses on-premises servers to host its applications. The company is running out of storage capacity. The applications use both block storage and NFS storage. The company needs a high-performing solution that supports local caching without re-architecting its existing applications.

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

- A. Mount Amazon S3 as a file system to the on-premises servers.
- B. Deploy an AWS Storage Gateway file gateway to replace NFS storage.
- C. Deploy AWS Snowball Edge to provision NFS mounts to on-premises servers.
- D. Deploy an AWS Storage Gateway volume gateway to replace the block storage
- E. Deploy Amazon Elastic File System (Amazon EFS) volumes and mount them to on- premises servers.

Answer: BD

Explanation:

<https://aws.amazon.com/storagegateway/file/>

File Gateway provides a seamless way to connect to the cloud in order to store application data files and backup images as durable objects in Amazon S3 cloud storage. File Gateway offers SMB or NFS-based access to data in Amazon S3 with local caching. It can be used for on-premises applications, and for Amazon EC2-based applications that need file protocol access to S3 object storage.

<https://aws.amazon.com/storagegateway/volume/>

Volume Gateway presents cloud-backed iSCSI block storage volumes to your on-premises applications. Volume Gateway stores and manages on-premises data in Amazon S3 on your behalf and operates in either cache mode or stored mode. In the cached Volume Gateway mode, your primary data is stored in Amazon S3, while retaining your frequently accessed data locally in the cache for low latency access.

NEW QUESTION 9

- (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 10

- (Topic 4)

A company wants to securely exchange data between its software as a service (SaaS) application Salesforce account and Amazon S3. The company must encrypt the data at rest by using AWS Key Management Service (AWS KMS) customer managed keys (CMKs). The company must also encrypt the data in transit. The company has enabled API access for the Salesforce account.

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

- A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3.

- B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3.
- C. Create Amazon AppFlow flows to transfer the data securely from Salesforce to Amazon S3.
- D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3.

Answer: C

Explanation:

Amazon AppFlow is a fully managed integration service that enables users to transfer data securely between SaaS applications and AWS services. It supports Salesforce as a source and Amazon S3 as a destination. It also supports encryption of data at rest using AWS KMS CMKs and encryption of data in transit using SSL/TLS1. By using Amazon AppFlow, the solution can meet the requirements with the least development effort.

- * A. Create AWS Lambda functions to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves writing custom code to interact with Salesforce and Amazon S3 APIs, handle authentication, encryption, error handling, and monitoring2.
- * B. Create an AWS Step Functions workflow Define the task to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves creating a state machine definition to orchestrate the data transfer task, and invoking Lambda functions or other services to perform the actual data transfer3.
- * D. Create a custom connector for Salesforce to transfer the data securely from Salesforce to Amazon S3. This solution will not meet the requirement of the least development effort, as it involves using the Amazon AppFlow Custom Connector SDK to build and deploy a custom connector for Salesforce, which requires additional configuration and management. Reference URL: <https://aws.amazon.com/appflow/>

NEW QUESTION 10

- (Topic 4)

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

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

Answer: A

Explanation:

A Network Load Balancer can be assigned one Elastic IP address for each Availability Zone it uses1. This allows the clients to reach the load balancer using a static IP address that can be authorized on their firewalls. An Application Load Balancer cannot be assigned an Elastic IP address2. An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address would not work because the load balancer would still use its own IP address as the source of the forwarded requests to the web service. An EC2 instance with a public IP address running as a proxy in front of the load balancer would add unnecessary complexity and cost, and would not provide the same scalability and availability as a Network Load Balancer. References: 1: Network Load Balancers - Elastic Load Balancing3, IP address type section2: How to assign Elastic IP to Application Load Balancer in AWS?4, answer section.

NEW QUESTION 15

- (Topic 4)

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

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

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

A solutions architect must ensure that the order collection process and the order fulfillment process can both scale properly during peak traffic hours. The solution must optimize utilization of the company's AWS resources. Which solution meets these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 19

- (Topic 4)

A company has established a new AWS account. The account is newly provisioned and no changes have been made to the default settings. The company is concerned about the security of the AWS account root user. What should be done to secure the root user?

- A. Create 1AM users for daily administrative task
- B. Disable the root user.
- C. Create 1AM users for daily administrative task

- D. Enable multi-factor authentication on the root user.
- E. Generate an access key for the root user Use the access key for daily administration tasks instead of the AWS Management Console.
- F. Provide the root user credentials to the most senior solutions architect
- G. Have the solutions architect use the root user for daily administration tasks.

Answer: B

Explanation:

This answer is the most secure and recommended option for securing the root user of a new AWS account. The root user is the identity that has complete access to all AWS services and resources in the account. It is accessed by signing in with the email address and password that were used to create the account. To protect the root user credentials from unauthorized use, AWS advises the following best practices:

? Create IAM users for daily administrative tasks. IAM users are identities that you create in your account that have specific permissions to access AWS resources. You can create individual IAM users for yourself and for others who need access to your account. You can also assign IAM users to IAM groups that have a set of policies that grant permissions to perform common tasks. By using IAM users instead of the root user, you can follow the principle of least privilege and reduce the risk of compromising your account.

? Enable multi-factor authentication (MFA) on the root user. MFA is a security feature that requires users to prove their identity by providing two pieces of information: their password and a code from a device that only they have access to. By enabling MFA on the root user, you can add an extra layer of protection to your account and prevent unauthorized access even if your password is compromised.

? Limit the tasks you perform with the root user account. You should use the root user only for tasks that require root user credentials, such as changing your account settings, closing your account, or managing consolidated billing. For a complete list of tasks that require root user credentials, see Tasks that require root user credentials. For all other tasks, you should use IAM users or roles that have the appropriate permissions.

References:

- ? AWS account root user
- ? Root user best practices for your AWS account
- ? Tasks that require root user credentials

NEW QUESTION 20

- (Topic 4)

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

handle millions of UDP internet traffic requests each second.

Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

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

NEW QUESTION 21

- (Topic 4)

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

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 26

- (Topic 4)

A company wants to migrate its three-tier application from on premises to AWS. The web tier and the application tier are running on third-party virtual machines (VMs). The database tier is running on MySQL. The company needs to migrate the application by making the fewest possible changes to the architecture. The company also needs a database solution that can restore data to a specific point in time. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 28

- (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 31

- (Topic 4)

To meet security requirements, a company needs to encrypt all of its application data in transit while communicating with an Amazon RDS MySQL DB instance. A recent security audit revealed that encryption at rest is enabled using AWS Key Management Service (AWS KMS), but data in transit is not enabled. What should a solutions architect do to satisfy the security requirements?

- A. Enable IAM database authentication on the database.
- B. Provide self-signed certificate
- C. Use the certificates in all connections to the RDS instance.
- D. Take a snapshot of the RDS instance
- E. Restore the snapshot to a new instance with encryption enabled.
- F. Download AWS-provided root certificate
- G. Provide the certificates in all connections to the RDS instance.

Answer: D

Explanation:

To satisfy the security requirements, the solutions architect should download AWS-provided root certificates and provide the certificates in all connections to the RDS instance. This will enable SSL/TLS encryption for data in transit between the application and the RDS instance. SSL/TLS encryption provides a layer of security by encrypting data that moves between the client and the server. Amazon RDS creates an SSL certificate and installs the certificate on the DB instance when the instance is provisioned. The application can use the AWS-provided root certificates to verify the identity of the DB instance and establish a secure connection¹.

The other options are not correct because they do not enable encryption for data in transit or are not relevant for the use case. Enabling IAM database authentication on the database is not correct because this option only provides a method of authentication, not encryption. IAM database authentication allows users to use AWS Identity and Access Management (IAM) users and roles to access a database, instead of using a database user name and password². Providing self-signed certificates is not correct because this option is not secure or reliable. Self-signed certificates are certificates that are signed by the same entity that issued them, instead of by a trusted certificate authority (CA). Self-signed certificates can be easily forged or compromised, and are not recognized by most browsers and applications³. Taking a snapshot of the RDS instance and restoring it to a new instance with encryption enabled is not correct because this option only enables encryption at rest, not encryption in transit. Encryption at rest protects data that is stored on disk, but does not protect data that is moving between the client and the server⁴.

References:

- ? Using SSL/TLS to encrypt a connection to a DB instance - Amazon Relational Database Service
- ? IAM database authentication for MySQL and PostgreSQL - Amazon Relational Database Service
- ? What are self-signed certificates?
- ? Encrypting Amazon RDS resources - Amazon Relational Database Service

NEW QUESTION 34

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

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

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

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

- ? Encrypting secrets used in Amazon EKS
- ? What Is AWS Key Management Service?
- ? What Is AWS Secrets Manager?
- ? Amazon EBS CSI driver
- ? Encryption at rest

NEW QUESTION 36

- (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 40

- (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 41

- (Topic 4)

A company runs a container application by using Amazon Elastic Kubernetes Service (Amazon EKS). The application includes microservices that manage customers and place orders. The company needs to route incoming requests to the appropriate microservices.

Which solution will meet this requirement MOST cost-effectively?

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

Answer: B

Explanation:

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

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

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

? You can route incoming requests to the appropriate microservices based on the rules you define in your Ingress resource. For example, you can route requests with different host names or paths to different microservices that handle customers and orders².

? You can improve the performance and availability of your container applications by distributing the load across multiple targets and enabling health checks and automatic scaling¹.

? You can reduce the cost and complexity of managing your load balancers by using a single controller that integrates with Amazon EKS and Kubernetes. You do not need to manually create or configure load balancers or update them when your cluster changes².

NEW QUESTION 45

- (Topic 4)

A company needs to minimize the cost of its 1 Gbps AWS Direct Connect connection. The company's average connection utilization is less than 10%. A solutions architect must recommend a solution that will reduce the cost without compromising security.

Which solution will meet these requirements?

- A. Set up a new 1 Gbps Direct Connect connectio
- B. Share the connection with another AWS account.
- C. Set up a new 200 Mbps Direct Connect connection in the AWS Management Console.
- D. Contact an AWS Direct Connect Partner to order a 1 Gbps connectio
- E. Share the connection with another AWS account.
- F. Contact an AWS Direct Connect Partner to order a 200 Mbps hosted connection for an existing AWS account.

Answer: D

Explanation:

company need to setup a cheaper connection (200 M) but B is incorrect because you can only order port speeds of 1, 10, or 100 Gbps for more flexibility you can go with hosted connection, You can order port speeds between 50 Mbps and 10 Gbps. <https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect.html>

NEW QUESTION 46

- (Topic 4)

A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will continue to grow and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

Explanation:

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

NEW QUESTION 48

- (Topic 4)

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

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

Answer: BE

Explanation:

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

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

NEW QUESTION 50

- (Topic 4)

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

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

Answer: D

Explanation:

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

NEW QUESTION 53

- (Topic 4)

A company runs a website that uses a content management system (CMS) on Amazon EC2. The CMS runs on a single EC2 instance and uses an Amazon Aurora MySQL Multi-AZ DB instance for the data tier. Website images are stored on an Amazon Elastic Block Store (Amazon EBS) volume that is mounted inside the EC2 instance.

Which combination of actions should a solutions architect take to improve the performance and resilience of the website? (Select TWO.)

- A. Move the website images into an Amazon S3 bucket that is mounted on every EC2 instance.
- B. Share the website images by using an NFS share from the primary EC2 instance
- C. Mount this share on the other EC2 instances.
- D. Move the website images onto an Amazon Elastic File System (Amazon EFS) file system that is mounted on every EC2 instance.
- E. Create an Amazon Machine Image (AMI) from the existing EC2 instance Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling group

- F. Configure the Auto Scaling group to maintain a minimum of two instance
- G. Configure an accelerator in AWS Global Accelerator for the website.
- H. Create an Amazon Machine Image (AMI) from the existing EC2 instanc
- I. Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling grou
- J. Configure the Auto Scaling group to maintain a minimum of two instance
- K. Configure an Amazon CloudFront distribution for the website.

Answer: CE

Explanation:

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

NEW QUESTION 56

- (Topic 4)

A global marketing company has applications that run in the ap-southeast-2 Region and the eu-west-1 Region. Applications that run in a VPC in eu-west-1 need to communicate securely with databases that run in a VPC in ap-southeast-2.

Which network design will meet these requirements?

- A. Create a VPC peering connection between the eu-west-1 VPC and the ap-southeast-2 VP
- B. Create an inbound rule in the eu-west-1 application security group that allows traffic from the database server IP addresses in the ap-southeast-2 security group.
- C. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VP
- D. Update the subnet route table
- E. Create an inbound rule in the ap-southeast-2 database security group that references the security group ID of the application servers in eu-west-1.
- F. Configure a VPC peering connection between the ap-southeast-2 VPC and the eu-west- 1 VP
- G. Update the subnet route tables Create an inbound rule in the ap-southeast-2 database security group that allows traffic from the eu-west-1 application server IP addresses.
- H. Create a transit gateway with a peering attachment between the eu-west-1 VPC and the ap-southeast-2 VP
- I. After the transit gateways are properly peered and routing is configured, create an inbound rule in the database security group that references the security group ID of the application servers in eu-west-1.

Answer: C

Explanation:

"You cannot reference the security group of a peer VPC that's in a different Region. Instead, use the CIDR block of the peer VPC."

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>

NEW QUESTION 58

- (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 1AM role.
- B. Create interface VPC endpoints to allow nodes to access the control plane.
- C. Recreate nodes in the public subnet Restrict security groups for EC2 nodes
- D. Allow outbound traffic in the security group of the nodes.

Answer: B

Explanation:

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

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

NEW QUESTION 60

- (Topic 4)

A social media company wants to allow its users to upload images in an application that is hosted in the AWS Cloud. The company needs a solution that automatically resizes the images so that the images can be displayed on multiple device types. The application experiences unpredictable traffic patterns throughout the day. The company is seeking a highly available solution that maximizes scalability.

What should a solutions architect do to meet these requirements?

- A. Create a static website hosted in Amazon S3 that invokes AWS Lambda functions to resize the images and store the images in an Amazon S3 bucket.
- B. Create a static website hosted in Amazon CloudFront that invokes AWS Step Functions to resize the images and store the images in an Amazon RDS database.
- C. Create a dynamic website hosted on a web server that runs on an Amazon EC2 instance Configure a process that runs on the EC2 instance to resize the images and store the images in an Amazon S3 bucket.
- D. Create a dynamic website hosted on an automatically scaling Amazon Elastic Container Service (Amazon ECS) cluster that creates a resize job in Amazon Simple Queue Service (Amazon SQS). Set up an image-resizing program that runs on an Amazon EC2 instance to process the resize jobs

Answer: A

Explanation:

By using Amazon S3 and AWS Lambda together, you can create a serverless architecture that provides highly scalable and available image resizing capabilities.

Here's how the solution would work: Set up an Amazon S3 bucket to store the original images uploaded by users. Configure an event trigger on the S3 bucket to invoke an AWS Lambda function whenever a new image is uploaded. The Lambda function can be designed to retrieve the uploaded image, perform the necessary resizing operations based on device requirements, and store the resized images back in the S3 bucket or a different bucket designated for resized images. Configure the Amazon S3 bucket to make the resized images publicly accessible for serving to users.

NEW QUESTION 64

- (Topic 4)

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

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

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

Answer: B

Explanation:

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

? 1 explains how to create a scheduled action for an Auto Scaling group.

? 2 describes the concept and benefits of an Auto Scaling group.

NEW QUESTION 66

- (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 71

- (Topic 4)

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

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

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

Answer: C

Explanation:

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

References:

? What Is VPC Peering?

? VPC Peering Pricing

NEW QUESTION 76

- (Topic 4)

A solutions architect must provide an automated solution for a company's compliance policy that states security groups cannot include a rule that allows SSH from

0.0.0.0/0. The company needs to be notified if there is any breach in the policy. A solution is needed as soon as possible. What should the solutions architect do to meet these requirements with the LEAST operational overhead?

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

Answer: B

Explanation:

The most suitable solution for the company's compliance policy is to enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created. This solution has the least operational overhead because it uses a predefined rule that is already available in AWS Config, which is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. The restricted-ssh rule checks whether security groups that are in use have inbound rules that allow SSH from 0.0.0.0/0 addresses, and reports them as noncompliant¹. Users can configure the rule to send notifications to an Amazon SNS topic when a noncompliant change occurs, and subscribe to the topic to receive alerts via email, SMS, or other methods².

The other options are not correct because they either have more operational overhead or do not meet the requirements. Writing an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one is not correct because it requires custom code development and maintenance, which adds complexity and cost to the solution. Creating an IAM role with permissions to globally open security groups and network ACLs, and creating an Amazon SNS topic to generate a notification every time the role is assumed by a user is not correct because it does not prevent or detect the creation of noncompliant rules by other users or roles, and it does not address the existing rules that may violate the policy. Configuring a service control policy (SCP) that prevents non-administrative users from creating or editing security groups, and creating a notification in the ticketing system when a user requests a rule that needs administrator permissions is not correct because it does not provide an automated solution for the policy enforcement and notification, and it may limit the flexibility and productivity of the users.

References:

- ? restricted-ssh - AWS Config
- ? Getting Notifications When Your Resources Change - AWS Config

NEW QUESTION 80

- (Topic 4)

A company runs a real-time data ingestion solution on AWS. The solution consists of the most recent version of Amazon Managed Streaming for Apache Kafka (Amazon MSK). The solution is deployed in a VPC in private subnets across three Availability Zones.

A solutions architect needs to redesign the data ingestion solution to be publicly available over the internet. The data in transit must also be encrypted. Which solution will meet these requirements with the MOST operational efficiency?

- A. Configure public subnets in the existing VPC
- B. Deploy an MSK cluster in the public subnet
- C. Update the MSK cluster security settings to enable mutual TLS authentication.
- D. Create a new VPC that has public subnet
- E. Deploy an MSK cluster in the public subnet
- F. Update the MSK cluster security settings to enable mutual TLS authentication.
- G. Deploy an Application Load Balancer (ALB) that uses private subnet
- H. Configure an ALB security group inbound rule to allow inbound traffic from the VPC CIDR block for HTTPS protocol.
- I. Deploy a Network Load Balancer (NLB) that uses private subnet
- J. Configure an NLB listener for HTTPS communication over the internet.

Answer: A

Explanation:

The solution that meets the requirements with the most operational efficiency is to configure public subnets in the existing VPC and deploy an MSK cluster in the public subnets. This solution allows the data ingestion solution to be publicly available over the internet without creating a new VPC or deploying a load balancer. The solution also ensures that the data in transit is encrypted by enabling mutual TLS authentication, which requires both the client and the server to present certificates for verification. This solution leverages the public access feature of Amazon MSK, which is available for clusters running Apache Kafka 2.6.0 or later versions¹.

The other solutions are not as efficient as the first one because they either create unnecessary resources or do not encrypt the data in transit. Creating a new VPC with public subnets would incur additional costs and complexity for managing network resources and routing. Deploying an ALB or an NLB would also add more costs and latency for the data ingestion solution. Moreover, an ALB or an NLB would not encrypt the data in transit by itself, unless they are configured with HTTPS listeners and certificates, which would require additional steps and maintenance. Therefore, these solutions are not optimal for the given requirements.

References:

- ? Public access - Amazon Managed Streaming for Apache Kafka

NEW QUESTION 84

- (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 89

- (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 91

- (Topic 4)

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

Which solution meets these requirements?

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

Answer: D

Explanation:

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

References :=

? Amazon CloudFront

? Application Load Balancer

? Amazon Route 53

? AWS WAF

NEW QUESTION 96

- (Topic 4)

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

Which solution will meet these requirements?

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

H. Increase the memory according to AWS Compute Optimizer recommendations.

Answer: D

Explanation:

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

? Configure provisioned concurrency for the Lambda functions. Provisioned

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

? Increase the memory according to AWS Compute Optimizer recommendations.

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

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

References:

? Provisioned Concurrency

? AWS Compute Optimizer

NEW QUESTION 99

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

References:

? Using Cost Allocation Tags

? Creating an AWS Budget

? Creating an Alert for an AWS Budget

NEW QUESTION 101

- (Topic 4)

A company runs a three-tier web application in the AWS Cloud that operates across three Availability Zones. The application architecture has an Application Load Balancer, an Amazon EC2 web server that hosts user session states, and a MySQL database that runs on an EC2 instance. The company expects sudden increases in application traffic. The company wants to be able to scale to meet future application capacity demands and to ensure high availability across all three Availability Zones.

Which solution will meet these requirements?

- A. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment
- B. Use Amazon ElastiCache for Redis with high availability to store session data and to cache read
- C. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
- D. Migrate the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment
- E. Use Amazon ElastiCache for Memcached with high availability to store session data and to cache read
- F. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
- G. Migrate the MySQL database to Amazon DynamoD
- H. Use DynamoDB Accelerator (DAX) to cache read
- I. Store the session data in DynamoD
- J. Migrate the web server to an Auto Scaling group that is in three Availability Zones.
- K. Migrate the MySQL database to Amazon RDS for MySQL in a single Availability Zone
- L. Use Amazon ElastiCache for Redis with high availability to store session data and to cache read
- M. Migrate the web server to an Auto Scaling group that is in three Availability Zones.

Answer: A

Explanation:

This answer is correct because it meets the requirements of scaling to meet future application capacity demands and ensuring high availability across all three Availability Zones. By migrating the MySQL database to Amazon RDS for MySQL with a Multi-AZ DB cluster deployment, the company can benefit from automatic failover, backup, and patching of the database across multiple Availability Zones. By using Amazon ElastiCache for Redis with high availability, the company can store session data and cache reads in a fast, in-memory data store that can also fail over across Availability Zones. By migrating the web server to an Auto Scaling group that is in three Availability Zones, the company can automatically scale the web server capacity based on the demand and traffic patterns. References:

? <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>
? <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html>
? <https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html>

NEW QUESTION 104

- (Topic 4)

A company uses Amazon Elastic Kubernetes Service (Amazon EKS) to run a container application. The EKS cluster stores sensitive information in the Kubernetes secrets object. The company wants to ensure that the information is encrypted. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the container application to encrypt the information by using AWS Key Management Service (AWS KMS).
- B. Enable secrets encryption in the EKS cluster by using AWS Key Management Service (AWS KMS).
- C. Implement an AWS Lambda function to encrypt the information by using AWS Key Management Service (AWS KMS).
- D. Use AWS Systems Manager Parameter Store to encrypt the information by using AWS Key Management Service (AWS KMS).

Answer: B

Explanation:

It allows the company to encrypt the Kubernetes secrets object in the EKS cluster with the least operational overhead. By enabling secrets encryption in the EKS cluster, the company can use AWS Key Management Service (AWS KMS) to generate and manage encryption keys for encrypting and decrypting secrets at rest. This is a simple and secure way to protect sensitive information in EKS clusters. References:

- ? [Encrypting Kubernetes secrets with AWS KMS](#)
- ? [Kubernetes Secrets](#)

NEW QUESTION 105

- (Topic 4)

A company's web application that is hosted in the AWS Cloud recently increased in popularity. The web application currently exists on a single Amazon EC2 instance in a single public subnet. The web application has not been able to meet the demand of the increased web traffic. The company needs a solution that will provide high availability and scalability to meet the increased user demand without rewriting the web application. Which combination of steps will meet these requirements? (Select TWO.)

- A. Replace the EC2 instance with a larger compute optimized instance.
- B. Configure Amazon EC2 Auto Scaling with multiple Availability Zones in private subnets.
- C. Configure a NAT gateway in a public subnet to handle web requests.
- D. Replace the EC2 instance with a larger memory optimized instance.
- E. Configure an Application Load Balancer in a public subnet to distribute web traffic.

Answer: BE

Explanation:

These two steps will meet the requirements because they will provide high availability and scalability for the web application without rewriting it. Amazon EC2 Auto Scaling allows you to automatically adjust the number of EC2 instances in response to changes in demand. By configuring Auto Scaling with multiple Availability Zones in private subnets, you can ensure that your web application is distributed across isolated and fault-tolerant locations, and that your instances are not directly exposed to the internet. An Application Load Balancer operates at the application layer and distributes incoming web traffic across multiple targets, such as EC2 instances, containers, or Lambda functions. By configuring an Application Load Balancer in a public subnet, you can enable your web application to handle requests from the internet and route them to the appropriate targets in the private subnets.

References:

- ? [What is Amazon EC2 Auto Scaling?](#)
- ? [What is an Application Load Balancer?](#)

NEW QUESTION 107

- (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 110

- (Topic 4)

A medical research lab produces data that is related to a new study. The lab wants to make the data available with minimum latency to clinics across the country for their on-premises, file-based applications. The data files are stored in an Amazon S3 bucket that has read-only permissions for each clinic. What should a solutions architect recommend to meet these requirements?

- A. Deploy an AWS Storage Gateway file gateway as a virtual machine (VM) on premises at each clinic
- B. Migrate the files to each clinic's on-premises applications by using AWS DataSync for processing.
- C. Deploy an AWS Storage Gateway volume gateway as a virtual machine (VM) on premises at each clinic.
- D. Attach an Amazon Elastic File System (Amazon EFS) file system to each clinic's on-premises servers.

Answer: A

Explanation:

AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. By deploying a file gateway as a virtual machine on each clinic's premises, the medical research lab can provide low-latency access to the data stored in the S3 bucket while maintaining read-only permissions for each clinic. This solution allows the clinics to access the data files directly from their on-premises file-based applications without the need for data transfer or migration.

NEW QUESTION 113

- (Topic 4)

A company is deploying a new application on Amazon EC2 instances. The application writes data to Amazon Elastic Block Store (Amazon EBS) volumes. The company needs to ensure that all data that is written to the EBS volumes is encrypted at rest. Which solution will meet this requirement?

- A. Create an IAM role that specifies EBS encryption
- B. Attach the role to the EC2 instances.
- C. Create the EBS volumes as encrypted volume

- D. Attach the EBS volumes to the EC2 instances
- E. Create an EC2 instance tag that has a key of Encrypt and a value of True
- F. Tag all instances that require encryption at the EBS level.
- G. Create an AWS Key Management Service (AWS KMS) key policy that enforces EBS encryption in the account
- H. Ensure that the key policy is active

Answer: B

Explanation:

The solution that will meet the requirement of ensuring that all data that is written to the EBS volumes is encrypted at rest is B. Create the EBS volumes as encrypted volumes and attach the encrypted EBS volumes to the EC2 instances. When you create an EBS volume, you can specify whether to encrypt the volume. If you choose to encrypt the volume, all data written to the volume is automatically encrypted at rest using AWS-managed keys. You can also use customer-managed keys (CMKs) stored in AWS KMS to encrypt and protect your EBS volumes. You can create encrypted EBS volumes and attach them to EC2 instances to ensure that all data written to the volumes is encrypted at rest.

NEW QUESTION 117

- (Topic 4)

A company is preparing a new data platform that will ingest real-time streaming data from multiple sources. The company needs to transform the data before writing the data to Amazon S3. The company needs the ability to use SQL to query the transformed data. Which solutions will meet these requirements? (Choose two.)

- A. Use Amazon Kinesis Data Streams to stream the data
- B. Use Amazon Kinesis Data Analytics to transform the data
- C. Use Amazon Kinesis Data Firehose to write the data to Amazon S3. Use Amazon Athena to query the transformed data from Amazon S3.
- D. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to stream the data
- E. Use AWS Glue to transform the data and to write the data to Amazon S3. Use Amazon Athena to query the transformed data from Amazon S3.
- F. Use AWS Database Migration Service (AWS DMS) to ingest the data
- G. Use Amazon EMR to transform the data and to write the data to Amazon S3. Use Amazon Athena to query the transformed data from Amazon S3.
- H. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to stream the data
- I. Use Amazon Kinesis Data Analytics to transform the data and to write the data to Amazon S3. Use the Amazon RDS query editor to query the transformed data from Amazon S3.
- J. Use Amazon Kinesis Data Streams to stream the data
- K. Use AWS Glue to transform the data
- L. Use Amazon Kinesis Data Firehose to write the data to Amazon S3. Use the Amazon RDS query editor to query the transformed data from Amazon S3.

Answer: AB

Explanation:

To ingest, transform, and query real-time streaming data from multiple sources, Amazon Kinesis and Amazon MSK are suitable solutions. Amazon Kinesis Data Streams can stream the data from various sources and integrate with other AWS services. Amazon Kinesis Data Analytics can transform the data using SQL or Apache Flink.

Amazon Kinesis Data Firehose can write the data to Amazon S3 or other destinations. Amazon Athena can query the transformed data from Amazon S3 using standard SQL. Amazon MSK can stream the data using Apache Kafka, which is a popular open-source platform for streaming data. AWS Glue can transform the data using Apache Spark or Python scripts and write the data to Amazon S3 or other destinations. Amazon Athena can also query the transformed data from Amazon S3 using standard SQL.

References:

- ? What Is Amazon Kinesis Data Streams?
- ? What Is Amazon Kinesis Data Analytics?
- ? What Is Amazon Kinesis Data Firehose?
- ? What Is Amazon Athena?
- ? What Is Amazon MSK?
- ? What Is AWS Glue?

NEW QUESTION 122

- (Topic 4)

A company is using a centralized AWS account to store log data in various Amazon S3 buckets. A solutions architect needs to ensure that the data is encrypted at rest before the data is uploaded to the S3 buckets. The data also must be encrypted in transit. Which solution meets these requirements?

- A. Use client-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- B. Use server-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- C. Create bucket policies that require the use of server-side encryption with S3 managed encryption keys (SSE-S3) for S3 uploads.
- D. Enable the security option to encrypt the S3 buckets through the use of a default AWS Key Management Service (AWS KMS) key.

Answer: A

Explanation:

Client-side encryption is a method of encrypting data before uploading it to Amazon S3. It allows users to manage the encryption process, encryption keys, and related tools. By using client-side encryption, the solution can ensure that the data is encrypted at rest and in transit, as Amazon S3 will not have access to the encryption keys or the unencrypted data.

NEW QUESTION 124

- (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 129

- (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 132

- (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 135

- (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 137

- (Topic 4)

A company copies 200 TB of data from a recent ocean survey onto AWS Snowball Edge Storage Optimized devices. The company has a high performance computing (HPC) cluster that is hosted on AWS to look for oil and gas deposits. A solutions architect must provide the cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices. The company is sending the devices back to AWS. Which solution will meet these requirements?

- A. Create an Amazon S3 bucket
- B. Import the data into the S3 bucket
- C. Configure an AWS Storage Gateway file gateway to use the S3 bucket
- D. Access the file gateway from the HPC cluster instances.
- E. Create an Amazon S3 bucket
- F. Import the data into the S3 bucket
- G. Configure an Amazon FSx for Lustre file system, and integrate it with the S3 bucket
- H. Access the FSx for Lustre file system from the HPC cluster instances.
- I. Create an Amazon S3 bucket and an Amazon Elastic File System (Amazon EFS) file system
- J. Import the data into the S3 bucket
- K. Copy the data from the S3 bucket to the EFS file system
- L. Access the EFS file system from the HPC cluster instances.
- M. Create an Amazon FSx for Lustre file system
- N. Import the data directly into the FSx for Lustre file system
- O. Access the FSx for Lustre file system from the HPC cluster instances.

Answer: B

Explanation:

To provide the HPC cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices, a solutions architect should configure an Amazon FSx for Lustre file system, and integrate it with an Amazon S3 bucket. This solution has the following benefits:

- ? It allows the HPC cluster to access the data on the Snowball Edge devices using a POSIX-compliant file system that is optimized for fast processing of large datasets¹.
- ? It enables the data to be imported from the Snowball Edge devices into the S3 bucket using the AWS Snow Family Console or the AWS CLI². The data can then be accessed from the FSx for Lustre file system using the S3 integration feature³.
- ? It supports high availability and durability of the data, as the FSx for Lustre file system can automatically copy the data to and from the S3 bucket³. The data can also be accessed from other AWS services or applications using the S3 API⁴.

References:

- ? 1: <https://aws.amazon.com/fsx/lustre/>
- ? 2: <https://docs.aws.amazon.com/snowball/latest/developer-guide/using-adapter.html>
- ? 3: <https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html>
- ? 4: <https://docs.aws.amazon.com/fsx/latest/LustreGuide/export-data-repo.html>

NEW QUESTION 142

- (Topic 4)

A company needs to store contract documents. A contract lasts for 5 years. During the 5-year period, the company must ensure that the documents cannot be overwritten or deleted. The company needs to encrypt the documents at rest and rotate the encryption keys automatically every year.

Which combination of steps should a solutions architect take to meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. Store the documents in Amazon S3. Use S3 Object Lock in governance mode.
- B. Store the documents in Amazon S3. Use S3 Object Lock in compliance mode.
- C. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Configure key rotation.
- D. Use server-side encryption with AWS Key Management Service (AWS KMS) customer managed key
- E. Configure key rotation.
- F. Use server-side encryption with AWS Key Management Service (AWS KMS) customer provided (imported) key
- G. Configure key rotation.

Answer: BD

Explanation:

Consider using the default aws/s3 KMS key if: You're uploading or accessing S3 objects using AWS Identity and Access Management (IAM) principals that are in the same AWS account as the AWS KMS key. You don't want to manage policies for the KMS key. Consider using a customer managed key if: You want to create, rotate, disable, or define access controls for the key. You want to grant cross-account access to your S3 objects. You can configure the policy of a customer managed key to allow access from another account. <https://repost.aws/knowledge-center/s3-object-encryption-keys>

NEW QUESTION 145

- (Topic 4)

A company wants to share accounting data with an external auditor. The data is stored in an Amazon RDS DB instance that resides in a private subnet. The auditor has its own AWS account and requires its own copy of the database.

What is the MOST secure way for the company to share the database with the auditor?

- A. Create a read replica of the databases
- B. Configure IAM standard database authentication to grant the auditor access.
- C. Export the database contents to text file
- D. Store the files in an Amazon S3 bucket
- E. Create a new IAM user for the auditor
- F. Grant the user access to the S3 bucket.
- G. Copy a snapshot of the database to an Amazon S3 bucket
- H. Create an IAM user
- I. Share the user's keys with the auditor to grant access to the object in the S3 bucket.
- J. Create an encrypted snapshot of the database
- K. Share the snapshot with the auditor
- L. Allow access to the AWS Key Management Service (AWS KMS) encryption key.

Answer: D

Explanation:

This answer is correct because it meets the requirements of sharing the database with the auditor in a secure way. You can create an encrypted snapshot of the database by using AWS Key Management Service (AWS KMS) to encrypt the snapshot with a customer managed key. You can share the snapshot with the auditor by modifying the permissions of the snapshot and specifying the AWS account ID of the auditor. You can also allow access to the AWS KMS encryption key by adding a key policy statement that grants permissions to the auditor's account. This way, you can ensure that only the auditor can access and restore the snapshot in their own AWS account.

References:

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

? <https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html#key-policy-default-allow-root-enable-iam>

NEW QUESTION 150

- (Topic 4)

A company wants to migrate 100 GB of historical data from an on-premises location to an Amazon S3 bucket. The company has a 100 megabits per second (Mbps) internet connection on premises. The company needs to encrypt the data in transit to the S3 bucket. The company will store new data directly in Amazon S3.

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

- A. Use the s3 sync command in the AWS CLI to move the data directly to an S3 bucket.
- B. Use AWS DataSync to migrate the data from the on-premises location to an S3 bucket.
- C. Use AWS Snowball to move the data to an S3 bucket.
- D. Set up an IPsec VPN from the on-premises location to AWS
- E. Use the s3 cp command in the AWS CLI to move the data directly to an S3 bucket.

Answer: B

Explanation:

AWS DataSync is a data transfer service that makes it easy for you to move large amounts of data online between on-premises storage and AWS storage services over the internet or AWS Direct Connect. DataSync automatically encrypts your data in transit using TLS encryption, and verifies data integrity during transfer using checksums. DataSync can transfer data up to 10 times faster than open-source tools, and reduces operational overhead by simplifying and automating tasks such as scheduling, monitoring, and resuming transfers. References: <https://aws.amazon.com/datasync/>

NEW QUESTION 153

- (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 154

- (Topic 4)

A company's infrastructure consists of Amazon EC2 instances and an Amazon RDS DB instance in a single AWS Region. The company wants to back up its data in a separate Region.

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

- A. Use AWS Backup to copy EC2 backups and RDS backups to the separate Region.
- B. Use Amazon Data Lifecycle Manager (Amazon DLM) to copy EC2 backups and RDS backups to the separate Region.
- C. Create Amazon Machine Images (AMIs) of the EC2 instance
- D. Copy the AMIs to the separate Region
- E. Create a read replica for the RDS DB instance in the separate Region.
- F. Create Amazon Elastic Block Store (Amazon EBS) snapshot
- G. Copy the EBS snapshots to the separate Region
- H. Create RDS snapshot
- I. Export the RDS snapshots to Amazon S3. Configure S3 Cross-Region Replication (CRR) to the separate Region.

Answer: A

Explanation:

To back up EC2 instances and RDS DB instances in a separate Region with the least operational overhead, AWS Backup is a simple and cost-effective solution. AWS Backup can copy EC2 backups and RDS backups to another Region automatically and securely. AWS Backup also supports backup policies, retention rules, and monitoring features.

References:

? What Is AWS Backup?

? Cross-Region Backup

NEW QUESTION 156

- (Topic 4)

A company has deployed a Java Spring Boot application as a pod that runs on Amazon Elastic Kubernetes Service (Amazon EKS) in private subnets. The application needs to write data to an Amazon DynamoDB table. A solutions architect must ensure that the application can interact with the DynamoDB table without exposing traffic to the internet.

Which combination of steps should the solutions architect take to accomplish this goal? (Choose two.)

- A. Attach an IAM role that has sufficient privileges to the EKS pod.
- B. Attach an IAM user that has sufficient privileges to the EKS pod.
- C. Allow outbound connectivity to the DynamoDB table through the private subnets' network ACLs.
- D. Create a VPC endpoint for DynamoDB.
- E. Embed the access keys in the Java Spring Boot code.

Answer: AD

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

<https://aws.amazon.com/about-aws/whats-new/2019/09/amazon-eks-adds-support-to-assign-iam-permissions-to-kubernetes-service-accounts/>

NEW QUESTION 160

- (Topic 4)

A company uses Amazon API Gateway to run a private gateway with two REST APIs in the same VPC. The BuyStock RESTful web service calls the CheckFunds RESTful

web service to ensure that enough funds are available before a stock can be purchased.

The company has noticed in the VPC flow logs that the BuyStock RESTful web

service calls the CheckFunds RESTful web service over the internet instead of through the VPC. A solutions architect must implement a solution so that the APIs communicate through the VPC.

Which solution will meet these requirements with the FEWEST changes to the code? (Select Correct Option/s and give detailed explanation from AWS Certified Solutions

Architect - Associate (SAA-C03) Study Manual or documents)

- A. Add an X-API-Key header in the HTTP header for authorization.
- B. Use an interface endpoint.
- C. Use a gateway endpoint.
- D. Add an Amazon Simple Queue Service (Amazon SQS) queue between the two REST APIs.

Answer: B

Explanation:

Using an interface endpoint will allow the BuyStock RESTful web service and the CheckFunds RESTful web service to communicate through the VPC without any changes to the code. An interface endpoint creates an elastic network interface (ENI) in the customer's VPC, and then configures the route tables to route traffic from the APIs to the ENI. This will ensure that the two APIs will communicate through the VPC without any changes to the code.

NEW QUESTION 162

- (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 166

- (Topic 4)

A research company uses on-premises devices to generate data for analysis. The company wants to use the AWS Cloud to analyze the data. The devices generate .csv files and support writing the data to SMB file share. Company analysts must be able to use SQL commands to query the data. The analysts will run queries periodically throughout the day.

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

- A. Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode.
- B. Deploy an AWS Storage Gateway on premises in Amazon FSx File Gateway mode.
- C. Set up an AWS Glue crawler to create a table based on the data that is in Amazon S3.
- D. Set up an Amazon EMR cluster with EMR File System (EMRFS) to query the data that is in Amazon S3. Provide access to analysts.
- E. Set up an Amazon Redshift cluster to query the data that is in Amazon S3. Provide access to analysts.
- F. Set up Amazon Athena to query the data that is in Amazon S3. Provide access to analysts.

Answer: ACF

Explanation:

To meet the requirements of the use case in a cost-effective way, the following steps are recommended:

? Deploy an AWS Storage Gateway on premises in Amazon S3 File Gateway mode. This will allow the company to write the .csv files generated by the devices to an SMB file share, which will be stored as objects in Amazon S3 buckets. AWS Storage Gateway is a hybrid cloud storage service that integrates on-premises environments with AWS storage. Amazon S3 File Gateway mode provides a seamless way to connect to Amazon S3 and access a virtually unlimited amount of cloud storage¹.
 ? Set up an AWS Glue crawler to create a table based on the data that is in Amazon S3. This will enable the company to use standard SQL to query the data stored in Amazon S3 buckets. AWS Glue is a serverless data integration service that simplifies data preparation and analysis. AWS Glue crawlers can automatically discover and classify data from various sources, and create metadata tables in the AWS Glue Data Catalog². The Data Catalog is a central repository that stores information about data sources and how to access them³.
 ? Set up Amazon Athena to query the data that is in Amazon S3. This will provide the company analysts with a serverless and interactive query service that can analyze data directly in Amazon S3 using standard SQL. Amazon Athena is integrated with the AWS Glue Data Catalog, so users can easily point Athena at the data source tables defined by the crawlers. Amazon Athena charges only for the queries that are run, and offers a pay-per-query pricing model, which makes it a cost-effective option for periodic queries⁴.
 The other options are not correct because they are either not cost-effective or not suitable for the use case. Deploying an AWS Storage Gateway on premises in Amazon FSx File Gateway mode is not correct because this mode provides low-latency access to fully managed Windows file shares in AWS, which is not required for the use case. Setting up an Amazon EMR cluster with EMR File System (EMRFS) to query the data that is in Amazon S3 is not correct because this option involves setting up and managing a cluster of EC2 instances, which adds complexity and cost to the solution. Setting up an Amazon Redshift cluster to query the data that is in Amazon S3 is not correct because this option also involves provisioning and managing a cluster of nodes, which adds overhead and cost to the solution.

References:

- ? What is AWS Storage Gateway?
- ? What is AWS Glue?
- ? AWS Glue Data Catalog
- ? What is Amazon Athena?

NEW QUESTION 171

- (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 176

- (Topic 4)

A company is developing a marketing communications service that targets mobile app users. The company needs to send confirmation messages with Short Message Service (SMS) to its users. The users must be able to reply to the SMS messages. The company must store the responses for a year for analysis. What should a solutions architect do to meet these requirements?

- A. Create an Amazon Connect contact flow to send the SMS message
- B. Use AWS Lambda to process the responses.
- C. Build an Amazon Pinpoint journey
- D. Configure Amazon Pinpoint to send events to an Amazon Kinesis data stream for analysis and archiving.
- E. Use Amazon Simple Queue Service (Amazon SQS) to distribute the SMS message
- F. Use AWS Lambda to process the responses.
- G. Create an Amazon Simple Notification Service (Amazon SNS) FIFO topic
- H. Subscribe an Amazon Kinesis data stream to the SNS topic for analysis and archiving.

Answer: B

Explanation:

<https://aws.amazon.com/pinpoint/product-details/sms/> Two-Way Messaging: Receive SMS messages from your customers and reply back to them in a chat-like interactive experience. With Amazon Pinpoint, you can create automatic responses when customers send you messages that contain certain keywords. You can even use Amazon Lex to create conversational bots. A majority of mobile phone users read incoming SMS messages almost immediately after receiving them. If you need to be able to provide your customers with urgent or important information, SMS messaging may be the right solution for you. You can use Amazon Pinpoint to create targeted groups of customers, and then send them campaign-based messages. You can also use Amazon Pinpoint to send direct messages, such as appointment confirmations, order updates, and one-time passwords.

NEW QUESTION 181

.....

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](#)