

Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate



NEW QUESTION 1

- (Topic 4)

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

Which solution meets these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:

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

References:

? AWS Fargate

? Amazon Elastic File System

? Using Amazon EFS file systems with Amazon ECS

NEW QUESTION 2

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

Amazon Athena is a service that allows you to run SQL queries on data stored in Amazon S3. It is serverless, meaning you do not need to provision or manage any infrastructure. You only pay for the queries you run and the amount of data scanned.

By using Amazon Athena to query your data in Amazon S3, you can achieve the following benefits:

? You can run queries for your report without affecting the performance of your

Amazon RDS for MySQL DB instance. You can export your data from your DB instance to an S3 bucket and use Athena to query the data in the bucket. This way, you can avoid the overhead and contention of running queries on your DB instance.

? You can reduce the cost and complexity of running queries for your report. You do

not need to create a read replica or a backup of your DB instance, which would incur additional charges and require maintenance. You also do not need to resize your DB instance to accommodate the additional workload, which would increase your operational overhead.

? You can leverage the scalability and flexibility of Amazon S3 and Athena. You can

store large amounts of data in S3 and query them with Athena without worrying about capacity or performance limitations. You can also use different formats, compression methods, and partitioning schemes to optimize your data storage and query performance.

NEW QUESTION 3

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

A gateway VPC endpoint for Amazon S3 enables private connections between the VPC and Amazon S3 that do not require an internet gateway or NAT device. This minimizes costs and prevents traffic from traversing the internet. A gateway VPC endpoint uses a prefix list as the route target in a VPC route table to route traffic privately to Amazon S3. Associating the endpoint with all route tables in the VPC ensures that all subnets can access Amazon S3 through the endpoint.

Option A is incorrect because S3 Intelligent-Tiering is a storage class that optimizes storage costs by automatically moving objects between two access tiers based on changing access patterns. It does not affect the network traffic between the VPC and Amazon S3.

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

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

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

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

NEW QUESTION 4

- (Topic 4)

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 5

- (Topic 4)

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: A

Explanation:

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

NEW QUESTION 6

- (Topic 4)

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

Which solution will migrate the database MOST cost-effectively?

- A. Order an AWS Snowball Edge Storage Optimized device
- B. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing change
- C. Send the Snowball Edge device to AWS to finish the migration and continue the ongoing replication.
- D. Order an AWS Snowmobile vehicle
- E. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
- F. Send the Snowmobile vehicle back to AWS to finish the migration and continue the ongoing replication.
- G. Order an AWS Snowball Edge Compute Optimized with GPU device
- H. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with ongoing change
- I. Send the Snowball device to AWS to finish the migration and continue the ongoing replication.

- J. Order a 1 GB dedicated AWS Direct Connect connection to establish a connection with the data center
K. Use AWS Database Migration Service (AWS DMS) with AWS Schema Conversion Tool (AWS SCT) to migrate the database with replication of ongoing changes.

Answer: A

Explanation:

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

References:

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

NEW QUESTION 7

- (Topic 4)

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

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

Answer: D

Explanation:

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

* A. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) customer managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. Although the company can create and use its own customer managed key in AWS KMS, the key is still stored and managed by AWS KMS, which is a service within the AWS Cloud. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

* B. Encrypt the data in the S3 bucket with server-side encryption (SSE) that uses an AWS Key Management Service (AWS KMS) AWS managed key. This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default AWS managed key in AWS KMS, which is created and managed by AWS on behalf of the company. The company has no control over the key rotation, deletion, or recovery policies. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards.

* C. Encrypt the data in the S3 bucket with the default server-side encryption (SSE). This option does not meet the requirements because it does not allow the company to manage the encryption keys outside the AWS Cloud. In this option, the company uses the default server-side encryption with Amazon S3 managed keys (SSE-S3), which is applied to every bucket in Amazon S3. The company has no visibility or control over the encryption keys, which are managed by Amazon S3. Moreover, the company still needs to use the AWS encryption features and APIs to encrypt and decrypt the data in the S3 bucket, which may not be compatible with the company's security policies or compliance standards. References:

- ? 1 Protecting data with encryption - Amazon Simple Storage Service
- ? 2 Protecting data with server-side encryption - Amazon Simple Storage Service
- ? 3 Protecting data by using client-side encryption - Amazon Simple Storage Service
- ? 4 AWS Key Management Service Concepts - AWS Key Management Service

NEW QUESTION 8

- (Topic 4)

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 monitoring.

* 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 transfer.

* 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 9

- (Topic 4)

A company needs to use its on-premises LDAP directory service to authenticate its users to the AWS Management Console. The directory service is not compatible with Security Assertion Markup Language (SAML).

Which solution meets these requirements?

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

Answer: D

Explanation:

The solution that meets the requirements is to develop an on-premises custom identity broker application or process that uses AWS Security Token Service (AWS STS) to get short-lived credentials. This solution allows the company to use its existing LDAP directory service to authenticate its users to the AWS Management Console, without requiring SAML compatibility. The custom identity broker application or process can act as a proxy between the LDAP directory service and AWS STS, and can request temporary security credentials for the users based on their LDAP attributes and roles. The users can then use these credentials to access the AWS Management Console via a sign-in URL generated by the identity broker. This solution also enhances security by using short-lived credentials that expire after a specified duration.

The other solutions do not meet the requirements because they either require SAML compatibility or do not provide access to the AWS Management Console. Enabling AWS IAM Identity Center (AWS Single Sign-On) between AWS and the on-premises LDAP would require the LDAP directory service to support SAML 2.0, which is not the case for this scenario. Creating an IAM policy that uses AWS credentials and integrating the policy into LDAP would not provide access to the AWS Management Console, but only to the AWS APIs. Setting up a process that rotates the IAM credentials whenever LDAP credentials are updated would also not provide access to the AWS Management Console, but only to the AWS CLI. Therefore, these solutions are not suitable for the given requirements.

NEW QUESTION 10

- (Topic 4)

A company runs an application on AWS. The application receives inconsistent amounts of usage. The application uses AWS Direct Connect to connect to an on-premises MySQL-compatible database. The on-premises database consistently uses a minimum of 2 GiB of memory.

The company wants to migrate the on-premises database to a managed AWS service. The company wants to use auto scaling capabilities to manage unexpected workload increases.

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

- A. Provision an Amazon DynamoDB database with default read and write capacity settings.
- B. Provision an Amazon Aurora database with a minimum capacity of 1 Aurora capacity unit (ACU).
- C. Provision an Amazon Aurora Serverless v2 database with a minimum capacity of 1 Aurora capacity unit (ACU).
- D. Provision an Amazon RDS for MySQL database with 2 GiB of memory.

Answer: C

Explanation:

it allows the company to migrate the on-premises database to a managed AWS service that supports auto scaling capabilities and has the least administrative overhead. Amazon Aurora Serverless v2 is a configuration of Amazon Aurora that automatically scales compute capacity based on workload demand. It can scale from hundreds to hundreds of thousands of transactions in a fraction of a second. Amazon Aurora Serverless v2 also supports MySQL-compatible databases and AWS Direct Connect connectivity. References:

? Amazon Aurora Serverless v2

? Connecting to an Amazon Aurora DB Cluster

NEW QUESTION 10

- (Topic 4)

A recent analysis of a company's IT expenses highlights the need to reduce backup costs. The company's chief information officer wants to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. The company must preserve the existing investment in the on-premises backup applications and workflows.

What should a solutions architect recommend?

- A. Set up AWS Storage Gateway to connect with the backup applications using the NFS interface.
- B. Set up an Amazon EFS file system that connects with the backup applications using the NFS interface.
- C. Set up an Amazon EFS file system that connects with the backup applications using the iSCSI interface.
- D. Set up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface.

Answer: D

Explanation:

it allows the company to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. By setting up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface, the company can store backup data on virtual tapes in S3 or Glacier. This preserves the existing investment in the on-premises backup applications and workflows while leveraging AWS storage services.

References:

? AWS Storage Gateway

? Tape Gateway

NEW QUESTION 13

- (Topic 4)

A company wants to migrate an on-premises legacy application to AWS. The application ingests customer order files from an on-premises enterprise resource planning (ERP) system. The application then uploads the files to an SFTP server. The application uses a scheduled job that checks for order files every hour. The company already has an AWS account that has connectivity to the on-premises network. The new application on AWS must support integration with the

existing ERP system. The new application must be secure and resilient and must use the SFTP protocol to process orders from the ERP system immediately. Which solution will meet these requirements?

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

Answer: D

Explanation:

This solution meets the requirements because it uses the following components and features:

? AWS Transfer Family SFTP internal server: This allows the application to securely transfer order files from the on-premises ERP system to AWS using the SFTP protocol over a private connection. The internal server is deployed in two Availability Zones for high availability and fault tolerance.

? Amazon S3 storage: This provides scalable, durable, and cost-effective object storage for the order files. Amazon S3 also supports encryption at rest and in transit, as well as lifecycle policies and versioning for data protection and compliance.

? AWS Lambda function: This enables the application to process the order files in a serverless manner, without provisioning or managing servers. The Lambda function can perform any custom logic or transformation on the order files, such as validating, parsing, or enriching the data.

? Transfer Family managed workflow: This simplifies the orchestration of the file processing tasks by triggering the Lambda function as soon as a file is uploaded to the SFTP server. The managed workflow also provides error handling, retry policies, and logging capabilities.

NEW QUESTION 17

- (Topic 4)

A city has deployed a web application running on Amazon EC2 instances behind an Application Load Balancer (ALB). The application's users have reported sporadic performance, which appears to be related to DDoS attacks originating from random IP addresses. The city needs a solution that requires minimal configuration changes and provides an audit trail for the DDoS sources.

Which solution meets these requirements?

- A. Enable an AWS WAF web ACL on the ALB, and configure rules to block traffic from unknown sources.
- B. Subscribe to Amazon Inspector
- C. Engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service.
- D. Subscribe to AWS Shield Advance
- E. Engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service.
- F. Create an Amazon CloudFront distribution for the application, and set the ALB as the origi
- G. Enable an AWS WAF web ACL on the distribution, and configure rules to block traffic from unknown sources.

Answer: C

Explanation:

To protect the web application from DDoS attacks originating from random IP addresses, a solutions architect should subscribe to AWS Shield Advanced and engage the AWS DDoS Response Team (DRT) to integrate mitigating controls into the service. AWS Shield Advanced is a managed service that provides protection against large and sophisticated DDoS attacks, with access to 24/7 support and response from the DRT. The DRT can help the city configure proactive and reactive safeguards, such as AWS WAF rules, rate-based rules, and network ACLs, to block malicious traffic and improve the application's resilience. The service also provides an audit trail for the DDoS sources through detailed attack reports and Amazon CloudWatch metrics.

NEW QUESTION 21

- (Topic 4)

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

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

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

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 23

- (Topic 4)

A company built an application with Docker containers and needs to run the application in the AWS Cloud. The company wants to use a managed service to host the application.

The solution must scale in and out appropriately according to demand on the individual container services. The solution also must not result in additional operational overhead or infrastructure to manage.

Which solutions will meet these requirements? (Select TWO)

- A. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.
- B. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate.
- C. Provision an Amazon API Gateway API. Connect the API to AWS Lambda to run the containers.
- D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes.
- E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes.

Answer: AB

Explanation:

These options are the best solutions because they allow the company to run the application with Docker containers in the AWS Cloud using a managed service that scales automatically and does not require any infrastructure to manage. By using AWS Fargate, the company can launch and run containers without having to provision, configure, or scale clusters of EC2 instances. Fargate allocates the right amount of compute resources for each container and scales them up or down as needed. By using Amazon ECS or Amazon EKS, the company can choose the container orchestration platform that suits its needs. Amazon ECS is a fully managed service that integrates with other AWS services and simplifies the deployment and management of containers. Amazon EKS is a managed service that runs Kubernetes on AWS and provides compatibility with existing Kubernetes tools and plugins.

* C. Provision an Amazon API Gateway API. Connect the API to AWS Lambda to run the containers. This option is not feasible because AWS Lambda does not support running Docker containers directly. Lambda functions are executed in a sandboxed environment that is isolated from other functions and resources. To run Docker containers on Lambda, the company would need to use a custom runtime or a wrapper library that emulates the Docker API, which can introduce additional complexity and overhead.

* D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes. This option is not optimal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.

* E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes. This option is not ideal because it requires the company to manage the EC2 instances that host the containers. The company would need to provision, configure, scale, patch, and monitor the EC2 instances, which can increase the operational overhead and infrastructure costs.

References:

- ? 1 AWS Fargate - Amazon Web Services
- ? 2 Amazon Elastic Container Service - Amazon Web Services
- ? 3 Amazon Elastic Kubernetes Service - Amazon Web Services
- ? 4 AWS Lambda FAQs - Amazon Web Services

NEW QUESTION 25

- (Topic 4)

A solutions architect needs to copy files from an Amazon S3 bucket to an Amazon Elastic File System (Amazon EFS) file system and another S3 bucket. The files must be copied continuously. New files are added to the original S3 bucket consistently. The copied files should be overwritten only if the source file changes.

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

- A. Create an AWS DataSync location for both the destination S3 bucket and the EFS file system.
- B. Create a task for the destination S3 bucket and the EFS file system.
- C. Set the transfer mode to transfer only data that has changed.
- D. Create an AWS Lambda function.
- E. Mount the file system to the function.
- F. Set up an S3 event notification to invoke the function when files are created and changed in Amazon S3. Configure the function to copy files to the file system and the destination S3 bucket.
- G. Create an AWS DataSync location for both the destination S3 bucket and the EFS file system.
- H. Create a task for the destination S3 bucket and the EFS file system.
- I. Set the transfer mode to transfer all data.
- J. Launch an Amazon EC2 instance in the same VPC as the file system.
- K. Mount the file system.
- L. Create a script to routinely synchronize all objects that changed in the origin S3 bucket to the destination S3 bucket and the mounted file system.

Answer: A

Explanation:

AWS DataSync is a service that makes it easy to move large amounts of data between AWS storage services and on-premises storage systems. AWS DataSync can copy files from an S3 bucket to an EFS file system and another S3 bucket continuously, as well as overwrite only the files that have changed in the source. This solution will meet the requirements with the least operational overhead, as it does not require any code development or manual intervention.

References:

- ? 4 explains how to create AWS DataSync locations for different storage services.

- ? 5 describes how to create and configure AWS DataSync tasks for data transfer.
- ? 6 discusses the different transfer modes that AWS DataSync supports.

NEW QUESTION 27

- (Topic 4)

A company hosts a data lake on Amazon S3. The data lake ingests data in Apache Parquet format from various data sources. The company uses multiple transformation steps to prepare the ingested data. The steps include filtering of anomalies, normalizing of data to standard date and time values, and generation of aggregates for analyses.

The company must store the transformed data in S3 buckets that data analysts access. The company needs a prebuilt solution for data transformation that does not require code. The solution must provide data lineage and data profiling. The company needs to share the data transformation steps with employees throughout the company.

Which solution will meet these requirements?

- A. Configure an AWS Glue Studio visual canvas to transform the data
- B. Share the transformation steps with employees by using AWS Glue jobs.
- C. Configure Amazon EMR Serverless to transform the data
- D. Share the transformation steps with employees by using EMR Serverless jobs.
- E. Configure AWS Glue DataBrew to transform the data
- F. Share the transformation steps with employees by using DataBrew recipes.
- G. Create Amazon Athena tables for the data
- H. Write Athena SQL queries to transform the data
- I. Share the Athena SQL queries with employees.

Answer: C

Explanation:

The most suitable solution for the company's requirements is to configure AWS Glue DataBrew to transform the data and share the transformation steps with employees by using DataBrew recipes. This solution will provide a prebuilt solution for data transformation that does not require code, and will also provide data lineage and data profiling. The company can easily share the data transformation steps with employees throughout the company by using DataBrew recipes. AWS Glue DataBrew is a visual data preparation tool that makes it easy for data analysts and data scientists to clean and normalize data for analytics or machine learning by up to 80% faster. Users can upload their data from various sources, such as Amazon S3, Amazon RDS, Amazon Redshift, Amazon Aurora, or Glue Data Catalog, and use a point-and-click interface to apply over 250 built-in transformations. Users can also preview the results of each transformation step and see how it affects the quality and distribution of the data¹.

A DataBrew recipe is a reusable set of transformation steps that can be applied to one or more datasets. Users can create recipes from scratch or use existing ones from the DataBrew recipe library. Users can also export, import, or share recipes with other users or groups within their AWS account or organization². DataBrew also provides data lineage and data profiling features that help users understand and improve their data quality. Data lineage shows the source and destination of each dataset and how it is transformed by each recipe step. Data profiling shows various statistics and metrics about each dataset, such as column

NEW QUESTION 28

- (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 31

- (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 connection
- 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 connection
- 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 35

- (Topic 4)

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:

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

? Lambda SnapStart for Java 11 Functions

? Lambda SnapStart FAQs

NEW QUESTION 40

- (Topic 4)

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

What should the solutions architect recommend?

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

Answer: B

Explanation:

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

? Using AWS Lambda with Amazon S3

? AWS CloudTrail FAQs

? What Is AWS AppSync?

? [What Is Amazon Kinesis Data Streams?]

? [What Is Amazon Simple Notification Service?]

NEW QUESTION 45

- (Topic 4)

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

? Viewing Events with CloudTrail Event History

? Querying AWS CloudTrail Logs

? Amazon Athena

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

Answer: BE

Explanation:

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

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

NEW QUESTION 51

- (Topic 4)

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

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

- A. Move the website images into an Amazon S3 bucket that is mounted on every EC2 instance.
- B. Share the website images by using an NFS share from the primary EC2 instanc
- C. Mounththis share on the other EC2 instances.
- D. Move the website images onto an Amazon Elastic File System (Amazon EFS) file system that is mounted on every EC2 instance.
- E. Create an Amazon Machine Image (AMI) from the existing EC2 instance Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling grou
- F. Configure the Auto Scaling group to maintain a minimum of two instance
- G. Configure an accelerator in AWS Global Accelerator for the website.
- H. Create an Amazon Machine Image (AMI) from the existing EC2 instanc
- I. Use the AMI to provision new instances behind an Application Load Balancer as part of an Auto Scaling group
- 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 53

- (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 54

- (Topic 4)

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

Which solution will allow the node to join the cluster?

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

Answer: B

Explanation:

Kubernetes API requests within your cluster's VPC (such as node to control plane communication) use the private VPC endpoint.
<https://docs.aws.amazon.com/eks/latest/userguide/cluster-endpoint.html>

NEW QUESTION 57

- (Topic 4)

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

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

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

Answer: B

Explanation:

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

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

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

References:

? Table classes - Amazon DynamoDB

? Read/write capacity mode - Amazon DynamoDB

NEW QUESTION 61

- (Topic 4)

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

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

NEW QUESTION 63

- (Topic 4)

A company hosts a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website serves static content. Website traffic is increasing.

and the company is concerned about a potential increase in cost.
What should a solutions architect do to reduce the cost of the website?

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

Answer: A

Explanation:

Amazon CloudFront is a content delivery network (CDN) that can improve the performance and reduce the cost of serving static content from a website. CloudFront

can cache static files at edge locations closer to the users, reducing the latency and data transfer costs. CloudFront can also integrate with Amazon S3 as the origin for the static content, eliminating the need for EC2 instances to host the website. CloudFront meets all the requirements of the question, while the other options do not. References:

? <https://aws.amazon.com/blogs/architecture/architecting-a-low-cost-web-content-publishing-system/>

? <https://nodeployfriday.com/posts/static-website-hosting/>

? <https://aws.amazon.com/cloudfront/>

NEW QUESTION 68

- (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 72

- (Topic 4)

A company has a popular gaming platform running on AWS. The application is sensitive to latency because latency can impact the user experience and introduce unfair advantages to some players. The application is deployed in every AWS Region. It runs on Amazon EC2 instances that are part of Auto Scaling groups configured behind Application Load Balancers (ALBs). A solutions architect needs to implement a mechanism to monitor the health of the application and redirect traffic to healthy endpoints.

Which solution meets these requirements?

- A. Configure an accelerator in AWS Global Accelerator.
- B. Add a listener for the port that the application listens on, and attach it to a Regional endpoint in each Region.
- C. Add the ALB as the endpoint.
- D. Create an Amazon CloudFront distribution and specify the ALB as the origin server.
- E. Configure the cache behavior to use origin cache header.
- F. Use AWS Lambda functions to optimize the traffic.
- G. Create an Amazon CloudFront distribution and specify Amazon S3 as the origin server.
- H. Configure the cache behavior to use origin cache header.
- I. Use AWS Lambda functions to optimize the traffic.
- J. Configure an Amazon DynamoDB database to serve as the data store for the application.
- K. Create a DynamoDB Accelerator (DAX) cluster to act as the in-memory cache for DynamoDB hosting the application data.

Answer: A

Explanation:

AWS Global Accelerator directs traffic to the optimal healthy endpoint based on health checks, it can also route traffic to the closest healthy endpoint based on geographic location of the client. By configuring an accelerator and attaching it to a Regional endpoint in each Region, and adding the ALB as the endpoint, the solution will redirect traffic to healthy endpoints, improving the user experience by reducing latency and ensuring that the application is running optimally. This solution will ensure that traffic is directed to the closest healthy endpoint and will help to improve the overall user experience.

NEW QUESTION 75

- (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 80

- (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 85

- (Topic 4)

A company has hired a solutions architect to design a reliable architecture for its application. The application consists of one Amazon RDS DB instance and two manually provisioned Amazon EC2 instances that run web servers. The EC2 instances are located in a single Availability Zone.

An employee recently deleted the DB instance, and the application was unavailable for 24 hours as a result. The company is concerned with the overall reliability of its environment.

What should the solutions architect do to maximize reliability of the application's infrastructure?

- A. Delete one EC2 instance and enable termination protection on the other EC2 instance
- B. Update the DB instance to be Multi-AZ, and enable deletion protection.
- C. Update the DB instance to be Multi-AZ, and enable deletion protection
- D. Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling group across multiple Availability Zones.
- E. Create an additional DB instance along with an Amazon API Gateway and an AWS Lambda function
- F. Configure the application to invoke the Lambda function through API Gateway
- G. Have the Lambda function write the data to the two DB instances.
- H. Place the EC2 instances in an EC2 Auto Scaling group that has multiple subnets located in multiple Availability Zone
- I. Use Spot Instances instead of On-Demand Instance
- J. Set up Amazon CloudWatch alarms to monitor the health of the instance
- K. Update the DB instance to be Multi-AZ, and enable deletion protection.

Answer: B

Explanation:

This answer is correct because it meets the requirements of maximizing the reliability of the application's infrastructure. You can update the DB instance to be Multi-AZ, which means that Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance. It can also help protect your databases

against DB instance failure and Availability Zone disruption. You can also enable deletion protection on the DB instance, which prevents the DB instance from being deleted by any user. You can place the EC2 instances behind an Application Load Balancer, which distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This increases the availability and fault tolerance of your applications. You can run the EC2 instances in an EC2 Auto Scaling group across multiple Availability Zones, which ensures that you have the correct number of EC2 instances available to handle the load for your application. You can use scaling policies to adjust the number of instances in your Auto Scaling group in response to changing demand.

References:

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

? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html#USER_DeleteInstance.DeletionProtection

? <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html>

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

NEW QUESTION 88

- (Topic 4)

A company runs a Java-based job on an Amazon EC2 instance. The job runs every hour and takes 10 seconds to run. The job runs on a scheduled interval and consumes 1 GB of memory. The CPU utilization of the instance is low except for short surges during which the job uses the maximum CPU available. The company wants to optimize the costs to run the job.

Which solution will meet these requirements?

- A. Use AWS App2Container (A2C) to containerize the job
- B. Run the job as an Amazon Elastic Container Service (Amazon ECS) task on AWS Fargate with 0.5 virtual CPU (vCPU) and 1 GB of memory.
- C. Copy the code into an AWS Lambda function that has 1 GB of memory
- D. Create an Amazon EventBridge scheduled rule to run the code each hour.
- E. Use AWS App2Container (A2C) to containerize the job
- F. Install the container in the existing Amazon Machine Image (AMI). Ensure that the schedule stops the container when the task finishes.
- G. Configure the existing schedule to stop the EC2 instance at the completion of the job and restart the EC2 instance when the next job starts.

Answer: B

Explanation:

AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers. You can create Lambda functions using various languages, including Java, and specify the amount of memory and CPU allocated to your function. Lambda charges you only for the compute time you consume, which is calculated based on the number of requests and the duration of your code execution. You can use Amazon EventBridge to trigger your Lambda function on a schedule, such as every hour, using cron or rate expressions. This solution will optimize the costs to run the job, as you will not pay for any idle time or unused resources, unlike running the job on an EC2 instance. References: 1: AWS Lambda - FAQs2, General Information section2: Tutorial: Schedule AWS Lambda functions using EventBridge3, Introduction section3: Schedule expressions using rate or cron - AWS Lambda4, Introduction section.

NEW QUESTION 92

- (Topic 4)

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

The permissions will be used by multiple IAM users and must be split between the developer and administrator teams. Each team requires different permissions. The company wants a solution that includes new users that are hired on both teams.

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

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

Answer: C

Explanation:

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

References:

? [Managing access to AWS accounts and applications](#)

? [Managing permissions sets](#)

? [Managing groups](#)

NEW QUESTION 95

- (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 99

- (Topic 4)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 103

- (Topic 4)

A serverless application uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The Lambda function needs permissions to read and write to the DynamoDB table.

Which solution will give the Lambda function access to the DynamoDB table MOST securely?

- A. Create an IAM user with programmatic access to the Lambda function
- B. Attach a policy to the user that allows read and write access to the DynamoDB table
- C. Store the access_key_id and secret_access_key parameters as part of the Lambda environment variable
- D. Ensure that other AWS users do not have read and write access to the Lambda function configuration
- E. Create an IAM role that includes Lambda as a trusted service
- F. Attach a policy to the role that allows read and write access to the DynamoDB table
- G. Update the configuration of the Lambda function to use the new role as the execution role.
- H. Create an IAM user with programmatic access to the Lambda function
- I. Attach a policy to the user that allows read and write access to the DynamoDB table
- J. Store the access_key_id and secret_access_key parameters in AWS Systems Manager Parameter Store as a secure string parameter
- K. Update the Lambda function code to retrieve the secure string parameters before connecting to the DynamoDB table.
- L. Create an IAM role that includes DynamoDB as a trusted service
- M. Attach a policy to the role that allows read and write access from the Lambda function
- N. Update the code of the Lambda function to attach to the new role as an execution role.

Answer: B

Explanation:

Option B suggests creating an IAM role that includes Lambda as a trusted service, meaning the role is specifically designed for Lambda functions. The role should have a policy attached to it that grants the required read and write access to the DynamoDB table.

NEW QUESTION 104

- (Topic 4)

A company has multiple AWS accounts with applications deployed in the us-west-2 Region. Application logs are stored within Amazon S3 buckets in each account. The company wants

to build a centralized log analysis solution that uses a single S3 bucket. Logs must not leave us-west-2, and the company wants to incur minimal operational overhead.

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

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

Answer: B

Explanation:

This solution meets the following requirements:

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

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

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

References:

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

NEW QUESTION 109

- (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 113

- (Topic 4)

A company wants to use artificial intelligence (AI) to determine the quality of its customer service calls. The company currently manages calls in four different languages, including English. The company will offer new languages in the future. The company does not have the resources to regularly maintain machine learning (ML) models.

The company needs to create written sentiment analysis reports from the customer service call recordings. The customer service call recording text must be translated into English.

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

- A. Use Amazon Comprehend to translate the audio recordings into English.
- B. Use Amazon Lex to create the written sentiment analysis reports.
- C. Use Amazon Polly to convert the audio recordings into text.
- D. Use Amazon Transcribe to convert the audio recordings in any language into text.
- E. Use Amazon Translate to translate text in any language to English.
- F. Use Amazon Comprehend to create the sentiment analysis reports.

Answer: DEF

Explanation:

These answers are correct because they meet the requirements of creating written sentiment analysis reports from the customer service call recordings in any language and translating them into English. Amazon Transcribe is a service that uses advanced machine learning technologies to recognize speech in audio files and transcribe them into text. You can use Amazon Transcribe to convert the audio recordings in any language into text, and specify the language code of the source audio. Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. You can use Amazon Translate to translate text in any language to English, and specify the source and target language codes. Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. You can use Amazon Comprehend to create the sentiment analysis reports, which determine if the text is positive, negative, neutral, or mixed.

References:

- ? <https://docs.aws.amazon.com/transcribe/latest/dg/what-is-transcribe.html>
- ? <https://docs.aws.amazon.com/translate/latest/dg/what-is.html>
- ? <https://docs.aws.amazon.com/comprehend/latest/dg/how-sentiment.html>

NEW QUESTION 118

- (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 121

- (Topic 4)

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

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

- A. Use Cost Explorer to create a daily report of costs by service
- B. Filter the report by EC2 instance
- C. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.
- D. Use Cost Explorer to create a monthly report of costs by service
- E. Filter the report by EC2 instance
- F. Configure Cost Explorer to send an Amazon Simple Email Service (Amazon SES) notification when a threshold is exceeded.

- G. Use AWS Budgets to create a cost budget for each account
- H. Set the period to monthl
- I. Set the scope to EC2 instance
- J. Set an alert threshold for the budge
- K. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.
- L. Use AWS Cost and Usage Reports to create a report with hourly granularit
- M. Integrate the report data with Amazon Athen
- N. Use Amazon EventBridge to schedule an Athena quer
- O. Configure an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification when a threshold is exceeded.

Answer: C

Explanation:

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

NEW QUESTION 126

- (Topic 4)

A company hosts multiple applications on AWS for different product lines. The applications use different compute resources, including Amazon EC2 instances and Application Load Balancers. The applications run in different AWS accounts under the same organization in AWS Organizations across multiple AWS Regions. Teams for each product line have tagged each compute resource in the individual accounts. The company wants more details about the cost for each product line from the consolidated billing feature in Organizations. Which combination of steps will meet these requirements? (Select TWO.)

- A. Select a specific AWS generated tag in the AWS Billing console.
- B. Select a specific user-defined tag in the AWS Billing console.
- C. Select a specific user-defined tag in the AWS Resource Groups console.
- D. Activate the selected tag from each AWS account.
- E. Activate the selected tag from the Organizations management account.

Answer: BE

Explanation:

User-defined tags are key-value pairs that can be applied to AWS resources to categorize and track them. User-defined tags can also be used to allocate costs and create detailed billing reports in the AWS Billing console. To use user-defined tags for cost allocation, the tags must be activated from the Organizations management account, which is the root account that has full control over all the member accounts in the organization. Once activated, the user-defined tags will appear as columns in the cost allocation report, and can be used to filter and group costs by product line. This solution will meet the requirements with the least operational overhead, as it leverages the existing tagging strategy and does not require any code development or manual intervention.

References:

- ? 1 explains how to use user-defined tags for cost allocation.
- ? 2 describes how to access and manage member accounts from the Organizations management account.
- ? 3 discusses how to create and view cost allocation reports in the AWS Billing console.

NEW QUESTION 131

- (Topic 4)

A company is building a three-tier application on AWS. The presentation tier will serve a static website. The logic tier is a containerized application. This application will store data in a relational database. The company wants to simplify deployment and to reduce operational costs. Which solution will meet these requirements?

- A. Use Amazon S3 to host static conten
- B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute powe
- C. Use a managed Amazon RDS cluster for the database.
- D. Use Amazon CloudFront to host static conten
- E. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 for compute powe
- F. Use a managed Amazon RDS cluster for the database.
- G. Use Amazon S3 to host static conten
- H. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute powe
- I. Use a managed Amazon RDS cluster for the database.
- J. Use Amazon EC2 Reserved Instances to host static conten
- K. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 for compute powe
- L. Use a managed Amazon RDS cluster for the database.

Answer: A

Explanation:

Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to host static content for your website, such as HTML files, images, videos, etc. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that allows you to run and scale containerized applications on AWS. AWS Fargate is a serverless compute engine for containers that works with both Amazon ECS and Amazon EKS. Fargate makes it easy for you to focus on building your applications by removing the need to provision and manage servers. You can use Amazon ECS with AWS Fargate for compute power for your containerized application logic tier. Amazon RDS is a managed relational database service that makes it easy to set up, operate, and scale a relational database in the cloud. You can use a managed Amazon RDS cluster for the database tier of your application. This solution will simplify deployment and reduce operational costs for your three-tier application. References:

- <https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>
- <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>
- <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 136

- (Topic 4)

A company is building a RESTful serverless web application on AWS by using Amazon API Gateway and AWS Lambda. The users of this web application will be geographically distributed, and the company wants to reduce the latency of API requests to these users Which type of endpoint should a solutions architect use to

meet these requirements?

- A. Private endpoint
- B. Regional endpoint
- C. Interface VPC endpoint
- D. Edge-optimized endpoint

Answer: D

Explanation:

An edge-optimized API endpoint is best for geographically distributed clients, as it routes the API requests to the nearest CloudFront Point of Presence (POP). This reduces the latency and improves the performance of the API. Edge-optimized endpoints are the default type for API Gateway REST APIs¹. A regional API endpoint is intended for clients in the same region as the API, and it does not use CloudFront to route the requests. A private API endpoint is an API endpoint that can only be accessed from a VPC using an interface VPC endpoint. A regional or private endpoint would not meet the requirement of reducing the latency for geographically distributed users¹.

NEW QUESTION 137

- (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 139

- (Topic 4)

A company runs container applications by using Amazon Elastic Kubernetes Service (Amazon EKS) and the Kubernetes Horizontal Pod Autoscaler. The workload is not

consistent throughout the day. A solutions architect notices that the number of nodes does not automatically scale out when the existing nodes have reached maximum capacity in the cluster, which causes performance issues

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

- A. Scale out the nodes by tracking the memory usage
- B. Use the Kubernetes Cluster Autoscaler to manage the number of nodes in the cluster.
- C. Use an AWS Lambda function to resize the EKS cluster automatically.
- D. Use an Amazon EC2 Auto Scaling group to distribute the workload.

Answer: B

Explanation:

The Kubernetes Cluster Autoscaler is a component that automatically adjusts the number of nodes in your cluster when pods fail or are rescheduled onto other nodes. It uses Auto Scaling groups to scale up or down the nodes according to the demand and capacity of your cluster¹.

By using the Kubernetes Cluster Autoscaler in your Amazon EKS cluster, you can achieve the following benefits:

- ? You can improve the performance and availability of your container applications by ensuring that there are enough nodes to run your pods and that there are no idle nodes wasting resources.
- ? You can reduce the administrative overhead of managing your cluster size manually or using custom scripts. The Cluster Autoscaler handles the scaling decisions and actions for you based on the metrics and events from your cluster.
- ? You can leverage the integration of Amazon EKS and AWS Auto Scaling to optimize the cost and efficiency of your cluster. You can use features such as launch templates, mixed instances policies, and spot instances to customize your node configuration and save up to 90% on compute costs²

NEW QUESTION 142

- (Topic 4)

A company is designing a tightly coupled high performance computing (HPC) environment in the AWS Cloud The company needs to include features that will optimize the HPC environment for networking and storage.

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

- A. Create an accelerator in AWS Global Accelerator
- B. Configure custom routing for the accelerator.
- C. Create an Amazon FSx for Lustre file system
- D. Configure the file system with scratch storage.

- E. Create an Amazon CloudFront distributio
- F. Configure the viewer protocol policy to be HTTP and HTTPS.
- G. Launch Amazon EC2 instance
- H. Attach an Elastic Fabric Adapter (EFA) to the instances.
- I. Create an AWS Elastic Beanstalk deployment to manage the environment.

Answer: BD

Explanation:

These two solutions will optimize the HPC environment for networking and storage. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. It is built on the world's most popular high-performance file system, Lustre, which is designed for applications that require fast storage, such as HPC and machine learning. By configuring the file system with scratch storage, you can achieve sub-millisecond latencies, up to hundreds of GBs/s of throughput, and millions of IOPS. Scratch file systems are ideal for temporary storage and shorter-term processing of data. Data is not replicated and does not persist if a file server fails. For more information, see Amazon FSx for Lustre.

Elastic Fabric Adapter (EFA) is a network interface for Amazon EC2 instances that enables customers to run applications requiring high levels of inter-node communications at scale on AWS. Its custom-built operating system (OS) bypass hardware interface enhances the performance of inter-instance communications, which is critical to scaling HPC and machine learning applications. EFA provides a low-latency, low-jitter channel for inter- instance communications, enabling your tightly-coupled HPC or distributed machine learning applications to scale to thousands of cores. EFA uses libfabric interface and libfabric APIs for communications, which are supported by most HPC programming models. For more information, see Elastic Fabric Adapter.

The other solutions are not suitable for optimizing the HPC environment for networking and storage. AWS Global Accelerator is a networking service that helps you improve the availability, performance, and security of your public applications by using the AWS global network. It provides two global static public IPs, deterministic routing, fast failover, and TCP termination at the edge for your application endpoints. However, it does not support OS- bypass capabilities or high-performance file systems that are required for HPC and machine learning applications. For more information, see AWS Global Accelerator. Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. CloudFront is integrated with AWS services such as Amazon S3, Amazon EC2, AWS Elemental Media Services, AWS Shield, AWS WAF, and AWS Lambda@Edge. However, CloudFront is not designed for HPC and machine learning applications that require high levels of inter-node communications and fast storage. For more information, see [Amazon CloudFront].

AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS. You can simply upload your code and Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring. However, Elastic Beanstalk is not optimized for HPC and machine learning applications that require OS-bypass capabilities and high-performance file systems. For more information, see [AWS Elastic Beanstalk].

References: Amazon FSx for Lustre, Elastic Fabric Adapter, AWS Global Accelerator, [Amazon CloudFront], [AWS Elastic Beanstalk].

NEW QUESTION 145

- (Topic 4)

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

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

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

Answer: BE

Explanation:

To predict the resources needed for manufacturing processes each month using historical values that are currently stored in an Amazon S3 bucket, a solutions architect should use Amazon SageMaker to train a model by using the historical data in the S3 bucket, and deploy an Amazon SageMaker model and create a SageMaker endpoint for inference. Amazon SageMaker is a fully managed service that provides an easy way to build, train, and deploy machine learning (ML) models. The solutions architect can use the built-in algorithms or frameworks provided by SageMaker, or bring their own custom code, to train a model using the historical data in the S3 bucket as input. The trained model can then be deployed to a SageMaker endpoint, which is a scalable and secure web service that can handle requests for predictions from the application. The solutions architect does not need to have any ML experience or manage any infrastructure to use SageMaker.

NEW QUESTION 146

- (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 151

- (Topic 4)

A company has deployed its application on Amazon EC2 instances with an Amazon RDS database. The company used the principle of least privilege to configure the database access credentials. The company's security team wants to protect the application and the database from SQL injection and other web-based attacks. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use security groups and network ACLs to secure the database and application servers.
- B. Use AWS WAF to protect the applicatio
- C. Use RDS parameter groups to configure thesecurity settings.
- D. Use AWS Network Firewall to protect the application and the database.
- E. Use different database accounts in the application code for different function
- F. Avoid granting excessive privileges to the database users.

Answer: B

Explanation:

AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF allows users to create rules that block, allow, or count web requests based on customizable web security rules. One of the types of rules that can be created is an SQL injection rule, which allows users to specify a list of IP addresses or IP address ranges that they want to allow or block. By using AWS WAF to protect the application, the company can prevent SQL injection and other web-based attacks from reaching the application and the database.

RDS parameter groups are collections of parameters that define how a database instance operates. Users can modify the parameters in a parameter group to change the behavior and performance of the database. By using RDS parameter groups to configure the security settings, the company can enforce best practices such as disabling remote root login, requiring SSL connections, and limiting the maximum number of connections.

The other options are not correct because they do not effectively protect the application and the database from SQL injection and other web-based attacks. Using security groups and network ACLs to secure the database and application servers is not sufficient because they only filter traffic at the network layer, not at the application layer. Using AWS Network Firewall to protect the application and the database is not necessary because it is a stateful firewall service that provides network protection for VPCs, not for individual applications or databases. Using different database accounts in the application code for different functions is a good practice, but it does not prevent SQL injection attacks from exploiting vulnerabilities in the application code.

References:

- ? AWS WAF
- ? How AWS WAF works
- ? Working with IP match conditions
- ? Working with DB parameter groups
- ? Amazon RDS security best practices

NEW QUESTION 152

- (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 Regio
- 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 Regio
- 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 155

- (Topic 4)

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

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

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

Answer: A

Explanation:

This answer is correct because it meets the requirements of accommodating the larger workload without adding infrastructure and minimizing the cost. Reserved DB instances are a billing discount applied to the use of certain on-demand DB instances in your account. Reserved DB instances provide you with a significant discount compared to on-demand DB instance pricing. You can buy reserved DB instances for the total workload and choose between three payment options: No

Upfront, Partial Upfront, or All Upfront. You can make the Amazon RDS for PostgreSQL DB instance larger by modifying its instance type to a higher performance class. This way, you can increase the CPU, memory, and network capacity of your DB instance and handle the increased workload. References:

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

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

NEW QUESTION 156

- (Topic 4)

A social media company is building a feature for its website. The feature will give users the ability to upload photos. The company expects significant increases in demand during large events and must ensure that the website can handle the upload traffic from users.

Which solution meets these requirements with the MOST scalability?

- A. Upload files from the user's browser to the application server
- B. Transfer the files to an Amazon S3 bucket.
- C. Provision an AWS Storage Gateway file gateway
- D. Upload files directly from the user's browser to the file gateway.
- E. Generate Amazon S3 presigned URLs in the application
- F. Upload files directly from the user's browser into an S3 bucket.
- G. Provision an Amazon Elastic File System (Amazon EFS) file system Upload files directly from the user's browser to the file system

Answer: C

Explanation:

This approach allows users to upload files directly to S3 without passing through the application servers, reducing the load on the application and improving scalability. It leverages the client-side capabilities to handle the file uploads and offloads the processing to S3.

NEW QUESTION 158

- (Topic 4)

A company provides an API interface to customers so the customers can retrieve their financial information. The company expects a larger number of requests during peak usage times of the year.

The company requires the API to respond consistently with low latency to ensure customer satisfaction. The company needs to provide a compute host for the API.

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

- A. Use an Application Load Balancer and Amazon Elastic Container Service (Amazon ECS).
- B. Use Amazon API Gateway and AWS Lambda functions with provisioned concurrency.
- C. Use an Application Load Balancer and an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- D. Use Amazon API Gateway and AWS Lambda functions with reserved concurrency.

Answer: B

Explanation:

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda scales automatically based on the incoming requests, but it may take some time to initialize new instances of your function if there is a sudden increase in demand. This may result in high latency or cold starts for your API. To avoid this, you can use provisioned concurrency, which ensures that your function is initialized and ready to respond at any time.

Provisioned concurrency also helps you achieve consistent low latency for your API by reducing the impact of scaling on performance. References:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-develop-integrations-lambda.html>

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html>

NEW QUESTION 163

- (Topic 4)

A company has deployed a database in Amazon RDS for MySQL. Due to increased transactions, the database support team is reporting slow reads against the DB instance and recommends adding a read replica.

Which combination of actions should a solutions architect take before implementing this change? (Choose two.)

- A. Enable binlog replication on the RDS primary node.
- B. Choose a failover priority for the source DB instance.
- C. Allow long-running transactions to complete on the source DB instance.
- D. Create a global table and specify the AWS Regions where the table will be available.
- E. Enable automatic backups on the source instance by setting the backup retention period to a value other than 0.

Answer: CE

Explanation:

"An active, long-running transaction can slow the process of creating the read replica. We recommend that you wait for long-running transactions to complete before creating a read replica. If you create multiple read replicas in parallel from the same source DB instance, Amazon RDS takes only one snapshot at the start of the first create action. When creating a read replica, there are a few things to consider. First, you must enable automatic backups on the source DB instance by setting the backup retention period to a value other than 0. This requirement also applies to a read replica that is the source DB instance for another read replica"

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

NEW QUESTION 166

- (Topic 4)

A company hosts multiple production applications. One of the applications consists of resources from Amazon EC2, AWS Lambda, Amazon RDS, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) across multiple AWS Regions. All company resources are tagged with a tag name of "application" and a value that corresponds to each application. A solutions architect must provide the quickest solution for identifying all of the tagged components.

Which solution meets these requirements?

- A. Use AWS CloudTrail to generate a list of resources with the application tag.
- B. Use the AWS CLI to query each service across all Regions to report the tagged components.
- C. Run a query in Amazon CloudWatch Logs Insights to report on the components with the application tag.
- D. Run a query with the AWS Resource Groups Tag Editor to report on the resources globally with the application tag.

Answer: D

Explanation:

<https://docs.aws.amazon.com/tag-editor/latest/userguide/tagging.html>

NEW QUESTION 171

- (Topic 4)

A company wants to analyze and troubleshoot Access Denied errors and Unauthorized errors that are related to IAM permissions. The company has AWS CloudTrail turned on. Which solution will meet these requirements with the LEAST effort?

- A. Use AWS Glue and write custom scripts to query CloudTrail logs for the errors.
- B. Use AWS Batch and write custom scripts to query CloudTrail logs for the errors.
- C. Search CloudTrail logs with Amazon Athena queries to identify the errors.
- D. Search CloudTrail logs with Amazon QuickSight.
- E. Create a dashboard to identify the errors.

Answer: C

Explanation:

This solution meets the following requirements:

? It is the least effort, as it does not require any additional AWS services, custom scripts, or data processing steps. Amazon Athena is a serverless interactive query service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to query CloudTrail logs directly from the S3 bucket where they are stored, without any data loading or transformation. You can also use the AWS Management Console, the AWS CLI, or the Athena API to run and manage your queries.

? It is effective, as it allows you to filter, aggregate, and join CloudTrail log data using SQL syntax. You can use various SQL functions and operators to specify the criteria for identifying Access Denied and Unauthorized errors, such as the error code, the user identity, the event source, the event name, the event time, and the resource ARN. You can also use subqueries, views, and common table expressions to simplify and optimize your queries.

? It is flexible, as it allows you to customize and save your queries for future use.

You can also export the query results to other formats, such as CSV or JSON, or integrate them with other AWS services, such as Amazon QuickSight, for further analysis and visualization.

References:

? Querying AWS CloudTrail Logs - Amazon Athena

? Analyzing Data in S3 using Amazon Athena | AWS Big Data Blog

? Troubleshoot IAM permission access denied or unauthorized errors | AWS re:Post

NEW QUESTION 175

.....

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](#)