

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

A company hosts a containerized web application on a fleet of on-premises servers that process incoming requests. The number of requests is growing quickly. The on-premises servers cannot handle the increased number of requests. The company wants to move the application to AWS with minimum code changes and minimum development effort.

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

- A. Use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling
- B. Use an Application Load Balancer to distribute the incoming requests.
- C. Use two Amazon EC2 instances to host the containerized web application
- D. Use an Application Load Balancer to distribute the incoming requests
- E. Use AWS Lambda with a new code that uses one of the supported languages
- F. Create multiple Lambda functions to support the load
- G. Use Amazon API Gateway as an entry point to the Lambda functions.
- H. Use a high performance computing (HPC) solution such as AWS ParallelCluster to establish an HPC cluster that can process the incoming requests at the appropriate scale.

Answer: A

Explanation:

AWS Fargate is a serverless compute engine that lets users run containers without having to manage servers or clusters of Amazon EC2 instances¹. Users can use AWS Fargate on Amazon Elastic Container Service (Amazon ECS) to run the containerized web application with Service Auto Scaling. Amazon ECS is a fully managed container orchestration service that supports both Docker and Kubernetes². Service Auto Scaling is a feature that allows users to adjust the desired number of tasks in an ECS service based on CloudWatch metrics, such as CPU utilization or request count³. Users can use AWS Fargate on Amazon ECS to migrate the application to AWS with minimum code changes and minimum development effort, as they only need to package their application in containers and specify the CPU and memory requirements.

Users can also use an Application Load Balancer to distribute the incoming requests. An Application Load Balancer is a load balancer that operates at the application layer and routes traffic to targets based on the content of the request. Users can register their ECS tasks as targets for an Application Load Balancer and configure listener rules to route requests to different target groups based on path or host headers. Users can use an Application Load Balancer to improve the availability and performance of their web application.

NEW QUESTION 2

- (Topic 1)

A company runs an online marketplace web application on AWS. The application serves hundreds of thousands of users during peak hours. The company needs a scalable, near-real-time solution to share the details of millions of financial transactions with several other internal applications. Transactions also need to be processed to remove sensitive data before being stored in a document database for low-latency retrieval.

What should a solutions architect recommend to meet these requirements?

- A. Store the transactions data into Amazon DynamoDB. Set up a rule in DynamoDB to remove sensitive data from every transaction upon write. Use DynamoDB Streams to share the transactions data with other applications.
- B. Stream the transactions data into Amazon Kinesis Data Firehose to store data in Amazon DynamoDB and Amazon S3. Use AWS Lambda integration with Kinesis Data Firehose to remove sensitive data.
- C. Other applications can consume the data stored in Amazon S3.
- D. Stream the transactions data into Amazon Kinesis Data Streams. Use AWS Lambda integration to remove sensitive data from every transaction and then store the transactions data in Amazon DynamoDB. Other applications can consume the transactions data off the Kinesis data stream.
- E. Store the batched transactions data in Amazon S3 as files.
- F. Use AWS Lambda to process every file and remove sensitive data before updating the files in Amazon S3. The Lambda function then stores the data in Amazon DynamoDB. Other applications can consume transaction files stored in Amazon S3.

Answer: C

Explanation:

The destination of your Kinesis Data Firehose delivery stream. Kinesis Data Firehose can send data records to various destinations, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, and any HTTP endpoint that is owned by you or any of your third-party service providers. The following are the supported destinations:

- * Amazon OpenSearch Service
- * Amazon S3
- * Datadog
- * Dynatrace
- * Honeycomb
- * HTTP Endpoint
- * Logic Monitor
- * MongoDB Cloud
- * New Relic
- * Splunk
- * Sumo Logic <https://docs.aws.amazon.com/firehose/latest/dev/create-name.html>

<https://aws.amazon.com/kinesis/data-streams/>

Amazon Kinesis Data Streams (KDS) is a massively scalable and durable real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events.

NEW QUESTION 3

- (Topic 1)

A company's application integrates with multiple software-as-a-service (SaaS) sources for data collection. The company runs Amazon EC2 instances to receive the data and to upload the data to an Amazon S3 bucket for analysis. The same EC2 instance that receives and uploads the data also sends a notification to the user when an upload is complete. The company has noticed slow application performance and wants to improve the performance as much as possible.

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

- A. Create an Auto Scaling group so that EC2 instances can scale out
- B. Configure an S3 event notification to send events to an Amazon Simple Notification Service (Amazon SNS) topic when the upload to the S3 bucket is complete.
- C. Create an Amazon AppFlow flow to transfer data between each SaaS source and the S3 bucket
- D. Configure an S3 event notification to send events to an Amazon Simple Notification Service (Amazon SNS) topic when the upload to the S3 bucket is complete.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for each SaaS source to send output data
- F. Configure the S3 bucket as the rule's target
- G. Create a second EventBridge (CloudWatch Events) rule to send events when the upload to the S3 bucket is complete
- H. Configure an Amazon Simple Notification Service (Amazon SNS) topic as the second rule's target.
- I. Create a Docker container to use instead of an EC2 instance
- J. Host the containerized application on Amazon Elastic Container Service (Amazon ECS). Configure Amazon CloudWatch Container Insights to send events to an Amazon Simple Notification Service (Amazon SNS) topic when the upload to the S3 bucket is complete.

Answer: B

Explanation:

Amazon AppFlow is a fully managed integration service that enables you to securely transfer data between Software-as-a-Service (SaaS) applications like Salesforce, SAP, Zendesk, Slack, and ServiceNow, and AWS services like Amazon S3 and Amazon Redshift, in just a few clicks.
<https://aws.amazon.com/appflow/>

NEW QUESTION 4

- (Topic 1)

A company hosts more than 300 global websites and applications. The company requires a platform to analyze more than 30 TB of clickstream data each day. What should a solutions architect do to transmit and process the clickstream data?

- A. Design an AWS Data Pipeline to archive the data to an Amazon S3 bucket and run an Amazon EMR cluster with the data to generate analytics
- B. Create an Auto Scaling group of Amazon EC2 instances to process the data and send it to an Amazon S3 data lake for Amazon Redshift to use for analysis
- C. Cache the data to Amazon CloudFront: Store the data in an Amazon S3 bucket. When an object is added to the S3 bucket, run an AWS Lambda function to process the data for analysis.
- D. Collect the data from Amazon Kinesis Data Stream
- E. Use Amazon Kinesis Data Firehose to transmit the data to an Amazon S3 data lake. Load the data in Amazon Redshift for analysis

Answer: D

Explanation:

<https://aws.amazon.com/es/blogs/big-data/real-time-analytics-with-amazon-redshift-streaming-ingestion/>

NEW QUESTION 5

- (Topic 1)

A development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.

Which solution meets these requirements MOST cost-effectively?

- A. Stop the DB instance when tests are complete
- B. Restart the DB instance when required.
- C. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.
- D. Create a snapshot when tests are complete
- E. Terminate the DB instance and restore the snapshot when required.
- F. Modify the DB instance to a low-capacity instance when tests are complete
- G. Modify the DB instance again when required.

Answer: A

Explanation:

To reduce the cost of running the tests without reducing the compute and memory attributes of the Amazon RDS for MySQL DB instance, the development team can stop the instance when tests are completed and restart it when required. Stopping the DB instance when not in use can help save costs because customers are only charged for storage while the DB instance is stopped. During this time, automated backups and automated DB instance maintenance are suspended. When the instance is restarted, it retains the same configurations, security groups, and DB parameter groups as when it was stopped.

Reference:

Amazon RDS Documentation: Stopping and Starting a DB instance (https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html)

NEW QUESTION 6

- (Topic 1)

A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins. Configure Route 53 to route traffic to the CloudFront distribution.
- B. Create an Amazon CloudFront distribution that has the ALB as an origin. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint.
- C. Configure Route 53 to route traffic to the CloudFront distribution.
- D. Create an Amazon CloudFront distribution that has the S3 bucket as an origin. Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints. Create a custom domain name that points to the accelerator DNS name. Use the custom domain name as an endpoint for the web application.
- E. Create an Amazon CloudFront distribution that has the ALB as an origin.
- F. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint. Create two domain names.
- G. Point one domain name to the CloudFront DNS name for dynamic content, Point the other domain name to the accelerator DNS name for static content. Use the domain names as endpoints for the web application.

Answer: C

Explanation:

Static content can be cached at Cloud front Edge locations from S3 and dynamic content EC2 behind the ALB whose performance can be improved by Global Accelerator whose one endpoint is ALB and other Cloud front. So with regards to custom domain name endpoint is web application is R53 alias records for the custom domain point to web application <https://aws.amazon.com/blogs/networking-and-content-delivery/improving-availability-and-performance-for-application-load-balancers-using-one-click-integration-with-aws-global-accelerator/>

NEW QUESTION 7

- (Topic 1)

A company is migrating a distributed application to AWS The application serves variable workloads The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes The company wants to modernize the application with a solution that maximizes resiliency and scalability. How should a solutions architect design the architecture to meet these requirements?

- A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group
- B. Configure EC2 Auto Scaling to use scheduled scaling
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs Implement the compute nodes with Amazon EC2 Instances that are managed in an Auto Scaling group Configure EC2 Auto Scaling based on the size of the queue
- D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed In an Auto Scaling group
- E. Configure AWS CloudTrail as a destination for the jobs Configure EC2 Auto Scaling based on the load on the primary server
- F. implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group Configure Amazon EventBridge (Amazon CloudWatch Events) as a destination for the jobs Configure EC2 Auto Scaling based on the load on the compute nodes

Answer: B

Explanation:

To maximize resiliency and scalability, the best solution is to use an Amazon SQS queue as a destination for the jobs. This decouples the primary server from the compute nodes, allowing them to scale independently. This also helps to prevent job loss in the event of a failure. Using an Auto Scaling group of Amazon EC2 instances for the compute nodes allows for automatic scaling based on the workload. In this case, it's recommended to configure the Auto Scaling group based on the size of the Amazon SQS queue, which is a better indicator of the actual workload than the load on the primary server or compute nodes. This approach ensures that the application can handle variable workloads, while also minimizing costs by automatically scaling up or down the compute nodes as needed.

NEW QUESTION 8

- (Topic 1)

A company stores call transcript files on a monthly basis. Users access the files randomly within 1 year of the call, but users access the files infrequently after 1 year. The company wants to optimize its solution by giving users the ability to query and retrieve files that are less than 1-year-old as quickly as possible. A delay in retrieving older files is acceptable.

Which solution will meet these requirements MOST cost-effectively?

- A. Store individual files with tags in Amazon S3 Glacier Instant Retrieval
- B. Query the tags to retrieve the files from S3 Glacier Instant Retrieval.
- C. Store individual files in Amazon S3 Intelligent-Tiering
- D. Use S3 Lifecycle policies to move the files to S3 Glacier Flexible Retrieval after 1 year
- E. Query and retrieve the files that are in Amazon S3 by using Amazon Athena
- F. Query and retrieve the files that are in S3 Glacier by using S3 Glacier Select.
- G. Store individual files with tags in Amazon S3 Standard storage
- H. Store search metadata for each archive in Amazon S3 Standard storage
- I. Use S3 Lifecycle policies to move the files to S3 Glacier Instant Retrieval after 1 year
- J. Query and retrieve the files by searching for metadata from Amazon S3.
- K. Store individual files in Amazon S3 Standard storage
- L. Use S3 Lifecycle policies to move the files to S3 Glacier Deep Archive after 1 year
- M. Store search metadata in Amazon RD
- N. Query the files from Amazon RD
- O. Retrieve the files from S3 Glacier Deep Archive.

Answer: B

Explanation:

"For archive data that needs immediate access, such as medical images, news media assets, or genomics data, choose the S3 Glacier Instant Retrieval storage class, an archive storage class that delivers the lowest cost storage with milliseconds retrieval. For archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, such as backup or disaster recovery use cases, choose S3 Glacier Flexible Retrieval (formerly S3 Glacier), with retrieval in minutes or free bulk retrievals in 5- 12 hours." <https://aws.amazon.com/about-aws/whats-new/2021/11/amazon-s3-glacier-instant-retrieval-storage-class/>

NEW QUESTION 9

- (Topic 1)

A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability. What should a solutions architect do to meet these requirements?

- A. Create an AWS Lambda function to apply the patch to all EC2 instances.
- B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.
- C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.
- D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: B

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/about-windows-app-patching.html>

NEW QUESTION 10

- (Topic 1)

A company hosts its multi-tier applications on AWS. For compliance, governance, auditing, and security, the company must track configuration changes on its AWS resources and record a history of API calls made to these resources.

What should a solutions architect do to meet these requirements?

- A. Use AWS CloudTrail to track configuration changes and AWS Config to record API calls
- B. Use AWS Config to track configuration changes and AWS CloudTrail to record API calls
- C. Use AWS Config to track configuration changes and Amazon CloudWatch to record API calls
- D. Use AWS CloudTrail to track configuration changes and Amazon CloudWatch to record API calls

Answer: B

Explanation:

AWS Config is a fully managed service that allows the company to assess, audit, and evaluate the configurations of its AWS resources. It provides a detailed inventory of the resources in use and tracks changes to resource configurations. AWS Config can detect configuration changes and alert the company when changes occur. It also provides a historical view of changes, which is essential for compliance and governance purposes. AWS CloudTrail is a fully managed service that provides a detailed history of API calls made to the company's AWS resources. It records all API activity in the AWS account, including who made the API call, when the call was made, and what resources were affected by the call. This information is critical for security and auditing purposes, as it allows the company to investigate any suspicious activity that might occur on its AWS resources.

NEW QUESTION 10

- (Topic 1)

A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning.

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

- A. Configure the application to send the data to Amazon Kinesis Data Firehose.
- B. Use Amazon Simple Email Service (Amazon SES) to format the data and to send the report by email.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Glue job to query the application's API for the data.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data.
- E. Store the application data in Amazon S3. Create an Amazon Simple Notification Service (Amazon SNS) topic as an S3 event destination to send the report by

Answer: BD

Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/send-email-formatted.html>

* D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data. This step can be done using AWS Lambda to extract the shipping statistics and organize the data into an HTML format.

* B. Use Amazon Simple Email Service (Amazon SES) to format the data and send the report by email. This step can be done by using Amazon SES to send the report to multiple email addresses at the same time every morning.

Therefore, options D and B are the correct choices for this question. Option A is incorrect because Kinesis Data Firehose is not necessary for this use case. Option C is incorrect because AWS Glue is not required to query the application's API. Option E is incorrect because S3 event notifications cannot be used to send the report by email.

NEW QUESTION 11

- (Topic 1)

A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images. Which method is the MOST cost-effective for hosting the website?

- A. Containerize the website and host it in AWS Fargate.
- B. Create an Amazon S3 bucket and host the website there.
- C. Deploy a web server on an Amazon EC2 instance to host the website.
- D. Configure an Application Load Balancer with an AWS Lambda target that uses the Express.js framework.

Answer: B

Explanation:

In Static Websites, Web pages are returned by the server which are prebuilt. They use simple languages such as HTML, CSS, or JavaScript.

There is no processing of content on the server (according to the user) in Static Websites. Web pages are returned by the server with no change therefore, static Websites are fast.

There is no interaction with databases.

Also, they are less costly as the host does not need to support server-side processing with different languages.

=====

In Dynamic Websites, Web pages are returned by the server which are processed during runtime means they are not prebuilt web pages but they are built during runtime according to the user's demand.

These use server-side scripting languages such as PHP, Node.js, ASP.NET and many more supported by the server.

So, they are slower than static websites but updates and interaction with databases are possible.

NEW QUESTION 13

- (Topic 1)

A company provides a Voice over Internet Protocol (VoIP) service that uses UDP connections. The service consists of Amazon EC2 instances that run in an Auto Scaling group. The company has deployments across multiple AWS Regions.

The company needs to route users to the Region with the lowest latency. The company also needs automated failover between Regions.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) and an associated target group
- B. Associate the target group with the Auto Scaling group

- C. Use the NLB as an AWS Global Accelerator endpoint in each Region.
- D. Deploy an Application Load Balancer (ALB) and an associated target group
- E. Associate the target group with the Auto Scaling group
- F. Use the ALB as an AWS Global Accelerator endpoint in each Region.
- G. Deploy a Network Load Balancer (NLB) and an associated target group
- H. Associate the target group with the Auto Scaling group
- I. Create an Amazon Route 53 latency record that points to aliases for each NL
- J. Create an Amazon CloudFront distribution that uses the latency record as an origin.
- K. Deploy an Application Load Balancer (ALB) and an associated target group
- L. Associate the target group with the Auto Scaling group
- M. Create an Amazon Route 53 weighted record that points to aliases for each AL
- N. Deploy an Amazon CloudFront distribution that uses the weighted record as an origin.

Answer: D

Explanation:

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

HTTP /HTTPS - ALB ; TCP and UDP - NLB; Lowest latency routing and more throughput. Also supports failover, uses Anycast Ip addressing - Global Accelerator Caching at Edge Locations – Cloudfont

WS Global Accelerator automatically checks the health of your applications and routes user traffic only to healthy application endpoints. If the health status changes or you make configuration updates, AWS Global Accelerator reacts instantaneously to route your users to the next available endpoint..

NEW QUESTION 14

- (Topic 1)

A company wants to run its critical applications in containers to meet requirements for scalability and availability The company prefers to focus on maintenance of the critical applications The company does not want to be responsible for provisioning and managing the underlying infrastructure that runs the containerized workload

What should a solutions architect do to meet those requirements?

- A. Use Amazon EC2 Instances, and Install Docker on the Instances
- B. Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 worker nodes
- C. Use Amazon Elastic Container Service (Amazon ECS) on AWS Fargate
- D. Use Amazon EC2 instances from an Amazon Elastic Container Service (Amazon ECS)- optimized Amazon Machine Image (AMI).

Answer: C

Explanation:

using AWS ECS on AWS Fargate since they requirements are for scalability and availability without having to provision and manage the underlying infrastructure to run the containerized workload. <https://docs.aws.amazon.com/AmazonECS/latest/userguide/what-is-fargate.html>

NEW QUESTION 18

- (Topic 1)

A company's containerized application runs on an Amazon EC2 instance. The application needs to download security certificates before it can communicate with other business applications. The company wants a highly secure solution to encrypt and decrypt the certificates in near real time. The solution also needs to store data in highly available storage after the data is encrypted.

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

- A. Create AWS Secrets Manager secrets for encrypted certificate
- B. Manually update the certificates as needed
- C. Control access to the data by using fine-grained IAM access.
- D. Create an AWS Lambda function that uses the Python cryptography library to receive and perform encryption operation
- E. Store the function in an Amazon S3 bucket.
- F. Create an AWS Key Management Service (AWS KMS) customer managed key
- G. Allow the EC2 role to use the KMS key for encryption operation
- H. Store the encrypted data on Amazon S3.
- I. Create an AWS Key Management Service (AWS KMS) customer managed key
- J. Allow the EC2 role to use the KMS key for encryption operation
- K. Store the encrypted data on Amazon Elastic Block Store (Amazon EBS) volumes.

Answer: D

NEW QUESTION 21

- (Topic 1)

A company runs an ecommerce application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales based on CPU utilization metrics. The ecommerce application stores the transaction data in a MySQL 8.0 database that is hosted on a large EC2 instance.

The database's performance degrades quickly as application load increases. The application handles more read requests than write transactions. The company wants a solution that will automatically scale the database to meet the demand of unpredictable read workloads while maintaining high availability.

Which solution will meet these requirements?

- A. Use Amazon Redshift with a single node for leader and compute functionality.
- B. Use Amazon RDS with a Single-AZ deployment Configure Amazon RDS to add reader instances in a different Availability Zone.
- C. Use Amazon Aurora with a Multi-AZ deployment
- D. Configure Aurora Auto Scaling with Aurora Replicas.
- E. Use Amazon ElastiCache for Memcached with EC2 Spot Instances.

Answer: C

Explanation:

AURORA is 5x performance improvement over MySQL on RDS and handles more read requests than write,; maintaining high availability = Multi-AZ deployment

NEW QUESTION 22

- (Topic 1)

A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege.

Which solution will meet these requirements?

- A. Share the dashboard from the CloudWatch console
- B. Enter the product manager's email address, and complete the sharing step
- C. Provide a shareable link for the dashboard to the product manager.
- D. Create an IAM user specifically for the product manager
- E. Attach the CloudWatch Read Only Access managed policy to the user
- F. Share the new login credential with the product manager
- G. Share the browser URL of the correct dashboard with the product manager.
- H. Create an IAM user for the company's employees, Attach the View Only Access AWS managed policy to the IAM user
- I. Share the new login credentials with the product manager
- J. Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.
- K. Deploy a bastion server in a public subnet
- L. When the product manager requires access to the dashboard, start the server and share the RDP credential
- M. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Answer: B

Explanation:

To provide the product manager access to the Amazon CloudWatch dashboard while following the principle of least privilege, a solution architect should create an IAM user specifically for the product manager and attach the CloudWatch Read Only Access managed policy to the user. This policy allows the user to view the dashboard without being able to make any changes to it. The solution architect should then share the new login credential with the product manager and provide them with the browser URL of the correct dashboard.

NEW QUESTION 26

- (Topic 1)

A company has a large Microsoft SharePoint deployment running on-premises that requires Microsoft Windows shared file storage. The company wants to migrate this workload to the AWS Cloud and is considering various storage options. The storage solution must be highly available and integrated with Active Directory for access control.

Which solution will satisfy these requirements?

- A. Configure Amazon EFS storage and set the Active Directory domain for authentication
- B. Create an SMB Me share on an AWS Storage Gateway file gateway in two Availability Zones
- C. Create an Amazon S3 bucket and configure Microsoft Windows Server to mount it as a volume
- D. Create an Amazon FSx for Windows File Server file system on AWS and set the Active Directory domain for authentication

Answer: D

NEW QUESTION 31

- (Topic 1)

A company has created an image analysis application in which users can upload photos and add photo frames to their images. The users upload images and metadata to indicate which photo frames they want to add to their images. The application uses a single Amazon EC2 instance and Amazon DynamoDB to store the metadata.

The application is becoming more popular, and the number of users is increasing. The company expects the number of concurrent users to vary significantly depending on the time of day and day of week. The company must ensure that the application can scale to meet the needs of the growing user base.

Which solution meets these requirements?

- A. Use AWS Lambda to process the photo
- B. Store the photos and metadata in DynamoDB.
- C. Use Amazon Kinesis Data Firehose to process the photos and to store the photos and metadata.
- D. Use AWS Lambda to process the photo
- E. Store the photos in Amazon S3. Retain DynamoDB to store the metadata.
- F. Increase the number of EC2 instances to three
- G. Use Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes to store the photos and metadata.

Answer: C

Explanation:

<https://www.quora.com/How-can-I-use-DynamoDB-for-storing-metadata-for-Amazon-S3-objects>

This solution meets the requirements of scalability, performance, and availability. AWS Lambda can process the photos in parallel and scale up or down automatically depending on the demand. Amazon S3 can store the photos and metadata reliably and durably, and provide high availability and low latency. DynamoDB can store the metadata efficiently and provide consistent performance. This solution also reduces the cost and complexity of managing EC2 instances and EBS volumes.

Option A is incorrect because storing the photos in DynamoDB is not a good practice, as it can increase the storage cost and limit the throughput. Option B is incorrect because Kinesis Data Firehose is not designed for processing photos, but for streaming data to destinations such as S3 or Redshift. Option D is incorrect because increasing the number of EC2 instances and using Provisioned IOPS SSD volumes does not guarantee scalability, as it depends on the load balancer and the application code. It also increases the cost and complexity of managing the infrastructure.

References:

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

? <https://www.examtips.com/discussions/amazon/view/7193-exam-aws-certified-solutions-architect-professional-topic-1/>

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

NEW QUESTION 35

- (Topic 1)

A company wants to migrate an on-premises data center to AWS. The data center hosts an SFTP server that stores its data on an NFS-based file system. The server holds 200 GB of data that needs to be transferred. The server must be hosted on an Amazon EC2 instance that uses an Amazon Elastic File System (Amazon EFS) file system

When combination of steps should a solutions architect take to automate this task? (Select TWO)

- A. Launch the EC2 instance into the same Availability Zone as the EFS file system
- B. Install an AWS DataSync agent in the on-premises data center
- C. Create a secondary Amazon Elastic Block Store (Amazon EBS) volume on the EC2 instance for the data
- D. Manually use an operating system copy command to push the data to the EC2 instance
- E. Use AWS DataSync to create a suitable location configuration for the on-premises SFTP server

Answer: BE

Explanation:

AWS DataSync is an online data movement and discovery service that simplifies data migration and helps users quickly, easily, and securely move their file or object data to, from, and between AWS storage services¹. Users can use AWS DataSync to transfer data between on-premises and AWS storage services. To use AWS DataSync, users need to install an AWS DataSync agent in the on-premises data center. The agent is a software appliance that connects to the source or destination storage system and handles the data transfer to or from AWS over the network². Users also need to use AWS DataSync to create a suitable location configuration for the on-premises SFTP server. A location is a logical representation of a storage system that contains files or objects that users want to transfer using DataSync. Users can create locations for NFS shares, SMB shares, HDFS file systems, self-managed object storage, Amazon S3 buckets, Amazon EFS file systems, Amazon FSx for Windows File Server file systems, Amazon FSx for Lustre file systems, Amazon FSx for OpenZFS file systems, Amazon FSx for NetApp ONTAP file systems, and AWS Snowcone devices³.

NEW QUESTION 36

- (Topic 1)

A company has a data ingestion workflow that consists of the following:

- ? An Amazon Simple Notification Service (Amazon SNS) topic for notifications about new data deliveries
- ? An AWS Lambda function to process the data and record metadata

The company observes that the ingestion workflow fails occasionally because of network connectivity issues. When such a failure occurs, the Lambda function does not ingest the corresponding data unless the company manually reruns the job.

Which combination of actions should a solutions architect take to ensure that the Lambda function ingests all data in the future? (Select TWO.)

- A. Configure the Lambda function in multiple Availability Zones.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue, and subscribe it to the SNS topic.
- C. Increase the CPU and memory that are allocated to the Lambda function.
- D. Increase the provisioned throughput for the Lambda function.
- E. Modify the Lambda function to read from an Amazon Simple Queue Service (Amazon SQS) queue

Answer: BE

Explanation:

To ensure that the Lambda function ingests all data in the future despite occasional network connectivity issues, the following actions should be taken:

- ? Create an Amazon Simple Queue Service (SQS) queue and subscribe it to the SNS topic. This allows for decoupling of the notification and processing, so that even if the processing Lambda function fails, the message remains in the queue for further processing later.
- ? Modify the Lambda function to read from the SQS queue instead of directly from SNS. This decoupling allows for retries and fault tolerance and ensures that all messages are processed by the Lambda function.

Reference:

AWS SNS documentation: <https://aws.amazon.com/sns/> AWS SQS documentation: <https://aws.amazon.com/sqs/>

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

NEW QUESTION 41

- (Topic 1)

A company needs to keep user transaction data in an Amazon DynamoDB table. The company must retain the data for 7 years.

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

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Use AWS Backup to create backup schedules and retention policies for the table.
- C. Create an on-demand backup of the table by using the DynamoDB console
- D. Store the backup in an Amazon S3 bucket
- E. Set an S3 Lifecycle configuration for the S3 bucket.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function
- G. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket
- H. Set an S3 Lifecycle configuration for the S3 bucket.

Answer: C

NEW QUESTION 45

- (Topic 1)

An e-commerce company wants to launch a one-deal-a-day website on AWS. Each day will feature exactly one product on sale for a period of 24 hours. The company wants to be able to handle millions of requests each hour with millisecond latency during peak hours.

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

- A. Use Amazon S3 to host the full website in different S3 buckets. Add Amazon CloudFront distributions. Set the S3 buckets as origins for the distributions. Store the order data in Amazon S3
- B. Deploy the full website on Amazon EC2 instances that run in Auto Scaling groups across multiple Availability Zones. Add an Application Load Balancer (ALB) to distribute the website traffic. Add another ALB for the backend APIs. Store the data in Amazon RDS for MySQL
- C. Migrate the full application to run in containers. Host the containers on Amazon Elastic Kubernetes Service (Amazon EKS). Use the Kubernetes Cluster Autoscaler to increase and decrease the number of pods to process bursts in traffic. Store the data in Amazon RDS for MySQL

- D. Use an Amazon S3 bucket to host the website's static content Deploy an Amazon CloudFront distributio
- E. Set the S3 bucket as the origin Use Amazon API Gateway and AWS Lambda functions for the backend APIs Store the data in Amazon DynamoDB

Answer: D

Explanation:

To launch a one-deal-a-day website on AWS with millisecond latency during peak hours and with the least operational overhead, the best option is to use an Amazon S3 bucket to host the website's static content, deploy an Amazon CloudFront distribution, set the S3 bucket as the origin, use Amazon API Gateway and AWS Lambda functions for the backend APIs, and store the data in Amazon DynamoDB. This option requires minimal operational overhead and can handle millions of requests each hour with millisecond latency during peak hours. Therefore, option D is the correct answer.

Reference: <https://aws.amazon.com/blogs/compute/building-a-serverless-multi-player-game-with-aws-lambda-and-amazon-dynamodb/>

NEW QUESTION 50

- (Topic 1)

A company has registered its domain name with Amazon Route 53. The company uses Amazon API Gateway in the ca-central-1 Region as a public interface for its backend microservice APIs. Third-party services consume the APIs securely. The company wants to design its API Gateway URL with the company's domain name and corresponding certificate so that the third-party services can use HTTPS.

Which solution will meet these requirements?

- A. Create stage variables in API Gateway with Name="Endpoint-URL" and Value="Company Domain Name" to overwrite the default UR
- B. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM).
- C. Create Route 53 DNS records with the company's domain nam
- D. Point the alias record to the Regional API Gateway stage endpoint
- E. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Region.
- F. Create a Regional API Gateway endpoint
- G. Associate the API Gateway endpoint with the company's domain nam
- H. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Regio
- I. Attach the certificate to the API Gateway endpoint
- J. Configure Route 53 to route traffic to the API Gateway endpoint.
- K. Create a Regional API Gateway endpoint
- L. Associate the API Gateway endpoint with the company's domain nam
- M. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the us-east-1 Regio
- N. Attach the certificate to the API Gateway API
- O. Create Route 53 DNS records with the company's domain nam
- P. Point an A record to the company's domain name.

Answer: C

Explanation:

To design the API Gateway URL with the company's domain name and corresponding certificate, the company needs to do the following: 1. Create a Regional API Gateway endpoint: This will allow the company to create an endpoint that is specific to a region. 2. Associate the API Gateway endpoint with the company's domain name: This will allow the company to use its own domain name for the API Gateway URL. 3. Import the public certificate associated with the company's domain name into AWS Certificate Manager (ACM) in the same Region: This will allow the company to use HTTPS for secure communication with its APIs. 4. Attach the certificate to the API Gateway endpoint: This will allow the company to use the certificate for securing the API Gateway URL. 5. Configure Route 53 to route traffic to the API Gateway endpoint: This will allow the company to use Route 53 to route traffic to the API Gateway URL using the company's domain name.

NEW QUESTION 51

- (Topic 1)

A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.

The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.

Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to extend the company's storage spac
- C. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
- D. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
- E. Install a utility on each user's computer to access Amazon S3. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: B

Explanation:

Amazon S3 File Gateway is a hybrid cloud storage service that enables on- premises applications to seamlessly use Amazon S3 cloud storage. It provides a file interface to Amazon S3 and supports SMB and NFS protocols. It also supports S3 Lifecycle policies that can automatically transition data from S3 Standard to S3 Glacier Deep Archive after a specified period of time. This solution will meet the requirements of increasing the company's available storage space without losing low-latency access to the most recently accessed files and providing file lifecycle management to avoid future storage issues.

Reference:

<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

NEW QUESTION 55

- (Topic 1)

A company hosts an application on multiple Amazon EC2 instances The application processes messages from an Amazon SQS queue writes to an Amazon RDS table and deletes the message from the queue Occasional duplicate records are found in the RDS table. The SQS queue does not contain any duplicate messages.

What should a solutions architect do to ensure messages are being processed once only?

- A. Use the CreateQueue API call to create a new queue
- B. Use the Add Permission API call to add appropriate permissions

- C. Use the ReceiveMessage API call to set an appropriate wait time
- D. Use the ChangeMessageVisibility API call to increase the visibility timeout

Answer: D

Explanation:

The visibility timeout begins when Amazon SQS returns a message. During this time, the consumer processes and deletes the message. However, if the consumer fails before deleting the message and your system doesn't call the DeleteMessage action for that message before the visibility timeout expires, the message becomes visible to other consumers and the message is received again. If a message must be received only once, your consumer should delete it within the duration of the visibility timeout. <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html>

Keyword: SQS queue writes to an Amazon RDS From this, Option D best suite & other Options ruled out [Option A - You can't intruduce one more Queue in the existing one; Option B - only Permission & Option C - Only Retrieves Messages] FIFO queues are designed to never introduce duplicate messages. However, your message producer might introduce duplicates in certain scenarios: for example, if the producer sends a message, does not receive a response, and then resends the same message. Amazon SQS APIs provide deduplication functionality that prevents your message producer from sending duplicates. Any duplicates introduced by the message producer are removed within a 5-minute deduplication interval. For standard queues, you might occasionally receive a duplicate copy of a message (at-least- once delivery). If you use a standard queue, you must design your applications to be idempotent (that is, they must not be affected adversely when processing the same message more than once).

NEW QUESTION 56

- (Topic 1)

A solutions architect must design a highly available infrastructure for a website. The website is powered by Windows web servers that run on Amazon EC2 instances. The solutions architect must implement a solution that can mitigate a large-scale DDoS attack that originates from thousands of IP addresses. Downtime is not acceptable for the website.

Which actions should the solutions architect take to protect the website from such an attack? (Select TWO.)

- A. Use AWS Shield Advanced to stop the DDoS attack.
- B. Configure Amazon GuardDuty to automatically block the attackers.
- C. Configure the website to use Amazon CloudFront for both static and dynamic content.
- D. Use an AWS Lambda function to automatically add attacker IP addresses to VPC network ACLs.
- E. Use EC2 Spot Instances in an Auto Scaling group with a target tracking scaling policy that is set to 80% CPU utilization

Answer: AC

Explanation:

(<https://aws.amazon.com/cloudfront>)

NEW QUESTION 60

- (Topic 1)

A company is planning to use an Amazon DynamoDB table for data storage. The company is concerned about cost optimization. The table will not be used on most mornings. In the evenings, the read and write traffic will often be unpredictable. When traffic spikes occur, they will happen very quickly. What should a solutions architect recommend?

- A. Create a DynamoDB table in on-demand capacity mode.
- B. Create a DynamoDB table with a global secondary index.
- C. Create a DynamoDB table with provisioned capacity and auto scaling.
- D. Create a DynamoDB table in provisioned capacity mode, and configure it as a global table.

Answer: A

NEW QUESTION 64

- (Topic 1)

A company has more than 5 TB of file data on Windows file servers that run on premises. Users and applications interact with the data each day. The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency. The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS. What should a solutions architect do to meet these requirements?

- A. Deploy and configure Amazon FSx for Windows File Server on AWS
- B. Move the on-premises file data to FSx for Windows File Server
- C. Reconfigure the workloads to use FSx for Windows File Server on AWS.
- D. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to the S3 File Gateway. Reconfigure the on-premises workloads and the cloud workloads to use the S3 File Gateway
- E. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to Amazon S3. Reconfigure the workloads to use either Amazon S3 directly or the S3 File Gateway, depending on each workload's location
- F. Deploy and configure Amazon FSx for Windows File Server on AWS. Deploy and configure an Amazon FSx File Gateway on premises. Move the on-premises file data to the FSx File Gateway. Configure the cloud workloads to use FSx for Windows File Server on AWS. Configure the on-premises workloads to use the FSx File Gateway

Answer: D

Explanation:

<https://docs.aws.amazon.com/filegateway/latest/filefsxw/what-is-file-fsxw.html>

To meet the requirements of the company to have access to both AWS and on-premises file storage with minimum latency, a hybrid cloud architecture can be used. One solution is to deploy and configure Amazon FSx for Windows File Server on AWS, which provides fully managed Windows file servers. The on-premises file data can be moved to the FSx File Gateway, which can act as a bridge between on-premises and AWS file storage. The cloud workloads can be configured to use FSx for Windows File Server on AWS, while the on-premises workloads can be configured to use the FSx File Gateway. This solution minimizes operational overhead and requires no significant changes to the existing file access patterns. The connectivity between on-premises and AWS can be established using an AWS Site-to-Site VPN connection.

Reference:

AWS FSx for Windows File Server: <https://aws.amazon.com/fsx/windows/> AWS FSx File Gateway: <https://aws.amazon.com/fsx/file-gateway/>

AWS Site-to-Site VPN: <https://aws.amazon.com/vpn/site-to-site-vpn/>

NEW QUESTION 65

- (Topic 1)

A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.

What should the solutions architect recommend?

- A. Launch an Amazon EC2 instance in us-east-1 and migrate the site to it.
- B. Move the website to Amazon S3. Use cross-Region replication between Regions.
- C. Use Amazon CloudFront with a custom origin pointing to the on-premises servers.
- D. Use an Amazon Route 53 geo-proximity routing policy pointing to on-premises servers.

Answer: C

Explanation:

<https://aws.amazon.com/pt/blogs/aws/amazon-cloudfront-support-for-custom-origins/>

You can now create a CloudFront distribution using a custom origin. Each distribution can point to an S3 or to a custom origin. This could be another storage service, or it could be something more interesting and more dynamic, such as an EC2 instance or even an Elastic Load Balancer

NEW QUESTION 68

- (Topic 1)

A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes. The application data must be stored in a standard file system structure. The company wants a solution that scales automatically, is highly available, and requires minimum operational overhead.

Which solution will meet these requirements?

- A. Migrate the application to run as containers on Amazon Elastic Container Service (Amazon ECS). Use Amazon S3 for storage.
- B. Migrate the application to run as containers on Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon Elastic Block Store (Amazon EBS) for storage.
- C. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling group.
- D. Use Amazon Elastic File System (Amazon EFS) for storage.
- E. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling group.
- F. Use Amazon Elastic Block Store (Amazon EBS) for storage.

Answer: C

Explanation:

EFS is a standard file system, it scales automatically and is highly available.

NEW QUESTION 69

- (Topic 1)

A company is developing a two-tier web application on AWS. The company's developers have deployed the application on an Amazon EC2 instance that connects directly to a backend Amazon RDS database. The company must not hardcode database credentials in the application. The company must also implement a solution to automatically rotate the database credentials on a regular basis.

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

- A. Store the database credentials in the instance metadata.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) rules to run a scheduled AWS Lambda function that updates the RDS credentials and instance metadata at the same time.
- C. Store the database credentials in a configuration file in an encrypted Amazon S3 bucket.
- D. Use Amazon EventBridge (Amazon CloudWatch Events) rules to run a scheduled AWS Lambda function that updates the RDS credentials and the credentials in the configuration file at the same time.
- E. Use S3 Versioning to ensure the ability to fall back to previous values.
- F. Store the database credentials as a secret in AWS Secrets Manager.
- G. Turn on automatic rotation for the secret.
- H. Attach the required permission to the EC2 role to grant access to the secret.
- I. Store the database credentials as encrypted parameters in AWS Systems Manager Parameter Store.
- J. Turn on automatic rotation for the encrypted parameter.
- K. Attach the required permission to the EC2 role to grant access to the encrypted parameters.

Answer: C

Explanation:

https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_database_secret.html

NEW QUESTION 72

- (Topic 1)

A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website.

The company needs to reduce coupling within the application and improve website performance. A solutions architect must design the most operationally efficient process for image uploads.

Which combination of actions should the solutions architect take to meet these requirements? (Choose two.)

- A. Configure the application to upload images to S3 Glacier.
- B. Configure the web server to upload the original images to Amazon S3.
- C. Configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL.
- D. Configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded.
- E. Use the function to resize the image.

F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function on a schedule to resize uploaded images.

Answer: CD

Explanation:

Amazon S3 is a highly scalable and durable object storage service that can store and retrieve any amount of data from anywhere on the web¹. Users can configure the application to upload images directly from each user's browser to Amazon S3 through the use of a presigned URL. A presigned URL is a URL that gives access to an object in an S3 bucket for a limited time and with a specific action, such as uploading an object². Users can generate a presigned URL programmatically using the AWS SDKs or AWS CLI. By using a presigned URL, users can reduce coupling within the application and improve website performance, as they do not need to send the images to the web server first. AWS Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources³. Users can configure S3 Event Notifications to invoke an AWS Lambda function when an image is uploaded. S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion. Users can configure S3 Event Notifications to invoke a Lambda function that resizes the image and stores it back in the same or a different S3 bucket. This way, users can offload the image resizing task from the web server to Lambda.

NEW QUESTION 75

- (Topic 1)

A company receives 10 TB of instrumentation data each day from several machines located at a single factory. The data consists of JSON files stored on a storage area network (SAN) in an on-premises data center located within the factory. The company wants to send this data to Amazon S3 where it can be accessed by several additional systems that provide critical near-real-time analytics. A secure transfer is important because the data is considered sensitive. Which solution offers the MOST reliable data transfer?

- A. AWS DataSync over public internet
- B. AWS DataSync over AWS Direct Connect
- C. AWS Database Migration Service (AWS DMS) over public internet
- D. AWS Database Migration Service (AWS DMS) over AWS Direct Connect

Answer: B

Explanation:

These are some of the main use cases for AWS DataSync: • Data migration

– Move active datasets rapidly over the network into Amazon S3, Amazon EFS, or FSx for Windows File Server. DataSync includes automatic encryption and data integrity validation to help make sure that your data arrives securely, intact, and ready to use.

"DataSync includes encryption and integrity validation to help make sure your data arrives securely, intact, and ready to use."

<https://aws.amazon.com/datasync/faqs/>

NEW QUESTION 78

- (Topic 1)

A company is designing an application where users upload small files into Amazon S3. After a user uploads a file, the file requires one-time simple processing to transform the data and save the data in JSON format for later analysis.

Each file must be processed as quickly as possible after it is uploaded. Demand will vary. On some days, users will upload a high number of files. On other days, users will upload a few files or no files.

Which solution meets these requirements with the LEAST operational overhead?

- A. Configure Amazon EMR to read text files from Amazon S3. Run processing scripts to transform the data
- B. Store the resulting JSON file in an Amazon Aurora DB cluster.
- C. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue
- D. Use Amazon EC2 instances to read from the queue and process the data
- E. Store the resulting JSON file in Amazon DynamoDB.
- F. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue
- G. Use an AWS Lambda function to read from the queue and process the data
- H. Store the resulting JSON file in Amazon DynamoDB
- I. Most Voted
- J. Configure Amazon EventBridge (Amazon CloudWatch Events) to send an event to Amazon Kinesis Data Streams when a new file is uploaded
- K. Use an AWS Lambda function to consume the event from the stream and process the data
- L. Store the resulting JSON file in Amazon Aurora DB cluster.

Answer: C

Explanation:

Amazon S3 sends event notifications about S3 buckets (for example, object created, object removed, or object restored) to an SNS topic in the same Region. The SNS topic publishes the event to an SQS queue in the central Region.

The SQS queue is configured as the event source for your Lambda function and buffers the event messages for the Lambda function.

The Lambda function polls the SQS queue for messages and processes the Amazon S3 event notifications according to your application's requirements.

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/subscribe-a-lambda-function-to-event-notifications-from-s3-buckets-in-different-aws-regions.html>

NEW QUESTION 81

- (Topic 1)

A company is preparing to store confidential data in Amazon S3. For compliance reasons the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys must be rotated every year.

Which solution meets these requirements and is the MOST operationally efficient?

- A. Server-side encryption with customer-provided keys (SSE-C)
- B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
- D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automatic rotation

Answer: D

Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html> When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.

Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

NEW QUESTION 83

- (Topic 1)

An application allows users at a company's headquarters to access product data. The product data is stored in an Amazon RDS MySQL DB instance. The operations team has isolated an application performance slowdown and wants to separate read traffic from write traffic. A solutions architect needs to optimize the application's performance quickly.

What should the solutions architect recommend?

- A. Change the existing database to a Multi-AZ deployment
- B. Serve the read requests from the primary Availability Zone.
- C. Change the existing database to a Multi-AZ deployment
- D. Serve the read requests from the secondary Availability Zone.
- E. Create read replicas for the database
- F. Configure the read replicas with half of the compute and storage resources as the source database.
- G. Create read replicas for the database
- H. Configure the read replicas with the same compute and storage resources as the source database.

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html

NEW QUESTION 84

- (Topic 1)

A company has an automobile sales website that stores its listings in a database on Amazon RDS. When an automobile is sold the listing needs to be removed from the website and the data must be sent to multiple target systems.

Which design should a solutions architect recommend?

- A. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) queue for the targets to consume
- B. Create an AWS Lambda function triggered when the database on Amazon RDS is updated to send the information to an Amazon Simple Queue Service (Amazon SQS) FIFO queue for the targets to consume
- C. Subscribe to an RDS event notification and send an Amazon Simple Queue Service (Amazon SQS) queue fanned out to multiple Amazon Simple Notification Service (Amazon SNS) topics. Use AWS Lambda functions to update the targets
- D. Subscribe to an RDS event notification and send an Amazon Simple Notification Service (Amazon SNS) topic fanned out to multiple Amazon Simple Queue Service (Amazon SQS) queues. Use AWS Lambda functions to update the targets

Answer: D

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html> <https://docs.aws.amazon.com/lambda/latest/dg/with-sns.html>

NEW QUESTION 87

- (Topic 1)

A company has a production web application in which users upload documents through a web interface or a mobile app. According to a new regulatory requirement, new documents cannot be modified or deleted after they are stored.

What should a solutions architect do to meet this requirement?

- A. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning and S3 Object Lock enabled
- B. Store the uploaded documents in an Amazon S3 bucket
- C. Configure an S3 Lifecycle policy to archive the documents periodically.
- D. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning enabled. Configure an ACL to restrict all access to read-only.
- E. Store the uploaded documents on an Amazon Elastic File System (Amazon EFS) volume
- F. Access the data by mounting the volume in read-only mode.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>

NEW QUESTION 88

- (Topic 1)

A company is designing an application. The application uses an AWS Lambda function to receive information through Amazon API Gateway and to store the information in an Amazon Aurora PostgreSQL database.

During the proof-of-concept stage, the company has to increase the Lambda quotas significantly to handle the high volumes of data that the company needs to load into the database. A solutions architect must recommend a new design to improve scalability and minimize the configuration effort.

Which solution will meet these requirements?

- A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instance
- B. Connect the database by using native Java Database Connectivity (JDBC) drivers.
- C. Change the platform from Aurora to Amazon DynamoDB
- D. Provision a DynamoDB Accelerator (DAX) cluster

- E. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
- F. Set up two Lambda function
- G. Configure one function to receive the informatio
- H. Configure the other function to load the information into the databas
- I. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
- J. Set up two Lambda function
- K. Configure one function to receive the informatio
- L. Configure the other function to load the information into the databas
- M. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Answer: B

Explanation:

bottlenecks can be avoided with queues (SQS).

NEW QUESTION 90

- (Topic 1)

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations. Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Add the aws:PrincipalOrgID global condition key with a reference to the organization ID to the S3 bucket policy.
- B. Create an organizational unit (OU) for each departmen
- C. Add the aws:PrincipalOrgPaths global condition key to the S3 bucket policy.
- D. Use AWS CloudTrail to monitor the CreateAccount, InviteAccountToOrganization, LeaveOrganization, and RemoveAccountFromOrganization event
- E. Update the S3 bucket policy accordingly.
- F. Tag each user that needs access to the S3 bucke
- G. Add the aws:PrincipalTag global condition key to the S3 bucket policy.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

The aws:PrincipalOrgID global key provides an alternative to listing all the account IDs for all AWS accounts in an organization. For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an object into the examtopics bucket.

```
{
  "Version": "2020-09-10",
  "Statement": {
    "Sid": "AllowPutObject", "Effect": "Allow",
    "Principal": "*",
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::examtopics/*", "Condition": {"StringEquals":
{"aws:PrincipalOrgID":["XXX"]}}}
```

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition_keys.html

NEW QUESTION 95

- (Topic 1)

A company wants to move a multi-tiered application from on premises to the AWS Cloud to improve the application's performance. The application consists of application tiers that communicate with each other by way of RESTful services. Transactions are dropped when one tier becomes overloaded. A solutions architect must design a solution that resolves these issues and modernizes the application. Which solution meets these requirements and is the MOST operationally efficient?

- A. Use Amazon API Gateway and direct transactions to the AWS Lambda functions as the application laye
- B. Use Amazon Simple Queue Service (Amazon SQS) as the communication layer between application services.
- C. Use Amazon CloudWatch metrics to analyze the application performance history to determine the server's peak utilization during the performance failure
- D. Increase the size of the application server's Amazon EC2 instances to meet the peak requirements.
- E. Use Amazon Simple Notification Service (Amazon SNS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling grou
- F. Use Amazon CloudWatch to monitor the SNS queue length and scale up and down as required.
- G. Use Amazon Simple Queue Service (Amazon SQS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling grou
- H. Use Amazon CloudWatch to monitor the SQS queue length and scale up when communication failures are detected.

Answer: A

Explanation:

<https://aws.amazon.com/getting-started/hands-on/build-serverless-web-app-lambda-apigateway-s3-dynamodb-cognito/module-4/>

Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito. This example showed similar setup as question: Build a Serverless Web Application with AWS Lambda, Amazon API Gateway, AWS Amplify, Amazon DynamoDB, and Amazon Cognito

NEW QUESTION 99

- (Topic 2)

A company sells ringtones created from clips of popular songs. The files containing the ringtones are stored in Amazon S3 Standard and are at least 128 KB in size. The company has millions of files, but downloads are infrequent for ringtones older than 90 days. The company needs to save money on storage while keeping the most accessed files readily available for its users.

Which action should the company take to meet these requirements MOST cost-effectively?

- A. Configure S3 Standard-Infrequent Access (S3 Standard-IA) storage for the initial storage tier of the objects.
- B. Move the files to S3 Intelligent-Tiering and configure it to move objects to a less expensive storage tier after 90 days.
- C. Configure S3 inventory to manage objects and move them to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.
- D. Implement an S3 Lifecycle policy that moves the objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.

Answer: D

Explanation:

This solution meets the requirements of saving money on storage while keeping the most accessed files readily available for the users. S3 Lifecycle policy can automatically move objects from one storage class to another based on predefined rules. S3 Standard-IA is a lower-cost storage class for data that is accessed less frequently, but requires rapid access when needed. It is suitable for ringtones older than 90 days that are downloaded infrequently. Option A is incorrect because configuring S3 Standard-IA for the initial storage tier of the objects can incur higher costs for frequent access and retrieval fees. Option B is incorrect because moving the files to S3 Intelligent-Tiering can incur additional monitoring and automation fees that may not be necessary for ringtones older than 90 days. Option C is incorrect because using S3 inventory to manage objects and move them to S3 Standard-IA can be complex and time-consuming, and it does not provide automatic cost savings. References:
? <https://aws.amazon.com/s3/storage-classes/>
? <https://aws.amazon.com/s3/cloud-storage-cost-optimization-ebook/>

NEW QUESTION 100

- (Topic 2)

A medical records company is hosting an application on Amazon EC2 instances. The application processes customer data files that are stored on Amazon S3. The EC2 instances are hosted in public subnets. The EC2 instances access Amazon S3 over the internet, but they do not require any other network access. A new requirement mandates that the network traffic for file transfers take a private route and not be sent over the internet. Which change to the network architecture should a solutions architect recommend to meet this requirement?

- A. Create a NAT gateway
- B. Configure the route table for the public subnets to send traffic to Amazon S3 through the NAT gateway.
- C. Configure the security group for the EC2 instances to restrict outbound traffic so that only traffic to the S3 prefix list is permitted.
- D. Move the EC2 instances to private subnet
- E. Create a VPC endpoint for Amazon S3, and link the endpoint to the route table for the private subnets
- F. Remove the internet gateway from the VP
- G. Set up an AWS Direct Connect connection, and route traffic to Amazon S3 over the Direct Connect connection.

Answer: C

Explanation:

To meet the new requirement of transferring files over a private route, the EC2 instances should be moved to private subnets, which do not have direct access to the internet. This ensures that the traffic for file transfers does not go over the internet. To enable the EC2 instances to access Amazon S3, a VPC endpoint for Amazon S3 can be created. VPC endpoints allow resources within a VPC to communicate with resources in other services without the traffic being sent over the internet. By linking the VPC endpoint to the route table for the private subnets, the EC2 instances can access Amazon S3 over a private connection within the VPC.

NEW QUESTION 102

- (Topic 2)

A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. What should a solutions architect do to meet this requirement?

- A. Configure AWS Storage Gateway in volume gateway mod
- B. Mount the volume to each Windows instance.
- C. Configure Amazon FSx for Windows File Serve
- D. Mount the Amazon FSx file system to each Windows instance.
- E. Configure a file system by using Amazon Elastic File System (Amazon EFS). Mount the EFS file system to each Windows instance.
- F. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required siz
- G. Attach each EC2 instance to the volum
- H. Mount the file system within the volume to each Windows instance.

Answer: B

Explanation:

This solution meets the requirement of migrating a Windows-based application that requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones. Amazon FSx for Windows File Server provides fully managed shared storage built on Windows Server, and delivers a wide range of data access, data management, and administrative capabilities. It supports the Server Message Block (SMB) protocol and can be mounted to EC2 Windows instances across multiple Availability Zones. Option A is incorrect because AWS Storage Gateway in volume gateway mode provides cloud-backed storage volumes that can be mounted as iSCSI devices from on-premises application servers, but it does not support SMB protocol or EC2 Windows instances. Option C is incorrect because Amazon Elastic File System (Amazon EFS) provides a scalable and elastic NFS file system for Linux-based workloads, but it does not support SMB protocol or EC2 Windows instances. Option D is incorrect because Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with EC2 instances, but it does not support SMB protocol or attaching multiple instances to the same volume. References:
? <https://aws.amazon.com/fsx/windows/>
? <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/using-file-shares.html>

NEW QUESTION 103

- (Topic 2)

A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet. Which solutions will meet these requirements? (Choose two.)

- A. Use AWS Control Tower to implement data residency guardrails to deny internet access and deny access to all AWS Regions except ap-northeast-3.
- B. Use rules in AWS WAF to prevent internet acces
- C. Deny access to all AWS Regions except ap-northeast-3 in the AWS account settings.
- D. Use AWS Organizations to configure service control policies (SCPS) that prevent VPCs from gaining internet acces
- E. Deny access to all AWS Regions except ap-northeast-3.

- F. Create an outbound rule for the network ACL in each VPC to deny all traffic from 0.0.0.0/0. Create an IAM policy for each user to prevent the use of any AWS Region other than ap-northeast-3.
- G. Use AWS Config to activate managed rules to detect and alert for internet gateways and to detect and alert for new resources deployed outside of ap-northeast-3.

Answer: AC

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_vpc.html#example_vpc_2

NEW QUESTION 108

- (Topic 2)

A company has a multi-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs to modify the infrastructure to be highly available without modifying the application. Which architecture should the solutions architect choose that provides high availability?

- A. Create an Auto Scaling group that uses three Instances across each of two Regions.
- B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.
- C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.
- D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

Answer: B

Explanation:

High availability can be enabled for this architecture quite simply by modifying the existing Auto Scaling group to use multiple availability zones. The ASG will automatically balance the load so you don't actually need to specify the instances per AZ.

NEW QUESTION 113

- (Topic 2)

A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution.

Which storage solution meets these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elasticsearch Service (Amazon ES)
- D. Amazon S3

Answer: D

Explanation:

Amazon S3 is cheapest and can be accessed from anywhere.

NEW QUESTION 117

- (Topic 2)

A company runs an application using Amazon ECS. The application creates resized versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3.

How can a solutions architect ensure that the application has permission to access Amazon S3?

- A. Update the S3 role in AWS IAM to allow read/write access from Amazon ECS, and then relaunch the container.
- B. Create an IAM role with S3 permissions, and then specify that role as the taskRoleArn in the task definition.
- C. Create a security group that allows access from Amazon ECS to Amazon S3, and update the launch configuration used by the ECS cluster.
- D. Create an IAM user with S3 permissions, and then relaunch the Amazon EC2 instances for the ECS cluster while logged in as this account.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ecs-taskdefinition.html>

NEW QUESTION 120

- (Topic 2)

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.

The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.

Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RDS
- B. Use RDS access controls to limit access.
- C. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawler
- D. Use Amazon Athena to query the data
- E. Use S3 policies to limit access.
- F. Create a data lake by using AWS Lake Formation
- G. Create an AWS Glue JDBC connection to Amazon RDS
- H. Register the S3 bucket in Lake Formation
- I. Use Lake Formation access controls to limit access.
- J. Create an Amazon Redshift cluster

- K. Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshift
- L. Use Amazon Redshift access controls to limit access.

Answer: C

Explanation:

To make all the data available to various teams and minimize operational overhead, the company can create a data lake by using AWS Lake Formation. This will allow the company to centralize all the data in one place and use fine-grained access controls to manage access to the data. To meet the requirements of the company, the solutions architect can create a data lake by using AWS Lake Formation, create an AWS Glue JDBC connection to Amazon RDS, and register the S3 bucket in Lake Formation. The solutions architect can then use Lake Formation access controls to limit access to the data. This solution will provide the ability to manage fine-grained permissions for the data and minimize operational overhead.

NEW QUESTION 121

- (Topic 2)

A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after it is deployed. The company is unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.

Which solution will meet these requirements?

- A. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository
- B. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the container
- C. Use target tracking to scale automatically based on demand.
- D. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository
- E. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the container
- F. Use target tracking to scale automatically based on demand.
- G. Store container images in a repository that runs on an Amazon EC2 instance
- H. Run the containers on EC2 instances that are spread across multiple Availability Zones
- I. Monitor the average CPU utilization in Amazon CloudWatch
- J. Launch new EC2 instances as needed
- K. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image. Launch EC2 instances in an Auto Scaling group across multiple Availability Zones
- L. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Answer: A

Explanation:

AWS Fargate is a serverless experience for user applications, allowing the user to concentrate on building applications instead of configuring and managing servers. Fargate also automates resource management, allowing users to easily scale their applications in response to demand.

NEW QUESTION 124

- (Topic 2)

A company is running an online transaction processing (OLTP) workload on AWS. This workload uses an unencrypted Amazon RDS DB instance in a Multi-AZ deployment. Daily database snapshots are taken from this instance.

What should a solutions architect do to ensure the database and snapshots are always encrypted moving forward?

- A. Encrypt a copy of the latest DB snapshot
- B. Replace existing DB instance by restoring the encrypted snapshot
- C. Create a new encrypted Amazon Elastic Block Store (Amazon EBS) volume and copy the snapshots to it. Enable encryption on the DB instance
- D. Copy the snapshots and enable encryption using AWS Key Management Service (AWS KMS). Restore encrypted snapshot to an existing DB instance
- E. Copy the snapshots to an Amazon S3 bucket that is encrypted using server-side encryption with AWS Key Management Service (AWS KMS) managed keys (SSE-KMS)

Answer: A

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#USER_RestoreFromSnapshot.CON
Under "Encrypt unencrypted resources" - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

NEW QUESTION 125

- (Topic 2)

A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year.

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

- A. Move the data to the S3 bucket
- B. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use the built-in key rotation behavior of SSE-S3 encryption keys.
- C. Create an AWS Key Management Service (AWS KMS) customer managed key
- D. Enable automatic key rotation
- E. Set the S3 bucket's default encryption behavior to use the customer managed KMS key
- F. Move the data to the S3 bucket.
- G. Create an AWS Key Management Service (AWS KMS) customer managed key
- H. Set the S3 bucket's default encryption behavior to use the customer managed KMS key
- I. Move the data to the S3 bucket
- J. Manually rotate the KMS key every year.
- K. Encrypt the data with customer key material before moving the data to the S3 bucket
- L. Create an AWS Key Management Service (AWS KMS) key without key material
- M. Import the customer key material into the KMS key
- N. Enable automatic key rotation.

Answer: B

Explanation:

SSE-S3 - is free and uses AWS owned CMKs (CMK = Customer Master Key). The encryption key is owned and managed by AWS, and is shared among many accounts. Its rotation is automatic with time that varies as shown in the table here. The time is not explicitly defined.

SSE-KMS - has two flavors:

AWS managed CMK. This is free CMK generated only for your account. You can only view its policies and audit usage, but not manage it. Rotation is automatic - once per 1095 days (3 years),

Customer managed CMK. This uses your own key that you create and can manage. Rotation is not enabled by default. But if you enable it, it will be automatically rotated every 1 year. This variant can also use an imported key material by you. If you create such key with an imported material, there is no automated rotation. Only manual rotation.

SSE-C - customer provided key. The encryption key is fully managed by you outside of AWS. AWS will not rotate it.

This solution meets the requirements of moving data to an Amazon S3 bucket, encrypting the data when it is stored in the S3 bucket, and automatically rotating the encryption key every year with the least operational overhead. AWS Key Management Service (AWS KMS) is a service that enables you to create and manage encryption keys for your data. A customer managed key is a symmetric encryption key that you create and manage in AWS KMS. You can enable automatic key rotation for a customer managed key, which means that AWS KMS generates new cryptographic material for the key every year. You can set the S3 bucket's default encryption behavior to use the customer managed KMS key, which means that any object that is uploaded to the bucket without specifying an encryption method will be encrypted with that key.

Option A is incorrect because using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not allow you to control or manage the encryption keys. SSE-S3 uses a unique key for each object, and encrypts that key with a master key that is regularly rotated by S3. However, you cannot enable or disable key rotation for SSE-S3 keys, or specify the rotation interval. Option C is incorrect because manually rotating the KMS key every year can increase the operational overhead and complexity, and it may not meet the requirement of rotating the key every year if you forget or delay the rotation process. Option D is incorrect because encrypting the data with customer key material before moving the data to the S3 bucket can increase the operational overhead and complexity, and it may not provide consistent encryption for all objects in the bucket. Creating a KMS key without key material and importing the customer key material into the KMS key can enable you to use your own source of random bits to generate your KMS keys, but it does not support automatic key rotation.

Option A is incorrect because using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) does not allow you to control or manage the encryption keys. SSE-S3 uses a unique key for each object, and encrypts that key with a master key that is regularly rotated by S3. However, you cannot enable or disable key rotation for SSE-S3 keys, or specify the rotation interval. Option C is incorrect because manually rotating the KMS key every year can increase the operational overhead and complexity, and it may not meet the requirement of rotating the key every year if you forget or delay the rotation process. Option D is incorrect because encrypting the data with customer key material before moving the data to the S3 bucket can increase the operational overhead and complexity, and it may not provide consistent encryption for all objects in the bucket. Creating a KMS key without key material and importing the customer key material into the KMS key can enable you to use your own source of random bits to generate your KMS keys, but it does not support automatic key rotation.

References:

? <https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html>

? <https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

? <https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html>

NEW QUESTION 129

- (Topic 2)

A media company is evaluating the possibility of moving its systems to the AWS Cloud. The company needs at least 10 TB of storage with the maximum possible I/O performance for video processing, 300 TB of very durable storage for storing media content, and 900 TB of storage to meet requirements for archival media that is not in use anymore.

Which set of services should a solutions architect recommend to meet these requirements?

- A. Amazon EBS for maximum performance, Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage
- B. Amazon EBS for maximum performance, Amazon EFS for durable data storage and Amazon S3 Glacier for archival storage
- C. Amazon EC2 instance store for maximum performance
- D. Amazon EFS for durable data storage and Amazon S3 for archival storage
- E. Amazon EC2 Instance store for maximum performance
- F. Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html>

NEW QUESTION 134

- (Topic 2)

A company runs its two-tier e-commerce website on AWS. The web tier consists of a load balancer that sends traffic to Amazon EC2 instances. The database tier uses an Amazon RDS DB instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available.

Which combination of configuration options will meet these requirements? (Choose two.)

- A. Use an Auto Scaling group to launch the EC2 instances in private subnet
- B. Deploy an RDS Multi-AZ DB instance in private subnets.
- C. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones
- D. Deploy an Application Load Balancer in the private subnets.
- E. Use an Auto Scaling group to launch the EC2 instances in public subnets across two Availability Zones
- F. Deploy an RDS Multi-AZ DB instance in private subnets.
- G. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zones
- H. Deploy an Application Load Balancer in the public subnet.
- I. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zones
- J. Deploy an Application Load Balancer in the public subnets.

Answer: AE

Explanation:

Before you begin: Decide which two Availability Zones you will use for your EC2 instances. Configure your virtual private cloud (VPC) with at least one public subnet in each of these Availability Zones. These public subnets are used to configure the load balancer. You can launch your EC2 instances in other subnets of these Availability Zones instead.

NEW QUESTION 135

- (Topic 2)

A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.

Which solution meets these requirements MOST cost-effectively?

- A. Use Spot Instances exclusively to handle the maximum capacity required.
- B. Use Reserved Instances exclusively to handle the maximum capacity required.
- C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
- D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

Answer: D

Explanation:

We recommend that you use On-Demand Instances for applications with short-term, irregular workloads that cannot be interrupted.
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-on-demand-instances.html>

NEW QUESTION 140

- (Topic 2)

A company has a highly dynamic batch processing job that uses many Amazon EC2 instances to complete it. The job is stateless in nature, can be started and stopped at any given time with no negative impact, and typically takes upwards of 60 minutes total to complete. The company has asked a solutions architect to design a scalable and cost-effective solution that meets the requirements of the job. What should the solutions architect recommend?

- A. Implement EC2 Spot Instances
- B. Purchase EC2 Reserved Instances
- C. Implement EC2 On-Demand Instances
- D. Implement the processing on AWS Lambda

Answer: A

Explanation:

EC2 Spot Instances allow users to bid on spare Amazon EC2 computing capacity and can be a cost-effective solution for stateless, interruptible workloads that can be started and stopped at any time. Since the batch processing job is stateless, can be started and stopped at any time, and typically takes upwards of 60 minutes to complete, EC2 Spot Instances would be a good fit for this workload.

NEW QUESTION 141

- (Topic 2)

A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using a mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible. How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

- A. Deploy the application stack in a single AWS Region
- B. Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
- C. Deploy the application stack in two AWS Region
- D. Use an Amazon Route 53 latency routing policy to serve all content from the ALB in the closest Region.
- E. Deploy the application stack in a single AWS Region
- F. Use Amazon CloudFront to serve the static content
- G. Serve the dynamic content directly from the ALB.
- H. Deploy the application stack in two AWS Region
- I. Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deliver-your-apps-dynamic-content-using-amazon-cloudfront-getting-started-template/>

NEW QUESTION 144

- (Topic 2)

A solutions architect must design a solution that uses Amazon CloudFront with an Amazon S3 origin to store a static website. The company's security policy requires that all website traffic be inspected by AWS WAF. How should the solutions architect comply with these requirements?

- A. Configure an S3 bucket policy to accept requests coming from the AWS WAF Amazon Resource Name (ARN) only.
- B. Configure Amazon CloudFront to forward all incoming requests to AWS WAF before requesting content from the S3 origin.
- C. Configure a security group that allows Amazon CloudFront IP addresses to access Amazon S3 only.
- D. Associate AWS WAF to CloudFront.
- E. Configure Amazon CloudFront and Amazon S3 to use an origin access identity (OAI) to restrict access to the S3 bucket.
- F. Enable AWS WAF on the distribution.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-aws-waf.html>

NEW QUESTION 149

- (Topic 2)

A solutions architect is optimizing a website for an upcoming musical event. Videos of the performances will be streamed in real time and then will be available on demand. The event is expected to attract a global online audience. Which service will improve the performance of both the real-time and on-demand streaming?

- A. Amazon CloudFront
- B. AWS Global Accelerator

- C. Amazon Route 53
- D. Amazon S3 Transfer Acceleration

Answer: A

Explanation:

You can use CloudFront to deliver video on demand (VOD) or live streaming video using any HTTP origin. One way you can set up video workflows in the cloud is by using CloudFront together with AWS Media Services. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand-streaming-video.html>

NEW QUESTION 153

- (Topic 2)

A company owns an asynchronous API that is used to ingest user requests and, based on the request type, dispatch requests to the appropriate microservice for processing. The company is using Amazon API Gateway to deploy the API front end, and an AWS Lambda function that invokes Amazon DynamoDB to store user requests before dispatching them to the processing microservices.

The company provisioned as much DynamoDB throughput as its budget allows, but the company is still experiencing availability issues and is losing user requests. What should a solutions architect do to address this issue without impacting existing users?

- A. Add throttling on the API Gateway with server-side throttling limits.
- B. Use DynamoDB Accelerator (DAX) and Lambda to buffer writes to DynamoDB.
- C. Create a secondary index in DynamoDB for the table with the user requests.
- D. Use the Amazon Simple Queue Service (Amazon SQS) queue and Lambda to buffer writes to DynamoDB.

Answer: D

Explanation:

By using an SQS queue and Lambda, the solutions architect can decouple the API front end from the processing microservices and improve the overall scalability and availability of the system. The SQS queue acts as a buffer, allowing the API front end to continue accepting user requests even if the processing microservices are experiencing high workloads or are temporarily unavailable. The Lambda function can then retrieve requests from the SQS queue and write them to DynamoDB, ensuring that all user requests are stored and processed. This approach allows the company to scale the processing microservices independently from the API front end, ensuring that the API remains available to users even during periods of high demand.

NEW QUESTION 156

- (Topic 2)

A company's web application is running on Amazon EC2 instances behind an Application Load Balancer. The company recently changed its policy, which now requires the application to be accessed from one specific country only. Which configuration will meet this requirement?

- A. Configure the security group for the EC2 instances.
- B. Configure the security group on the Application Load Balancer.
- C. Configure AWS WAF on the Application Load Balancer in a VPC.
- D. Configure the network ACL for the subnet that contains the EC2 instances.

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/10/aws-waf-now-supports-geographic-match/>

NEW QUESTION 159

- (Topic 2)

A company has a data ingestion workflow that includes the following components:

- An Amazon Simple Notification Service (Amazon SNS) topic that receives notifications about new data deliveries
- An AWS Lambda function that processes and stores the data

The ingestion workflow occasionally fails because of network connectivity issues. When a failure occurs the corresponding data is not ingested unless the company manually reruns the job. What should a solutions architect do to ensure that all notifications are eventually processed?

- A. Configure the Lambda function (or deployment across multiple Availability Zones)
- B. Modify the Lambda functions configuration to increase the CPU and memory allocations for the function
- C. Configure the SNS topic's retry strategy to increase both the number of retries and the wait time between retries
- D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as the on-failure destination. Modify the Lambda function to process messages in the queue

Answer: D

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>

NEW QUESTION 160

- (Topic 2)

A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics software is written in PHP and uses a MySQL database. The analytics software, the web server that provides PHP, and the database server are all hosted on the EC2 instance. The application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly. Which solution will meet these requirements MOST cost-effectively?

- A. Migrate the database to an Amazon RDS for MySQL DB instance
- B. Create an AMI of the web application
- C. Use the AMI to launch a second EC2 On-Demand Instance
- D. Use an Application Load Balancer to distribute the load to each EC2 instance.
- E. Migrate the database to an Amazon RDS for MySQL DB instance
- F. Create an AMI of the web application

- G. Use the AMI to launch a second EC2 On-Demand Instance
- H. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
- I. Migrate the database to an Amazon Aurora MySQL DB instance
- J. Create an AWS Lambda function to stop the EC2 instance and change the instance type
- K. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
- L. Migrate the database to an Amazon Aurora MySQL DB instance
- M. Create an AMI of the web application
- N. Apply the AMI to a launch template
- O. Create an Auto Scaling group with the launch template. Configure the launch template to use a Spot Fleet
- P. Attach an Application Load Balancer to the Auto Scaling group.

Answer: D

Explanation:

Migrate the database to Amazon Aurora MySQL - this will let the DB scale on its own; it'll scale automatically without needing adjustment. Create AMI of the web app and using a launch template - this will make the creating of any future instances of the app seamless. They can then be added to the auto scaling group which will save them money as it will scale up and down based on demand. Using a spot fleet to launch instances- This solves the "MOST cost-effective" portion of the question as spot instances come at a huge discount at the cost of being terminated at any time Amazon deems fit. I think this is why there's a bit of disagreement on this. While it's the most cost effective, it would be a terrible choice if Amazon were to terminate that spot instance during a busy period.

NEW QUESTION 162

- (Topic 2)

A global company is using Amazon API Gateway to design REST APIs for its loyalty club users in the us-east-1 Region and the ap-southeast-2 Region. A solutions architect must design a solution to protect these API Gateway managed REST APIs across multiple accounts from SQL injection and cross-site scripting attacks. Which solution will meet these requirements with the LEAST amount of administrative effort?

- A. Set up AWS WAF in both Region
- B. Associate Regional web ACLs with an API stage.
- C. Set up AWS Firewall Manager in both Region
- D. Centrally configure AWS WAF rules.
- E. Set up AWS Shield in both Region
- F. Associate Regional web ACLs with an API stage.
- G. Set up AWS Shield in one of the Region
- H. Associate Regional web ACLs with an API stage.

Answer: A

Explanation:

Using AWS WAF has several benefits. Additional protection against web attacks using criteria that you specify. You can define criteria using characteristics of web requests such as the following: Presence of SQL code that is likely to be malicious (known as SQL injection). Presence of a script that is likely to be malicious (known as cross-site scripting). AWS Firewall Manager simplifies your administration and maintenance tasks across multiple accounts and resources for a variety of protections. <https://docs.aws.amazon.com/waf/latest/developer/guide/what-is-aws-waf.html>

NEW QUESTION 166

- (Topic 2)

A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone. The company needs to redesign its architecture to provide the highest availability with the least operational overhead. What should a solutions architect do to meet these requirements?

- A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
- B. Create a Multi-AZ Auto Scaling group (or EC2 instances that host the applicatio
- C. Create another Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.
- D. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon M
- E. Create a Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
- F. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- G. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queu
- H. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- I. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue. Create another Multi-AZ Auto Scaling group for EC2 instances that host the applicatio
- J. Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Answer: B

Explanation:

Migrating to Amazon MQ reduces the overhead on the queue management. C and D are dismissed. Deciding between A and B means deciding to go for an Auto Scaling group for EC2 or an RDS for PostgreSQL (both multi-AZ). The RDS option has less operational impact, as provided as a service the tools and software required. Consider for instance, the effort to add an additional node like a read replica, to the DB. <https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/active-standby-broker-deployment.html> <https://aws.amazon.com/rds/postgresql/>

NEW QUESTION 169

- (Topic 2)

A company runs a stateless web application in production on a group of Amazon EC2 On-Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight. Application usage is low during weekends. The company wants to minimize its EC2 costs without affecting the availability of the application. Which solution will meet these requirements?

- A. Use Spot Instances for the entire workload.
- B. Use Reserved instances for the baseline level of usage Use Spot Instances for any additional capacity that the application needs.
- C. Use On-Demand Instances for the baseline level of usage
- D. Use Spot Instances for any additional capacity that the application needs
- E. Use Dedicated Instances for the baseline level of usage
- F. Use On-Demand Instances for any additional capacity that the application needs

Answer: B

Explanation:

Reserved is cheaper than on demand the company has. And it's meet the availability (HA) requirement as to spot instance that can be disrupted at any time. PRICING BELOW. On- Demand: 0% There's no commitment from you. You pay the most with this option. Reserved : 40%-60% 1-year or 3-year commitment from you. You save money from that commitment. Spot 50%-90% Ridiculously inexpensive because there's no commitment from the AWS side.

NEW QUESTION 172

- (Topic 2)

A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month. Which storage option meets these requirements MOST cost-effectively?

- A. Store the logs in Amazon S3 Use AWS Backup to move logs more than 1 month old to S3 Glacier Deep Archive
- B. Store the logs in Amazon S3 Use S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive
- C. Store the logs in Amazon CloudWatch Logs Use AWS Backup to move logs more than 1 month old to S3 Glacier Deep Archive
- D. Store the logs in Amazon CloudWatch Logs Use Amazon S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive

Answer: B

Explanation:

You need S3 to be able to archive the logs after one month. Cannot do that with CloudWatch Logs.

NEW QUESTION 175

- (Topic 2)

A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable. Which solution will meet these requirements?

- A. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive immediately.
- B. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 2 years.
- C. Use S3 Intelligent-Tiering
- D. Activate the archiving option to ensure that data is archived in S3 Glacier Deep Archive.
- E. Set up an S3 Lifecycle policy to transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately and to S3 Glacier Deep Archive after 2 years.

Answer: B

Explanation:

https://aws.amazon.com/about-aws/whats-new/2018/04/announcing-s3-one-zone-infrequent-access-a-new-amazon-s3-storage-class/?nc1=h_ls

NEW QUESTION 179

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 182

- (Topic 2)

A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours. Which solution meets these requirements?

- A. Use Amazon DynamoDB with auto scaling Use on-demand backups and Amazon DynamoDB Streams
- B. Use Amazon Redshift
- C. Configure concurrency scaling

- D. Activate audit loggin
- E. Perform database snapshots every 4 hours.
- F. Use Amazon RDS with Provisioned IOPS Activate the database auditing parameter Perform database snapshots every 5 hours
- G. Use Amazon Aurora MySQL with auto scalin
- H. Activate the database auditing parameter

Answer: A

Explanation:

This solution meets the requirements of a customer-facing application that has a clearly defined access pattern throughout the year and a variable number of reads and writes that depend on the time of year. Amazon DynamoDB is a fully managed NoSQL database service that can handle any level of request traffic and data size. DynamoDB auto scaling can automatically adjust the provisioned read and write capacity based on the actual workload. DynamoDB on-demand backups can create full backups of the tables for data protection and archival purposes. DynamoDB Streams can capture a time-ordered sequence of item-level modifications in the tables for audit purposes.

Option B is incorrect because Amazon Redshift is a data warehouse service that is designed for analytical workloads, not for customer-facing applications. Option C is incorrect because Amazon RDS with Provisioned IOPS can provide consistent performance for relational databases, but it may not be able to handle unpredictable spikes in traffic and data size. Option D is incorrect because Amazon Aurora MySQL with auto scaling can provide high performance and availability for relational databases, but it does not support audit logging as a parameter.

References:

- ? <https://aws.amazon.com/dynamodb/>
- ? <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>
- ? <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BackupRestore.html>
- ? <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html>

NEW QUESTION 185

- (Topic 2)

A company wants to run applications in containers in the AWS Cloud. These applications are stateless and can tolerate disruptions within the underlying infrastructure. The company needs a solution that minimizes cost and operational overhead.

What should a solutions architect do to meet these requirements?

- A. Use Spot Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- B. Use Spot Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.
- C. Use On-Demand Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- D. Use On-Demand Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.

Answer: A

Explanation:

<https://aws.amazon.com/cn/blogs/compute/cost-optimization-and-resilience-eks-with-spot-instances/>

NEW QUESTION 189

- (Topic 2)

A gaming company is designing a highly available architecture. The application runs on a modified Linux kernel and supports only UDP-based traffic. The company needs the front-end tier to provide the best possible user experience. That tier must have low latency, route traffic to the nearest edge location, and provide static IP addresses for entry into the application endpoints.

What should a solutions architect do to meet these requirements?

- A. Configure Amazon Route 53 to forward requests to an Application Load Balance
- B. Use AWS Lambda for the application in AWS Application Auto Scaling.
- C. Configure Amazon CloudFront to forward requests to a Network Load Balance
- D. Use AWS Lambda for the application in an AWS Application Auto Scaling group.
- E. Configure AWS Global Accelerator to forward requests to a Network Load Balance
- F. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.
- G. Configure Amazon API Gateway to forward requests to an Application Load Balance
- H. Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.

Answer: C

Explanation:

AWS Global Accelerator and Amazon CloudFront are separate services that use the AWS global network and its edge locations around the world. CloudFront improves performance for both cacheable content (such as images and videos) and dynamic content (such as API acceleration and dynamic site delivery). Global Accelerator improves performance for a wide range of applications over TCP or UDP by proxying packets at the edge to applications running in one or more AWS Regions. Global Accelerator is a good fit for non-HTTP use cases, such as gaming (UDP), IoT (MQTT), or Voice over IP, as well as for HTTP use cases that specifically require static IP addresses or deterministic, fast regional failover. Both services integrate with AWS Shield for DDoS protection.

NEW QUESTION 191

- (Topic 2)

A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.

What should a solutions architect do to meet these requirements?

- A. Configure an IAM policy for AWS Systems Manager Session Manager
- B. Create an IAM role for the polic
- C. Update the trust relationship of the rol
- D. Set up automatic start and stop for the DB instance.
- E. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stoppe
- F. Invalidate the cache after the DB instance is started.
- G. Launch an Amazon EC2 instanc
- H. Create an IAM role that grants access to Amazon RD
- I. Attach the role to the EC2 instanc

- J. Configure a cron job to start and stop the EC2 instance on the desired schedule.
- K. Create AWS Lambda functions to start and stop the DB instance
- L. Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda function
- M. Configure the Lambda functions as event targets for the rules

Answer: D

Explanation:

In a typical development environment, dev and test databases are mostly utilized for 8 hours a day and sit idle when not in use. However, the databases are billed for the compute and storage costs during this idle time. To reduce the overall cost, Amazon RDS allows instances to be stopped temporarily. While the instance is stopped, you're charged for storage and backups, but not for the DB instance hours. Please note that a stopped instance will automatically be started after 7 days. This post presents a solution using AWS Lambda and Amazon EventBridge that allows you to schedule a Lambda function to stop and start the idle databases with specific tags to save on compute costs. The second post presents a solution that accomplishes stop and start of the idle Amazon RDS databases using AWS Systems Manager.

NEW QUESTION 195

- (Topic 2)

A company needs to move data from an Amazon EC2 instance to an Amazon S3 bucket. The company must ensure that no API calls and no data are routed through public internet routes. Only the EC2 instance can have access to upload data to the S3 bucket. Which solution will meet these requirements?

- A. Create an interface VPC endpoint for Amazon S3 in the subnet where the EC2 instance is located
- B. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
- C. Create a gateway VPC endpoint for Amazon S3 in the Availability Zone where the EC2 instance is located
- D. Attach appropriate security groups to the endpoint
- E. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
- F. Run the nslookup tool from inside the EC2 instance to obtain the private IP address of the S3 bucket's service API endpoint
- G. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket
- H. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.
- I. Use the AWS provided, publicly available ip-ranges.json file to obtain the private IP address of the S3 bucket's service API endpoint
- J. Create a route in the VPC route table to provide the EC2 instance with access to the S3 bucket
- K. Attach a resource policy to the S3 bucket to only allow the EC2 instance's IAM role for access.

Answer: A

Explanation:

(<https://aws.amazon.com/blogs/security/how-to-restrict-amazon-s3-bucket-access-to-a-specific-iam-role/>)

NEW QUESTION 196

.....

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

SAA-C03 Practice Exam Features:

- * SAA-C03 Questions and Answers Updated Frequently
- * SAA-C03 Practice Questions Verified by Expert Senior Certified Staff
- * SAA-C03 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAA-C03 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAA-C03 Practice Test Here](#)