

Amazon-Web-Services

Exam Questions SAA-C03

AWS Certified Solutions Architect - Associate (SAA-C03)



NEW QUESTION 1

- (Topic 1)

A company has an application that generates a large number of files, each approximately 5 MB in size. The files are stored in Amazon S3. Company policy requires the files to be stored for 4 years before they can be deleted. Immediate accessibility is always required as the files contain critical business data that is not easy to reproduce. The files are frequently accessed in the first 30 days of the object creation but are rarely accessed after the first 30 days. Which storage solution is MOST cost-effective?

- A. Create an S3 bucket lifecycle policy to move files from S3 Standard to S3 Glacier 30 days from object creation. Delete the files 4 years after object creation.
- B. Create an S3 bucket lifecycle policy to move files from S3 Standard to S3 One Zone- Infrequent Access (S3 One Zone-IA) 30 days from object creation. Delete the files 4 years after object creation.
- C. Delete the files 4 years after object creation.
- D. Create an S3 bucket lifecycle policy to move files from S3 Standard-Infrequent Access (S3 Standard-IA) 30 days from object creation. Delete the files 4 years after object creation.
- E. Delete the files 4 years after object creation.
- F. Create an S3 bucket lifecycle policy to move files from S3 Standard to S3 Standard- Infrequent Access (S3 Standard-IA) 30 days from object creation. Move the files to S3 Glacier 4 years after object creation.

Answer: B

Explanation:

https://aws.amazon.com/s3/storage-classes/?trk=66264cd8-3b73-416c-9693-ea7cf4fe846a&sc_channel=ps&s_kwid=AL14422!3!536452716950!p!g!aws%20s3%20pricing&ef_id=Cj0KCQjwnbmaBhD-ARIsAGTPcfVHUZN5_BMrzI5zBcaC8KnqpnNZvbZzqPkH6k7q4JcYO5KFLx0YYgaAm6nEALw_wcB:s&s_kwid=AL14422!3!536452716950!p!g!aws%20s3%20pricing

NEW QUESTION 2

- (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 3

- (Topic 1)

A company is implementing a new business application. The application runs on two Amazon EC2 instances and uses an Amazon S3 bucket for document storage. A solutions architect needs to ensure that the EC2 instances can access the S3 bucket. What should the solutions architect do to meet this requirement?

- A. Create an IAM role that grants access to the S3 bucket.
- B. Attach the role to the EC2 instances.
- C. Create an IAM policy that grants access to the S3 bucket.
- D. Attach the policy to the EC2 instances.
- E. Create an IAM group that grants access to the S3 bucket.
- F. Attach the group to the EC2 instances.
- G. Create an IAM user that grants access to the S3 bucket.
- H. Attach the user account to the EC2 instances.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/ec2-instance-access-s3-bucket/>

NEW QUESTION 4

- (Topic 1)

A company needs the ability to analyze the log files of its proprietary application. The logs are stored in JSON format in an Amazon S3 bucket. Queries will be simple and will run on-demand. A solutions architect needs to perform the analysis with minimal changes to the existing architecture. What should the solutions architect do to meet these requirements with the LEAST amount of operational overhead?

- A. Use Amazon Redshift to load all the content into one place and run the SQL queries as needed.
- B. Use Amazon CloudWatch Logs to store the logs. Run SQL queries as needed from the Amazon CloudWatch console.
- C. Use Amazon Athena directly with Amazon S3 to run the queries as needed.
- D. Use AWS Glue to catalog the logs. Use a transient Apache Spark cluster on Amazon EMR to run the SQL queries as needed.

Answer: C

Explanation:

Amazon Athena can be used to query JSON in S3

NEW QUESTION 5

- (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 6

- (Topic 1)

A company is hosting a static website on Amazon S3 and is using Amazon Route 53 for DNS. The website is experiencing increased demand from around the world. The company must decrease latency for users who access the website.

Which solution meets these requirements MOST cost-effectively?

- A. Replicate the S3 bucket that contains the website to all AWS Region
- B. Add Route 53 geolocation routing entries.
- C. Provision accelerators in AWS Global Accelerator
- D. Associate the supplied IP addresses with the S3 bucket
- E. Edit the Route 53 entries to point to the IP addresses of the accelerators.
- F. Add an Amazon CloudFront distribution in front of the S3 bucket
- G. Edit the Route 53 entries to point to the CloudFront distribution.
- H. Enable S3 Transfer Acceleration on the bucket
- I. Edit the Route 53 entries to point to the new endpoint.

Answer: C

Explanation:

Amazon CloudFront is a content delivery network (CDN) that caches content at edge locations around the world, providing low latency and high transfer speeds to users accessing the content. Adding a CloudFront distribution in front of the S3 bucket will cache the static website's content at edge locations around the world, decreasing latency for users accessing the website. This solution is also cost-effective as it only charges for the data transfer and requests made by users accessing the content from the CloudFront edge locations. Additionally, this solution provides scalability and reliability benefits as CloudFront can automatically scale to handle increased demand and provide high availability for the website.

NEW QUESTION 7

- (Topic 1)

A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database table that contains more than 10 million rows. The database has 2 TB of General Purpose SSD storage. There are millions of updates against this data every day through the company's website. The company has noticed that some insert operations are taking 10 seconds or longer. The company has determined that the database storage performance is the problem.

Which solution addresses this performance issue?

- A. Change the storage type to Provisioned IOPS SSD
- B. Change the DB instance to a memory optimized instance class
- C. Change the DB instance to a burstable performance instance class
- D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Answer: A

Explanation:

<https://aws.amazon.com/ebs/features/>

"Provisioned IOPS volumes are backed by solid-state drives (SSDs) and are the highest performance EBS volumes designed for your critical, I/O intensive database applications.

These volumes are ideal for both IOPS-intensive and throughput-intensive workloads that require extremely low latency."

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

NEW QUESTION 8

- (Topic 1)

A company recently migrated a message processing system to AWS. The system receives messages into an ActiveMQ queue running on an Amazon EC2 instance. Messages are processed by a consumer application running on Amazon EC2. The consumer application processes the messages and writes results to a MySQL database running on Amazon EC2. The company wants this application to be highly available with low operational complexity.

Which architecture offers the HIGHEST availability?

- A. Add a second ActiveMQ server to another Availability Zone. Add an additional consumer EC2 instance in another Availability Zone.
- B. Replicate the MySQL database to another Availability Zone.

- C. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an additional consumer EC2 instance in another Availability Zone
- D. Replicate the MySQL database to another Availability Zone.
- E. Use Amazon MQ with active/standby brokers configured across two Availability Zone
- F. Add an additional consumer EC2 instance in another Availability Zone
- G. Use Amazon RDS for MySQL with Multi-AZ enabled.
- H. Use Amazon MQ with active/standby brokers configured across two Availability Zones Add an Auto Scaling group for the consumer EC2 instances across two Availability Zone
- I. Use Amazon RDS for MySQL with Multi-AZ enabled.

Answer: D

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 9

- (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 10

- (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.examtactics.com/discussions/amazon/view/7193-exam-aws-certified-solutions-architect-professional-topic-1/>

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

NEW QUESTION 10

- (Topic 1)

A company has applications that run on Amazon EC2 instances in a VPC. One of the applications needs to call the Amazon S3 API to store and read objects. According to the company's security regulations, no traffic from the applications is allowed to travel across the internet.

Which solution will meet these requirements?

- A. Configure an S3 interface endpoint.
- B. Configure an S3 gateway endpoint.
- C. Create an S3 bucket in a private subnet.
- D. Create an S3 bucket in the same Region as the EC2 instance.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3>
<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>

NEW QUESTION 11

- (Topic 1)

A company recently launched Linux-based application instances on Amazon EC2 in a private subnet and launched a Linux-based bastion host on an Amazon EC2 instance in a public subnet of a VPC. A solutions architect needs to connect from the on-premises network, through the company's internet connection to the bastion host and to the application servers. The solutions architect must make sure that the security groups of all the EC2 instances will allow that access.

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

- A. Replace the current security group of the bastion host with one that only allows inbound access from the application instances.
- B. Replace the current security group of the bastion host with one that only allows inbound access from the internal IP range for the company.
- C. Replace the current security group of the bastion host with one that only allows inbound access from the external IP range for the company.
- D. Replace the current security group of the application instances with one that allows inbound SSH access from only the private IP address of the bastion host.
- E. Replace the current security group of the application instances with one that allows inbound SSH access from only the public IP address of the bastion host.

Answer: CD

Explanation:

<https://digitalcloud.training/ssh-into-ec2-in-private-subnet/>

NEW QUESTION 13

- (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 space.
- 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 16

- (Topic 1)

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.

What is the FASTEST way to aggregate data from all of these global sites?

- A. Enable Amazon S3 Transfer Acceleration on the destination bucket.
- B. Use multipart uploads to directly upload site data to the destination bucket.
- C. Upload site data to an Amazon S3 bucket in the closest AWS Region.
- D. Use S3 cross-Region replication to copy objects to the destination bucket.
- E. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region.
- F. Use S3 cross-Region replication to copy objects to the destination bucket.
- G. Upload the data to an Amazon EC2 instance in the closest Region.
- H. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume.
- I. Once a day take an EBS snapshot and copy it to the centralized Region.
- J. Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Answer: A

Explanation:

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:

You have customers that upload to a centralized bucket from all over the world. You transfer gigabytes to terabytes of data on a regular basis across continents.

You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html>

[https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications:](https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications:)

"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects.

Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet."

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

"Improved throughput - You can upload parts in parallel to improve throughput."

NEW QUESTION 19

- (Topic 1)

A company is preparing to deploy a new serverless workload. A solutions architect must use the principle of least privilege to configure permissions that will be used to run an AWS Lambda function. An Amazon EventBridge (Amazon CloudWatch Events) rule will invoke the function. Which solution meets these requirements?

- A. Add an execution role to the function with `lambda: InvokeFunction` as the action and `*` as the principal.
- B. Add an execution role to the function with `lambda: InvokeFunction` as the action and `Service:amazonaws.com` as the principal.
- C. Add a resource-based policy to the function with `lambda:*` as the action and `Service:events.amazonaws.com` as the principal.
- D. Add a resource-based policy to the function with `lambda: InvokeFunction` as the action and `Service:events.amazonaws.com` as the principal.

Answer: D

Explanation:

<https://docs.aws.amazon.com/eventbridge/latest/userguide/resource-based-policies-eventbridge.html#lambda-permissions>

NEW QUESTION 20

- (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 23

- (Topic 1)

A company runs a photo processing application that needs to frequently upload and download pictures from Amazon S3 buckets that are located in the same AWS Region. A solutions architect has noticed an increased cost in data transfer fees and needs to implement a solution to reduce these costs. How can the solutions architect meet this requirement?

- A. Deploy Amazon API Gateway into a public subnet and adjust the route table to route S3 calls through it.
- B. Deploy a NAT gateway into a public subnet and attach an endpoint policy that allows access to the S3 buckets.
- C. Deploy the application into a public subnet and allow it to route through an internet gateway to access the S3 Buckets
- D. Deploy an S3 VPC gateway endpoint into the VPC and attach an endpoint policy that allows access to the S3 buckets.

Answer: D

Explanation:

The correct answer is Option D. Deploy an S3 VPC gateway endpoint into the VPC and attach an endpoint policy that allows access to the S3 buckets. By deploying an S3 VPC gateway endpoint, the application can access the S3 buckets over a private network connection within the VPC, eliminating the need for data transfer over the internet. This can help reduce data transfer fees as well as improve the performance of the application. The endpoint policy can be used to specify which S3 buckets the application has access to.

NEW QUESTION 24

- (Topic 1)

A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance. A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours. The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue. Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the DB instance to be a Multi-AZ deployment
- B. Create a read replica of the database. Configure the script to query only the read replica.
- C. Instruct the development team to manually export the entries in the database at the end of each day.
- D. Use Amazon ElastiCache to cache the common queries that the script runs against the database.

Answer: B

NEW QUESTION 27

- (Topic 1)

A company runs a shopping application that uses Amazon DynamoDB to store customer information. In case of data corruption, a solutions architect needs to design a solution that meets a recovery point objective (RPO) of 15 minutes and a recovery time objective (RTO) of 1 hour.

What should the solutions architect recommend to meet these requirements?

- A. Configure DynamoDB global table
- B. For RPO recovery, point the application to a different AWS Region.
- C. Configure DynamoDB point-in-time recover
- D. For RPO recovery, restore to the desired point in time.
- E. Export the DynamoDB data to Amazon S3 Glacier on a daily basi
- F. For RPO recovery, import the data from S3 Glacier to DynamoDB.
- G. Schedule Amazon Elastic Block Store (Amazon EBS) snapshots for the DynamoDB table every 15 minute
- H. For RPO recovery, restore the DynamoDB table by using the EBS snapshot.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/PointInTimeRecovery.html>

NEW QUESTION 31

- (Topic 1)

A survey company has gathered data for several years from areas in the United States. The company hosts the data in an Amazon S3 bucket that is 3 TB in size and growing. The company has started to share the data with a European marketing firm that has S3 buckets. The company wants to ensure that its data transfer costs remain as low as possible.

Which solution will meet these requirements?

- A. Configure the Requester Pays feature on the company's S3 bucket
- B. Configure S3 Cross-Region Replication from the company's S3 bucket to one of the marketing firm's S3 buckets.
- C. Configure cross-account access for the marketing firm so that the marketing firm has access to the company's S3 bucket.
- D. Configure the company's S3 bucket to use S3 Intelligent-Tiering Sync the S3 bucket to one of the marketing firm's S3 buckets

Answer: A

Explanation:

"Typically, you configure buckets to be Requester Pays buckets when you want to share data but not incur charges associated with others accessing the data. For example, you might use Requester Pays buckets when making available large datasets, such as zip code directories, reference data, geospatial information, or web crawling data." <https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysBuckets.html>

NEW QUESTION 33

- (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 34

- (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 37

- (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 layer
- 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 group
- 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 group
- 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 39

- (Topic 2)

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.

Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instance
- B. Stop the RDS database when it is not in use.
- C. Purchase EC2 Instance Savings Plans to cover five EC2 instance
- D. Purchase an RDS Reserved DB Instance
- E. Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as needed
- F. Stop the RDS database when it is not in use.
- G. Purchase EC2 Instance Savings Plans to cover two EC2 instance
- H. Use up to three additional EC2 On-Demand Instances as needed
- I. Purchase an RDS Reserved DB Instance.

Answer: C

Explanation:

This solution meets the requirements of a two-tier application that has a variable demand based on the time of day and must be available at all times, while minimizing the overall cost. EC2 Reserved Instances can provide significant savings compared to On-Demand Instances for the baseline level of usage, and they can guarantee capacity reservation when needed. EC2 Spot Instances can provide up to 90% savings compared to On-Demand Instances for any additional capacity that the application needs during peak hours. Spot Instances are suitable for stateless applications that can tolerate interruptions and can be replaced by other instances. Stopping the RDS database when it is not in use can reduce the cost of running the database tier. Option A is incorrect because using all EC2 Spot Instances can affect the availability of the application if there are not enough spare capacity or if the Spot price exceeds the maximum price. Stopping the RDS database when it is not in use can reduce the cost of running the database tier, but it can also affect the availability of the application. Option B is incorrect because purchasing EC2 Instance Savings Plans to cover five EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Purchasing an RDS Reserved DB Instance can provide savings for the database tier, but it does not allow stopping the database when it is not in use. Option D is incorrect because purchasing EC2 Instance Savings Plans to cover two EC2 instances can lock in a fixed amount of compute usage per hour, which may not match the actual usage pattern of the application. Using up to three additional EC2 On-Demand Instances as needed can incur higher costs than using Spot Instances.

References:

- ? <https://aws.amazon.com/ec2/pricing/reserved-instances/>
- ? <https://aws.amazon.com/ec2/spot/>
- ? https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

NEW QUESTION 40

- (Topic 2)

A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date.

Which solution will meet these requirements?

- A. Use S3 Object Lock In governance mode with a legal hold of 1 year
- B. Use S3 Object Lock in compliance mode with a retention period of 365 days.
- C. Use an IAM role to restrict all users from deleting or changing objects in the S3 bucket Use an S3 bucket policy to only allow the IAM role
- D. Configure the S3 bucket to invoke an AWS Lambda function every time an object is added Configure the function to track the hash of the saved object to that modified objects can be marked accordingly

Answer: B

Explanation:

In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period. In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. In Governance mode, Objects can be deleted by some users with special permissions, this is against the requirement.

Compliance:

- Object versions can't be overwritten or deleted by any user, including the root user
- Objects retention modes can't be changed, and retention periods can't be shortened

Governance:

- Most users can't overwrite or delete an object version or alter its lock settings
- Some users have special permissions to change the retention or delete the object

NEW QUESTION 44

- (Topic 2)

An application runs on Amazon EC2 instances across multiple Availability Zones. The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer. The application performs best when the CPU utilization of the EC2 instances is at or near 40%. What should a solutions architect do to maintain the desired performance across all instances in the group?

- A. Use a simple scaling policy to dynamically scale the Auto Scaling group
- B. Use a target tracking policy to dynamically scale the Auto Scaling group
- C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.
- D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-target-tracking.html>

NEW QUESTION 48

- (Topic 2)

An ecommerce company hosts its analytics application in the AWS Cloud. The application generates about 300 MB of data each month. The data is stored in JSON format. The company is evaluating a disaster recovery solution to back up the data. The data must be accessible in milliseconds if it is needed, and the data must be kept for 30 days.

Which solution meets these requirements MOST cost-effectively?

- A. Amazon OpenSearch Service (Amazon Elasticsearch Service)
- B. Amazon S3 Glacier
- C. Amazon S3 Standard
- D. Amazon RDS for PostgreSQL

Answer: C

Explanation:

This solution meets the requirements of a disaster recovery solution to back up the data that is generated by an analytics application, stored in JSON format, and must be accessible in milliseconds if it is needed. Amazon S3 Standard is a durable and scalable storage class for frequently accessed data. It can store any amount of data and provide high availability and performance. It can also support millisecond access time for data retrieval.

Option A is incorrect because Amazon OpenSearch Service (Amazon Elasticsearch Service) is a search and analytics service that can index and query data, but it is not a backup solution for data stored in JSON format. Option B is incorrect because Amazon S3 Glacier is a low-cost storage class for data archiving and long-term backup, but it does not support millisecond access time for data retrieval. Option D is incorrect because Amazon RDS for PostgreSQL is a relational database service that can store and query structured data, but it is not a backup solution for data stored in JSON format.

References:

? <https://aws.amazon.com/s3/storage-classes/>

? https://aws.amazon.com/s3/faqs/#Durability_and_data_protection

NEW QUESTION 53

- (Topic 2)

A company uses a popular content management system (CMS) for its corporate website. However, the required patching and maintenance are burdensome. The company is redesigning its website and wants a new solution. The website will be updated four times a year and does not need to have any dynamic content available. The solution must provide high scalability and enhanced security.

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

- A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality
- B. Create and deploy an AWS Lambda function to manage and serve the website content
- C. Create the new website and an Amazon S3 bucket. Deploy the website on the S3 bucket with static website hosting enabled
- D. Create the new website
- E. Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer.

Answer: AD

Explanation:

A -> We can configure CloudFront to require HTTPS from clients (enhanced security)

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-viewers-to-cloudfront.html> D -> storing static website on S3 provides scalability and less operational overhead, then configuration of Application LB and EC2 instances (hence E is out)

NEW QUESTION 56

- (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 a 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.

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 58

- (Topic 2)

A company is planning to build a high performance computing (HPC) workload as a service solution that is hosted on AWS. A group of 16 Amazon EC2 Linux instances requires the lowest possible latency for node-to-node communication. The instances also need a shared block device volume for high-performance storage.

Which solution will meet these requirements?

- A. Use a distributed placement group
- B. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach
- C. Use a cluster placement group
- D. Create shared file systems across the instances by using Amazon Elastic File System (Amazon EFS)
- E. Use a partition placement group
- F. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
- G. Use a spread placement group
- H. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach

Answer: A

Explanation:

1. lowest possible latency + node to node ==> cluster placement (must be within one AZ), so C, D out

* 2. For EBS Multi-Attach, up to 16 instances can be attached to a single volume ==> we have 16 Linux instances ==> more close to A

* 3. "need a shared block device volume" ==> EBS Multi-attach is Block Storage whereas EFS is File Storage ==> B out

* 4. EFS automatically replicates data within and across 3 AZ ==> we use cluster placement

so all EC2 are within one AZ.

* 5. EBS Multi-attach volumes can be used for clients within a single AZ. <https://repost.aws/questions/QUK2RANw1QTKCwpDUwCCI72A/efs-vs-ebs-mult-attach>

NEW QUESTION 60

- (Topic 2)

A corporation has recruited a new cloud engineer who should not have access to the CompanyConfidential Amazon S3 bucket. The cloud engineer must have read and write permissions on an S3 bucket named AdminTools.

Which IAM policy will satisfy these criteria?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::AdminTools"
    },
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject" ],
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::CompanyConfidential/*",
        "arn:aws:s3:::CompanyConfidential"
      ]
    }
  ]
}
```

B.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": [
        "arn:aws:s3:::AdminTools",
        "arn:aws:s3:::CompanyConfidential/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject", "s3>DeleteObject" ],
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::CompanyConfidential"
    }
  ]
}
```

C.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject" ],
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::CompanyConfidential/*",
        "arn:aws:s3:::CompanyConfidential"
      ]
    }
  ]
}
```

D.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
      "Resource": "arn:aws:s3:::AdminTools/"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::CompanyConfidential",
        "arn:aws:s3:::CompanyConfidential/*",
        "arn:aws:s3:::AdminTools/*"
      ]
    }
  ]
}
```

A.

Answer: A

Explanation:

https://docs.amazonaws.cn/en_us/IAM/latest/UserGuide/reference_policies_examples_s3_rw-bucket.html

The policy is separated into two parts because the ListBucket action requires permissions on the bucket while the other actions require permissions on the objects in the bucket. You must use two different Amazon Resource Names (ARNs) to specify bucket-level and object-level permissions. The first Resource element specifies arn:aws:s3:::AdminTools for the ListBucket action so that applications can list all objects in the AdminTools bucket.

NEW QUESTION 62

- (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 67

- (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 70

- (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 72

- (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 74

- (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 76

- (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 policy
- C. Update the trust relationship of the role
- 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 stopped
- F. Invalidate the cache after the DB instance is started.
- G. Launch an Amazon EC2 instance
- H. Create an IAM role that grants access to Amazon RDS
- I. Attach the role to the EC2 instance
- 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 79

- (Topic 2)

A company wants to migrate its MySQL database from on-premises to AWS. The company recently experienced a database outage that significantly impacted the business. To ensure this does not happen again, the company wants a reliable database solution on AWS that minimizes data loss and stores every transaction on at least two nodes.

Which solution meets these requirements?

- A. Create an Amazon RDS DB instance with synchronous replication to three nodes in three Availability Zones.
- B. Create an Amazon RDS MySQL DB instance with Multi-AZ functionality enabled to synchronously replicate the data.
- C. Create an Amazon RDS MySQL DB instance and then create a read replica in a separate AWS Region that synchronously replicates the data.
- D. Create an Amazon EC2 instance with a MySQL engine installed that triggers an AWS Lambda function to synchronously replicate the data to an Amazon RDS MySQL DB instance.

Answer: B

Explanation:

Q: What does Amazon RDS manage on my behalf?

Amazon RDS manages the work involved in setting up a relational database: from provisioning the infrastructure capacity you request to installing the database software. Once your database is up and running, Amazon RDS automates common administrative tasks such as performing backups and patching the software that powers your database. With optional Multi-AZ deployments, Amazon RDS also manages synchronous data replication across Availability Zones with automatic failover. <https://aws.amazon.com/rds/faqs/>

NEW QUESTION 82

- (Topic 3)

A company is migrating an old application to AWS. The application runs a batch job every hour and is CPU intensive. The batch job takes 15 minutes on average with an on-premises server. The server has 64 virtual CPU (vCPU) and 512 GiB of memory.

Which solution will run the batch job within 15 minutes with the LEAST operational overhead?

- A. Use AWS Lambda with functional scaling
- B. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate
- C. Use Amazon Lightsail with AWS Auto Scaling
- D. Use AWS Batch on Amazon EC2

Answer: D

Explanation:

Use AWS Batch on Amazon EC2. AWS Batch is a fully managed batch processing service that can be used to easily run batch jobs on Amazon EC2 instances. It can scale the number of instances to match the workload, allowing the batch job to be completed in the desired time frame with minimal operational overhead.

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

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

NEW QUESTION 86

- (Topic 3)

An e-commerce company is experiencing an increase in user traffic. The company's store is deployed on Amazon EC2 instances as a two-tier web application consisting of a web tier and a separate database tier. As traffic increases, the company notices that the architecture is causing significant delays in sending timely marketing and order confirmation emails to users. The company wants to reduce the time it spends resolving complex email delivery issues and minimize operational overhead.

What should a solutions architect do to meet these requirements?

- A. Create a separate application tier using EC2 instances dedicated to email processing.
- B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES).
- C. Configure the web instance to send email through Amazon Simple Notification Service (Amazon SNS)
- D. Create a separate application tier using EC2 instances dedicated to email processing

E. Place the instances in an Auto Scaling group.

Answer: B

Explanation:

Amazon SES is a cost-effective and scalable email service that enables businesses to send and receive email using their own email addresses and domains. Configuring the web instance to send email through Amazon SES is a simple and effective solution that can reduce the time spent resolving complex email delivery issues and minimize operational overhead.

NEW QUESTION 89

- (Topic 3)

A company will deployed a web application on AWS. The company hosts the backend database on Amazon RDS for MySQL with a primary DB instance and five read replicas to support scaling needs. The read replicas must log no more than 1 second behind the primary DB Instance. The database routinely runs scheduled stored procedures.

As traffic on the website increases, the replicas experinces additional lag during periods of peak lead. A solutions architect must reduce the replication lag as much as possible. The solutions architect must minimize changes to the applicatin code and must minimize ongoing overhead.

Which solution will meet these requirements?

Migrate the database to Amazon Aurora MySQL. Replace the read replicas with Aurora Replicas, and configure Aurora Auto Scaling. Replace the stored procedures with Aurora MySQL native functions.

Deploy an Amazon ElastiCache for Redis cluser in front of the database. Modify the application to check the cache before the application queries the database. Repace the stored procedures with AWS Lambda funcions.

- A. Migrate the database to a MYSQL database that runs on Amazn EC2 instance
- B. Choose large, compute optimized for all replica node
- C. Maintain the stored procedures on the EC2 instances.
- D. Deploy an Amazon ElastiCache for Redis cluster in fornt of the databas
- E. Modify the application to check the cache before the application queries the databas
- F. Replace the stored procedures with AWS Lambda functions.
- G. Migrate the database to a MySQL database that runs on Amazon EC2 instance
- H. Choose large, compute optimized EC2 instances for all replica nodes, Maintain the stored procedures on the EC2 instances.
- I. Migrate the database to Amazon DynamoDB, Provision number of read capacity units (RCUs) to support the required throughput, and configure on-demand capacity scalin
- J. Replace the stored procedures with DynamoDB streams.

Answer: A

Explanation:

Option A is the most appropriate solution for reducing replication lag without significant changes to the application code and minimizing ongoing operational overhead. Migrating the database to Amazon Aurora MySQL allows for improved replication performance and higher scalability compared to Amazon RDS for MySQL. Aurora Replicas provide faster replication, reducing the replication lag, and Aurora Auto Scaling ensures that there are enough Aurora Replicas to handle the incoming traffic. Additionally, Aurora MySQL native functions can replace the stored procedures, reducing the load on the database and improving performance.

NEW QUESTION 91

- (Topic 3)

A company deploys an appliation on five Amazon EC2 instances. An Applicatin Load Balancer (ALB) distributes traffic to the instances by using a target group. The average CPU usage on each of the insatances is below 10% most of the time. With occasional surges to 65%.

A solution architect needs to implement a solution to automate the scalability of the application. The solution must optimize the cost of the architecture and must ensure that the application has enough CPU resources when surges occur.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm that enters the ALARM state when the CPUUtilization metric is less than 20%. Create an AWS Lambda function that the CloudWatch alarm invokes to terminate one of the EC2 instances in the ALB target group.
- B. Create an EC2 Auto Scalin
- C. Select the exisiting ALB as the load balancer and the existing target group as the target grou
- D. Set a target tracking scaling policy that is based on the ASGAverageCPUUtilization metri
- E. Set the minimum instances to 2, the desired capacity to 3, the desired capacity to 3, the maximum instances to 6, and the target value to 50%. And the EC2 instances to the Auto Scaling group.
- F. Create an EC2 Auto Scalin
- G. Select the exisiting ALB as the load balancer and the existing target grou
- H. Set the minimum instances to 2, the desired capacity to 3, and the maximum instances to 6 Add the EC2 instances to the Scaling group.
- I. Create two Amazon CloudWatch alarm
- J. Configure the first CloudWatch alarm to enter the ALARM satet when the average CPUTUtilization metric is below 20%. Configure the seconnd CloudWatch alarm to enter the ALARM state when the average CPUUtilization metric is aboove 50%. Configure the alarms to publish to an Amazon Simple Notification Service (Amazon SNS) topic to send an email messag
- K. After receiving the message, log in to decrease or increase the number of EC2 instances that are running

Answer: B

Explanation:

- An Auto Scaling group will automatically scale the EC2 instances to match changes in demand. This optimizes cost by only running as many instances as needed.
- A target tracking scaling policy monitors the ASGAverageCPUUtilization metric and scales to keep the average CPU around the 50% target value. This ensures there are enough resources during CPU surges.
- The ALB and target group are reused, so the application architecture does not change. The Auto Scaling group is associated to the existing load balancer setup.
- A minimum of 2 and maximum of 6 instances provides the ability to scale between 3 and 6 instances as needed based on demand.
- Costs are optimized by starting with only 3 instances (the desired capacity) and scaling up as needed. When CPU usage drops, instances are terminated to match the desired capacity.

NEW QUESTION 92

- (Topic 3)

A company experienced a breach that affected several applications in its on-premises data center. The attacker took advantage of vulnerabilities in the custom applications that were running on the servers. The company is now migrating its applications to run on Amazon EC2 instances. The company wants to implement a solution that actively scans for vulnerabilities on the EC2 instances and sends a report that details the findings. Which solution will meet these requirements?

- A. Deploy AWS Shield to scan the EC2 instances for vulnerabilities. Create an AWS Lambda function to log any findings to AWS CloudTrail.
- B. Deploy Amazon Macie and AWS Lambda functions to scan the EC2 instances for vulnerabilities. Log any findings to AWS CloudTrail.
- C. Turn on Amazon GuardDuty. Deploy the GuardDuty agents to the EC2 instances. Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings.
- D. Turn on Amazon Inspector. Deploy the Amazon Inspector agent to the EC2 instances. Configure an AWS Lambda function to automate the generation and distribution of reports that detail the findings.

Answer: D

Explanation:

Amazon Inspector:

- Performs active vulnerability scans of EC2 instances. It looks for software vulnerabilities, unintended network accessibility, and other security issues.
- Requires installing an agent on EC2 instances to perform scans. The agent must be deployed to each instance.
- Provides scheduled scan reports detailing any findings of security risks or vulnerabilities. These reports can be used to patch or remediate issues.
- Is best suited for proactively detecting security weaknesses and misconfigurations in your AWS environment.

NEW QUESTION 97

- (Topic 3)

A company has an application that runs on several Amazon EC2 instances. Each EC2 instance has multiple Amazon Elastic Block Store (Amazon EBS) data volumes attached to it. The application's EC2 instance configuration and data need to be backed up nightly. The application also needs to be recoverable in a different AWS Region.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Write an AWS Lambda function that schedules nightly snapshots of the application's EBS volumes and copies the snapshots to a different Region.
- B. Create a backup plan by using AWS Backup to perform nightly backup.
- C. Copy the backups to another Region. Add the application's EC2 instances as resources.
- D. Create a backup plan by using AWS Backup to perform nightly backups. Copy the backups to another Region. Add the application's EBS volumes as resources.
- E. Write an AWS Lambda function that schedules nightly snapshots of the application's EBS volumes and copies the snapshots to a different Availability Zone.

Answer: B

Explanation:

The most operationally efficient solution to meet these requirements would be to create a backup plan by using AWS Backup to perform nightly backups and copying the backups to another Region. Adding the application's EBS volumes as resources will ensure that the application's EC2 instance configuration and data are backed up, and copying the backups to another Region will ensure that the application is recoverable in a different AWS Region.

NEW QUESTION 100

- (Topic 3)

A company is hosting a web application from an Amazon S3 bucket. The application uses Amazon Cognito as an identity provider to authenticate users and return a JSON Web Token (JWT) that provides access to protected resources that are stored in another S3 bucket.

Upon deployment of the application, users report errors and are unable to access the protected content. A solutions architect must resolve this issue by providing proper permissions so that users can access the protected content.

Which solution meets these requirements?

- A. Update the Amazon Cognito identity pool to assume the proper IAM role for access to the protected content.
- B. Update the S3 ACL to allow the application to access the protected content.
- C. Redeploy the application to Amazon S3 to prevent eventually consistent reads in the S3 bucket from affecting the ability of users to access the protected content.
- D. Update the Amazon Cognito pool to use custom attribute mappings within the identity pool and grant users the proper permissions to access the protected content.

Answer: A

Explanation:

Amazon Cognito identity pools assign your authenticated users a set of temporary, limited-privilege credentials to access your AWS resources. The permissions for each user are controlled through IAM roles that you create. <https://docs.aws.amazon.com/cognito/latest/developerguide/role-based-access-control.html>

NEW QUESTION 103

- (Topic 3)

A solutions architect must migrate a Windows Internet Information Services (IIS) web application to AWS. The application currently relies on a file share hosted in the user's on-premises network-attached storage (NAS). The solutions architect has proposed migrating the MS web servers to Amazon EC2 instances in multiple Availability Zones that are connected to the storage solution, and configuring an Elastic Load Balancer attached to the instances.

Which replacement to the on-premises file share is MOST resilient and durable?

- A. Migrate the file share to Amazon RDS.
- B. Migrate the file share to AWS Storage Gateway.
- C. Migrate the file share to Amazon FSx for Windows File Server.
- D. Migrate the file share to Amazon Elastic File System (Amazon EFS).

Answer: C

Explanation:

This answer is correct because it provides a resilient and durable replacement for the on-premises file share that is compatible with Windows IIS web servers.

Amazon FSx for Windows File Server is a fully managed service that provides shared file storage built on Windows Server. It supports the SMB protocol and integrates with Microsoft Active Directory, which enables seamless access and authentication for Windows-based applications. Amazon FSx for Windows File Server also offers the following benefits:

? Resilience: Amazon FSx for Windows File Server can be deployed in multiple

Availability Zones, which provides high availability and failover protection. It also supports automatic backups and restores, as well as self-healing features that detect and correct issues.

? Durability: Amazon FSx for Windows File Server replicates data within and across

Availability Zones, and stores data on highly durable storage devices. It also supports encryption at rest and in transit, as well as file access auditing and data deduplication.

? Performance: Amazon FSx for Windows File Server delivers consistent sub-

millisecond latencies and high throughput for file operations. It also supports SSD storage, native Windows features such as Distributed File System (DFS) Namespaces and Replication, and user-driven performance scaling.

References:

? Amazon FSx for Windows File Server

? Using Microsoft Windows file shares

NEW QUESTION 107

- (Topic 3)

A company hosts its application on AWS. The company uses Amazon Cognito to manage users. When users log in to the application, the application fetches required data from Amazon DynamoDB by using a REST API that is hosted in Amazon API Gateway. The company wants an AWS managed solution that will control access to the REST API to reduce development efforts.

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

A. Configure an AWS Lambda function to be an authorizer in API Gateway to validate which user made the request.

B. For each user, create and assign an API key that must be sent with each request. Validate the key by using an AWS Lambda function.

C. Send the user's email address in the header with every request. Invoke an AWS Lambda function to validate that the user with that email address has proper access.

D. Configure an Amazon Cognito user pool authorizer in API Gateway to allow Amazon Cognito to validate each request.

Answer: D

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-integrate-with-cognito.html>

To control access to the REST API and reduce development efforts, the company can use an Amazon Cognito user pool authorizer in API Gateway. This will allow Amazon Cognito to validate each request and ensure that only authenticated users can access the API. This solution has the LEAST operational overhead, as it does not require the company to develop and maintain any additional infrastructure or code.

NEW QUESTION 109

- (Topic 3)

An IAM user made several configuration changes to AWS resources in their company's account during a production deployment last week. A solutions architect learned that a couple of security group rules are not configured as desired. The solutions architect wants to confirm which IAM user was responsible for making changes.

Which service should the solutions architect use to find the desired information?

A. Amazon GuardDuty

B. Amazon Inspector

C. AWS CloudTrail

D. AWS Config

Answer: C

Explanation:

The best option is to use AWS CloudTrail to find the desired information. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of AWS account activities. CloudTrail can be used to log all changes made to resources in an AWS account, including changes made by IAM users, EC2 instances, AWS management console, and other AWS services. By using CloudTrail, the solutions architect can identify the IAM user who made the configuration changes to the security group rules.

NEW QUESTION 114

- (Topic 3)

A company serves a dynamic website from a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The website needs to support multiple languages to serve customers around the world. The website's architecture is running in the us-west-1 Region and is exhibiting high request latency for users that are located in other parts of the world.

The website needs to serve requests quickly and efficiently regardless of a user's location. However, the company does not want to recreate the existing architecture across multiple Regions.

What should a solutions architect do to meet these requirements?

A. Replace the existing architecture with a website that is served from an Amazon S3 bucket. Configure an Amazon CloudFront distribution with the S3 bucket as the origin. Set the cache behavior settings to cache based on the Accept-Language request header.

B. Configure an Amazon CloudFront distribution with the ALB as the origin. Set the cache behavior settings to cache based on the Accept-Language request header.

C. Create an Amazon API Gateway API that is integrated with the ALB. Configure the API to use the HTTP integration type. Set up an API Gateway stage to enable the API cache based on the Accept-Language request header.

D. Launch an EC2 instance in each additional Region and configure NGINX to act as a cache server for that Region. Put all the EC2 instances and the ALB behind an Amazon Route 53 record set with a geolocation routing policy.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/header-caching.html> Configuring caching based on the language of the viewer: If you want CloudFront to cache different versions of your objects based on the language specified in the request, configure CloudFront to forward the Accept-Language

header to your origin.

NEW QUESTION 117

- (Topic 3)

A company has a custom application with embedded credentials that retrieves information from an Amazon RDS MySQL DB instance. Management says the application must be made more secure with the least amount of programming effort.

What should a solutions architect do to meet these requirements?

- A. Use AWS Key Management Service (AWS KMS) customer master keys (CMKs) to create key
- B. Configure the application to load the database credentials from AWS KM
- C. Enable automatic key rotation.
- D. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manage
- E. Configure the application to load the database credentials from Secrets Manage
- F. Create an AWS Lambda function that rotates the credentials in Secret Manager.
- G. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manage
- H. Configure the application to load the database credentials from Secrets Manage
- I. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Secrets Manager.
- J. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Systems Manager Parameter Stor
- K. Configure the application to load the database credentials from Parameter Stor
- L. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Parameter Store.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>

NEW QUESTION 121

- (Topic 3)

A company must migrate 20 TB of data from a data center to the AWS Cloud within 30 days. The company's network bandwidth is limited to 15 Mbps and cannot exceed 70% utilization. What should a solutions architect do to meet these requirements?

- A. Use AWS Snowball.
- B. Use AWS DataSync.
- C. Use a secure VPN connection.
- D. Use Amazon S3 Transfer Acceleration.

Answer: A

Explanation:

AWS Snowball is a secure data transport solution that accelerates moving large amounts of data into and out of the AWS cloud. It can move up to 80 TB of data at a time, and provides a network bandwidth of up to 50 Mbps, so it is well-suited for the task. Additionally, it is secure and easy to use, making it the ideal solution for this migration.

NEW QUESTION 123

- (Topic 3)

A company is running a publicly accessible serverless application that uses Amazon API Gateway and AWS Lambda. The application's traffic recently spiked due to fraudulent requests from botnets.

Which steps should a solutions architect take to block requests from unauthorized users? (Select TWO.)

- A. Create a usage plan with an API key that is shared with genuine users only.
- B. Integrate logic within the Lambda function to ignore the requests from fraudulent IP addresses.
- C. Implement an AWS WAF rule to target malicious requests and trigger actions to filter them out.
- D. Convert the existing public API to a private AP
- E. Update the DNS records to redirect users to the new API endpoint.
- F. Create an IAM role for each user attempting to access the AP
- G. A user will assume the role when making the API call.

Answer: AC

Explanation:

[https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html](https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html#:~:text=Don%27t%20rely%20on%20API%20keys%20as%20y our%20only%20means%20of%20authentication%20and%20authorization%20for%20your%20APIs)
<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

NEW QUESTION 127

- (Topic 3)

A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time.

The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy.

Which solution meets these requirements?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon S3 Glacier Deep Archive
- D. AWS Backup

Answer: A

Explanation:

Amazon EFS provides a simple, scalable, fully managed file system that can be simultaneously accessed from multiple EC2 instances and provides built-in redundancy. It is optimized for multiple EC2 instances to access the same files, and it is designed to be highly available, durable, and secure. It can scale up to petabytes of data and can handle thousands of concurrent connections, and is a cost-effective solution for storing and accessing large amounts of data.

NEW QUESTION 130

- (Topic 3)

A company provides an API to its users that automates inquiries for tax computations based on item prices. The company experiences a larger number of inquiries during the holiday season only that cause slower response times. A solutions architect needs to design a solution that is scalable and elastic. What should the solutions architect do to accomplish this?

- A. Provide an API hosted on an Amazon EC2 instance
- B. The EC2 instance performs the required computations when the API request is made.
- C. Design a REST API using Amazon API Gateway that accepts the item name
- D. API Gateway passes item names to AWS Lambda for tax computations.
- E. Create an Application Load Balancer that has two Amazon EC2 instances behind it
- F. The EC2 instances will compute the tax on the received item names.
- G. Design a REST API using Amazon API Gateway that connects with an API hosted on an Amazon EC2 instance
- H. API Gateway accepts and passes the item names to the EC2 instance for tax computations.

Answer: B

Explanation:

Lambda server-less is scalable and elastic than EC2 api gateway solution

NEW QUESTION 133

- (Topic 3)

An online learning company is migrating to the AWS Cloud. The company maintains its student records in a PostgreSQL database. The company needs a solution in which its data is available and online across multiple AWS Regions at all times.

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

- A. Migrate the PostgreSQL database to a PostgreSQL cluster on Amazon EC2 instances.
- B. Migrate the PostgreSQL database to an Amazon RDS for PostgreSQL DB instance with the Multi-AZ feature turned on.
- C. Migrate the PostgreSQL database to an Amazon RDS for PostgreSQL DB instance
- D. Create a read replica in another Region.
- E. Migrate the PostgreSQL database to an Amazon RDS for PostgreSQL DB instance
- F. Set up DB snapshots to be copied to another Region.

Answer: C

Explanation:

"online across multiple AWS Regions at all times". Currently only Read Replica supports cross-regions, Multi-AZ does not support cross-region (it works only in same region) <https://aws.amazon.com/about-aws/whats-new/2018/01/amazon-rds-read-replicas-now-support-multi-az-deployments/>

NEW QUESTION 136

- (Topic 3)

A company wants to migrate its 1 PB on-premises image repository to AWS. The images will be used by a serverless web application. Images stored in the repository are rarely accessed, but they must be immediately available. Additionally, the images must be encrypted at rest and protected from accidental deletion. Which solution meets these requirements?

- A. Implement client-side encryption and store the images in an Amazon S3 Glacier vault. Set a vault lock to prevent accidental deletion.
- B. Store the images in an Amazon S3 bucket in the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Enable versioning, default encryption, and MFA Delete on the S3 bucket.
- C. Store the images in an Amazon FSx for Windows File Server file share. Configure the Amazon FSx file share to use an AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the images in the file share. Use NTFS permission sets on the images to prevent accidental deletion.
- D. Store the images in an Amazon Elastic File System (Amazon EFS) file share in the Infrequent Access storage class. Configure the EFS file share to use an AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the images in the file share.
- E. Use NFS permission sets on the images to prevent accidental deletion.

Answer: B

Explanation:

This answer is correct because it provides a resilient and durable replacement for the on-premises file share that is compatible with a serverless web application. Amazon S3 is a fully managed object storage service that can store any amount of data and serve it over the internet. It supports the following features:

? Resilience: Amazon S3 stores data across multiple Availability Zones within a Region, and offers 99.999999999% (11 9's) of durability. It also supports cross-region replication, which enables automatic and asynchronous copying of objects across buckets in different AWS Regions.

? Durability: Amazon S3 encrypts data at rest using server-side encryption with either Amazon S3-managed keys (SSE-S3), AWS KMS keys (SSE-KMS), or customer-provided keys (SSE-C). It also supports encryption in transit using SSL/TLS. Amazon S3 also provides data protection features such as versioning, which keeps multiple versions of an object in the same bucket, and MFA Delete, which requires additional authentication for deleting an object version or changing the versioning state of a bucket.

? Performance: Amazon S3 delivers high performance and scalability for serving static and dynamic web content. It also supports features such as S3 Transfer Acceleration, which speeds up data transfers by routing requests to AWS edge locations, and S3 Select, which enables retrieving only a subset of data from an object by using simple SQL expressions.

The S3 Standard-Infrequent Access (S3 Standard-IA) storage class is suitable for storing images that are rarely accessed, but must be immediately available when needed. It offers the same high durability, throughput, and low latency as S3 Standard, but with a lower storage cost per GB and a higher per-request cost.

References:

? Amazon Simple Storage Service

? Storage classes - Amazon Simple Storage Service

NEW QUESTION 137

- (Topic 3)

A company has an application that is backed by an Amazon DynamoDB table. The company's compliance requirements specify that database backups must be taken every month, must be available for 6 months, and must be retained for 7 years. Which solution will meet these requirements?

- A. Create an AWS Backup plan to back up the DynamoDB table on the first day of each month
- B. Specify a lifecycle policy that transitions the backup to cold storage after 6 months
- C. Set the retention period for each backup to 7 years.
- D. Create a DynamoDB on-demand backup of the DynamoDB table on the first day of each month Transition the backup to Amazon S3 Glacier Flexible Retrieval after 6 months
- E. Create an S3 Lifecycle policy to delete backups that are older than 7 years.
- F. Use the AWS SDK to develop a script that creates an on-demand backup of the DynamoDB table
- G. Set up an Amazon EventBridge rule that runs the script on the first day of each month
- H. Create a second script that will run on the second day of each month to transition DynamoDB backups that are older than 6 months to cold storage and to delete backups that are older than 7 years.
- I. Use the AWS CLI to create an on-demand backup of the DynamoDB table Set up an Amazon EventBridge rule that runs the command on the first day of each month with a cron expression Specify in the command to transition the backups to cold storage after 6 months and to delete the backups after 7 years.

Answer: A

Explanation:

This solution satisfies the requirements in the following ways:

- AWS Backup will automatically take full backups of the DynamoDB table on the schedule defined in the backup plan (the first of each month).
- The lifecycle policy can transition backups to cold storage after 6 months, meeting that requirement.
- Setting a 7-year retention period in the backup plan will ensure each backup is retained for 7 years as required.
- AWS Backup manages the backup jobs and lifecycle policies, requiring no custom scripting or management.

NEW QUESTION 141

- (Topic 3)

A company's application runs on AWS. The application stores large documents in an Amazon S3 bucket that uses the S3 Standard-infrequent Access (S3 Standard-IA) storage class. The company will continue paying to store the data but wants to save on its total S3 costs. The company wants authorized external users to have the ability to access the documents in milliseconds.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure the S3 bucket to be a Requester Pays bucket
- B. Change the storage tier to S3 Standard for all existing and future objects.
- C. Turn on S3 Transfer Acceleration for the S3 bucket
- D. Use Amazon CloudFront to handle all the requests to the S3 bucket

Answer: D

Explanation:

This option is the most efficient because it uses Amazon CloudFront, which is a web service that speeds up distribution of your static and dynamic web content, such as .html,

.css, .js, and image files, to your users¹. It also uses CloudFront to handle all the requests to the S3 bucket, which reduces the S3 costs by caching the content at the edge locations and serving it from there. It also allows authorized external users to access the documents in milliseconds, as CloudFront delivers the content with low latency and high data transfer rates. This solution meets the requirement of continuing paying to store the data but saving on its total S3 costs. Option A is less efficient because it configures the S3 bucket to be a Requester Pays bucket, which is a way to shift the cost of data transfer and requests from the bucket owner to the requester². However, this does not reduce the total S3 costs, as the company still has to pay for storing the data and for any requests made by its own users. Option B is less efficient because it changes the storage tier to S3 Standard for all existing and future objects, which is a way to store frequently accessed data with high durability and availability³. However, this does not reduce the total S3 costs, as S3 Standard has higher storage costs than S3 Standard-IA. Option C is less efficient because it turns on S3 Transfer Acceleration for the S3 bucket, which is a way to speed up transfers into and out of an S3 bucket by routing requests through CloudFront edge locations⁴. However, this does not reduce the total S3 costs, as S3 Transfer Acceleration has additional charges for data transfer and requests.

NEW QUESTION 143

- (Topic 3)

A company hosts a three-tier web application that includes a PostgreSQL database The database stores the metadata from documents The company searches the metadata for key terms to retrieve documents that the company reviews in a report each month The documents are stored in Amazon S3 The documents are usually written only once, but they are updated frequently The reporting process takes a few hours with the use of relational queries The reporting process must not affect any document modifications or the addition of new documents.

What are the MOST operationally efficient solutions that meet these requirements? (Select TWO)

- A. Set up a new Amazon DocumentDB (with MongoDB compatibility) cluster that includes a read replica Scale the read replica to generate the reports.
- B. Set up a new Amazon RDS for PostgreSQL Reserved Instance and an On-Demand read replica Scale the read replica to generate the reports
- C. Set up a new Amazon Aurora PostgreSQL DB cluster that includes a Reserved Instance and an Aurora Replica issue queries to the Aurora Replica to generate the reports.
- D. Set up a new Amazon RDS for PostgreSQL Multi-AZ Reserved Instance Configure the reporting module to query the secondary RDS node so that the reporting module does not affect the primary node
- E. Set up a new Amazon DynamoDB table to store the documents Use a fixed write capacity to support new document entries Automatically scale the read capacity to support the reports

Answer: BC

Explanation:

These options are operationally efficient because they use Amazon RDS read replicas to offload the reporting workload from the primary DB instance and avoid affecting any document modifications or the addition of new documents¹. They also use Reserved Instances for the primary DB instance to reduce costs and On-Demand or Aurora Replicas for the read replicas to scale as needed. Option A is less efficient because it uses Amazon S3 Glacier Flexible Retrieval, which is a cold storage class that has higher retrieval costs and longer retrieval times than Amazon S3 Standard. It also uses EventBridge rules to invoke the job nightly, which does not meet the requirement of processing incoming data files as soon as possible. Option D is less efficient because it uses AWS Lambda to process the files, which has a maximum execution time of 15 minutes per invocation, which might not be enough for processing each file that needs 3-8 minutes. It also uses

S3 event notifications to invoke the Lambda function when the files arrive, which could cause concurrency issues if there are thousands of small data files arriving periodically. Option E is less efficient because it uses Amazon DynamoDB, which is a NoSQL database service that does not support relational queries, which are needed for generating the reports. It also uses fixed write capacity, which could cause throttling or underutilization depending on the incoming data files.

NEW QUESTION 144

- (Topic 3)

A company runs an application on a large fleet of Amazon EC2 instances. The application reads and write entries into an Amazon DynamoDB table. The size of the DynamoDB table continuously grows, but the application needs only data from the last 30 days. The company needs a solution that minimizes cost and development effort.

Which solution meets these requirements?

- A. Use an AWS CloudFormation template to deploy the complete solution
- B. Redeploy the CloudFormation stack every 30 days, and delete the original stack.
- C. Use an EC2 instance that runs a monitoring application from AWS Marketplace
- D. Configure the monitoring application to use Amazon DynamoDB Streams to store the timestamp when a new item is created in the table
- E. Use a script that runs on the EC2 instance to delete items that have a timestamp that is older than 30 days.
- F. Configure Amazon DynamoDB Streams to invoke an AWS Lambda function when a new item is created in the table
- G. Configure the Lambda function to delete items in the table that are older than 30 days.
- H. Extend the application to add an attribute that has a value of the current timestamp plus 30 days to each new item that is created in the table
- I. Configure DynamoDB to use the attribute as the TTL attribute.

Answer: D

Explanation:

Amazon DynamoDB Time to Live (TTL) allows you to define a per-item timestamp to determine when an item is no longer needed. Shortly after the date and time of the specified timestamp, DynamoDB deletes the item from your table without consuming any write throughput. TTL is provided at no extra cost as a means to reduce stored data volumes by retaining only the items that remain current for your workload's needs.

TTL is useful if you store items that lose relevance after a specific time. The following are example TTL use cases:

Remove user or sensor data after one year of inactivity in an application.

Archive expired items to an Amazon S3 data lake via Amazon DynamoDB Streams and AWS Lambda.

Retain sensitive data for a certain amount of time according to contractual or regulatory obligations.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html>

NEW QUESTION 149

- (Topic 3)

An ecommerce company is building a distributed application that involves several serverless functions and AWS services to complete order-processing tasks.

These tasks require manual approvals as part of the workflow. A solutions architect needs to design an

architecture for the order-processing application. The solution must be able to combine multiple AWS Lambda functions into responsive serverless applications. The

solution also must orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers.

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

- A. Use AWS Step Functions to build the application.
- B. Integrate all the application components in an AWS Glue job.
- C. Use Amazon Simple Queue Service (Amazon SQS) to build the application.
- D. Use AWS Lambda functions and Amazon EventBridge (Amazon CloudWatch Events) events to build the application.

Answer: A

Explanation:

AWS Step Functions is a fully managed service that makes it easy to build applications by coordinating the components of distributed applications and microservices using visual workflows. With Step Functions, you can combine multiple AWS Lambda functions into responsive serverless applications and orchestrate data and services that run on Amazon EC2 instances, containers, or on-premises servers. Step Functions also allows for manual approvals as part of the workflow. This solution meets all the requirements with the least operational overhead.

[https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20\(ML\)%20pipelines.](https://aws.amazon.com/step-functions/#:~:text=AWS%20Step%20Functions%20is%20a,machine%20learning%20(ML)%20pipelines.)

NEW QUESTION 153

- (Topic 3)

A company runs an application that receives data from thousands of geographically dispersed remote devices that use UDP. The application processes the data immediately and sends a message back to the device if necessary. No data is stored.

The company needs a solution that minimizes latency for the data transmission from the devices. The solution also must provide rapid failover to another AWS Region.

Which solution will meet these requirements?

- A. Configure an Amazon Route 53 failover routing policy. Create a Network Load Balancer (NLB) in each of the two Regions. Configure the NLB to invoke an AWS Lambda function to process the data.
- B. Use AWS Global Accelerator. Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the NLB. Process the data in Amazon ECS.
- D. Use AWS Global Accelerator. Create an Application Load Balancer (ALB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster.
- E. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.
- F. Configure an Amazon Route 53 failover routing policy. Create an Application Load Balancer (ALB) in each of the two Regions. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the ALB. Process the data in Amazon ECS.

Answer: B

Explanation:

To meet the requirements of minimizing latency for data transmission from the devices and providing rapid failover to another AWS Region, the best solution would be to use AWS Global Accelerator in combination with a Network Load Balancer (NLB) and Amazon Elastic Container Service (Amazon ECS). AWS Global

Accelerator is a service that improves the availability and performance of applications by using static IP addresses (Anycast) to route traffic to optimal AWS endpoints. With Global Accelerator, you can direct traffic to multiple Regions and endpoints, and provide automatic failover to another AWS Region.

NEW QUESTION 158

- (Topic 3)

A company needs to provide its employee with secure access to confidential and sensitive files. The company wants to ensure that the files can be accessed only by authorized users. The files must be downloaded securely to the employees' devices.

The files are stored in an on-premises Windows file server. However, due to an increase in remote usage, the file server is out of capacity.

Which solution will meet these requirements?

- A. Migrate the file server to an Amazon EC2 instance in a public subnet
- B. Configure the security group to limit inbound traffic to the employees' IP addresses.
- C. Migrate the files to an Amazon FSx for Windows File Server file system
- D. Integrate the Amazon FSx file system with the on-premises Active Directory. Configure AWS Client VPN.
- E. Migrate the files to Amazon S3, and create a private VPC endpoint
- F. Create a signed URL to allow download.
- G. Migrate the files to Amazon S3, and create a public VPC endpoint. Allow employees to sign on with AWS IAM Identity Center (AWS SSO).

Answer: B

Explanation:

Windows file server is on-premise and we need something to replicate the data to the cloud, the only option we have is AWS FSx for Windows File Server. Also, since the information is confidential and sensitive, we also want to make sure that the appropriate users have access to it in a secure manner.

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

NEW QUESTION 161

- (Topic 3)

A company needs to export its database once a day to Amazon S3 for other teams to access. The exported object size varies between 2 GB and 5 GB. The S3 access pattern for the data is variable and changes rapidly. The data must be immediately available and must remain accessible for up to 3 months. The company needs the most cost-effective solution that will not increase retrieval time.

Which S3 storage class should the company use to meet these requirements?

- A. S3 Intelligent-Tiering
- B. S3 Glacier Instant Retrieval
- C. S3 Standard
- D. S3 Standard-Infrequent Access (S3 Standard-IA)

Answer: D

Explanation:

S3 Intelligent-Tiering is a cost-optimized storage class that automatically moves data to the most cost-effective access tier based on changing access patterns. Although it offers cost savings, it also introduces additional latency and retrieval time into the data retrieval process, which may not meet the requirement of "immediately available" data. On the other hand, S3 Standard-Infrequent Access (S3 Standard-IA) provides low cost storage with low latency and high throughput performance. It is designed for infrequently accessed data that can be recreated if lost, and can be retrieved in a timely manner if required. It is a cost-effective solution that meets the requirement of immediately available data and remains accessible for up to 3 months.

NEW QUESTION 162

- (Topic 3)

An application that is hosted on Amazon EC2 instances needs to access an Amazon S3 bucket. Traffic must not traverse the internet. How should a solutions architect configure access to meet these requirements?

- A. Create a private hosted zone by using Amazon Route 53
- B. Set up a gateway VPC endpoint for Amazon S3 in the VPC
- C. Configure the EC2 instances to use a NAT gateway to access the S3 bucket
- D. Establish an AWS Site-to-Site VPN connection between the VPC and the S3 bucket

Answer: B

Explanation:

This option is the most efficient because it uses a gateway VPC endpoint for Amazon S3, which provides reliable connectivity to Amazon S3 without requiring an internet gateway or a NAT device for the VPC. A gateway VPC endpoint routes traffic from the VPC to Amazon S3 using a prefix list for the service and does not leave the AWS network. This meets the requirement of not traversing the internet. Option A is less efficient because it uses a private hosted zone by using Amazon Route 53, which is a DNS service that allows you to create custom domain names for your resources within your VPC. However, this does not provide connectivity to Amazon S3 without an internet gateway or a NAT device. Option C is less efficient because it uses a NAT gateway to access the S3 bucket, which is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. However, this does not meet the requirement of not traversing the internet. Option D is less efficient because it uses an AWS Site-to-Site VPN connection between the VPC and the S3 bucket, which is a secure and encrypted network connection between your on-premises network and your VPC. However, this does not meet the requirement of not traversing the internet.

NEW QUESTION 166

- (Topic 3)

A company hosts a multi-tier web application that uses an Amazon Aurora MySQL DB cluster for storage. The application tier is hosted on Amazon EC2 instances. The company's IT security guidelines mandate that the database credentials be encrypted and rotated every 14 days.

What should a solutions architect do to meet this requirement with the LEAST operational effort?

- A. Create a new AWS Key Management Service (AWS KMS) encryption key. Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Associate the secret with the Aurora DB cluster. Configure a custom rotation period of 14 days.
- B. Create two parameters in AWS Systems Manager Parameter Store: one for the username as a string parameter and one that uses the SecureString type for the password. Select AWS Key Management Service (AWS KMS) encryption for the password parameter, and load these parameters in the application tier. Implement

an AWS Lambda function that rotates the password every 14 days.

- C. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon Elastic File System (Amazon EFS) file system
- Mount the EFS file system in all EC2 instances of the application
- D. Restrict the access to the file on the file system so that the application can read the file and that only super users can modify the file
- Implement an AWS Lambda function that rotates the key in Aurora every 14 days and writes new credentials into the file
- E. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon S3 bucket that the application uses to load the credentials
- Download the file to the application regularly to ensure that the correct credentials are used
- Implement an AWS Lambda function that rotates the Aurora credentials every 14 days and uploads these credentials to the file in the S3 bucket

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types-oracle/>

NEW QUESTION 167

- (Topic 3)

An ecommerce company is running a multi-tier application on AWS. The front-end and backend tiers run on Amazon EC2, and the database runs on Amazon RDS for MySQL. The backend tier communicates with the RDS instance. There are frequent calls to return identical database from the database that are causing performance slowdowns.

Which action should be taken to improve the performance of the backend?

- A. Implement Amazon SNS to store the database calls.
- B. Implement Amazon ElastiCache to cache the large database.
- C. Implement an RDS for MySQL read replica to cache database calls.
- D. Implement Amazon Kinesis Data Firehose to stream the calls to the database.

Answer: B

Explanation:

the best solution is to implement Amazon ElastiCache to cache the large datasets, which will store the frequently accessed data in memory, allowing for faster retrieval times. This can help to alleviate the frequent calls to the database, reduce latency, and improve the overall performance of the backend tier.

NEW QUESTION 171

- (Topic 3)

A solution architect must create a disaster recovery (DR) plan for a high-volume software as a service (SaaS) platform. All data for the platform is stored in an Amazon Aurora MySQL DB cluster.

The DR plan must replicate data to a secondary AWS Region. Which solution will meet these requirements MOST cost-effectively? Use MySQL binary log replication to an Aurora cluster

- A. Use MySQL binary log replication to an Aurora cluster in the secondary Region Provision one DB instance for the Aurora cluster in the secondary Region.
- B. Set up an Aurora global database for the DB cluster
- C. When setup is complete, remove the DB instance from the secondary Region.
- D. Use AWS Database Migration Service (AWS DMS) to continuously replicate data to an Aurora cluster in the secondary Region Remove the DB instance from the secondary Region.
- E. Set up an Aurora global database for the DB cluster Specify a minimum of one DB instance in the secondary Region

Answer: D

Explanation:

"Replication from the primary DB cluster to all secondaries is handled by the Aurora storage layer rather than by the database engine, so lag time for replicating changes is minimal—typically, less than 1 second. Keeping the database engine out of the replication process means that the database engine is dedicated to processing workloads. It also means that you don't need to configure or manage the Aurora MySQL binlog (binary logging) replication."

NEW QUESTION 175

- (Topic 3)

A company is migrating its on-premises workload to the AWS Cloud. The company already uses several Amazon EC2 instances and Amazon RDS DB instances. The company wants a solution that automatically starts and stops the EC2 instances and D6 instances outside of business hours. The solution must minimize cost and infrastructure maintenance.

Which solution will meet these requirements?

- A. Scale the EC2 instances by using elastic resize Scale the DB instances to zero outside of business hours
- B. Explore AWS Marketplace for partner solutions that will automatically start and stop the EC2 Instances and DB instances on a schedule
- C. Launch another EC2 instance
- D. Configure a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule.
- E. Create an AWS Lambda function that will start and stop the EC2 instances and DB instances Configure Amazon EventBridge to invoke the Lambda function on a schedule

Answer: D

Explanation:

The most efficient solution for automatically starting and stopping EC2 instances and DB instances on a schedule while minimizing cost and infrastructure maintenance is to create an AWS Lambda function and configure Amazon EventBridge to invoke the function on a schedule.

Option A, scaling EC2 instances by using elastic resize and scaling DB instances to zero outside of business hours, is not feasible as DB instances cannot be scaled to zero.

Option B, exploring AWS Marketplace for partner solutions, may be an option, but it may not be the most efficient solution and could potentially add additional costs.

Option C, launching another EC2 instance and configuring a crontab schedule to run shell scripts that will start and stop the existing EC2 instances and DB instances on a schedule, adds unnecessary infrastructure and maintenance.

NEW QUESTION 176

- (Topic 3)

A company runs an application on Amazon EC2 Linux instances across multiple Availability Zones. The application needs a storage layer that is highly available and Portable Operating System Interface (POSIX) compliant. The storage layer must provide maximum data durability and must be shareable across the EC2 instances. The data in the storage layer will be accessed frequently for the first 30 days and will be accessed infrequently after that time. Which solution will meet these requirements MOST cost-effectively?

- A. Use the Amazon S3 Standard storage class. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Glacier.
- B. Use the Amazon S3 Standard storage class.
- C. Create an S3 Lifecycle policy to move infrequently accessed data to S3 Standard-Infrequent Access (S3 Standard-IA).
- D. Use the Amazon Elastic File System (Amazon EFS) Standard storage class.
- E. Create a Lifecycle management policy to move infrequently accessed data to EFS Standard- Infrequent Access (EFS Standard-IA).
- F. Use the Amazon Elastic File System (Amazon EFS) One Zone storage class.
- G. Create a Lifecycle management policy to move infrequently accessed data to EFS One Zone- Infrequent Access (EFS One Zone-IA).

Answer: C

Explanation:

<https://aws.amazon.com/efs/features/infrequent-access/>

NEW QUESTION 179

- (Topic 3)

A company is migrating a Linux-based web server group to AWS. The web servers must access files in a shared file store for some content. The company must not make any changes to the application.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 Standard bucket with access to the web servers.
- B. Configure an Amazon CloudFront distribution with an Amazon S3 bucket as the origin.
- C. Create an Amazon Elastic File System (Amazon EFS) file system.
- D. Mount the EFS file system on all web servers.
- E. Configure a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume.
- F. Mount the EBS volume to all web servers.

Answer: C

Explanation:

Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system on all web servers. To meet the requirements of providing a shared file store for Linux-based web servers without making changes to the application, using an Amazon EFS file system is the best solution. Amazon EFS is a managed NFS file system service that provides shared access to files across multiple Linux-based instances, which makes it suitable for this use case. Amazon S3 is not ideal for this scenario since it is an object storage service and not a file system, and it requires additional tools or libraries to mount the S3 bucket as a file system. Amazon CloudFront can be used to improve content delivery performance but is not necessary for this requirement. Additionally, Amazon EBS volumes can only be mounted to one instance at a time, so it is not suitable for sharing files across multiple instances.

NEW QUESTION 180

- (Topic 3)

A company is running a batch application on Amazon EC2 instances. The application consists of a backend with multiple Amazon RDS databases. The application is causing a high number of reads on the databases. A solutions architect must reduce the number of database reads while ensuring high availability.

What should the solutions architect do to meet this requirement?

- A. Add Amazon RDS read replicas.
- B. Use Amazon ElastiCache for Redis.
- C. Use Amazon Route 53 DNS caching.
- D. Use Amazon ElastiCache for Memcached.

Answer: A

Explanation:

This solution meets the requirement of reducing the number of database reads while ensuring high availability for a batch application that consists of a backend with multiple Amazon RDS databases. Amazon RDS read replicas are copies of the primary database instance that can serve read-only traffic. You can create one or more read replicas for a primary database instance and connect to them using a special endpoint. Read replicas can improve the performance and availability of your application by offloading read queries from the primary database instance.

Option B is incorrect because using Amazon ElastiCache for Redis can provide a fast, in-memory data store that can cache frequently accessed data, but it does not support replication from Amazon RDS databases. Option C is incorrect because using Amazon Route 53 DNS caching can improve the performance and availability of DNS queries, but it does not reduce the number of database reads. Option D is incorrect because using Amazon ElastiCache for Memcached can provide a fast, in-memory data store that can cache frequently accessed data, but it does not support replication from Amazon RDS databases.

References:

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

NEW QUESTION 181

- (Topic 3)

A company wants to implement a disaster recovery plan for its primary on-premises file storage volume. The file storage volume is mounted from an Internet Small Computer Systems Interface (iSCSI) device on a local storage server. The file storage volume holds hundreds of terabytes (TB) of data.

The company wants to ensure that end users retain immediate access to all file types from the on-premises systems without experiencing latency.

Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

- A. Provision an Amazon S3 File Gateway as a virtual machine (VM) that is hosted on-premise.
- B. Set the local cache to 10 TB.
- C. Modify existing applications to access the files through the NFS protocol.
- D. To recover from a disaster, provision an Amazon EC2 instance and mount the S3 bucket that contains the files.
- E. Provision an AWS Storage Gateway tape gateway.

- F. Use a data backup solution to back up all existing data to a virtual tape library
- G. Configure the data backup solution to run nightly after the initial backup is complete
- H. To recover from a disaster, provision an Amazon EC2 instance and restore the data to an Amazon Elastic Block Store (Amazon EBS) volume from the volumes in the virtual tape library.
- I. Provision an AWS Storage Gateway Volume Gateway cached volume
- J. Set the local cache to 10 TB
- K. Mount the Volume Gateway cached volume to the existing file server by using iSCSI
- L. and copy all files to the storage volume
- M. Configure scheduled snapshots of the storage volume
- N. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.
- O. Provision an AWS Storage Gateway Volume Gateway stored volume with the same amount of disk space as the existing file storage volume
- P. Mount the Volume Gateway stored volume to the existing file server by using iSCSI, and copy all files to the storage volume
- Q. Configure scheduled snapshots of the storage volume
- R. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.

Answer: D

Explanation:

"The company wants to ensure that end users retain immediate access to all file types from the on-premises systems" - Cached volumes: low latency access to most recent data - Stored volumes: entire dataset is on premise, scheduled backups to S3 Hence Volume Gateway stored volume is the apt choice.

NEW QUESTION 182

- (Topic 3)

A company needs a backup strategy for its three-tier stateless web application. The web application runs on Amazon EC2 instances in an Auto Scaling group with a dynamic scaling policy that is configured to respond to scaling events. The database tier runs on Amazon RDS for PostgreSQL. The web application does not require temporary local storage on the EC2 instances. The company's recovery point objective (RPO) is 2 hours. The backup strategy must maximize scalability and optimize resource utilization for this environment. Which solution will meet these requirements?

- A. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances and database every 2 hours to meet the RPO
- B. Configure a snapshot lifecycle policy to take Amazon Elastic Block Store (Amazon EBS) snapshots. Enable automated backups in Amazon RDS to meet the RPO
- C. Retain the latest Amazon Machine Images (AMIs) of the web and application tiers. Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO
- D. Take snapshots of Amazon Elastic Block Store (Amazon EBS) volumes of the EC2 instances every 2 hours. Enable automated backups in Amazon RDS and use point-in-time recovery to meet the RPO

Answer: C

Explanation:

Since the application has no local data on instances, AMIs alone can meet the RPO by restoring instances from the most recent AMI backup. When combined with automated RDS backups for the database, this provides a complete backup solution for this environment. The other options involving EBS snapshots would be unnecessary given the stateless nature of the instances. AMIs provide all the backup needed for the app tier. This uses native, automated AWS backup features that require minimal ongoing management: - AMI automated backups provide point-in-time recovery for the stateless app tier. - RDS automated backups provide point-in-time recovery for the database.

NEW QUESTION 187

- (Topic 3)

A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch runs. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application. What should a solution architect recommend to ensure the application is able to handle the workload and avoid downtime?

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

Answer: C

Explanation:

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

NEW QUESTION 188

.....

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](#)